Learning to Cooperate via Conditional Commitments

Jobst Heitzig, PIK RD4

RD4 seminar 21 April 2020

Overview

- Motivation, Inspiration, Rationale, Example

- Theoretical background
- Formal results

- Learning & Agent-based modeling
- Simulation results

Overview

- **Motivation, Inspiration, Rationale, Example**

- Theoretical background
- Formal results

- Learning & Agent-based modeling
- Simulation results

Motivation: International Climate Mitigation

- GHG reductions are a positive externality → free-riding → need for cooperation
- **How to establish cooperation?**
  - negotiate a **"grand" treaty** (UNFCCC/COP, Kyoto, Paris)
    - slow, not yet very successful, may lead to only unambitious treaties
    - but concept of INDCs contains **idea of conditional commitments**

Motivation: International Climate Mitigation

- GHG reductions are a positive externality → free-riding → need for cooperation
- **How to establish cooperation?**
  - negotiate a **"grand" treaty** (UNFCCC/COP, Kyoto, Paris)
    - slow, not yet very successful, may lead to only unambitious treaties
    - but concept of INDCs contains **idea of conditional commitments**
  - dynamically form **small then larger coalitions** "bottom-up"
    - ongoing process, not yet very successful, but may succeed eventually (Auer et al. Sci.Rep. 2015; Heitzig & Kornek, NCC 2018)
    - leads to a hierarchy of bi- or multilateral treaties

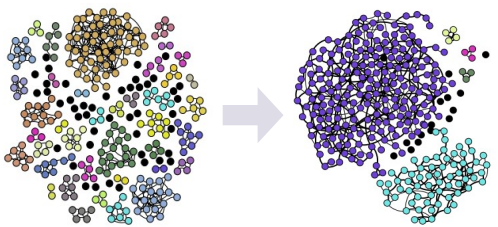

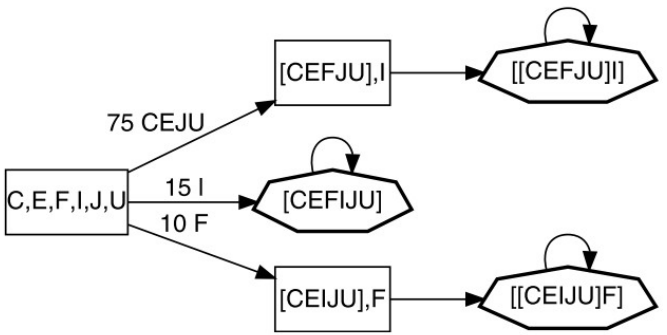

Motivation: International Climate Mitigation

- GHG reductions are a positive externality → free-riding → need for cooperation
- **How to establish cooperation?**
  - negotiate a **"grand" treaty** (UNFCCC/COP, Kyoto, Paris)
    - slow, not yet very successful, may lead to only unambitious treaties
    - but concept of INDCs contains **idea of conditional commitments**
  - dynamically form **small then larger coalitions** "bottom-up"
    - ongoing process, not yet very successful, but may succeed eventually (Auer et al. Sci.Rep. 2015; Heitzig & Kornek, NCC 2018)
    - leads to a hierarchy of bi- or multilateral treaties
  - unilateral approaches without formal treaties
    - e.g. some countries pioneer unconditionally & hope for others to follow
    - or: **use unilateral but binding, mutually conditional commitments**

Inspiration: The NPVIC

Scheme: Agents unilaterally (!) but bindingly commit to behave in certain way if others behave in certain ways.

Here: US federal states pass federal state laws

Internationally: Countries pass domestic laws?

National Popular Vote Interstate Compact

From Wikipedia, the free encyclopedia

The **National Popular Vote Interstate Compact** (**NPVIC**) is an agreement among a group of U.S. states and the District of Columbia to award all their electoral votes to whichever presidential candidate wins the overall popular vote in the 50 states and the District of Columbia. The compact is designed to ensure that the candidate who receives the most votes nationwide is elected president, and it would come into effect only when it would guarantee that outcome.[3][4] As of March 2020, it has been adopted by fifteen states and the District of Columbia, although it is suspended in Colorado.[2] Together, they have 196 electoral votes counting Colorado, which is 36% of the Electoral College and 73% of the 270 votes needed to give the compact legal force. Certain legal questions, however, may affect implementation of the compact.

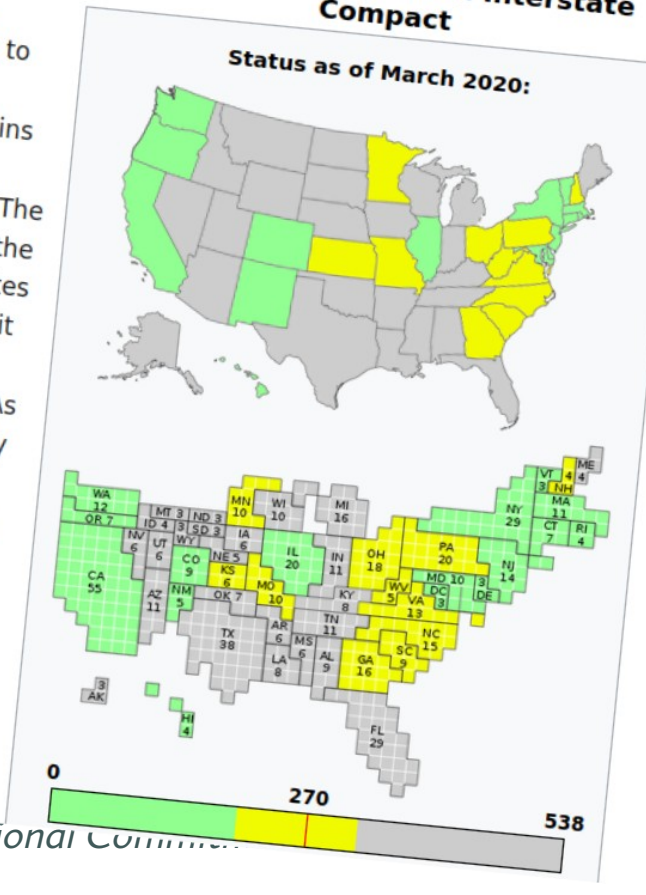

Rationale

- Without prior international negotations, a country could pass a domestic law that requires it to take specific climate protection measures as soon as (and as long as) certain other countries have passed similar laws that specify at least a certain amount of certain measures.
  - e.g.: I'll reduce emissions by 20% if you invest 1% of GDP into the Green Climate Fund

Rationale

- Without prior international negotations, a country could pass a domestic law
  that requires it to take specific climate protection measures
  as soon as (and as long as) certain other countries have passed similar laws
  that specify at least a certain amount of certain measures.
  - e.g.: I'll reduce emissions by 20% if you invest 1% of GDP into the Green Climate Fund

- If the ambition is low enough initially,
  this gives the other countries incentives to indeed pass similar laws.
- At each point in time the set of *laws currently in force*
  imply a set of *current obligations* for all participating countries.

Rationale

- Without prior international negotations, a country could pass a domestic law that requires it to take specific climate protection measures as soon as (and as long as) certain other countries have passed similar laws that specify at least a certain amount of certain measures.
  - e.g.: I'll reduce emissions by 20% if you invest 1% of GDP into the Green Climate Fund

- If the ambition is low enough initially, this gives the other countries incentives to indeed pass similar laws.
- At each point in time the set of *laws currently in force* imply a set of *current obligations* for all participating countries.
- These laws can be adjusted more easily than international treaties to react to circumstances and to increase ambition.
- **Hypothesis: over time, an "efficient" level of mitigation will arise!**

