METEOR Seminar, Maasticht University

Definitive solutions of strategic games

Sergei Artemov

June 16, 2011

Prisoner's Dilemma

	$cooperate_B$		$defect_B$
$cooperate_A$	$\mathbf{2,2}$	$\Rightarrow$	$\mathbf{0,3}$
	$\Downarrow$		$\Downarrow$
$defect_A$	$\mathbf{3,0}$	$\Rightarrow$	$\mathbf{1,1}$

One Nash equilibrium:

$$(defect_A, defect_B).$$

Strictly dominant strategies for both players.

Bach or Stravinsky

	Bach		*Stravinsky*
Bach	**2**,**1**	$\Leftarrow$	**0**,**0**
	$\Uparrow$		$\Downarrow$
Stravinsky	**0**,**0**	$\Rightarrow$	**1**,**2**

There are two symmetric Nash equilibria:

(Bach, Bach) and *(Stravinsky,Stravinsky)*.

No dominant strategies.

Matching Pennies

	Head		*Tail*
Head	**1**,**-1**	$\Rightarrow$	**-1**,**1**
	$\Uparrow$		$\Downarrow$
Tail	**-1**,**1**	$\Leftarrow$	**1**,**-1**

There are no Nash equilibria and no dominant strategies.

War and Peace Dilemma

	war_S		$peace_S$
war_B	**2,0**	$\Rightarrow$	**1,2**
	$\Uparrow$		$\Downarrow$
$peace_B$	**0,1**	$\Rightarrow$	**3,3**

S has the dominant strategy $peace_S$, but not *B*.

The Nash equilibrium $(peace_B, peace_S)$.

Strategic (= normal-form) Games

Strategic game: $\langle N,(A_i),(\succsim_i)\rangle$

1. **Players**: a finite set $N=(1,2,...,n)$.
2. **Strategies**: for each player i, a nonempty set A_i of *strategies* available to i.
   Strategy profile is $a =(a_1,a_2,...,a_n)$ such that $a_i \in A_i$.
3. **Preference relation**: for each player $i \in N$, a *preference relation* $\succsim_i$ on
   $A = A_1 \times A_2 \times ... \times A_n$, $\succsim_i$, is complete reflexive transitive.

****

This official definition gives **the ordinal payoff model**: only the order of the preferences matter, not their cardinal value.

Preference relation is often represented by ***payoff functions***

$$u_i: A \rightarrow \mathbf{R}, \quad u_i(a) \geqq u_i(b) \quad \Leftrightarrow \quad a \succsim_i b \,.$$

A game is ***generic*** if for any i, payoffs for i of different outcomes are different.

Ordinal vs. Cardinal Payoffs

Strategic game: $\langle N,(A_i),(\succsim_i)\rangle$ *or* $\langle N,(A_i),(u_i)\rangle$

Prisoner's Dilemma with ordinal payoffs (preferences) ⇒

	$cooperate_B$		$defect_B$
$cooperate_A$	**2,2**	⇒	**0,3**
	⇓		⇓
$defect_A$	**3,0**	⇒	**1,1**

	$cooperate_B$		$defect_B$
$cooperate_A$	**-1,-1**	⇒	**-10,0**
	⇓		⇓
$defect_A$	**0,-10**	⇒	**-7,-7**

⇐ Prisoner's Dilemma with cardinal payoffs (years in jail)

Basically the same analysis!

