Dual Reduction and Elementary Games with Senders and Receivers

See GEB 21:183-202 (1997): https://doi.org/10.1006/game.1997.0573

https://home.uchicago.edu/~rmyerson/research/eldual2023notes.pdf

A finite ***senders-receivers game*** is any $\Gamma = (I, (T_i)_{i\in I}, p, J, (C_j)_{j\in J}, (u_k)_{k\in I\cup J})$
with nonempty finite sets: I={senders}, T_i={i's types}, J={receivers}, C_j={j's actions},
$I\cap J=\varnothing$, $T=\times_{i\in I} T_i$, $C=\times_{j\in J} C_j$, $p\in\Delta(T)$ probabilities, $u_k{:}C\times T\to\mathbb{R}$ utility payoffs.
We may write $c = (c_j)_{j\in J} = (c_{-j},c_j) \in C$, $t = (t_i)_{i\in I} = (t_{-i},t_i) \in T$.

A direct ***coordination mechanism*** is any $\mu{:}T\to\Delta(C)$.
Let $U_j(\mu,c_j) = \sum_t p(t) \sum_{c-j} \mu(c|t)u_j(c,t)$, $\hat{U}_j(\mu,d_j,c_j) = \sum_t p(t) \sum_{c-j} \mu(c|t)u_j((c_{-j},d_j),t)$,
$U_i(\mu,t_i) = \sum_{t-i} p(t) \sum_c \mu(c|t)u_i(c,t)$, $\hat{U}_i(\mu,s_i,t_i) = \sum_{t-i} p(t) \sum_c \mu(c|t_{-i},s_i)u_i(c,t)$. *[proby-discounted]*
An ***incentive-compatible*** (IC) mechanism is any μ satisfying:
$\mu(c|t)\geq 0\ \ \forall c\in C, \forall t\in T$; $\sum_{c\in C} \mu(c|t) = 1\ \forall t\in T$; [probability constraints]
$U_j(\mu,c_j) \geq \hat{U}_j(\mu,d_j,c_j)\ \ \forall c_j\in C_j, \forall d_j\in C_j, \forall j\in J$; [moral-hazard constraints]
$U_i(\mu,t_i) \geq \hat{U}_i(\mu,s_i,t_i)\ \ \forall t_i\in T_i, \forall s_i\in T_i, \forall i\in I$. [adverse-selection constraints]

Γ is an ***elementary*** game iff there exists some $\mu^*{:}T\to\Delta(C)$ such that
$U_j(\mu^*,c_j) > \hat{U}_j(\mu^*,d_j,c_j)\ \forall c_j, \forall d_j\neq c_j, \forall j$; and $U_i(\mu^*,t_i) > \hat{U}_i(\mu^*,s_i,t_i)\ \forall t_i, \forall s_i\neq t_i, \forall i$.

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

Consider the following ***primal linear programming problem***:

minimize $\sum_i \sum_{ti} \pi_i(t_i) + \sum_j \sum_{cj} \pi_j(c_j)$ over $\mu \geq \mathbf{0}$ & π such that

$\pi_j(c_j) + \sum_t \sum_{c-j} p(t)\mu(c|t)(u_j(c,t) - u_j((c_{-j},d_j),t)) \geq 0 \quad \forall c_j \in C_j, \forall d_j \in C_j, \forall j \in J;$ $[\alpha_j(d_j|c_j)]$

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

$\sum_c \mu(c|t) = 1 \quad \forall t \in T.$ $[\beta(t)]$

The solutions to this LP are the incentive-compatible mechanisms, which exist (Nash) and yield optimal value 0 with $\pi=\mathbf{0}$. (Trivial constraints with $d_j=c_j$ and $s_i=t_i$ imply $\pi \geq \mathbf{0}$.)

