ARTIFICIAL INTELLIGENCE RESEARCH GROUP

ON NASH EQUILIBRIA IN MULTI-OBJECTIVE GAMES

Willem Röpke

Roxana Rădulescu

Diederik M. Roijers

Ann Nowé

willem.ropke@vub.be https://twitter.com/willem_ropke

OVERVIEW

- Background
- Game-theoretic work
- What's next
- Q&A

Multi-objective games present a natural framework for studying ***strategic interactions between rational individuals concerned with more than one objective***.

Strategic interactions between rational individuals

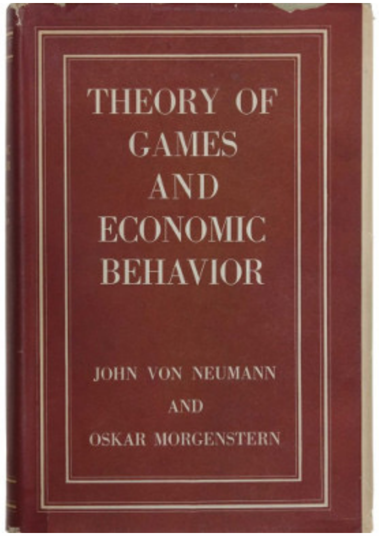


Game theory

Rational individuals concerned with more than one objective

Multi-objective decision making

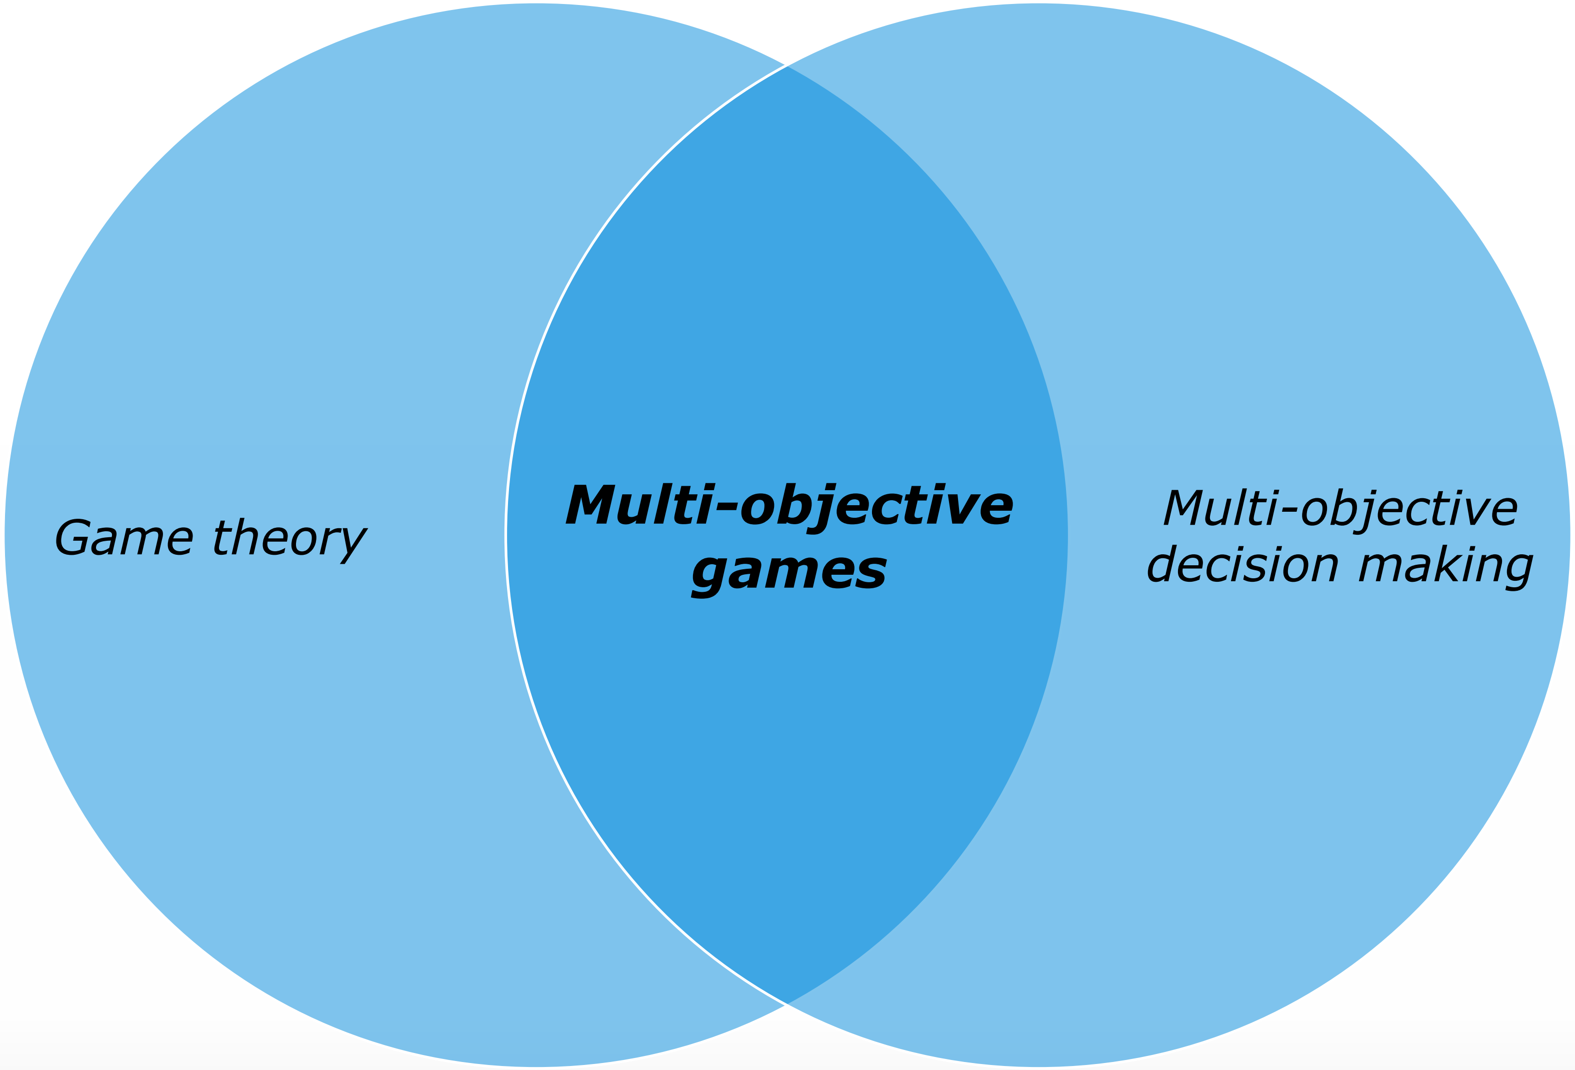
Game theory
Multi-objective games
Multi-objective decision making