Ordinal vs. Cardinal Payoffs

Strategic game: $\langle N,(A_i),(\succsim_i)\rangle$ *or* $\langle N,(A_i),(u_i)\rangle$

Bach-Stravinsky Dilemma with ordinal payoffs (preferences) ⇒

Mixed strategies solution: (2/3, 1/3)

	Bach		*Stravinsky*
Bach	**2**,1	⇐	0,0
	⇑		⇓
Stravinsky	0,0	⇒	1,**2**

	Bach		*Stravinsky*
Bach	**9**,1	⇐	0,0
	⇑		⇓
Stravinsky	0,0	⇒	1,**9**

Bach-Stravinsky Dilemma with cardinal payoffs (preferences) ⇐

Payoff cannot be measured in Euros, does not make much sense.
Mixed strategies solution: (9/10, 1/10)

Ordinal vs. Cardinal Payoffs

1. Mixed strategies can responsibly be applied ***only to cardinal payoffs***. In the ordinal payoffs domain, they can yield different results for the same game. **This makes ordinal payoff games the subject of pure strategies analysis.**

2. Making an ordinal payoff game a cardinal game requires **additional assumptions** (e.g., what is a $$ equivalent of your losses if you go to the wrong concert with your wife, how many roubles a victorious war would bring you, etc.) which are **not necessarily justified** by the game description.

Nash Equilibrium

Calculated from $\langle N, (A_i), (\succsim_i) \rangle$

Nash Equilibrium is a profile $a \in A$ such that no player can deviate profitably given the actions of other players, i.e., for each player $i \in N$,

$$(a_{-i}, a_i) \succsim_i (a_{-i}, \tilde{a}_i) \text{ for all actions } \tilde{a}_i \in A_i.$$

Here $a = (a_1, a_2, \ldots, a_n)$, $a_{-i} = (a_1, a_2, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$.

Should the players choose Nash equilibrium? The naive justification is circular:

A plays NE because B plays NE because A plays NE, etc.

Some games have many *NE* (e.g., BoS), some have none (e.g., MP).

Nash's own motivation for Equilibrium

In his dissertation of 1950, Nash explained that his concept of solution to a game was obtained by

*"...using the principles that **a rational prediction should be unique**, that the players should be able **to deduce** and make use of it..."*

and further:

*"... we need to assume **the players know the full structure of the game** in order to be able **to deduce** the prediction for themselves."*

We will show that under the usual game-theoretical assumptions, Nash's conditions are rarely met, and even if they are, the Nash Equilibrium is necessary, but not sufficient, for a definitive solution. Moreover, **many games just don't have definitive solutions**.

Aumann's account of rationality

> *Rationality of a player means that he is a habitual pay-off maximizer: that no matter where he finds himself – at which vertex – he will not knowingly continue with a strategy that yields him less than he could have gotten with a different strategy.*

This informal account has been made rigorous for specific classes of games, including strategic games. The intuitive formulation is quite sufficient for our talk. The key consequence:

rational players don't play dominated strategies

(think PD, WPD).

Definitive solutions and *NE*

CKGR - *Common Knowledge of the Game and Rationality*
(the standard assumption is Game Theory).

Rationality: we assume that any notion of rationality is at least as strong as Aumann's rationality at the same node:

Rational ⇒ *Aumann-rational.*

Definitive Solution, ***DS*** - strategy profile σ such that
each player knows that GAME RULES yields "all players choose σ."

DS finds natural formalization in the logic of knowledge.

As a convenient formalization feature, we assume that each strategy σ_i is completely specified by a corresponding logical sentence S_i stating that

player i has committed to strategy σ_i.

Definitive solutions and *NE*

A strategy profile $\sigma = \{\sigma_1, \sigma_2, \ldots, \sigma_n\}$ is a **definitive solution** of the game if each player knows that the choice of $\sigma_1, \sigma_2, \ldots, \sigma_n$ logically follows from the description of the game. We formalize the justified solution requirement by assuming

$$\mathbf{K}_i[GAME\ RULES \rightarrow (S_1 \wedge S_2 \wedge \ldots \wedge S_n)] \quad \text{for each } i = 1, 2, \ldots, n. \quad (1)$$

for sentences $S_1, S_2 \ldots S_n$ corresponding to $\sigma_1, \sigma_2, \ldots, \sigma_n$. Since different strategy profiles are incompatible, the definitive solution, if it exists, is unique.