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

$\beta(t) + \sum_j \sum_{dj} \alpha_j(d_j|c_j)p(t)(u_j(c,t) - u_j((c_{-j},d_j),t)) +$

$+ \sum_i \sum_{ti} (\alpha_i(s_i|t_i)p(t)u_i(c,t) - \alpha_i(t_i|s_i)p(t_{-i},s_i)u_i(c,(t_{-i},s_i))) \leq 0 \quad \forall t \in T, \forall c \in C;$ $[\mu(c|t)]$

$\sum_{dj} \alpha_j(d_j|c_j) = 1 \quad \forall c_j \in C_j, \quad \forall j \in J;$ $[\pi_j(c_j)]$

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

Nontrivial dual solutions exist, with some $\alpha_k(e_k|f_k)>0$ & $e_k \neq f_k$, ***iff Γ is not elementary***.

Let $D(c,t,\alpha) = \sum_j \sum_{dj} \alpha_j(d_j|c_j)p(t)(u_j((c_{-j},d_j),t) - u_j(c,t)) +$

$+ \sum_i \sum_{si} (\alpha_i(t_i|s_i)p(t_{-i},s_i)u_i(c,(t_{-i},s_i)) - \alpha_i(s_i|t_i)p(t)u_i(c,t)).$ [α-deviatn value @c,t]

Then the dual optimum has $\beta(t) = \min_{c \in C} D(c,t,\alpha) \; \forall t \in T$, and $\sum_{t \in T} \beta(t) = 0$.

Lemma: Given any dual solution α, for any mechanism $\mu:T \to \Delta(C)$:

$$\sum_{j \in J} \sum_{cj} \sum_{dj} \alpha_j(d_j|c_j)(\hat{U}_j(\mu,d_j,c_j) - U_j(\mu,c_j)) + \sum_{i \in I} \sum_{ti} \sum_{si} \alpha_i(s_i|t_i)(\hat{U}_i(\mu,s_i,t_i) - U_i(\mu,t_i))$$
$$= \sum_{t \in T} \sum_{c \in C} \mu(c|t)D(c,t,\alpha) \geq \sum_{t \in T} \beta(t) = 0.$$

(*Expected net gains of unilateral α-deviations from μ must have a nonnegative sum.*)

Given any Markov chain $\delta{:}X \to \Delta(X)$, we let X/δ denote the set of ***minimal nonempty*** ***δ-absorbing subsets*** of X. ($Y \subseteq X$ is δ-absorbing iff $\delta(x|y)=0$ $\forall y \in Y$, $\forall x \notin Y$.)

We define functions ψ and φ so that, for any absorbing set $R \in X/\delta$ and any $x \in X$:

$\varphi(x|R,\delta)$ is the probability of x in the unique ***δ-stationary distribution on R***, and

$\psi(R|x,\delta)$ is the ***probability of a δ-stochastic process reaching R from x***. That is:

$$\varphi(y|R,\delta) = \sum_{z \in R} \varphi(z|R,\delta)\delta(y|z)\ \forall y \in R;\ \ \sum_{y \in R} \varphi(y|R,\delta) = 1;\ \ \varphi(z|R,\delta)=0 \text{ if } z \notin R;$$

$$\psi(R|y,\delta) = 1 \text{ if } y \in R;\ \ \psi(R|y,\delta) = 0 \text{ if } y \in \hat{R} \in X/\delta\ \&\ \hat{R} \neq R;$$

$$\psi(R|y,\delta) = \sum_{z \in X} \delta(z|y)\psi(R|z,\delta)\ \forall y \in X.$$

Dual solutions α define Markov chains $\alpha_i{:}T_i \to \Delta(T_i)$ and $\alpha_j{:}C_j \to \Delta(C_j)$.

Given any dual solution α for the game Γ, let us define the ***α-reduced game***:

$\Gamma/\alpha = (I, (T_i/\alpha_i)_{i \in I}, q, J, (C_j/\alpha_j)_{j \in J}, (v_k)_{k \in I \cup J})$ where

$q(\tau) = \sum_{t \in T} p(t) \left(\prod_{i \in I} \psi(\tau_i|t_i,\alpha_i)\right)$, and

$v_k(\gamma,\tau) = \sum_{t \in T} \sum_{c \in C} p(t)\left(\prod_{i \in I} \psi(\tau_i|t_i,\alpha_i)\right)\left(\prod_{j \in J} \varphi(c_j|\gamma_j,\alpha_j)\right)u_k(c,t)/q(\tau)$.