WHAT
BACKGROUND

- Multi-Objective Normal-Form Games (MONFGs)
- Utility based approach
- Utility function $u_i \colon \mathbb{R}^d \rightarrow \mathbb{R}$

	A	B
A	(10, 2); (10, 2)	(2, 3); (2, 3)
B	(4, 2); (4, 2)	(6, 3); (6, 3)

$$u_1(p_1, p_2) = p_1 \cdot p_2$$

OPTIMISATION CRITERIA

BACKGROUND

- Two possible choices of optimisation criteria

- Expected Scalarised Returns (ESR)
  - Calculate the *expectation of your utility* from the payoffs
  - Utility of an individual policy execution

- Scalarised Expected Returns (SER)
  - Calculate the *utility of your expected payoff*
  - Utility of the average payoff from several executions of the policy

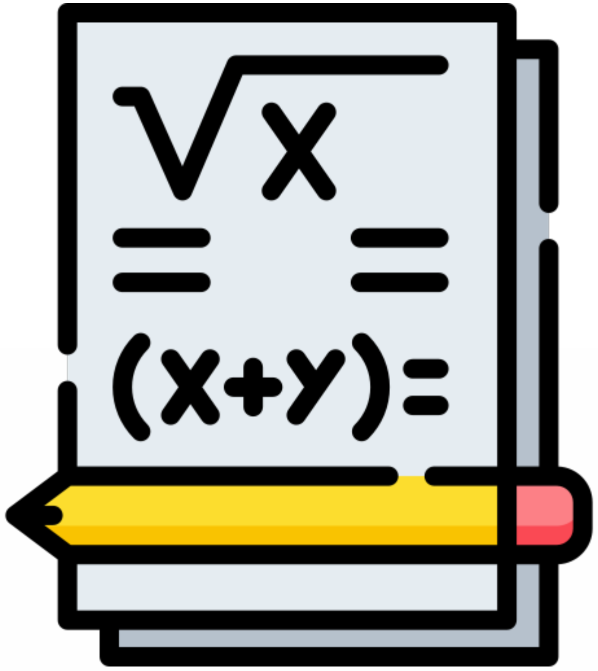

OPTIMISATION CRITERIA

EXAMPLE

$\boldsymbol{p} = (4, 2)$

Car

$\boldsymbol{p} = (2, 4)$

Bike

$$u(p_1, p_2) = p_1 \cdot p_2$$

Speed

Eco-friendliness

What happens when you take the car 50% of the time and the bike 50% of the time?