Definitive solutions and *NE*

Theorem 1. *For any notion of rationality and any game with CKGR, a definitive solution is a Nash equilibrium.*

Proof. Without loss of generality, we check that ***DS*** yields that player 1 chooses ***NE***.

$$\mathbf{K}_1[\textit{GAME RULES} \rightarrow (S_1 \wedge S_2 \wedge \ldots \wedge S_n)].$$

By standard epistemic reasoning,

$$\mathbf{K}_1(\textit{GAME RULES}) \rightarrow \mathbf{K}_1(S_1 \wedge S_2 \wedge \ldots \wedge S_n)$$

and

$$\mathbf{K}_1(\textit{GAME RULES}) \rightarrow \mathbf{K}_1(S_1) \wedge \mathbf{K}_1(S_2) \wedge \ldots \wedge \mathbf{K}_1(S_n).$$

Since the game is known to each player,

$$\mathbf{K}_1(\textit{GAME RULES})$$

holds, hence player 1 knows the strategies of all other players:

$$\mathbf{K}_1(S_1) \wedge \mathbf{K}_1(S_2) \wedge \ldots \wedge \mathbf{K}_1(S_n).$$

Since 1 knows others' strategies, 1 plays ***NE***.

Definitive solutions and *NE*

Theorem 1. *For any notion of rationality and any game with CKGR, a definitive solution is a Nash equilibrium.*

Corollary 1. *Games without Nash equilibria do not have definitive solutions under any notions of rationality, no matter how strong.*

Before. (R. Aumann and A. Brandenburger. *Epistemic conditions for Nash equilibrium*. 1995):

all players know that σ is played ⇒ σ is NE.

After. ***σ is DS ⇒ all players know that σ is played,***

hence ***DS ⇒ NE,***

hence ***no NE ⇒ no DS ⇒ no Nash's solution.***

A right definition of *DS* + epistemic logic argument capture Nash's

"*to be able to deduce the prediction*"

line and lead to the impossibility conclusion w.r.t. Nash's ideal solution.

A unique *NE* is not necessarily a solution

Consider the following game

$$\begin{pmatrix} \mathbf{1},\mathbf{2} & 1,0 & 0,1 \\ \\ 0,0 & 0,\mathbf{2} & \mathbf{1},1 \end{pmatrix}$$

It has a unique Nash equilibrium **(1,2)**, but no definitive solution within the scope of Aumann's rationality, even if the game and rationality are commonly known. Indeed, each strategy in this game is Aumann-rational and hence cannot be deleted by *IDSDS*.

Many Nash equilibria - no definitive solutions?

For Aumann's rationality, any strategic game with two or more Nash equilibria does not have a definitive solution.

This formulation is naive and admits a counterexample. Consider the payoff matrix of the generic version of the Coordination game:

$$\begin{pmatrix} 3,3 & 1,1 \\ \\ \\ 0,0 & 2,2 \end{pmatrix}$$

and define a new game *G* by adding some epistemic specifications to this game:

each player knows that his opponent will be playing **2**.

Then for each player, the only Aumann-rational strategy is to play **2**, too!
This solution singles out as a definitive solution one of the two Nash equilibria, which seemingly contradicts the "naive" formulation above.

The moral: control epistemic assumptions!

Stronger rationality can yield *DS* with many *NE*

Let us fix a game G with n players, for each player $i = 1, 2, \ldots, n$ a nonempty set A_i of actions (strategies) available to i, and a payoff function u_i. Let also strategy profile

$$\sigma = \{\sigma_1, \sigma_2, \ldots, \sigma_n\}$$

be a Nash equilibrium of game G. A **Bullet Nash Rationality** corresponding to σ is the predicate on strategies which accepts as rational for each player $i = 1, 2, \ldots, n$ only strategy σ_i from σ.

Lemma 1 *Bullet Nash Rationality yields Aumann's rationality.*

Proof. Easy, since σ is an Nash equilibrium.