A mechanism $\lambda{:}(\times_{i \in I} T_i/\alpha_i) \to \Delta(\times_{j \in J} C_j/\alpha_j)$ for Γ/α induces a mechanism μ^λ on Γ:

$\mu^\lambda(c|t) = \sum_\tau \sum_\gamma \left(\prod_{i \in I} \psi(\tau_i|t_i,\alpha_i)\right)\lambda(\gamma|\tau)\left(\prod_{j \in J} \varphi(c_j|\gamma_j,\alpha_j)\right)$.

Theorem: 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.

Fact: For any finite Γ, iterative dual reduction yields an elementary reduced game.

In the reduced game Γ/α, sender i's information is reduced to the absorbing set in T_i/α_i that was reached by an α_i stochastic process from i's true type in T_i, and receiver j's choices are reduced to the α_j-stationary distributions on absorbing sets in C_j/α_j.

To show that incentive-compatible mechanisms in Γ/α are still IC in Γ, we must show that, when all players are expected to act within this reduced strategic domain, one player could not gain by a deviation that uses his full information or choice set in Γ.

By the Lemma, for any mechanism in Γ, if each player considered deviating by α_i or α_j, the sum of their expected gains from unilaterally deviating must be nonnegative.

But for players who are acting according to their reduced strategic options in Γ/α, deviating by α_i or α_j would not change their probabilistic behavior (by construction).

Now suppose, *contrary to the Theorem*, that some j could expect to gain strictly by deviating to some c_j in C_j when γ_j in C_j/α_j is the λ-recommended action for j.

Consider the scenario that differs from λ only in that player j deviates to a uniform distribution over all of j's best-responses in C_j when γ_j is recommended.

Since α deviations from this scenario would not affect any other behavior, the Lemma implies that j cannot expect to lose by further deviating from this scenario by α_j when γ_j is recommended.

So if c_j is a best response to γ_j and $\alpha_j(d_j|c_j)>0$ then d_j is also a best response to γ_j.

So j's best responses in C_j are an α_j-absorbing set, and so an α_j-stationary option in C_j/α_j is also optimal for j, and (by λ IC) it is γ_j, *contradicting the supposition above.*

Now suppose, *contrary to the Theorem*, that some type t_i of some sender i could expect to gain strictly by deviating from λ by changing his report from τ_i to some other ρ_i.

Consider the scenario that differs from λ only in that all types t_i with $\psi(\tau_i|t_i,\alpha_i)>0$ which strictly prefer misreporting ρ_i do so when τ_i would be correct in λ, all those which would strictly lose by misreporting stay with τ_i, and the indifferent randomize equally.

Since α deviations from this scenario would not affect any other behavior, the Lemma implies that i could not expect to lose by a further α_i-deviation (each t_i acting like another s_i in this scenario with probability $\alpha_i(s_i|t_i)$).

But an $\alpha_i(s_i|t_i)>0$ probability of t_i further imitating an s_i with $\psi(\tau_i|s_i,\alpha_i)>0$ in this scenario would yield an expected strict loss when t_i strictly prefers misreporting & s_i does not, or when t_i strictly prefers to not misreport & s_i is willing to do so, and otherwise it would not make any difference for i.

So all of i's types in the α_i chain that leads into τ_i must be willing to misreport ρ_i, while some supposed types strictly prefer such misreporting.

But the incentive-compatibility of λ in Γ/α implies that the probability-weighted sum of these types' expected payoffs cannot be higher than what they get by reporting τ_i, *contradicting the supposition above.*

This intrinsic alignment of preferences among i's types in T_i that are pooled in a reduced-game type τ_i shows that the incentive-compatible mechanisms of the reduced game will not depend on the relative weighting of utility payoffs for these types.