Example

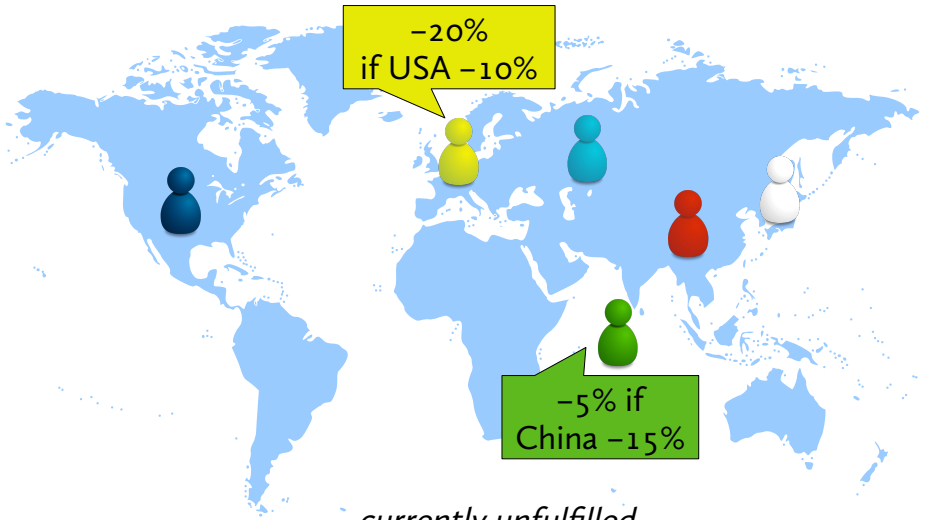

currently unfulfilled conditions

Example

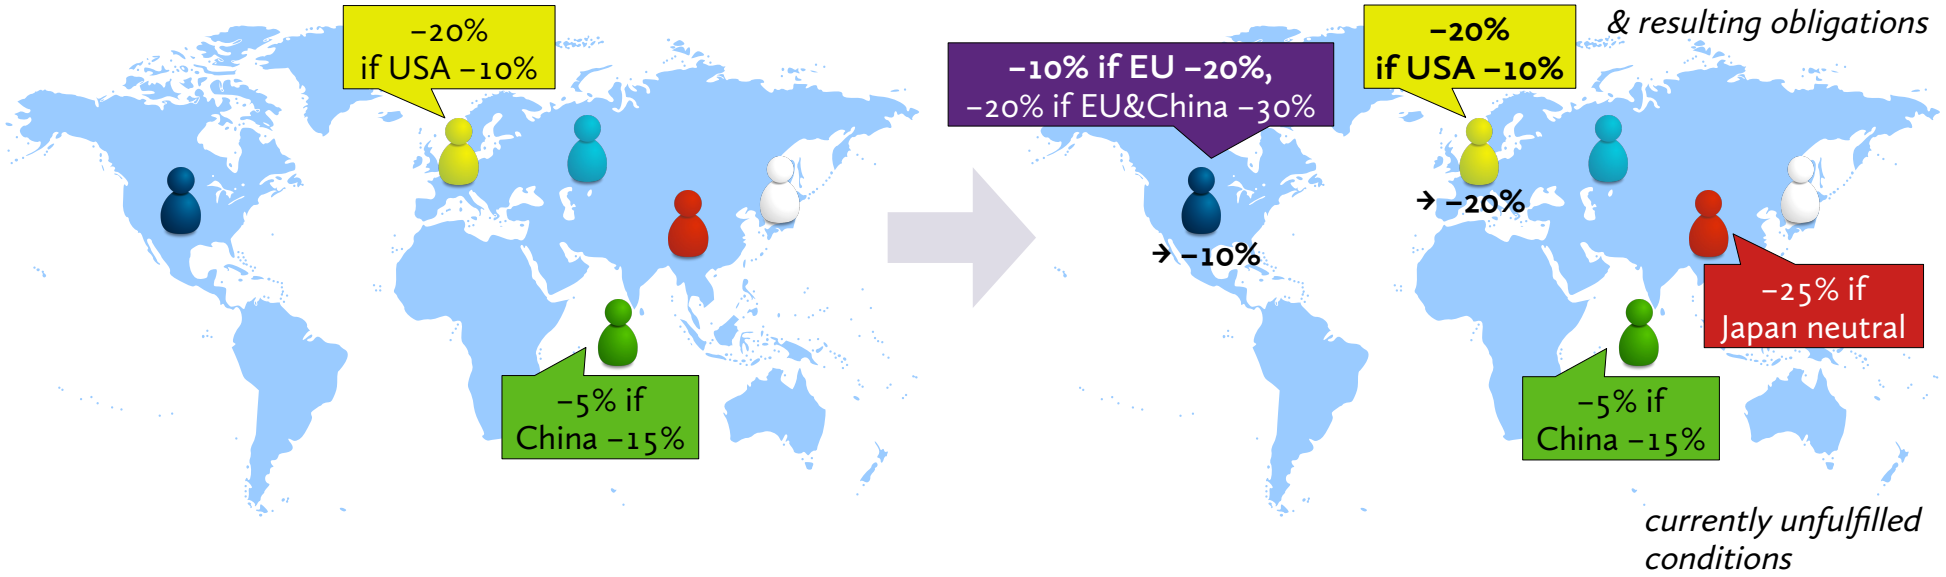
−20%
if USA −10%
−5% if
China −15%
−10% if EU −20%,
−20% if EU&China −30%
→ −10%
−20%
if USA −10%
→ −20%
−25% if
Japan neutral
−5% if
China −15%
bold: currently fulfilled conditions
& resulting obligations
currently unfulfilled conditions

Example

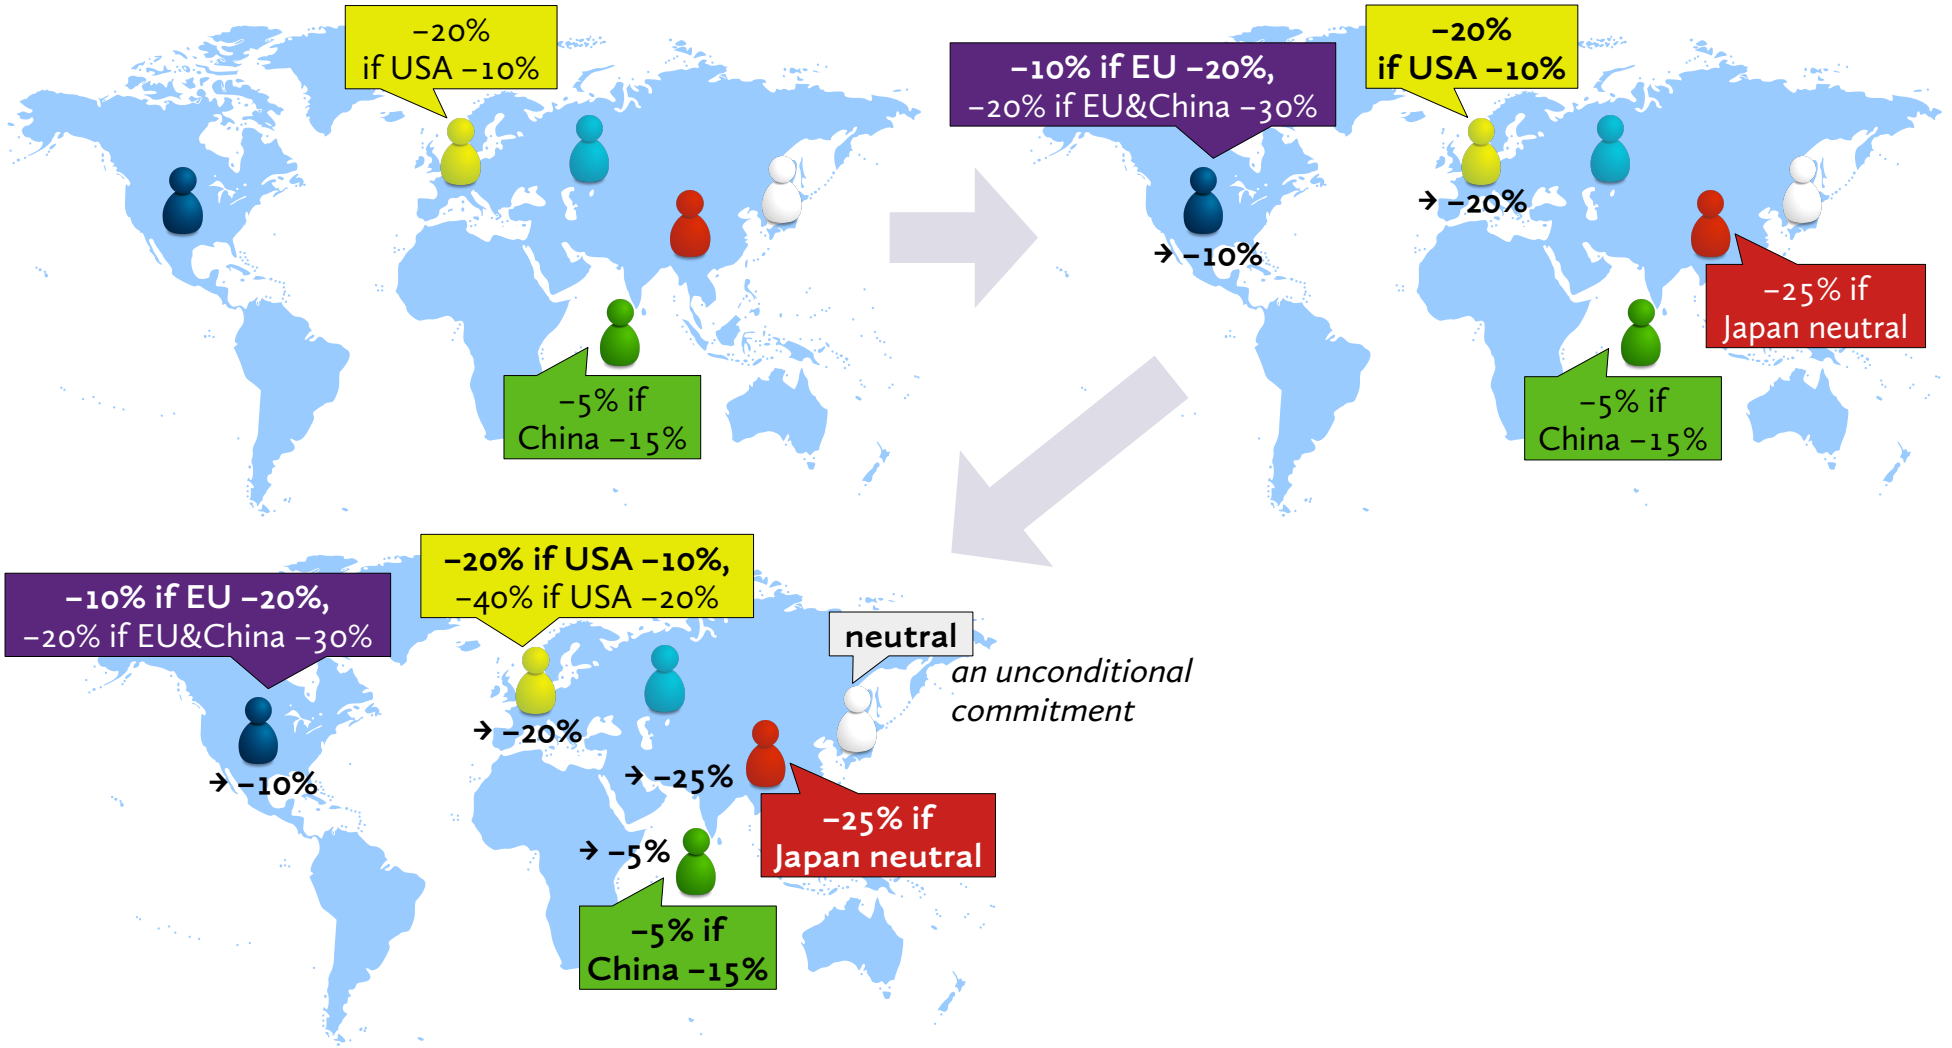
−20%
if USA −10%
−5% if
China −15%
−10% if EU −20%,
−20% if EU&China −30%
→ −10%
−20%
if USA −10%
→ −20%
−25% if
Japan neutral
−5% if
China −15%
−10% if EU −20%,
−20% if EU&China −30%
→ −10%
−20% if USA −10%,
−40% if USA −20%
→ −20%
→ −25%
→ −5%
−25% if
Japan neutral
−5% if
China −15%
neutral
an unconditional
commitment