OPTIMISATION CRITERIA

EXAMPLE

$\boldsymbol{p} = (4, 2)$

Car

$\boldsymbol{p} = (2, 4)$

Bike

$$u(p_1, p_2) = p_1 \cdot p_2$$

ESR

OPTIMISATION CRITERIA

EXAMPLE

$\boldsymbol{p} = (4, 2)$

Car

$\boldsymbol{p} = (2, 4)$

Bike

$u(p_1, p_2) = p_1 \cdot p_2$

$u(4, 2) = 8$

ESR

$u(2, 4) = 8$

OPTIMISATION CRITERIA

EXAMPLE

$\boldsymbol{p} = (4, 2)$

Car

$\boldsymbol{p} = (2, 4)$

Bike

$u(p_1, p_2) = p_1 \cdot p_2$

ESR

$u(4, 2) = 8$

$\frac{1}{2} \cdot 8 = 4$

$u(2, 4) = 8$

$\frac{1}{2} \cdot 8 = 4$

OPTIMISATION CRITERIA

EXAMPLE

$\boldsymbol{p} = (4, 2)$

Car

$\boldsymbol{p} = (2, 4)$

Bike

$u(p_1, p_2) = p_1 \cdot p_2$

$u(4, 2) = 8$

$\frac{1}{2} \cdot 8 = 4$

ESR

+ 8 +

$u(2, 4) = 8$

$\frac{1}{2} \cdot 8 = 4$

$$\mathbb{E}u(\boldsymbol{p}) = 8$$

VRIJE UNIVERSITEIT BRUSSEL

OPTIMISATION CRITERIA

EXAMPLE

$\boldsymbol{p} = (4, 2)$

Car

$\boldsymbol{p} = (2, 4)$

Bike

$$u(p_1, p_2) = p_1 \cdot p_2$$

SER

OPTIMISATION CRITERIA

EXAMPLE

$\boldsymbol{p} = (4, 2)$

Car

$\boldsymbol{p} = (2, 4)$

Bike

$$u(p_1, p_2) = p_1 \cdot p_2$$

SER

$$\frac{1}{2} \cdot (4, 2) = (2, 1)$$

$$\frac{1}{2} \cdot (2, 4) = (1, 2)$$

OPTIMISATION CRITERIA

EXAMPLE

$\boldsymbol{p} = (4, 2)$

Car

$\boldsymbol{p} = (2, 4)$

Bike

$u(p_1, p_2) = p_1 \cdot p_2$

$\frac{1}{2} \cdot (4, 2) = (2, 1)$

SER

$\frac{1}{2} \cdot (2, 4) = (1, 2)$

$+ \rightarrow (3, 3) \leftarrow +$

VRIJE UNIVERSITEIT BRUSSEL

OPTIMISATION CRITERIA

EXAMPLE

$$\boldsymbol{p} = (4, 2)$$

Car

$$\boldsymbol{p} = (2, 4)$$

Bike

$$u(p_1, p_2) = p_1 \cdot p_2$$

SER

$$\frac{1}{2} \cdot (4, 2) = (2, 1)$$

$$\frac{1}{2} \cdot (2, 4) = (1, 2)$$

$$(2,1) + (1,2) = (3, 3)$$

$$u(\mathbb{E}\boldsymbol{p}) = 9$$

OPTIMISATION CRITERIA

EXAMPLE

$\boldsymbol{p} = (4, 2)$

Car

$\boldsymbol{p} = (2, 4)$

Bike

$$u(p_1, p_2) = p_1 \cdot p_2$$

What happens when you take the car 50% of the time and the bike 50% of the time?

ESR = 8

SER = 9

SOLUTION CONCEPTS

BACKGROUND

- Nash equilibria
  - *No agent can improve their utility by unilatteraly deviating from the joint strategy*

- Nash equilibria
  - *No agent can improve their utility by unilatteraly deviating from the joint strategy*

$$u_1(p_1, p_2) = u_2(p_1, p_2) = p_1 \cdot p_2$$

	A	B
A	(10, 2); (10, 2)	(0, 0); (0, 0)
B	(0, 0); (0, 0)	(2, 10); (2, 10)

SOLUTION CONCEPTS

BACKGROUND

- Nash equilibria
  - *No agent can improve their utility by unilatteraly deviating from the joint strategy*

$$u_1(p_1, p_2) = u_2(p_1, p_2) = p_1 \cdot p_2$$

Nash equilibrium → (A, A)

$$u_1(10, 2) = 10 \cdot 2 = 20$$
$$u_2(10, 2) = 10 \cdot 2 = 20$$

	A	B
A	(10, 2); (10, 2)	(0, 0); (0, 0)
B	(0, 0); (0, 0)	(2, 10); (2, 10)

- Nash equilibria
  - *No agent can improve their utility by unilatteraly deviating from the joint strategy*

$$u_1(p_1, p_2) = u_2(p_1, p_2) = p_1 \cdot p_2$$

Nash equilibrium →

$$u_1(10, 2) = 10 \cdot 2 = 20$$
$$u_2(10, 2) = 10 \cdot 2 = 20$$

	A	B
A	(10, 2); (10, 2)	(0, 0); (0, 0)
B	(0, 0); (0, 0)	(2, 10); (2, 10)

Strictly worse to deviate to B

- Nash equilibria
  - *No agent can improve their utility by unilatteraly deviating from the joint strategy*

$$u_1(p_1, p_2) = u_2(p_1, p_2) = p_1 \cdot p_2$$

Nash equilibrium → (A, A)

Strictly worse to deviate to B

	A	**B**
A	(10, 2); (10, 2)	(0, 0); (0, 0)
B	(0, 0); (0, 0)	(2, 10); (2, 10)

$$u_1(10, 2) = 10 \cdot 2 = 20$$
$$u_2(10, 2) = 10 \cdot 2 = 20$$

GOAL

- Theoretical
  - Existence or non-existence guarantees
  - Algorithms
- Learning in these environments
  - Communication
  - Commitment

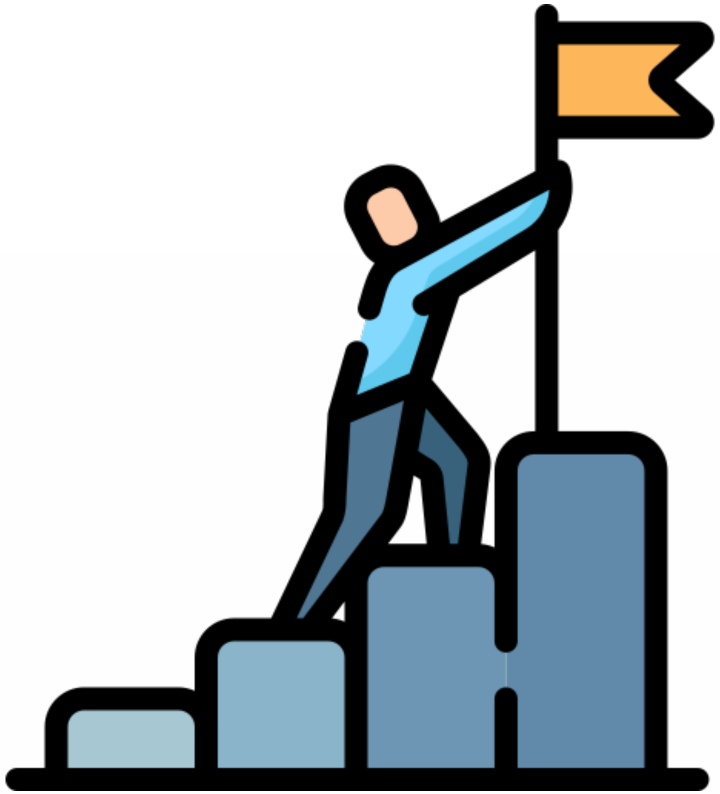

GOAL

- **Theoretical**
  - **Existence or non-existence guarantees**
  - **Algorithms**
- Learning in these environments
  - Communication
  - Commitment