Fact: Consider a dual solution α and $i \in I$, $\tau_i \in T_i/\alpha_i$, and $r_i \in \tau_i$. Let λ and η be any two mechanisms for Γ/α that do not depend on i's (reduced) type, and suppose that $U_i(\mu^\lambda,r_i) \geq U_i(\mu^\eta,r_i)$. Then $U_i(\mu^\lambda,t_i) \geq U_i(\mu^\eta,t_i)$ for every t_i such that $\psi(\tau_i|t_i,\alpha_i)>0$.

Suppose, *contrary to the Fact*, that some type s_i has $\psi(\tau_i|s_i,\alpha_i)>0$ & $U_i(\mu^\eta,s_i) > U_i(\mu^\lambda,s_i)$.

Consider the scenario which coincides with μ^λ except that it switches to μ^η iff i's type t_i satisfies $U_i(\mu^\eta,t_i) > U_i(\mu^\lambda,t_i)$.

This scenario satisfies all the strategic restrictions of the reduced game for all players other than sender i. Thus, since α deviations from this scenario would not affect any other behavior, the Lemma implies that player i could not expect to lose by an α_i-deviation (each t_i acting like another $\hat{t}_i$ in this scenario with probability $\alpha_i(\hat{t}_i|t_i)$).

An $\alpha_i(\hat{t}_i|t_i)>0$ probability of t_i imitating $\hat{t}_i$ can make a difference in this scenario only in two cases: (1) if $\hat{t}_i$ satisfies $U_i(\mu^\eta,\hat{t}_i)>U_i(\mu^\lambda,\hat{t}_i)$ but $U_i(\mu^\lambda,t_i)\geq U_i(\mu^\eta,t_i)$,
or (2) if t_i satisfies $U_i(\mu^\eta,t_i)>U_i(\mu^\lambda,t_i)$ but $U_i(\mu^\lambda,\hat{t}_i)\geq U_i(\mu^\eta,\hat{t}_i)$.

In case (1), the imitation cannot help player i in type t_i, as it would just substitute η for the weakly preferred λ. In case (2), the imitation strictly hurts player i in type t_i, as it substitutes λ for the strictly preferred η.

But $\psi(\tau_i|s_i,\alpha_i)>0$ implies the existence of a positive α_i-chain from s_i to $r_i \in \tau_i$, and so the strict-loss case (2) must happen at least once, *contradicting the above supposition*.

Corollary: If $\{r_i,s_i\} \subseteq \tau_i$ then $U_i(\mu^\lambda,r_i) \geq U_i(\mu^\eta,r_i) \iff U_i(\mu^\lambda,s_i) \geq U_i(\mu^\eta,s_i)$.

Examples with no senders (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 game with no senders:

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 we change the diagonal elements from (0,0) to $(\varepsilon,\varepsilon)$, for some $\varepsilon<1$.

If $\varepsilon<0$, then the game becomes elementary, and the above-described correlated equilibrium satisfies all incentive constraints strictly.

If $0<\varepsilon<1$ then no incentive constraint can be satisfied strictly, and there is a dual solution with all nontrivial incentive constraints having strictly positive dual variables:

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

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

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

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 (ε,ε), for some ε<1.
If ε<0, then the game becomes elementary, and the above-described incentive-compatible mechanism μ satisfies all incentive constraints strictly.
If 0<ε<1 then no incentive constraint can be satisfied strictly, and there is a dual solution with all nontrivial incentive constraints having strictly positive dual variables:
$\alpha_1(s_1|r_1)=\alpha_1(t_1|s_1)=\alpha_1(r_1|t_1)=\alpha_2(c_2|a_2)=\alpha_2(a_2|b_2)=\alpha_2(b_2|c_2)=\varepsilon$,
$\alpha_1(t_1|r_1)=\alpha_1(r_1|s_1)=\alpha_1(s_1|t_1)=\alpha_2(b_2|a_2)=\alpha_1(c_2|b_2)=\alpha_2(a_2|c_2)=1-\varepsilon$.)

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 components of α being 0.

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

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

p:	1's reduced type	$\{b_2\}$	$\{c_2\}$
4/9	$\{s_1\}$~0.75[s_1]+0.25[r_1]	6, 6.75	0, 0.75
5/9	$\{t_1\}$~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

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

All these reduced games are elementary, with strict mechanism ($\{s_1\}\rightarrow\{b_2\}$,$\{t_1\}\rightarrow\{c_2\}$).