Example

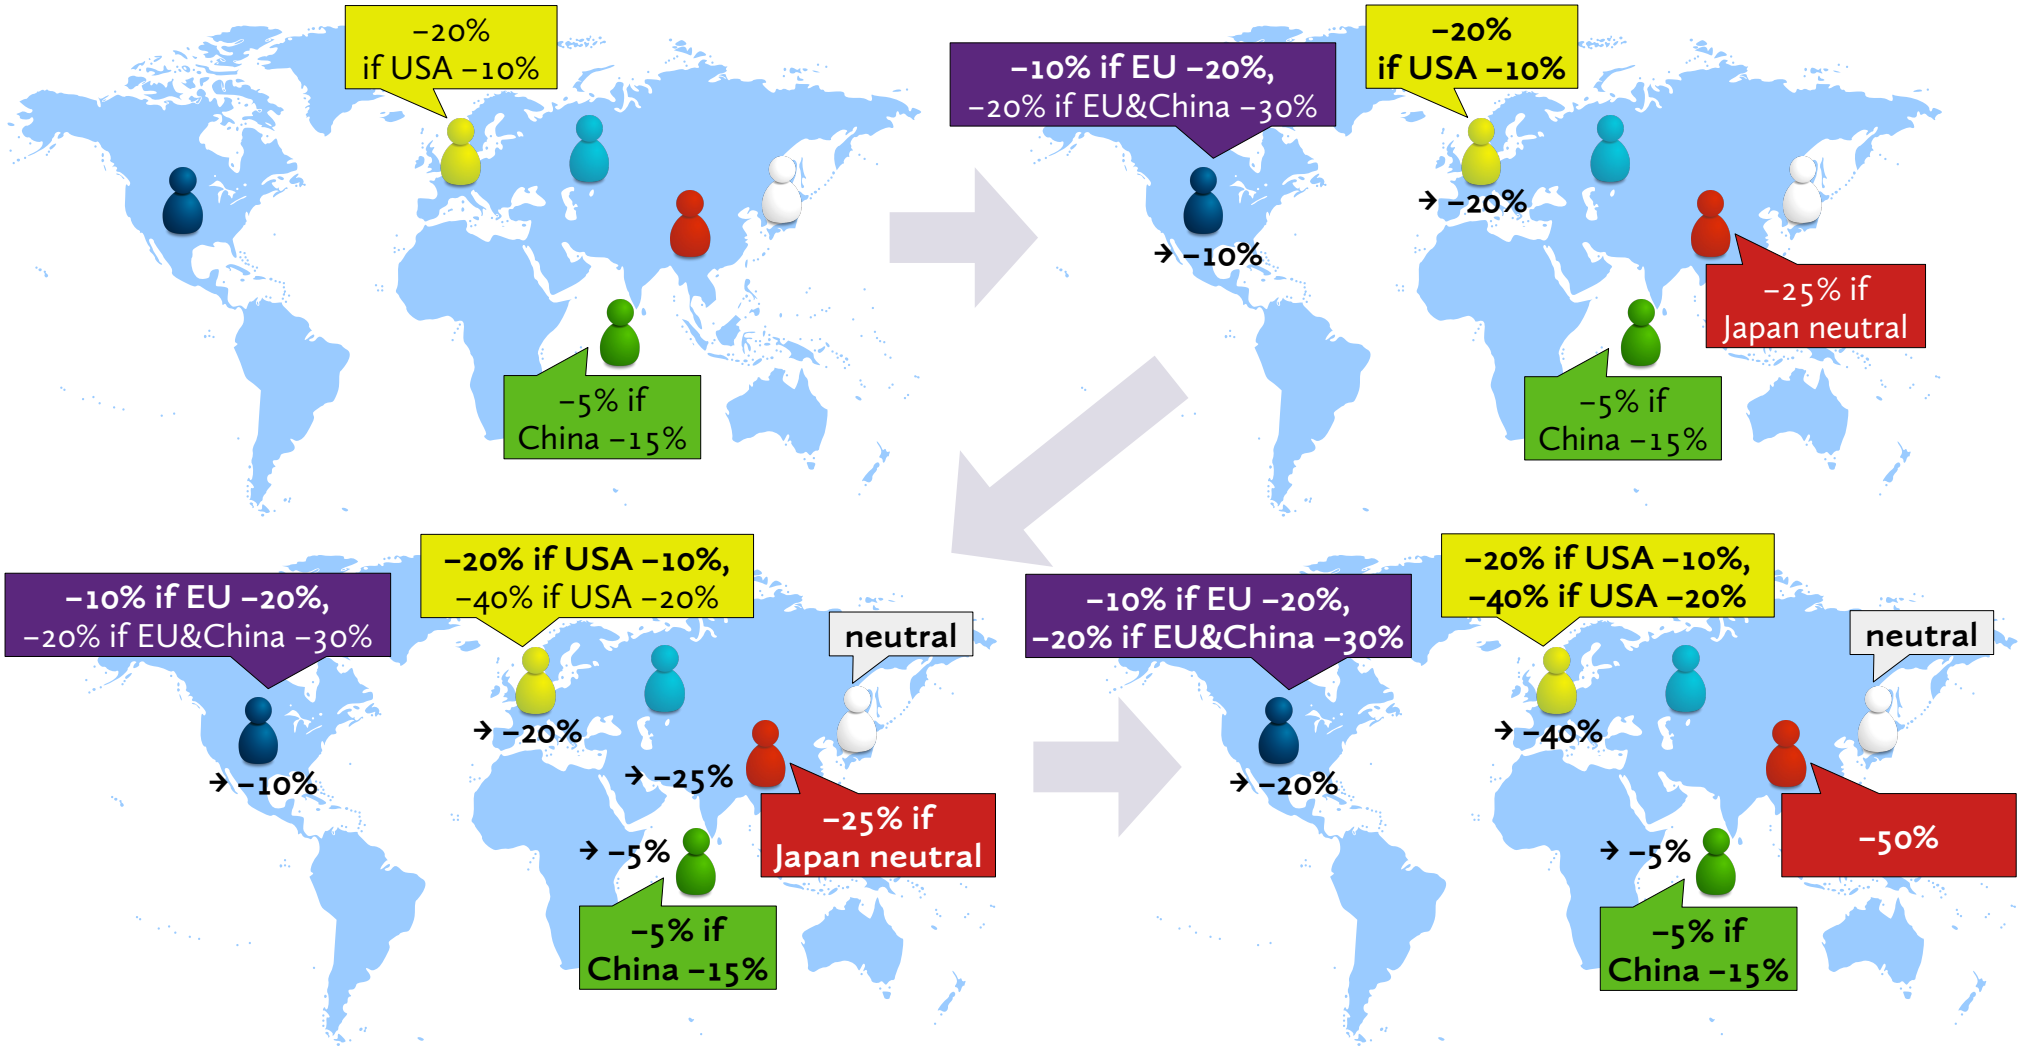
−20% if USA −10%
−5% if China −15%
−10% if EU −20%, −20% if EU&China −30%
→ −10%
−20% if USA −10%
→ −20%
−25% if Japan neutral
−5% if China −15%
−10% if EU −20%, −20% if EU&China −30%
→ −10%
−20% if USA −10%, −40% if USA −20%
neutral
→ −20%
→ −25%
→ −5%
−25% if Japan neutral
−5% if China −15%
−10% if EU −20%, −20% if EU&China −30%
→ −20%
−20% if USA −10%, −40% if USA −20%
neutral
→ −40%
→ −5%
−50%
−5% if China −15%