Stronger rationality can yield *DS* with many *NE*

Theorem 1 *Let G be a strategic game under Bullet Nash rationality corresponding to a Nash equilibrium σ. Then σ is a definitive solution to G.*

Proof. Since each player is Bullet Nash-rational with σ, each player i choses σ_i. If everybody knows that all players are Bullet Nash-rational with σ, everybody knows that each player i choses σ_i. □

Corollary 2 *If a game has a Nash equilibrium, then for an appropriate extension of Aumann's rationality, such a game has a definitive solution.*

Such "reverse engineered" Bullet Nash rationality is a theoretical notion which we do not offer as a viable practical notion. However, Bullet Nash rationality helps us to see the boundaries of the notion of definitive solution.

Theorem 1 also has a close predecessor in [3] where it was stated that for any Nash equilibrium σ there is a belief system in which each player assigns probability one to σ. Theorem 1 conveys basically the same message with the "probability one" replaced by strictly deductive notion of knowledge.

This is a theoretical space for refinement methods: consider stronger notions of rationality. However, on this path, one could steer to any Nash equilibrium...

Regular Strategic Games

We now return to the normal Aumann's notion of rationality, and we want to answer the question of whether Aumann's rationality suffices to resolve a game with multiple Nash equilibria. We can achieve definitive solution by either twisting a player's knowledge of the game, or the notion of rationality.

The question is whether there could be a definitive solution to a regular game with multiple Nash equilibria within the regular (Aumann's) notion of rationality. And the answer is NO.

We have to define a regular game though... In addition to their knowledge of the world - elementary logic, arithmetic, etc., players commonly know the parameters of the given game: # of players, possible strategies and payoffs, as well as players' rationality - all standard game-theoretical assumptions, ***but nothing more about each others' intentions.***

There is a tedious formal definition, of course...

Regular Strategic Games

a. Conditions on strategy propositions S_i^j stating '*player i chooses strategy j*.' These conditions express that each player i chooses one and only one strategy:

$$(S_i^1 \vee \ldots \vee S_i^{n_i}) \text{ and } \neg(S_i^j \wedge S_i^l) \text{ for each } j \neq l.$$

This part is determined **by the size of the game**: the number of players and number of strategies for each player.

b. A complete description of the preference relation for each player at each outcome. This part is determined **by the payoff matrix**.

c. Knowledge of one's own moves: $S_i^j \rightarrow \mathbf{K}_i(S_i^j)$ and $(\neg S_i^j) \rightarrow \mathbf{K}_i(\neg S_i^j)$ for all i, j's. This part is determined **by the size of the game**.

d. Aumann's rationality condition: *for each* $i = 1, \ldots, n$, *if player* i *knows that his strategy* j *is strictly dominated, then* $\neg S_i^j$. We may assume that this condition is determined **by the size of the game only** and covers all combinations of preferences on outcomes and all epistemic conditions of players, hence does not depend on particular set of preferences and epistemic conditions for this game.

e. Common knowledge of **a** – **d** above.

Knowing a Nash equilibrium is consistent

Lemma 2 *A regular strategic game is consistent with the knowledge of any of its Nash equilibria: for each player i and Nash equilibrium* $\{\sigma_1^{e_1}, \ldots, \sigma_n^{e_n}\}$,

$$\mathbf{K}_i[GAME\ RULES] \wedge \mathbf{K}_i(S_1^{e_1} \wedge \ldots \wedge S_n^{e_n})$$

is consistent.

This is the key lemma established by presenting an appropriate Kripke-style model of the game in the logic of knowledge.

The probabilistic version of Lemma 2 was found in R. Aumann and A. Brandenburger. *Epistemic conditions for Nash equilibrium*. 1995

In that paper, “know” means “ascribe probability 1 to.”
The A&B result states that for any *NE*, there is a belief system which assigns probability 1 to players’ beliefs in this *NE*.