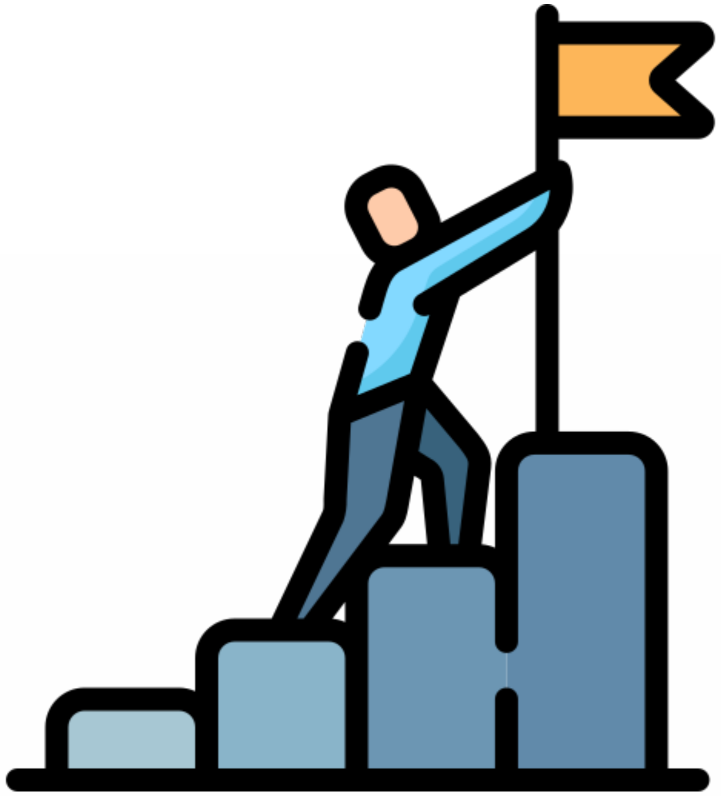

WHAT ARE MULTI-OBJECTIVE GAMES?

A NOVEL INTUITION

	A	B	C
A	(4, 1); (4, 1)	(1, 2); (4, 2)	(2, 1); (1, 2)
B	(3, 1); (2, 3)	(3, 2); (6, 3)	(1, 2); (2, 1)
C	(1, 2); (2, 1)	(2, 1); (1, 2)	(1, 3); (1, 3)

It turns out we can go from this

VRIJE UNIVERSITEIT BRUSSEL

WHAT ARE MULTI-OBJECTIVE GAMES?

A NOVEL INTUITION

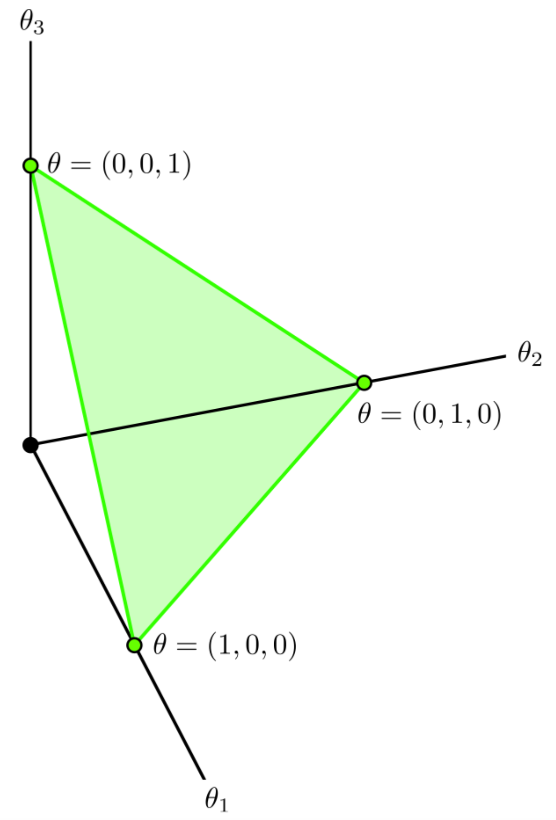

To this

Every MONFG with continuous utility functions can be reduced to a continuous game

Continuous game

- Single objective
- Infinite number of pure strategies
- Reuse utility functions

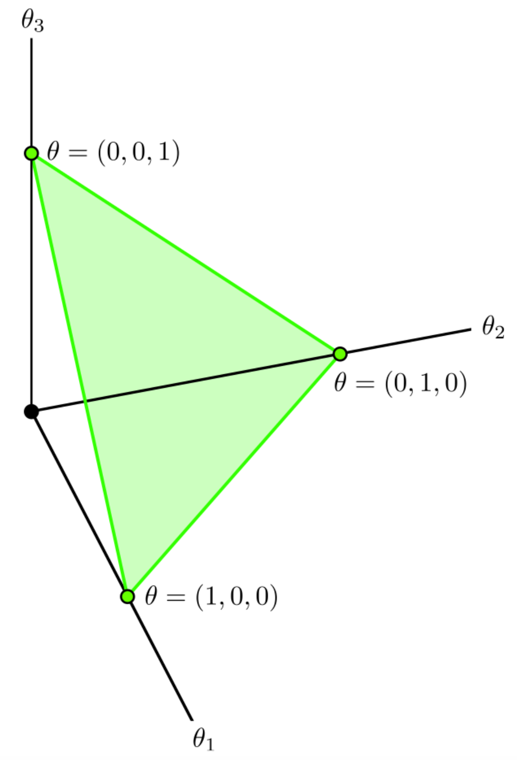

Make every ***mixed strategy*** in the MONFG a ***pure strategy*** in the continuous game

	A	B	C
A	(4, 1); (4, 1)	(1, 2); (4, 2)	(2, 1); (1, 2)
B	(3, 1); (2, 3)	(3, 2); (6, 3)	(1, 2); (2, 1)
C	(1, 2); (2, 1)	(2, 1); (1, 2)	(1, 3); (1, 3)