Overview

- Motivation, Inspiration, Rationale, Example

- **Theoretical background**
- **Formal results**

- Learning & Agent-based modeling
- Simulation results

Theoretical background

- *Cooperative* Game Theory, Efficiency, Bargaining Solutions
  → **the "core"** of a cooperative game

- *Non-cooperative* Game Theory & Forms of Strategic Equilibrium
  → **"strong" equilibria** of a non-cooperative game

- The Nash Program & Mechanism Design

Cooperative Game Theory

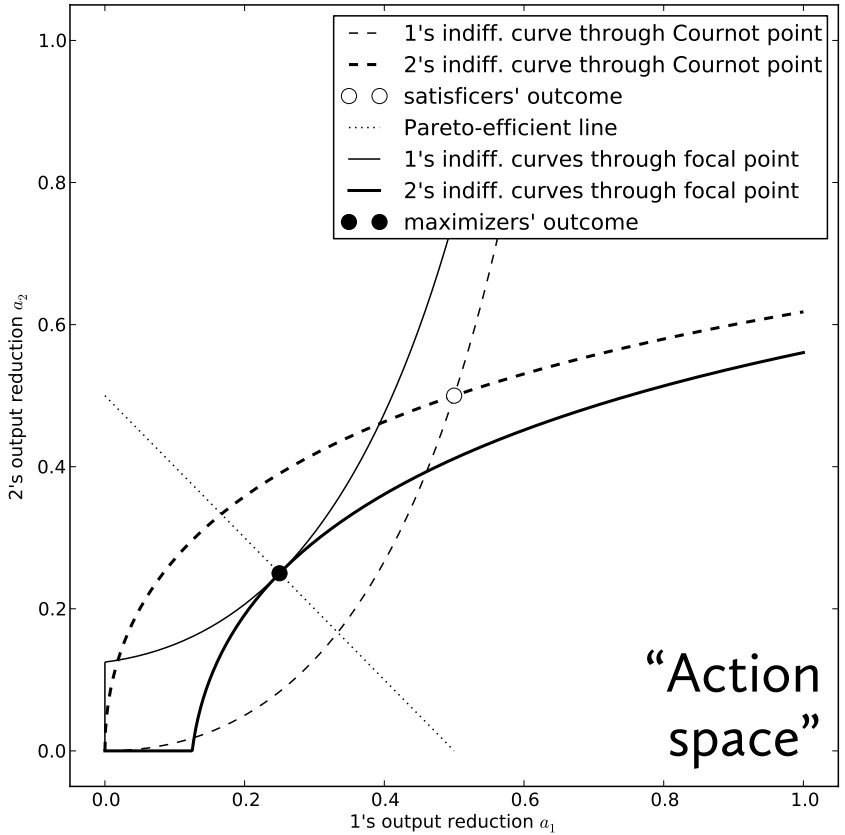

Example: Cournot duopoly
(e.g., two non-OPEC countries reducing output)

Cooperative Game Theory

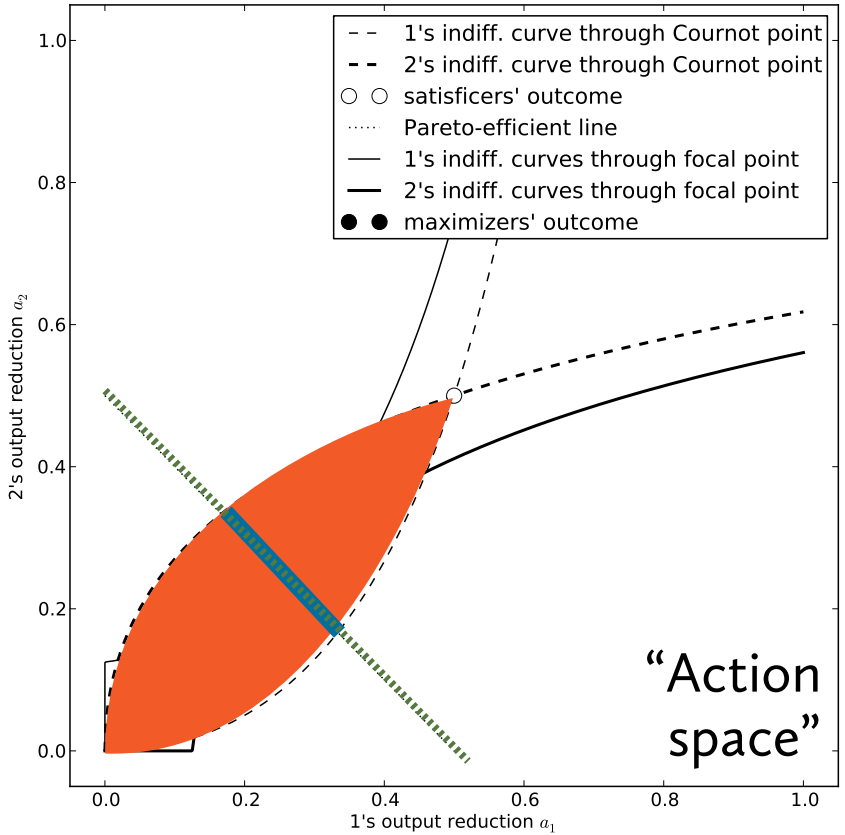

Example: Cournot duopoly
(e.g., two non-OPEC countries reducing output)

A combination of actions is ...

- (Pareto-)**efficient**: no other combination gives all players more payoff
- in the "**bargaining set**":
  all players get at least what they would get at the disagreement point (here: (0,0))
- in the "**core**" of the game: no group can get more by changing their actions, assuming all others will then react by doing nothing

Non-cooperative Game Theory

- No binding agreements are possible
- Players may use complex strategies rather than just plain actions
- A combination of strategies is a …
  - **Nash equilibrium**:
    no individual player has an incentive to deviate unilaterally
    → many games have too many Nash equilibria → concept too weak
  - **strong equilibrium**:
    no group of players has an incentive to deviate together
    → takes possibility to communicate into account, but may not exist…

The Nash Program & Mechanism Design

Nash (1953):

Reduce certain cooperative solutions *(e.g. the "Nash bargaining solution")*
to certain non-cooperative equilibria *(e.g. "Markov-perfect equilibrium)*
of suitable non-cooperative versions of a cooperative game!
(e.g. the Rubinstein bargaining protocol in case of bargaining)

Mechanism Design:

Construct a non-cooperative game form *(e.g. a type of auction)*
so that in certain types of equilibrium, behaviour will meet a given goal!
(e.g. revelation of preferences or maximal revenue for the auctioneer)

Our achievement will be:

Based on the idea of conditional commitments,
construct a game form for positive externality problems
so that **all strong equilibria correspond to core outcomes**
and hence **agents behave in a "jointly optimal" way**.