Playing a Nash equilibrium is consistent

Corollary 4 *A regular strategic game is consistent with playing any of its Nash equilibria: for any Nash equilibrium* $\{\sigma_1^{e_1}, \ldots, \sigma_n^{e_n}\}$,

$$GAME\ RULES \wedge (S_1^{e_1} \wedge \ldots \wedge S_n^{e_n}) \quad \textit{is consistent.}$$

Proof. In the proof of Lemma 2, $\sigma \Vdash GAME\ RULES \wedge (S_1^{e_1} \wedge \ldots \wedge S_n^{e_n})$. □

Impossibility Theorem: Mutiple *NE* - no *DS*

Theorem 3 *No regular strategic game with more than one Nash equilibium can have a definitive solution.*

Proof. Logician's trick: since knowledge of each of different Nash equilibria is consistent and these cases are exclusive, none follows from the game description.

Impossibility Theorems: findings

1. ***No NE*** - no *DS* for any notion of rationality, no matter how strong.
The requirement *DS*+No*NE* is just logically inconsistent.

2. ***One NE*** - can be either way.

3. ***Multiple NE*** - can be refined to any given equilibrium by assuming more and more from the game and/or from the notion of rationality. However, for the regular notion of rationality and a regular game with the standard assumption "common knowledge of the game and rationality," there are no definitive solutions. This invites studying *DS* for a range of "real" epistemic/rationality assumptions.

Moral: For strategic games with ordinal payoffs (the most general class of strategic games), Nash conditions for the equilibrium solution concepts are not met in cases (1) and (3). In case (2), Nash conditions sometimes hold and sometimes do not hold, but the Nash Equilibrium is not a criterion there.

Apparently, Nash's solution concept does not look properly principled here. The right approach to solving strategic games with ordinal payoffs is **iterated dominance**.

Iterated Deletion of Strictly Dominated Strategies

	war_S		$peace_S$
war_B	$\mathbf{2,0}$	$\Rightarrow$	$\mathbf{1,2}$
	$\Uparrow$		$\Downarrow$
$peace_B$	$\mathbf{0,1}$	$\Rightarrow$	$\mathbf{3,3}$

1. For all players, delete all strictly dominated strategies.
2. Iterate 1 as long as possible.

In WPD above:
First iteration: war_S is deleted, what remains is the right column only;
Second iteration: war_B is deleted, which leaves us with just one node
$peace_B, peace_S$.
Does not always converge to a single node, can be justified by *CKGR*.