Games with transitive dual solutions.

For any game Γ, we say that a dual solution α is *maximal* iff its support (the set of (k,x_k,y_k) such that $\alpha_k(x_k|y_k)>0$) is maximal among all dual solutions to the game.

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.

We say that the game Γ has *transitive* dual solutions iff, in any maximal dual solution α, for any player $k \in I \cup J$ and any three of k's alternative types or actions $\{x_k, y_k, z_k\}$, if $\alpha_k(x_k|y_k)>0$ and $\alpha_k(y_k|z_k)>0$ then $\alpha_k(x_k|z_k)>0$.

Any $(I, (T_i)_{i \in I}, p, J, (C_j)_{j \in J})$ as above may be called a *framework* for senders receivers games, which can then be defined by specifying a utility vector in $\mathbb{R}^{(I \cup J) \times C \times T}$.

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

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

We have seen 3x3 examples of games which do not have transitive dual solutions, but this failure of dual transitivity is very sensitive to small changes in the utilities.

Claim. Given any framework $(I, (T_i)_{i \in I}, p, J, (C_j)_{j \in J})$ for senders-receivers 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 $\mathbb{R}^{(I \cup J) \times C \times T}$.

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 that is minimal in the set. Then we ca show that this $\Gamma(u)$ with locally minimal dual magnitude will have transitive dual solutions.

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 [integer$\geq$0] that is minimal in the set.*

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

Let α be any maximal dual solution for $\Gamma(u)$.

$\Gamma(u)$ also has an IC mechanism μ that is *maximally* incentive compatible, satisfying with strict inequality every incentive constraint that has a zero component in α.

If $\Gamma(u)$ does not have transitive dual solutions then either

- (*case 1*) some receiver j has an action d_j such that some other actions e_i and c_j satisfy $\alpha_j(e_j|d_j)>0$ and $\alpha_j(d_j|c_j)>0$ but $\alpha_j(e_j|c_j)=0$; or
- (*case 2*) some sender i has a type s_i such that some other types r_i and t_i satisfy $\alpha_i(t_i|s_i)>0$ and $\alpha_i(s_i|r_i)>0$ but $\alpha_i(t_i|r_i)=0$.

In case 1, for any small $\varepsilon>0$, we can construct $\tilde{u}$ to be the same as u except that:

$\tilde{u}_j((c_{-j},d_j),t) = (1-\varepsilon)u_j((c_{-j},d_j),t) + \varepsilon\sum_{ej} \alpha_j(e_j|d_j)u_j((c_{-j},e_j),t),\ \forall c_{-j},\ \forall t.$

In case 2, for any small $\varepsilon>0$, we can construct $\tilde{u}$ to be the same as u except that

$p(t_{-i},s_i)\tilde{u}_i(c,t_{-i},s_i) = (1-\varepsilon)p(t_{-i},s_i)u_i(c,t_{-i},s_i) + \varepsilon\sum_{ri} \alpha_i(s_i|r_i)p(t_{-i},r_i)u_i(c,t_{-i},r_i)\ \forall t_{-i},\ \forall c.$

Then any μ that is maximally incentive compatible for $\Gamma(u)$ will also be incentive compatible for $\Gamma(\tilde{u})$, but its set of strictly satisfied constraints will expand to include the incentive constraint for the dual variable $\alpha_j(d_j|c_j)$ [in case 1] or $\alpha_i(t_i|s_i)$ [in case 2].

So in each case, 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.

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.

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.