Nash (1953):
Reduce certain coop. solutions to certain non-cooperative equilibria
of suitable non-cooperative versions of a cooperative game!

Mechanism Design:

Construct a non-cooperative game form
so that in certain types of equilibrium, behaviour will meet a given goal!

A GAME FORM BASED ON CONDITIONAL COMMITMENTS

Strategy spaces: Each player i chooses a
conditional commitment function (CCF) c_i

= a map from others' action combinations a_{-i} to max. own actions $c_i(a_{-i})$
(interpretation: "if they do at least a_{-i}, I do at least a_i")

Outcome: each player is commited to perform the action a_i given by the largest action profile a that meets all conditions $a_i \leq c_i(a_{-i})$.

Such a unique largest feasible action combination exists
if all action spaces are supremum-complete partially ordered sets.

Example: Output Reduction (Cournot Duopoly)

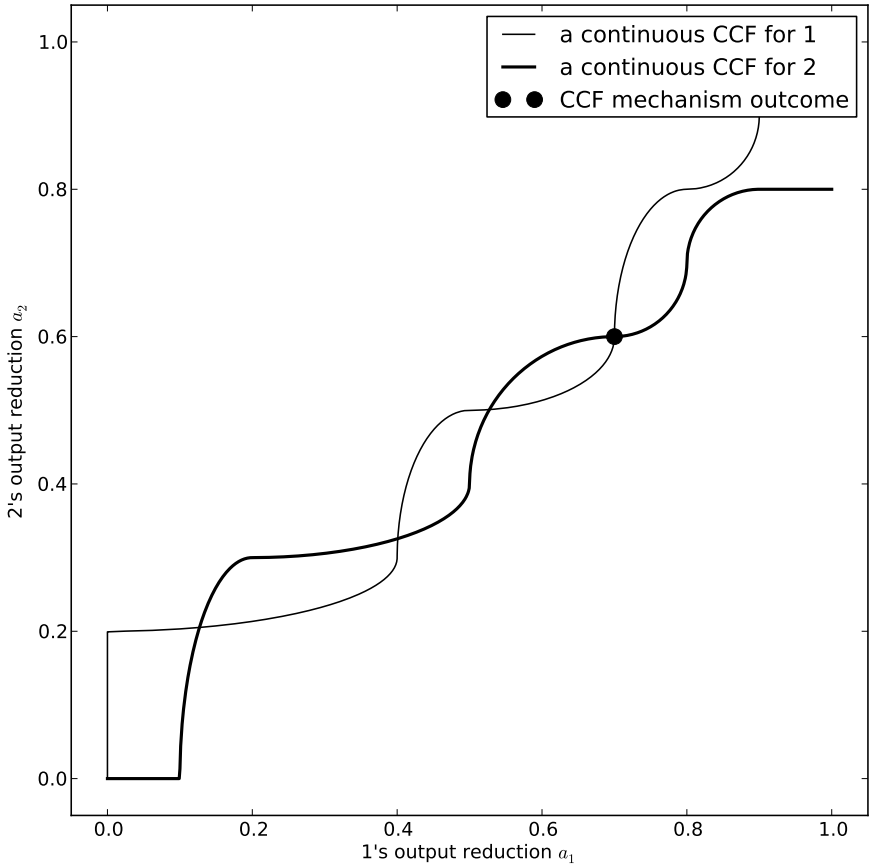

not a Nash equilibrium

Example: Output Reduction (Cournot Duopoly)

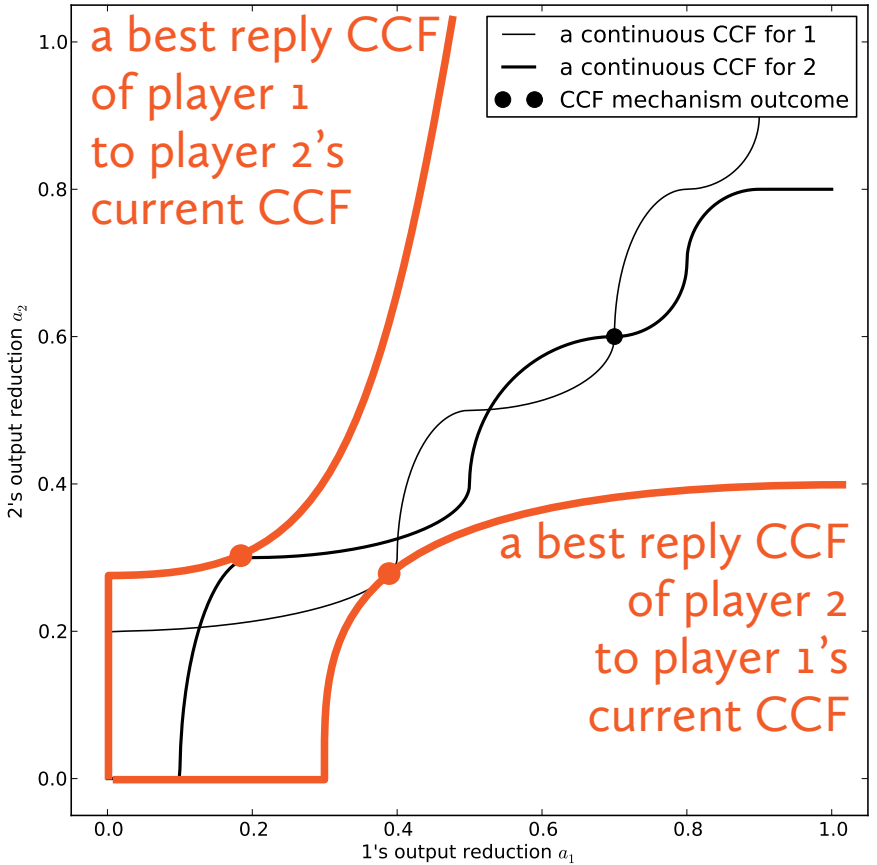

not a Nash equilibrium

each player has an incentive
to switch to a different CCF

Example: Output Reduction (Cournot Duopoly)