Iterated Deletion of Strictly Dominated Strategies

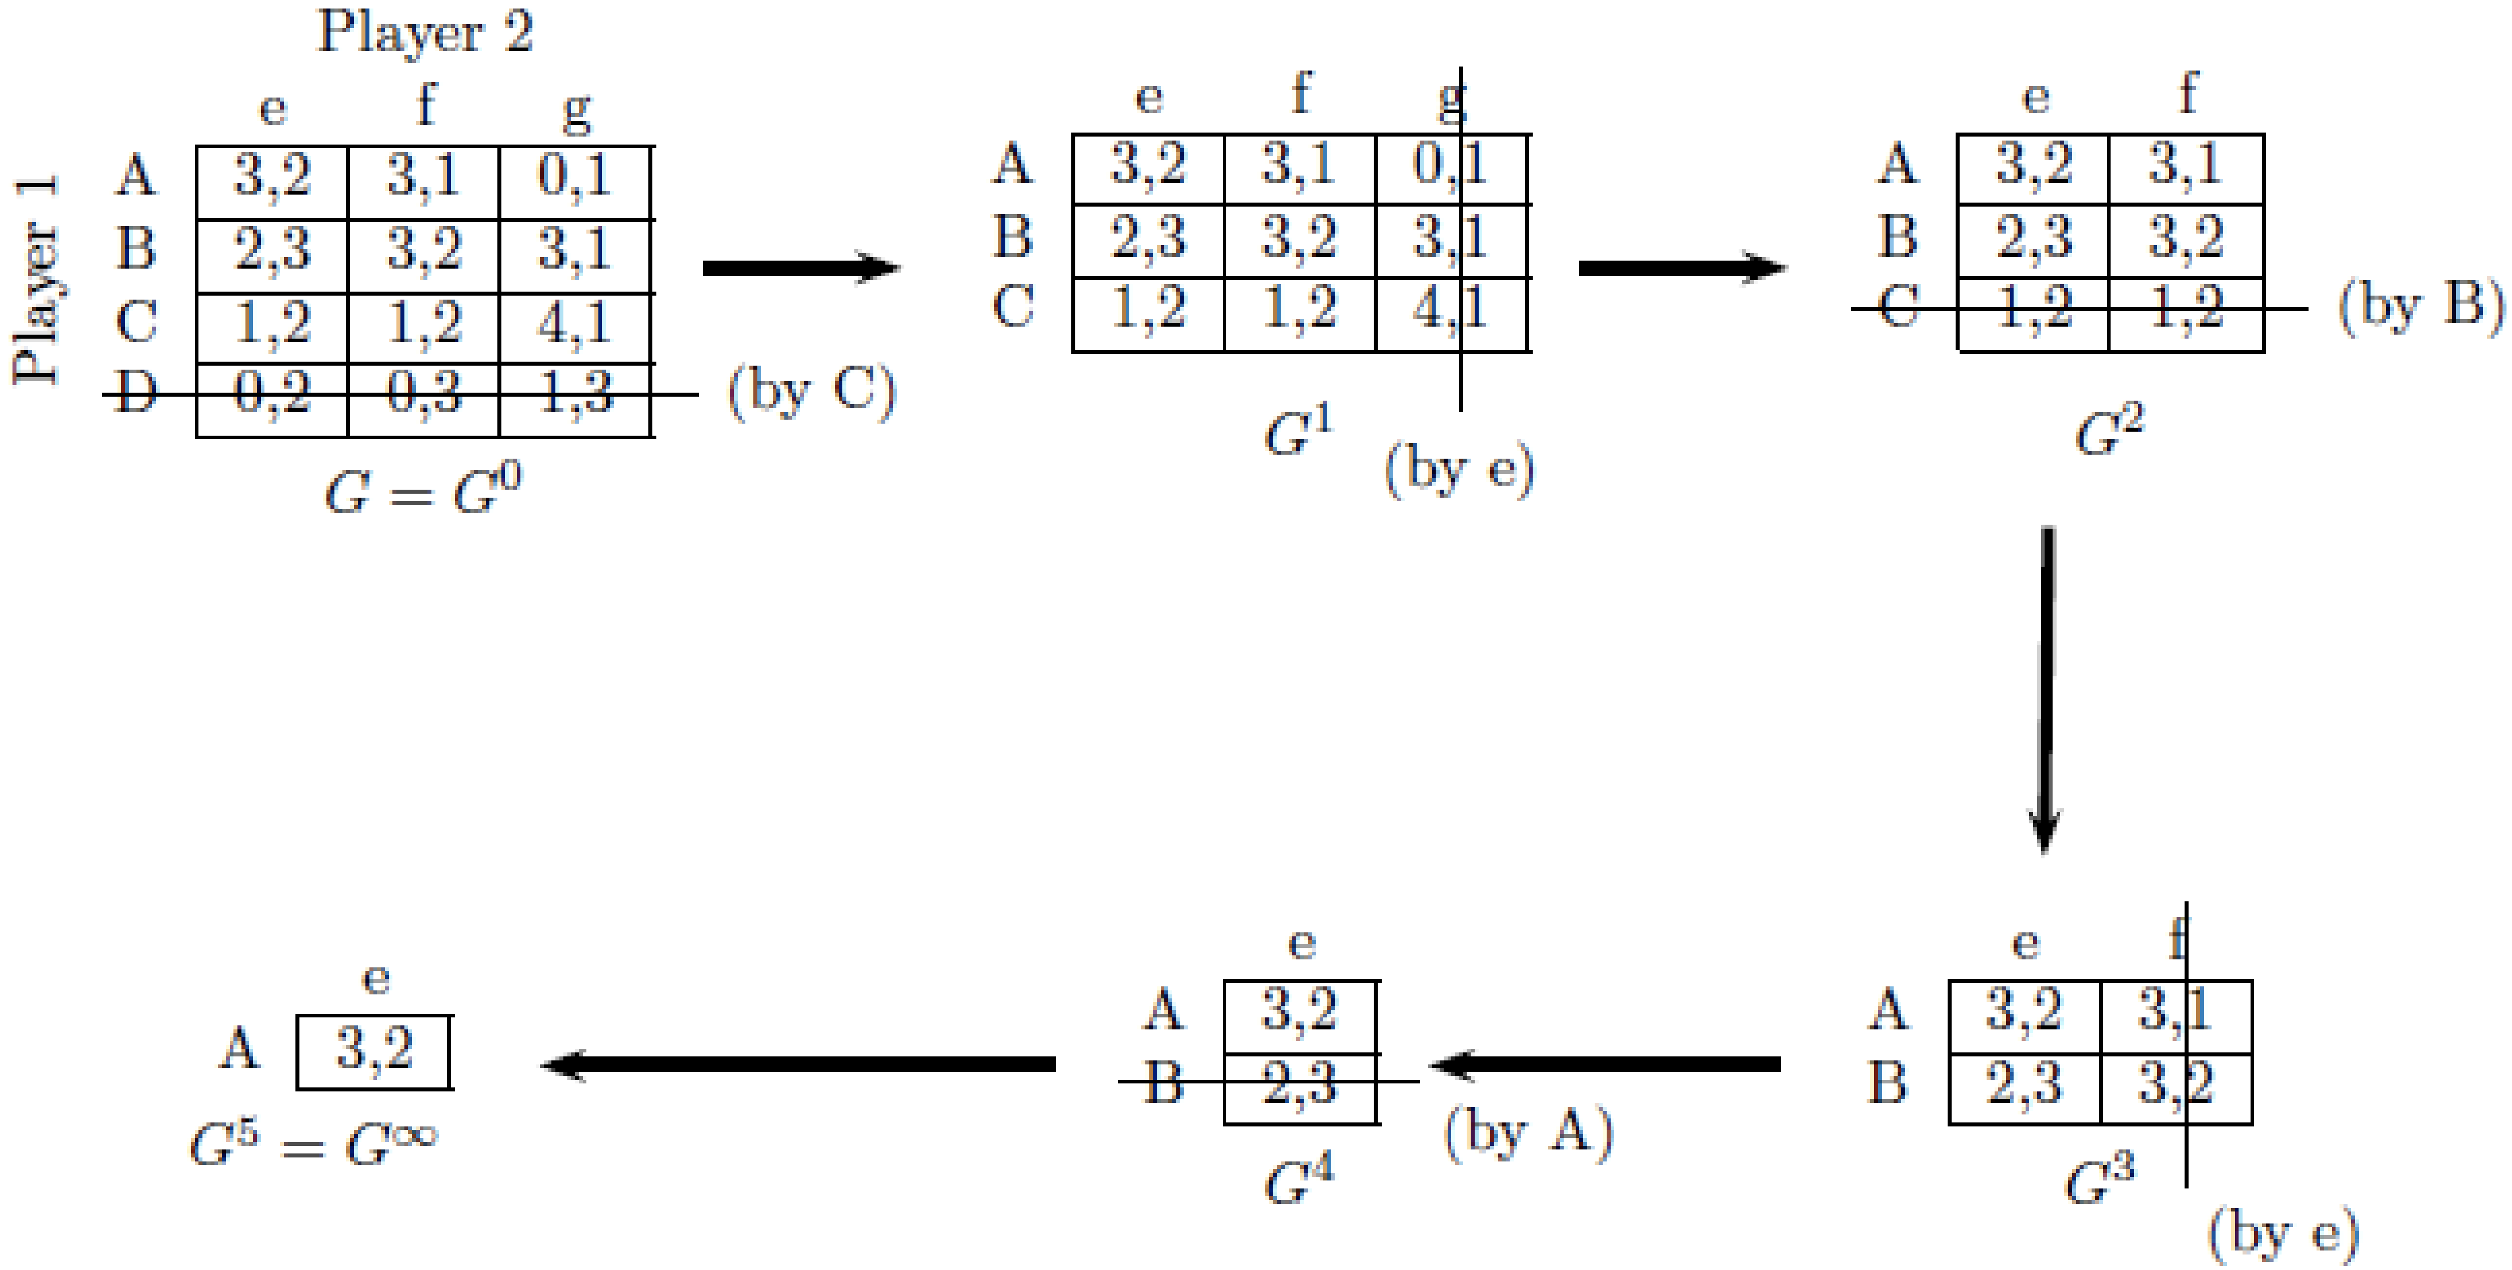

FIGURE 2. Illustration of the iterated deletion of strictly dominated strategies.

Figure from Giacomo Bonanno's paper of 2008.

Criterion of *DS* for strategic games

IDSDS *= Iterated Deletion of Strictly Dominated Strategies.*

Theorem [*IDSDS*=*DS*].
For a regular strategic game G with Aumann-rational players,
a) if IDSDS converges to a single profile σ, then σ is the DS to G;
b) if IDSDS does not converge to a single profile, then G does not have DS.

Corollary. *A regular strategic game with Aumann-rational players has DS iff IDSDS converges.*

Proof. Rather involved model-theoretical construction.
For (a), it suffices to prove that no solution can be deleted by *IDSDS* (easy).
For (b), suppose *IDSDS* stalls, leaving at least two strategy profiles *S1* and *S2* which hence do not dominate each other. We build models for

$$GAME\ RULES \cup \{S1\} \text{ and } GAME\ RULES \cup \{S2\}$$

to show that each of these sets is consistent, hence neither *S1* nor *S2* is derivable.
Giacomo Bonanno's paper of 2008 helps to find appropriate models,