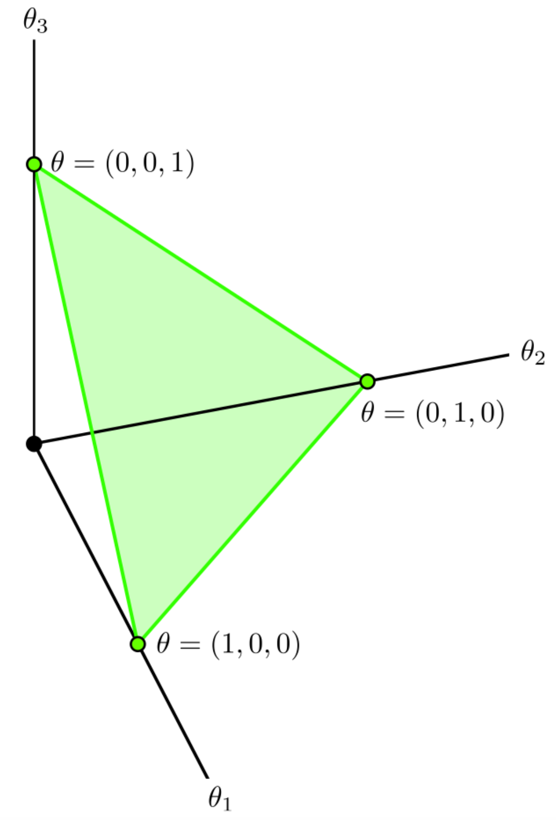

Make every ***mixed strategy*** in the MONFG a ***pure strategy*** in the continuous game

	A	B	C
A	(4, 1); (4, 1)	(1, 2); (4, 2)	(2, 1); (1, 2)
B	(3, 1); (2, 3)	(3, 2); (6, 3)	(1, 2); (2, 1)
C	(1, 2); (2, 1)	(2, 1); (1, 2)	(1, 3); (1, 3)

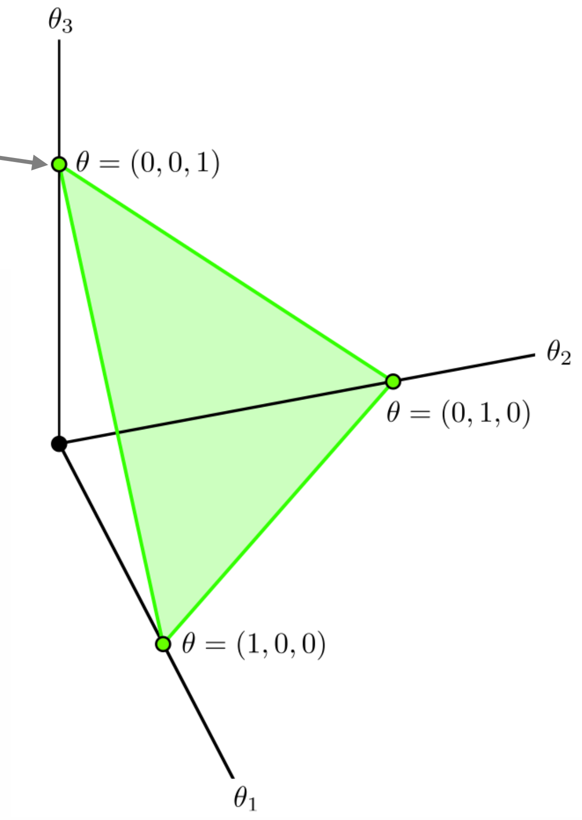

Make every ***mixed strategy*** in the MONFG a ***pure strategy*** in the continuous game

	A	B	C
A	(4, 1); (4, 1)	(1, 2); (4, 2)	(2, 1); (1, 2)
B	(3, 1); (2, 3)	(3, 2); (6, 3)	(1, 2); (2, 1)
C	(1, 2); (2, 1)	(2, 1); (1, 2)	(1, 3); (1, 3)

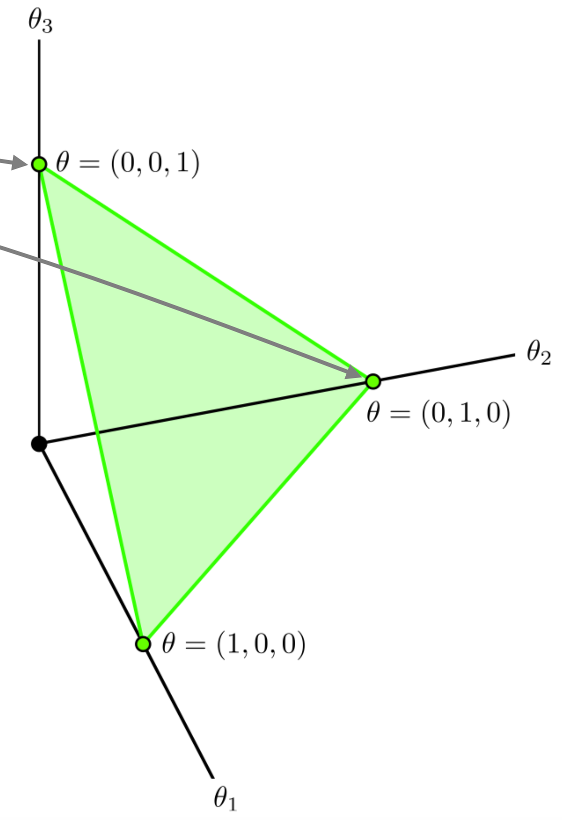

Make every ***mixed strategy*** in the MONFG a ***pure strategy*** in the continuous game

	A	B	C
A	(4, 1); (4, 1)	(1, 2); (4, 2)	(2, 1); (1, 2)
B	(3, 1); (2, 3)	(3, 2); (6, 3)	(1, 2); (2, 1)
C	(1, 2); (2, 1)	(2, 1); (1, 2)	(1, 3); (1, 3)