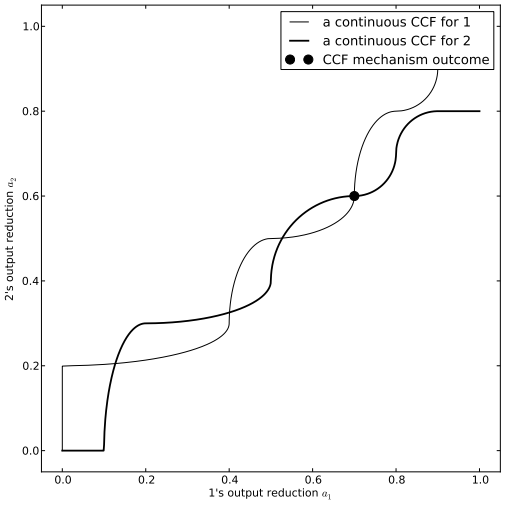

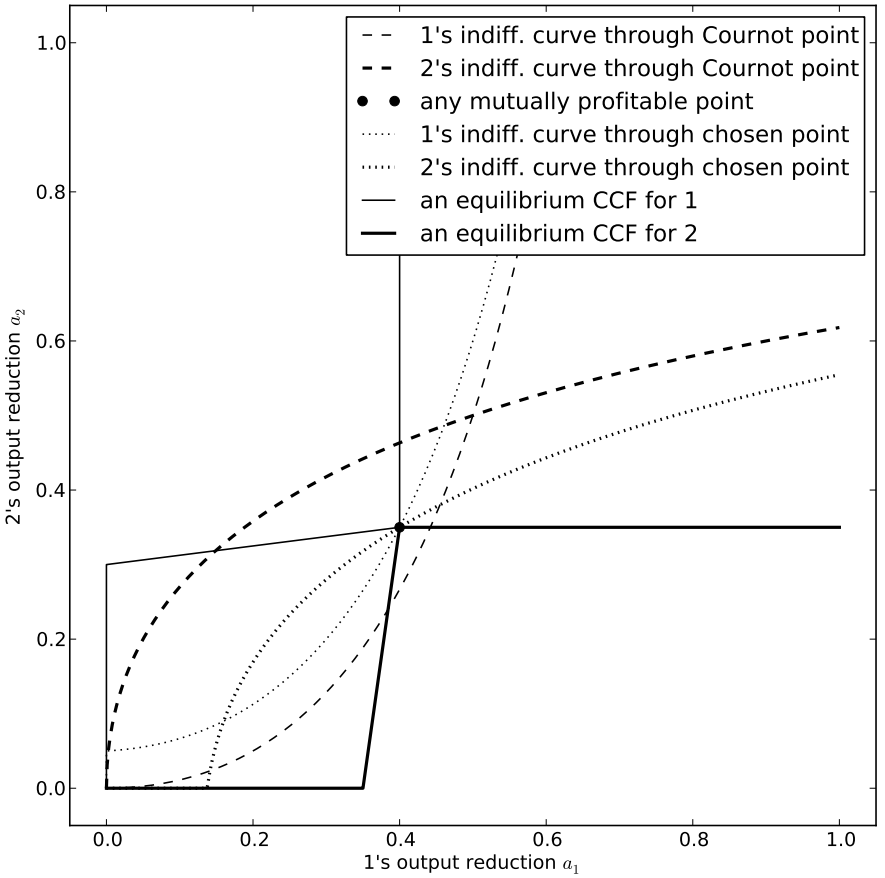

no Nash equilibrium

Nash equilibrium but *not strong*

Example: Output Reduction (Cournot Duopoly)

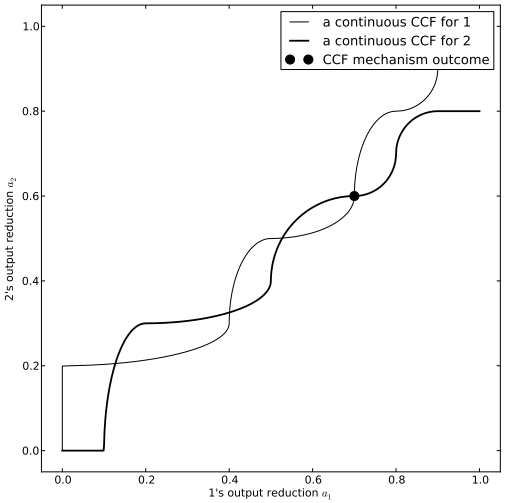

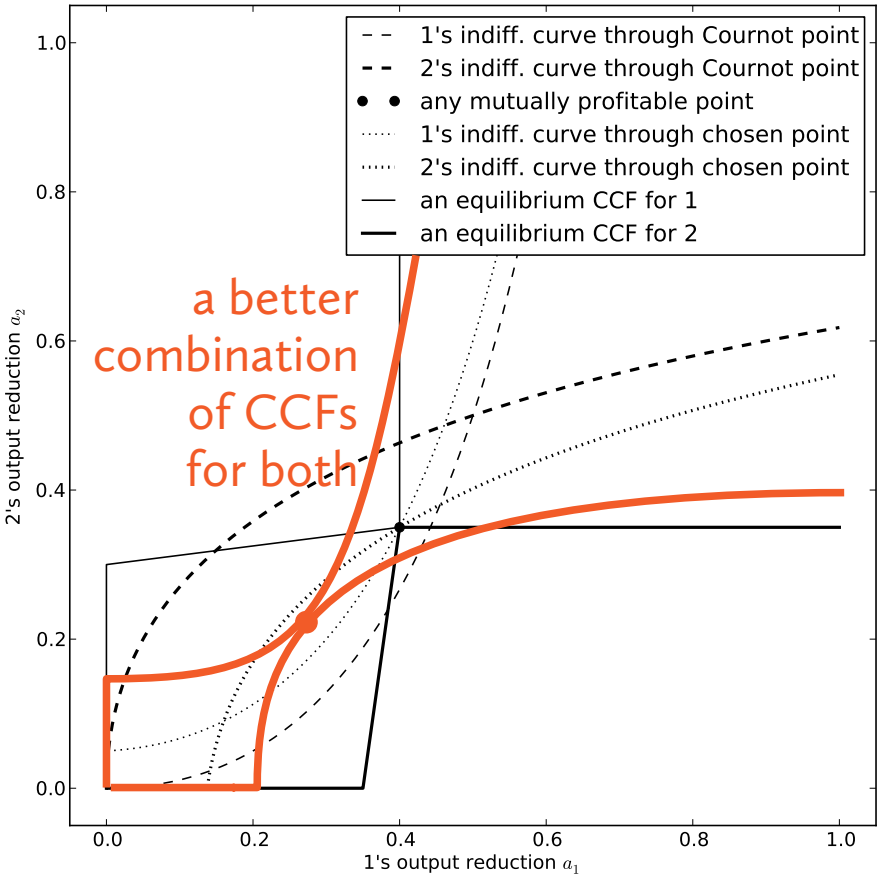

no Nash equilibrium

Nash equilibrium but *not strong*

the players have an incentive to switch *together* to different CCFs

Example: Output Reduction (Cournot Duopoly)

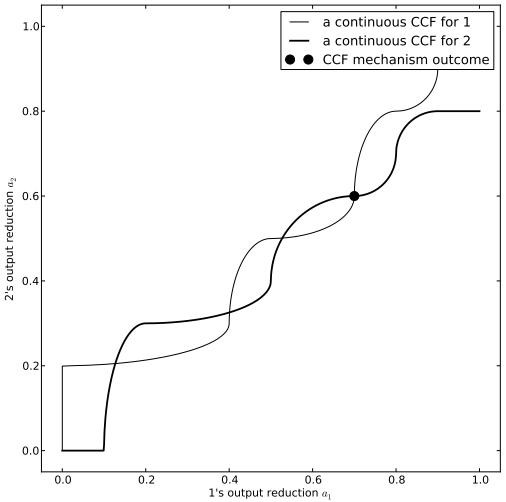

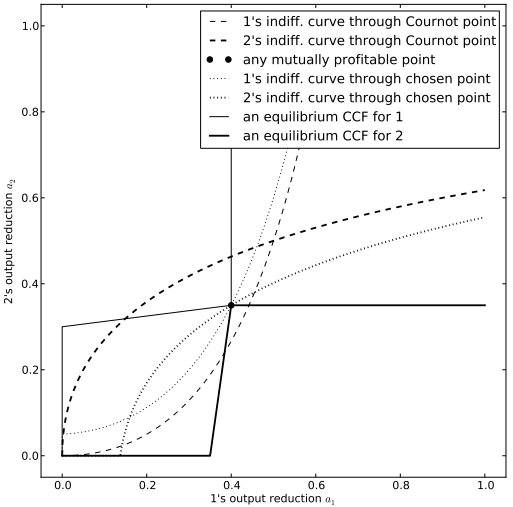

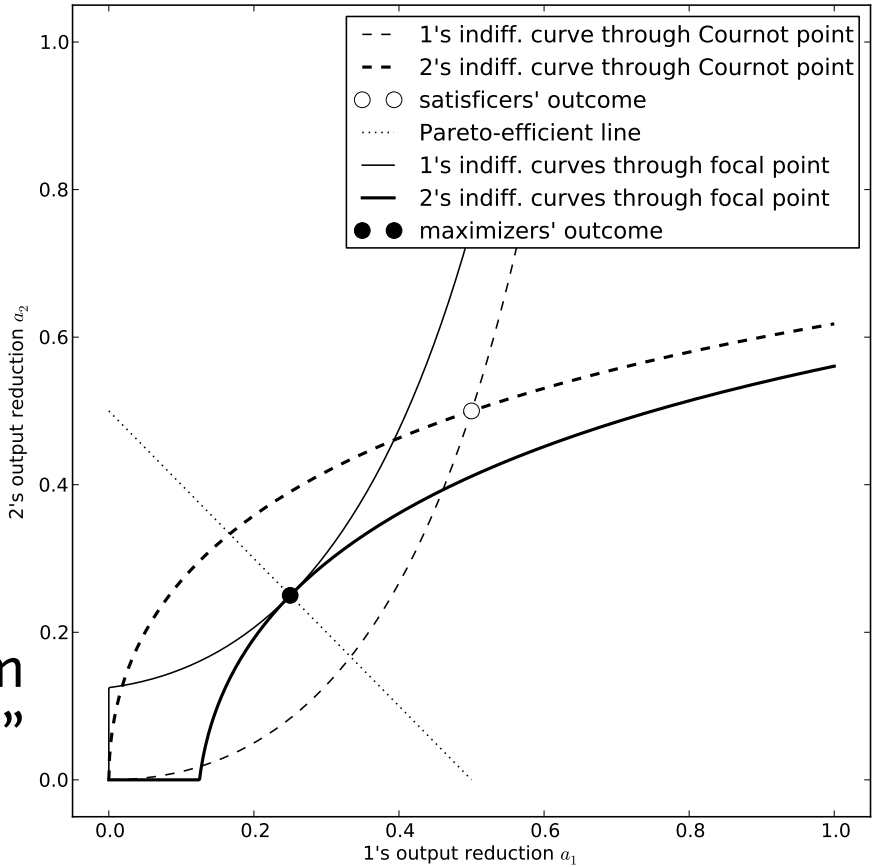

a strong equilibrium
that is even "focal"