Abstract: *Consider the incentive constraints that define the incentive-compatible mechanisms of a senders-receivers game. Duals of these linear constraints form Markov chains on the senders' type sets and the receivers' action sets. The minimal nonempty absorbing sets of these Markov chains can be interpreted as the types and actions in a dual-reduced game. Any incentive-compatible mechanism of a dual-reduced game induces an equivalent incentive-compatible mechanism for the original game. We say that a game is elementary if all nontrivial incentive constraints can be satisfied as strict inequalities in incentive-compatible mechanisms. Any senders-receivers game can be reduced to an elementary game by iterative dual reduction.*

https://home.uchicago.edu/~rmyerson/research/eldual2023notes.pdf

Dual Reduction and Elementary Games with Senders and Receivers
Addendum On the Density of Games with Transitive Dual Solutions

For any game Γ, we say that a dual solution α is *maximal* iff its support (the set of (k,x_k,y_k) such that $\alpha_k(x_k|y_k)>0$) is maximal among all dual solutions to the game. A maximal dual solution must exist, because the union of the supports of any two dual solutions will be the support of the average of the two given dual solutions, which is also a dual solution.
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 an elementary game would have dual magnitude equal to 0.)

An incentive-compatible mechanism μ is *maximal* iff its set of incentive constraints satisfied as strict inequalities is maximal. A maximal IC mechanism must also exist. An incentive constraint has a strictly positive dual variable in a maximal dual solution if and only if the constraint is binding with 0 slack in any maximal incentive-compatible mechanism.

We say that the game Γ has *transitive* dual solutions iff, in any maximal dual solution α, for any player $k\in I\cup J$ and any three of k's alternative types or actions $\{x_k, y_k, z_k\}$, if $\alpha_k(x_k|y_k)>0$ and $\alpha_k(y_k|z_k)>0$ then $\alpha_k(x_k|z_k)>0$.

For any given senders-receivers game $\Gamma = (I, (T_i)_{i\in I}, p, J, (C_j)_{j\in J}, (u_k)_{k\in I\cup J})$, for any players $i\in I$ and $j\in J$, and any $c_j\in C_j$ or $t_i\in T_i$, let

$A_j(c_j)=\{d_j\in C_j|$ a maximal dual solution α has $\alpha_j(d_j|c_j)>0\}$,

$Q_i(t_i)=\{s_i\in T_i|$ a maximal dual solution α has $\alpha_i(t_i|s_i)>0\}$.

So Γ has transitive dual solutions iff $\forall i\in I$, $\forall j\in J$:

$d_j\in A_j(c_j) \Rightarrow A_j(d_j)\subseteq A_j(c_j)$, and $s_i\in Q_i(t_i) \Rightarrow Q_i(s_i)\subseteq Q_i(t_i)$.

Any elementary game satisfies dual transitivity, because its $A_j(c_j)$ & $Q_i(t_i)$ are all empty sets.

Any $(I, (T_i)_{i\in I}, p, J, (C_j)_{j\in J})$ as above may be called a *framework* for senders receivers games, which can then be defined by specifying a utility vector in $\mathbb{R}^{(I\cup J)\times C\times T}$. Given this framework, the game defined by a utility vector u may then be denoted $\Gamma(u)$.
We say that the type distribution p has *full support* iff $p(t) > 0$ for all $t\in T$.

We have seen 3x3 examples of games which do not have transitive dual solutions, but this failure of dual transitivity is very sensitive to small changes in the utilities.
We can now show that the games with transitive dual solutions are dense in the set of games on any given framework with a full-support type distribution.
Claim. Given any framework $(I, (T_i)_{i\in I}, p, J, (C_j)_{j\in J})$ for senders-receivers 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 $\mathbb{R}^{(I\cup J)\times C\times T}$.

(That is, for any senders-receivers game $\Gamma(u)$ where the dual solutions are not transitive, for any $\varepsilon>0$, there exist other utility vectors $\tilde{u}$ such that $\Gamma(\tilde{u})$ has transitive dual solutions and no component of $\tilde{u}$ differs by more than ε from the corresponding component of u.)

Proof: Contrary to the claim, suppose that we can find some nonempty open subset of $\mathbb{R}^{(I\cup J)\times C\times T}$ such that all utility vectors in this set induce games where the dual solutions are not transitive. Let us select u to be any utility vector in this set such that $\Gamma(u)$ has a dual magnitude which is minimal among the dual magnitudes of all games with utilities in this open set. (Such a magnitude-minimizing u can be found, even on an open set, because there are only finitely many possible values for the dual magnitudes of these games.)
Then let α denote a maximal dual solution for $\Gamma(u)$.