What does logic add to Game Theory?

Before. (R. Aumann and A. Brandenburger. *Epistemic conditions for Nash equilibrium*. 1995):

all players know that σ is played ⇒ σ is NE.

Now: ***σ is DS ⇒ all players know that σ is played,***

hence ***DS ⇒ NE,***

hence ***no NE ⇒ no DS ⇒ no Nash's solution!***

What does logic add to Game Theory?

Before. *(Aumann & Brandenburger, 1995):*

for any Nash equilibrium σ, there is a belief system in which each player assigns probability one to σ.

After. ***for any NE σ there is a notion of rationality in which σ is DS,***

hence ***game has DS for some notion of rationality*** ⇔ ***game has NE .***

Furthermore, for Aumann's rationality,

for any NE σ, it is consistent with the rules of the game to assume that all players know that σ is played,

hence

two or more Nash Equilibria ⇒ ***no DS.***

What does logic add to Game Theory?

These impossibility observations need a proper (logical) setting to be set and proved formally. Comparison: Halting Problem in computability Theory. Before Turing's impossibility result the state of the art was

I don't know how to establish algorithmically whether a given program halts on a given input, and I don't see how this is at all possible.

Turing Theorem states that

it is impossible to establish algorithmically whether a given program halts on a given input.

The same happens here. Before:

I don't see how a regular game with two or more NE could possibly have Nash's definitive solution.

After: *A regular game with two or more NE does not have Nash's definitive solution.*

What does logic add to Game Theory?

Before. *(Nash, 1950)* ***A vision of a Definitive Solution.***

After. *Complete and efficient characterization of DS in strategic games with CKGR:*

IDSDS = DS .

Note:

- 75% of generic 2x2 games have *DS,*
- about 25% of generic 3x3 games have *DS*,
- the proportion of solvable games **quickly goes to 0** when the size of the game grows.

Further work

Mixed strategies: the same consistency lemma and the the same conclusion about absence of definitive solutions in mixed strategies holds here.

Other specific notions of rationality could be studied.

Arbitrary epistemic conditions should also be considered.