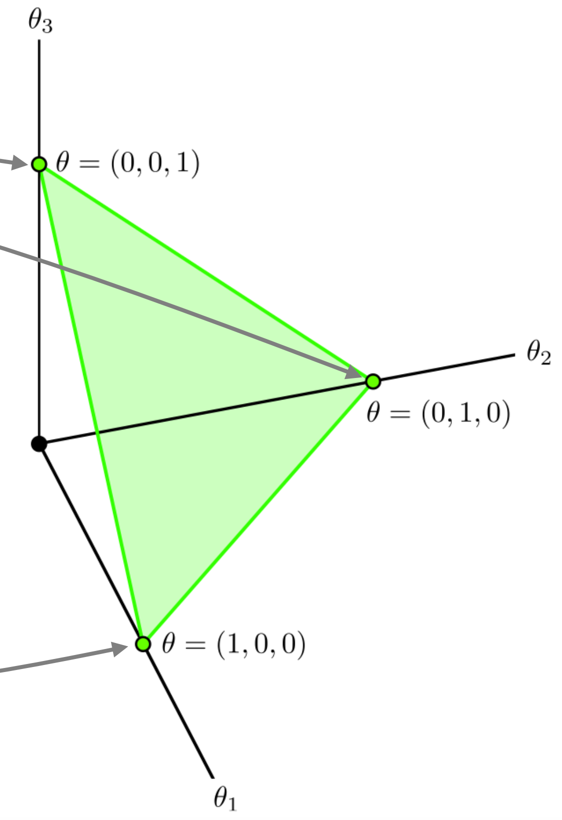

Make every ***mixed strategy*** in the MONFG a ***pure strategy*** in the continuous game

	A	B	C
A	(4, 1); (4, 1)	(1, 2); (4, 2)	(2, 1); (1, 2)
B	(3, 1); (2, 3)	(3, 2); (6, 3)	(1, 2); (2, 1)
C	(1, 2); (2, 1)	(2, 1); (1, 2)	(1, 3); (1, 3)

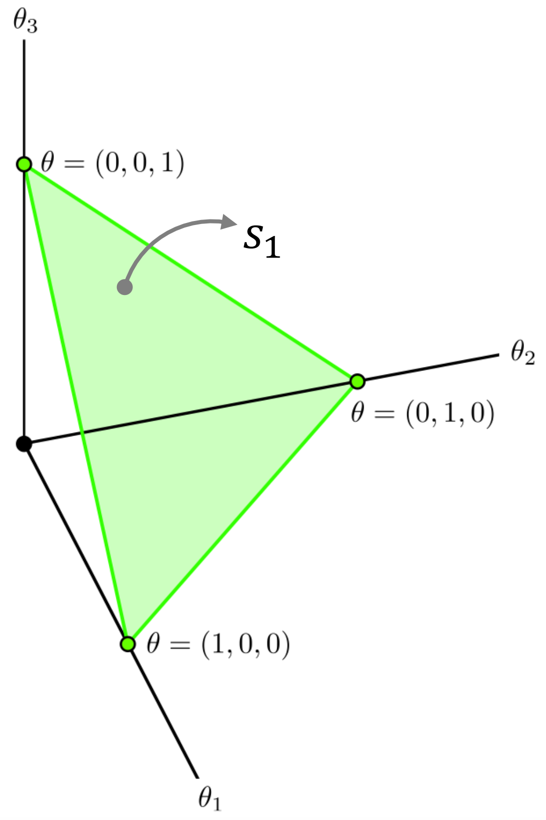

Make every ***mixed strategy*** in the MONFG a ***pure strategy*** in the continuous game

	A	B	C
A	(4, 1); (4, 1)	(1, 2); (4, 2)	(2, 1); (1, 2)
B	(3, 1); (2, 3)	(3, 2); (6, 3)	(1, 2); (2, 1)
C	(1, 2); (2, 1)	(2, 1); (1, 2)	(1, 3); (1, 3)

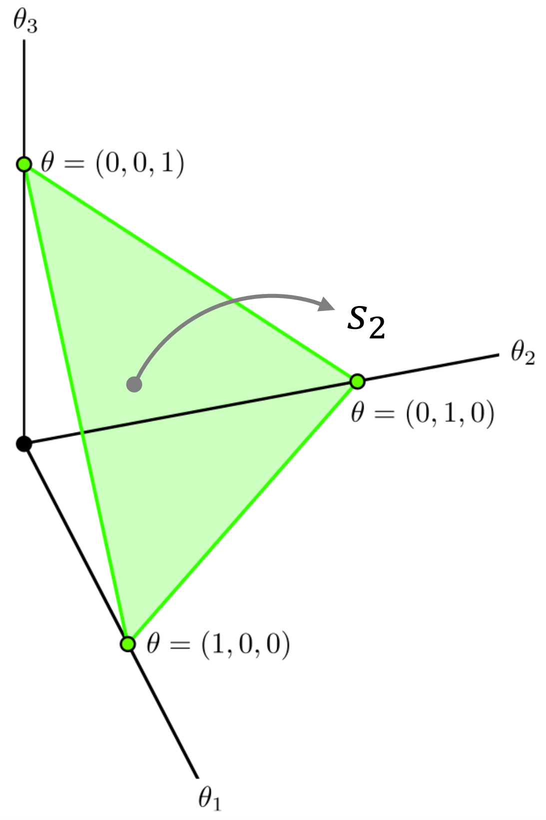

Make every ***mixed strategy*** in the MONFG a ***pure strategy*** in the continuous game

	A	B	C
A	(4, 1); (4, 1)	(1, 2); (4, 2)	(2, 1); (1, 2)
B	(3, 1); (2, 3)	(3, 2); (6, 3)	(1, 2); (2, 1)
C	(1, 2); (2, 1)	(2, 1); (1, 2)	(1, 3); (1, 3)

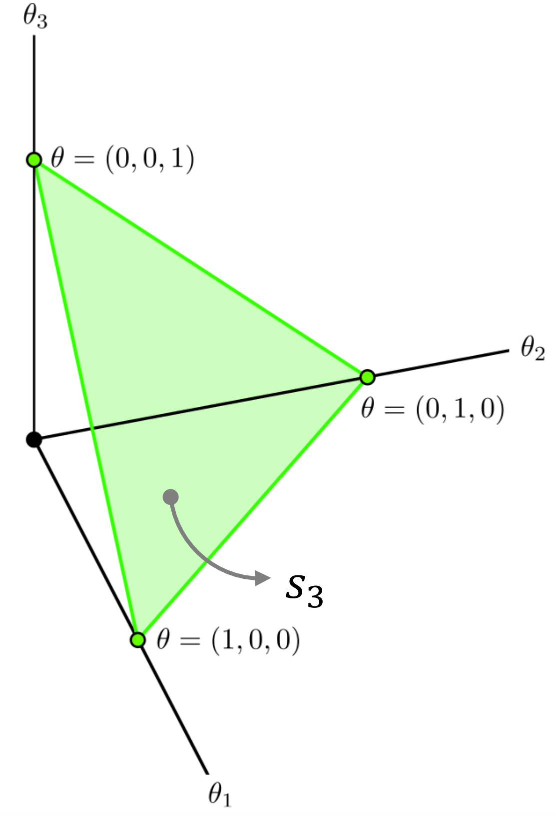

VRIJE UNIVERSITEIT BRUSSEL

WHY ARE NASH EQUILIBRIA NOT GUARANTEED?