Since the dual solutions of $\Gamma(u)$ are not transitive, either (*case 1*) some receiver j has some action d_j such that we can find other actions c_j, and e_j in C_j satisfying $d_j \in A_j(c_j)$, $e_j \in A_j(d_j)$, but $e_j \notin A_j(c_j)$; or (*case 2*) some sender i has some type s_i such that we can find other type t_i and r_i satisfying $s_i \in Q_i(t_i)$, $r_i \in Q_i(s_i)$, but $r_i \notin Q_i(t_i)$.

First consider case 1, where u has a violation of transitivity for some action d_j of some receiver j. For any small $\varepsilon>0$, we can construct $\tilde{u}$ to be the same as u except that player j's payoff from choosing action d_j is changed to:

$$\tilde{u}_j((c_{-j},d_j),t) = (1-\varepsilon)u_j((c_{-j},d_j),t) + \varepsilon\textstyle\sum_{ej\in Aj(dj)} \alpha_j(e_j|d_j)u_j((c_{-j},e_j),t),\ \forall c_{-j},\ \forall t.$$

That is, player j's expected $\tilde{u}$-payoff from d_j would as if j imagined that choosing d_j would involve an ε chance of (accidentally) deviating to another action in $A_j(d_j)$ according to the $\alpha_j(\bullet|d_j)$ distribution; but j's imagined deviation would not affect other players' payoffs.
We now show that any maximal incentive-compatible mechanism μ for the game $\Gamma(u)$ will also be incentive compatible for the new game $\Gamma(\tilde{u})$, and all the incentive constraints that μ satisfies with positive slack for $\Gamma(u)$ will also be satisfied with positive slack for $\Gamma(\tilde{u})$. We only need to consider incentive constraints involving d_j, because all other payoffs in $\tilde{u}$ remain the same as in u. In the game $\Gamma(u)$, by the complementary slackness of primal and dual solutions, j is indifferent between choosing d_j and any other e_j in $A_j(d_j)$ when dj is recommended by any μ that is incentive-compatible in $\Gamma(u)$; and so j's imagined chances of deviating from d_j to another action e_j in $A_j(d_j)$ would not affect j's expected payoffs from choosing d_j when it is recommended by μ in $\Gamma(\tilde{u})$. So for any mechanism μ that is incentive compatible in $\Gamma(u)$, the slack in primal incentive constraints corresponding to any $\alpha_j(c_j|d_j)$ would remain the same in $\Gamma(\tilde{u})$ as in $\Gamma(u)$. On the other hand, for primal incentive constraints corresponding to any $\alpha_j(d_j|c_j)$, the fact that j would gain by deviating in $\Gamma(u)$ from μ's recommendation of c_j to d_j or any other e_j implies that j would also not gain by deviating from any recommended c_j to the imagined mixture of d_j and other e_j that j gets from the action d_j in $\Gamma(\tilde{u})$. Thus, any μ that is incentive compatible in $\Gamma(u)$ is also incentive compatible in $\Gamma(\tilde{u})$, and any incentive constraint that μ satisfies with positive slack in $\Gamma(u)$ must also be satisfied with positive slack in $\Gamma(\tilde{u})$ (for all sufficiently small ε). So any dual variable which is 0 in the maximal dual solutions μ for $\Gamma(u)$ must remain 0 in all dual solutions for $\Gamma(\tilde{u})$.

But as u does not have transitive dual solutions (by supposition), the case-1 conditions $d_j \in A_j(c_j)$, $e_j \in A_j(d_j)$, & $e_j \notin A_j(c_j)$ for $\Gamma(u)$ together imply that the incentive constraint for the dual variable $\alpha_j(d_j|c_j)$ would now be satisfied strictly in $\Gamma(\tilde{u})$ (with $\varepsilon>0$) for any μ that is maximally incentive compatible in $\Gamma(u)$; and so the dual variable $\alpha_j(d_j|c_j)$ which is strictly positive in maximal dual solutions for $\Gamma(u)$ must be 0 in dual solutions for $\Gamma(\tilde{u})$. Thus, the dual magnitude of the game $\Gamma(\tilde{u})$ must be strictly smaller than the dual magnitude of the game with utilities u. But with ε sufficiently small, this $\tilde{u}$ will be in the given open set, contradicting the selection of u above.