More players → Cournot *Oligopoly*

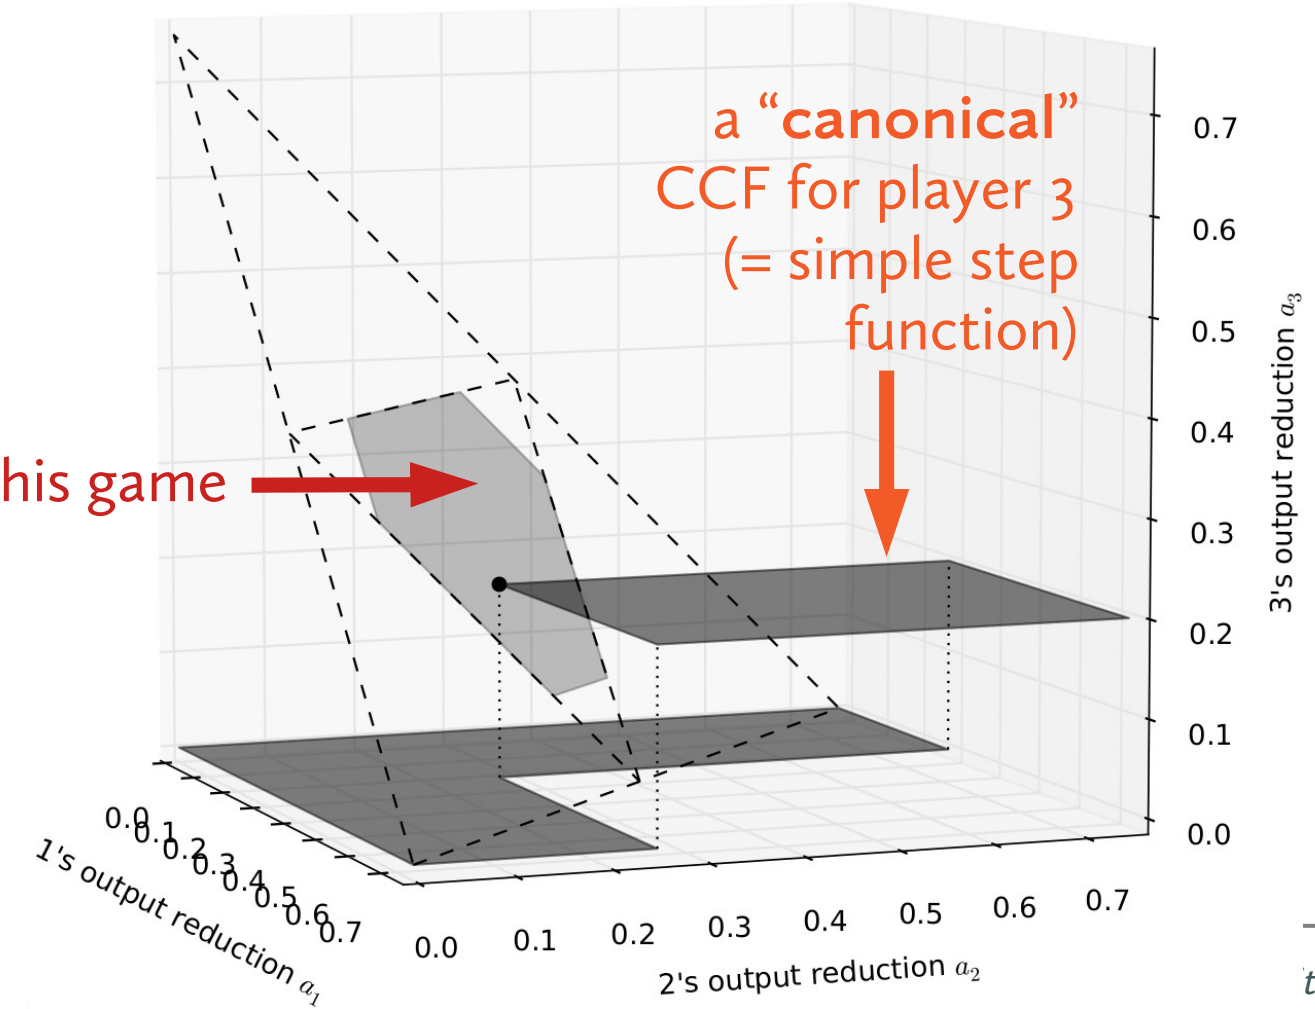

Theorem 1

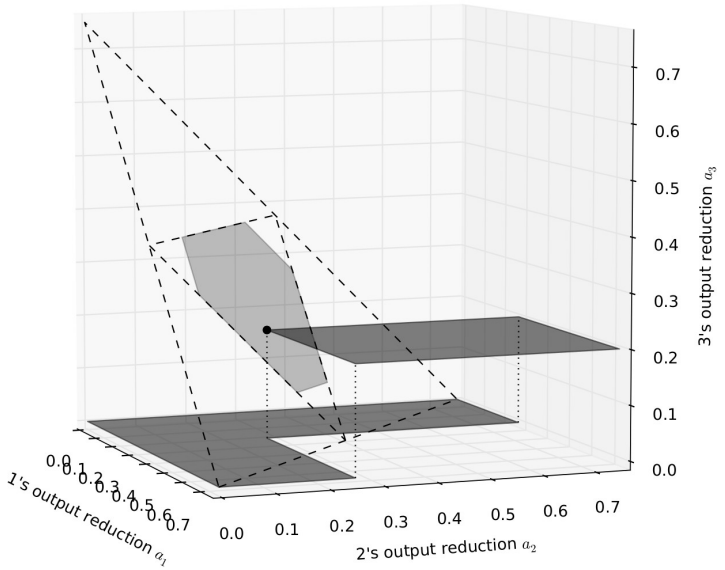

Assume the CCF mechanism is applied
to any Costly Positive Externality Problem
(CPEP, a certain fairly broad class of games)
and certain fairly weak conditions apply.

Then the outcomes
that result from strong equilibria
are exactly the "core" outcomes,

and to sustain them
it suffices to use "canonical" CCFs.

(Heitzig 2019,
seven times desk-rejected,
ssrn.com/abstract=3449004)

Other examples of CPEPs

- Public good provision, e.g. emission reduction
- Bilateral trade with a broker
- Political package deals like Helmut Schmidt's Bonn 1978 G7 deal
  (discrete but high-dimensional action spaces)
- Supply chains/networks with uncertain costs and capacities
  (multi-dimensional continuous action spaces)
- Commodity exchanges (e.g. electricity markets)
  (many players having only few information)
- …

(Heitzig 2019, ssrn.com/abstract=3449004)

Overview

- Motivation, Inspiration, Rationale, Example

- Theoretical background
- Formal results

- **Learning & Agent-based modeling**
- **Simulation results**

Some forms of learning

- **Individual** learning
  - based on own experiences
    (e.g. trial and error, regret matching, reinforcement learning)
  - based on other's behaviour (e.g. simple imitation)
  - based on other's experiences
    (e.g. "social" learning by observing others performance)

- **Collective** learning
  - cooperatively (e.g. by sharing experiences)
  - non-cooperatively (e.g. by **alternating trial and error**)

A simple agent-based model of collective learning using conditional commitment functions

All players start with a zero CCF

At random time points, a random player updates her CCF:

- she finds her "favourite" point x on the joint CCF of the other players
- she determines
  (1) her canonical CCF leading through x
  (2) her indifference curve through x
- she uses any curve lying between the two curves as her new CCF

(thereby she offers to do more if others do more without risking a loss)