A NOVEL INTUITION

- Nash equilibria are not guaranteed in MONFGs
  - They are guaranteed in single-objective NFGs, so why not here?
- ***Mixed strategy equilibria*** in the MONFG are ***pure strategy equilibria*** in the continuous game
- Continuous games are not guaranteed to have a ***pure strategy*** Nash equilibrium

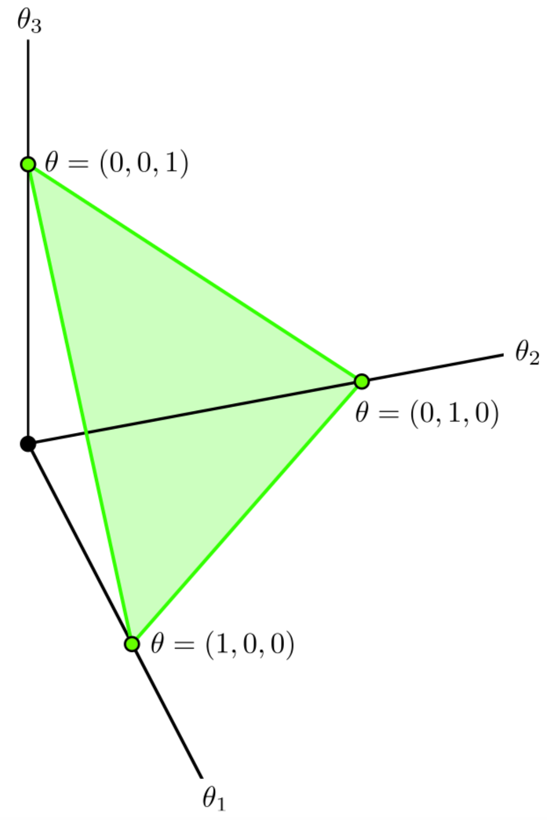

EXISTENCE GUARANTEE

- Existence is guaranteed with ***(quasi)concave*** utility functions
  - Used in economics as well
  - Represents "well-behaved" preferences
- Intuition
  - You can reduce an MONFG to a continuous game
  - In this game it is known that a pure strategy Nash equilibrium exists when assuming only quasiconcave utility functions
  - This equilibrium is also an equilibrium in the original MONFG

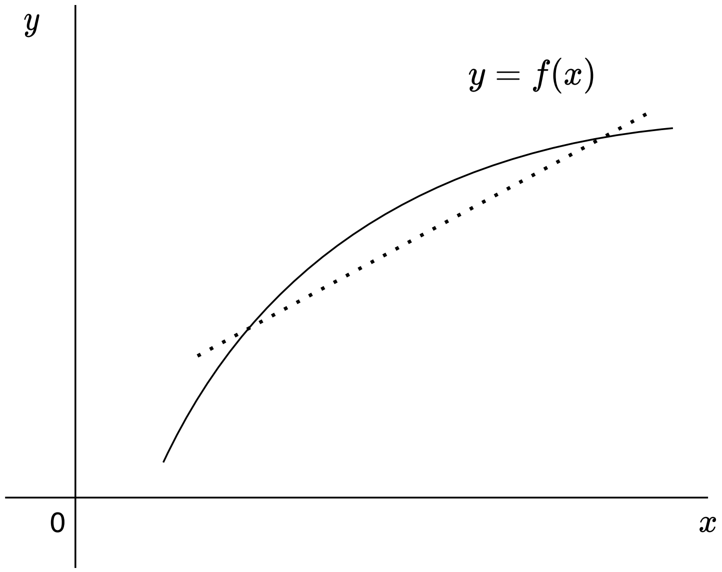

NON-EXISTENCE

- We can show that no Nash equilibrium exists in this game
  - With **strict convex** utility functions
- Saving grace
  - Techniques we developped are generally useful
  - Can use it to prove counterexamples for additional possible properties
  - Can use it for an efficient algoritmh (future work)

	A	B
A	(2, 0); (1, 0)	(1, 0); (0, 2)
B	(0, 1); (2, 0)	(0, 2); (0, 1)

$$u_1(p_1, p_2) = u_2(p_1, p_2) = p_1^2 + p_2^2$$

RELATIONS BETWEEN OPTIMISATION CRITERIA

MIXED STRATEGY EQUILIBRIA

- ***No relation*** between both optimisation criteria ***in general***

	A	B
A	(1, 0); (1, 0)	(0, 1); (0, 1)
B	(0, 1); (0, 1)	(-10, 0); (-10, 0)

Multi-objective reward vectors

	A	B
A	0.1; 0.1	0; 0
B	0; 0	-0.1; -0.1

Scalarised utility for both agents

No sharing of number of equilibria or equilibria themselves

RELATIONS BETWEEN OPTIMISATION CRITERIA

PURE STRATEGY EQUILIBRIA

- ***Relation*** when only considering ***pure strategy*** equilibria
  - Pure strategy equilibrium under SER is also one under ESR
  - Bidirectional when assuming (quasi)convex utility functions
- We can extend this to ***blended*** settings
  - Pure strategy equilibrium under SER is also one in any blended setting
  - Bidirectional when assuming (quasi)convex utility functions