Now consider case 2, where u has a violation of transitivity for some type s_i of some sender i. In case 2, for any small $\varepsilon>0$, we can construct $\tilde{u}$ to be the same as u except that, if player i's type is s_j then

$$p(t_{-i},s_i)\tilde{u}_i(c,t_{-i},s_i) = (1-\varepsilon)p(t_{-i},s_i)u_i(c,t_{-i},s_i) + \varepsilon\sum_{ri\in Qi(si)} \alpha_i(s_i|r_i)p(t_{-i},r_i)u_i(c,t_{-i},r_i).$$

That is, player i's $\tilde{u}$-payoff with type s_i would be as if s_i imagined that there was a small chance of i's having had another type r_i in $Q_j(s_j)$ which, with probability $\varepsilon\alpha_i(s_i|r_i)$, chose to identify itself as s_i and then forgot about this self-delusion. The payoffs of other players' or other types of i would not be affected by this imagined self-delusion of i's type s_i, however.
We now show that any maximal incentive-compatible mechanism μ for the game $\Gamma(u)$ will also be incentive compatible for this new game $\Gamma(\tilde{u})$, and all the incentive constraints that μ satisfies with positive slack for $\Gamma(u)$ will also be satisfied with positive slack for $\Gamma(\tilde{u})$. We only need to consider incentive constraints for type s_i of player i, because all other payoffs in $\tilde{u}$ remain the same as in u. In the game $\Gamma(u)$, by complementary slackness of primal and dual solutions, any type r_i in $Q_i(s_i)$ would be indifferent between reporting type s_i and reporting its true type, and weakly prefers reporting its true type over any other type t_i in the mechanism μ; so in $\Gamma(\tilde{u})$ where i is unsure whether he is getting u-payoffs for type s_i or for some other type r_i in $Q(s_i)$, participating honestly in the mechanism μ is still optimal for type s_i. Thus, any μ that is incentive compatible in $\Gamma(u)$ is also incentive compatible in $\Gamma(\tilde{u})$, and any incentive constraint that μ satisfies with positive slack in $\Gamma(u)$ must also be satisfied with positive slack in $\Gamma(\tilde{u})$ for all sufficiently small ε. So any dual variable which is 0 in the maximal dual solutions μ for $\Gamma(u)$ must remain 0 in all dual solutions for $\Gamma(\tilde{u})$.
Now, however, the case-2 dual-intransitivity conditions $s_i \in Q_i(t_i)$, $r_i \in Q_i(s_i)$, & $r_i \notin Q_i(t_i)$ for $\Gamma(u)$ would together imply that the incentive constaint for the dual variable $\alpha_i(t_i|s_i)$ would be satisfied strictly in $\Gamma(\tilde{u})$ (with $\varepsilon>0$) for any μ that is maximally incentive compatible in $\Gamma(u)$; and so the dual variable $\alpha_i(t_i|s_i)$ which is strictly positive in maximal dual solutions for $\Gamma(u)$ must be 0 in dual solutions for $\Gamma(\tilde{u})$. Thus, the dual magnitude of the game with payoffs $\tilde{u}$ must be strictly smaller than the dual magnitude of the game with utilities u. But with ε sufficiently small, this $\tilde{u}$ will also be in the given open set, again contradicting the magnitude-minimizing selection of u.

These contradictions for cases 1 and 2 together prove that, for utilities in any nonempty open subset of $\mathbb{R}^{(I\cup J)\times C\times T}$, the minimal dual magnitude can be achieved only at games that have transitive duals. Thus, utilities that yield games with transitive duals are dense in $\mathbb{R}^{(I\cup J)\times C\times T}$.