ALGORITHMIC IMPLICATIONS

- Algorithm for calculating ***all pure strategy equilibria*** in a given MONFG ***with quasiconvex utility functions***
- Shown to work because of our theoretical contributions

NASH EQUILIBRIA

RECENT WORK

Algorithm 1 Computing all PSNE in an MONFG

Input: an MONFG $G = (N, \mathcal{A}, \boldsymbol{p})$ and quasiconvex utility functions $\boldsymbol{u} = (u_1, \cdots, u_n)$

function REDUCE_MONFG(monfg, u)
    $N, \mathcal{A}, \boldsymbol{p} \leftarrow$ monfg
    $u_1, \cdots, u_n \leftarrow$ u
    $f \leftarrow (u_1 \circ \boldsymbol{p}_1, \cdots, u_n \circ \boldsymbol{p}_n)$
    $G' \leftarrow (N, \mathcal{A}, f)$ ▷ An induced normal-form game
    return G'
end function
function COMPUTE_ALL_PSNE(nfg)
    $S = \emptyset$
    for PS **in** nfg **do** ▷ Loop over all pure strategies
        if PS is a PSNE **then** ▷ If it is a PSNE add it to the set
            $S \leftarrow S \cup \{\text{PS}\}$
        end if
    end for
    return S
end function
nfg $\leftarrow$ REDUCE_MONFG(G, u)
PSNE $\leftarrow$ COMPUTE_ALL_PSNE(nfg)

NASH EQUILIBRIA

RECENT WORK

Algorithm 1 Computing all PSNE in an MONFG

Input: an MONFG $G = (N, \mathcal{A}, \boldsymbol{p})$ and quasiconvex utility functions $u = (u_1, \cdots, u_n)$

function REDUCE_MONFG(monfg, u)
    $N, \mathcal{A}, \boldsymbol{p} \leftarrow$ monfg
    $u_1, \cdots, u_n \leftarrow$ u
    $f \leftarrow (u_1 \circ \boldsymbol{p}_1, \cdots, u_n \circ \boldsymbol{p}_n)$
    $G' \leftarrow (N, \mathcal{A}, f)$ ▷ An induced normal-form game
    return G'
end function
function COMPUTE_ALL_PSNE(nfg)
    $S = \emptyset$
    for PS **in** nfg **do** ▷ Loop over all pure strategies
        if PS is a PSNE **then** ▷ If it is a PSNE add it to the set
            $S \leftarrow S \cup \{\text{PS}\}$
        end if
    end for
    return S
end function
nfg $\leftarrow$ REDUCE_MONFG(G, u)
PSNE $\leftarrow$ COMPUTE_ALL_PSNE(nfg)

Reduce the MONFG

NASH EQUILIBRIA

RECENT WORK

Algorithm 1 Computing all PSNE in an MONFG

Input: an MONFG $G = (N, \mathcal{A}, \boldsymbol{p})$ and quasiconvex utility functions $u = (u_1, \cdots, u_n)$

function REDUCE_MONFG(monfg, u)
    $N, \mathcal{A}, \boldsymbol{p} \leftarrow$ monfg
    $u_1, \cdots, u_n \leftarrow$ u
    $f \leftarrow (u_1 \circ \boldsymbol{p}_1, \cdots, u_n \circ \boldsymbol{p}_n)$
    $G' \leftarrow (N, \mathcal{A}, f)$ ▷ An induced normal-form game
    return G'
end function
function COMPUTE_ALL_PSNE(nfg)
    $S = \emptyset$
    for PS **in** nfg **do** ▷ Loop over all pure strategies
        if PS is a PSNE **then** ▷ If it is a PSNE add it to the set
            $S \leftarrow S \cup \{\text{PS}\}$
        end if
    end for
    return S
end function
nfg $\leftarrow$ REDUCE_MONFG(G, u)
PSNE $\leftarrow$ COMPUTE_ALL_PSNE(nfg)

Solve the trade-off game

CONCLUSION

- Lots of new theoretical insights
  - Nash equilibrium guarantees
  - Relation between optimisation criteria when only considering pure strategies
  - We can extend this to blended settings

CONCLUSION

- Lots of new theoretical insights
  - Nash equilibrium guarantees
  - Relation between optimisation criteria when only considering pure strategies
  - We can extend this to blended settings
- Incorporate everything into a novel algorithm

CONCLUSION

- Lots of new theoretical insights
  - Nash equilibrium guarantees
  - Relation between optimisation criteria when only considering pure strategies
  - We can extend this to blended settings
- Incorporate everything into a novel algorithm
- Additional guarantees for MONFGs
  - Zero-sum games
  - Exploit continuous game reduction

CONCLUSION

- Lots of new theoretical insights
  - Nash equilibrium guarantees
  - Relation between optimisation criteria when only considering pure strategies
  - We can extend this to blended settings
- Incorporate everything into a novel algorithm
- Additional guarantees for MONFGs
  - Zero-sum games
  - Exploit continuous game reduction
- More algorithmic work
  - Use theorems to find Nash equilibria efficiently

OTHER WORK

- Explored communication
  - Communication protocols
  - Commit to actions or policies
  - Evaluate in different settings

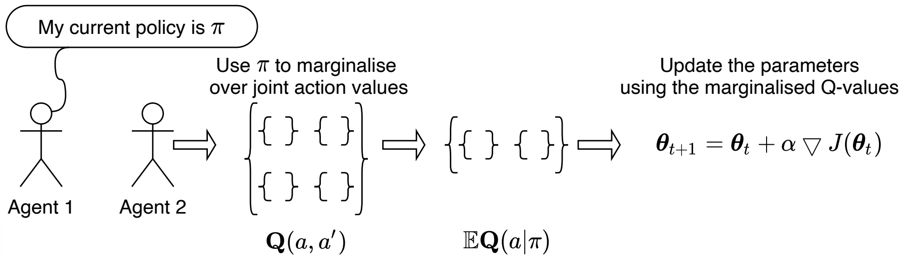

- Explored commitment
  - Theoretically
  - Evaluate using reinforcement learning