Example from before

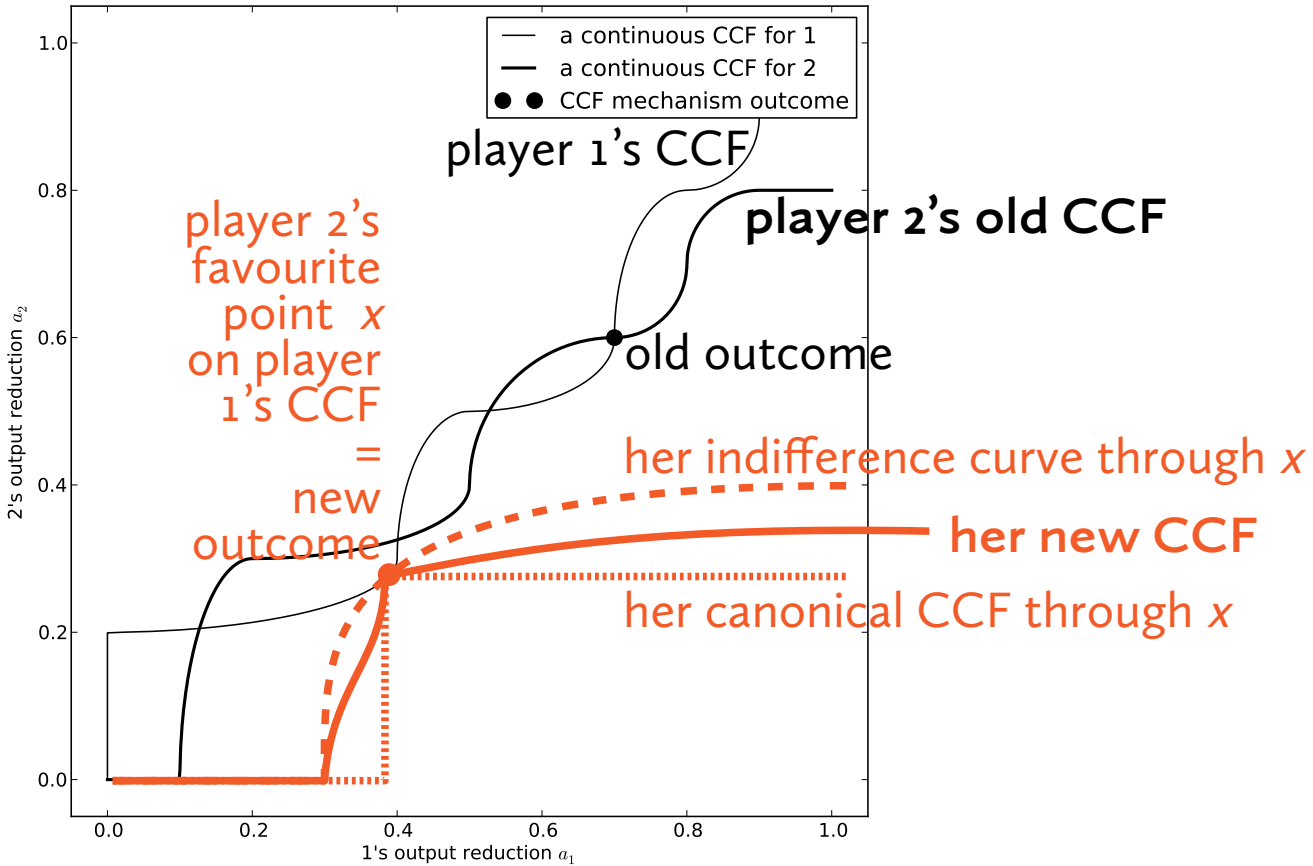

Simulation results for 3 countries' GHG emissions reductions

Almost all runs converge very fast to a "core" outcome

Very rarely the process seems to get stuck with some player mitigating nothing (might be a numerical error)

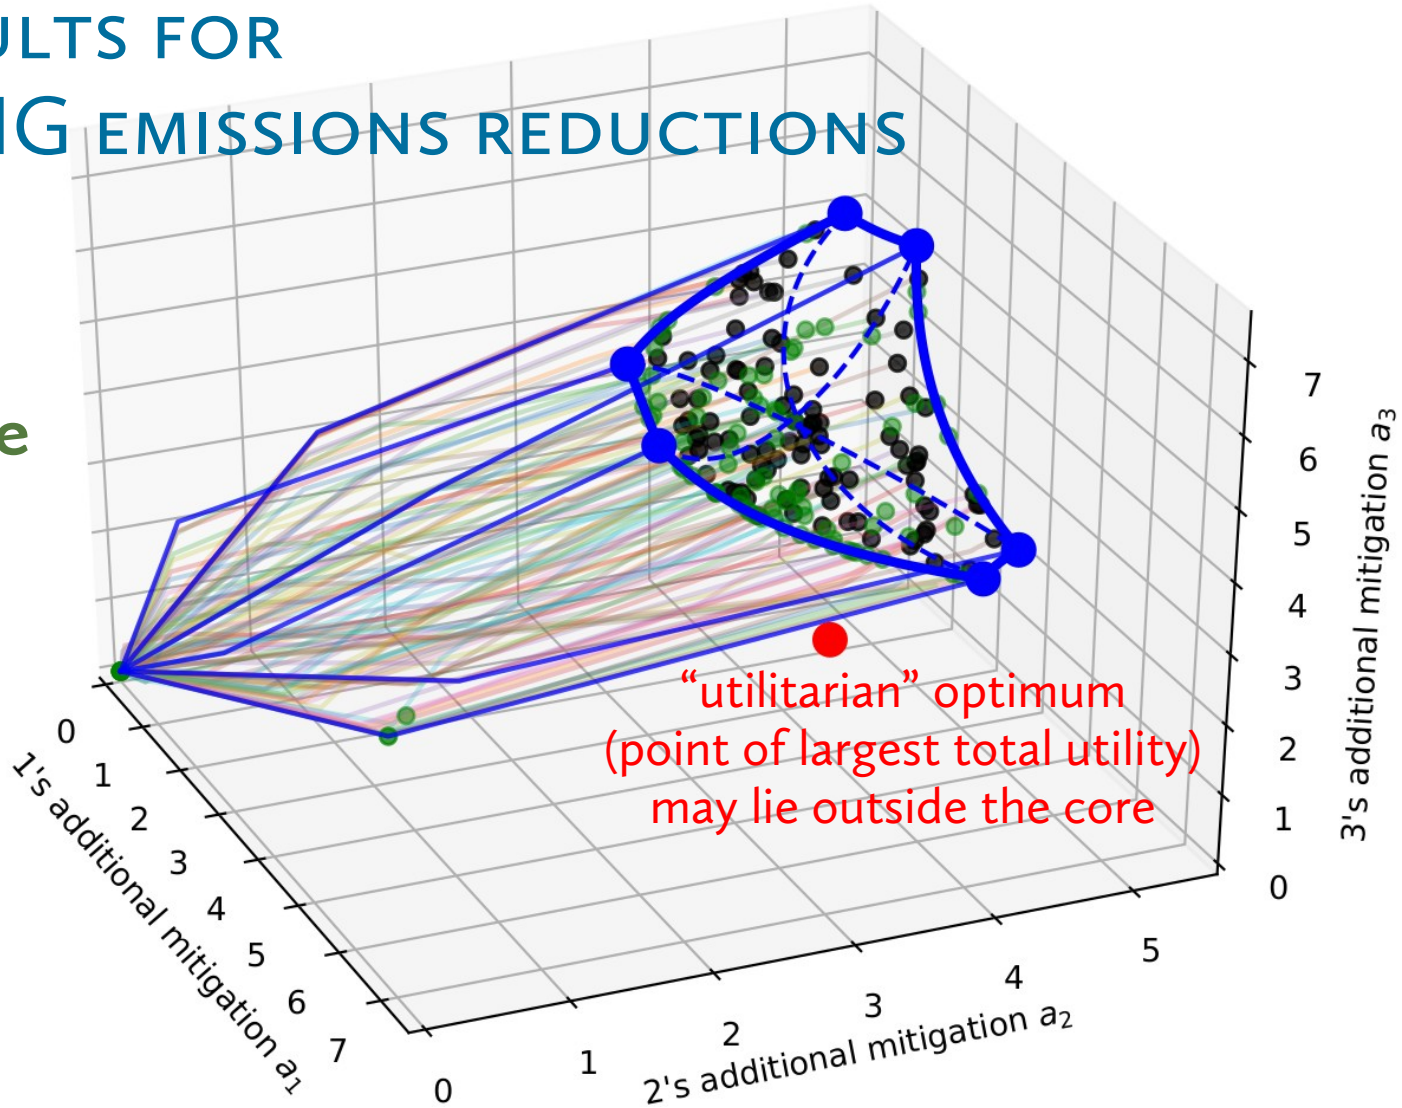

PIK

MCC

Theorem 2

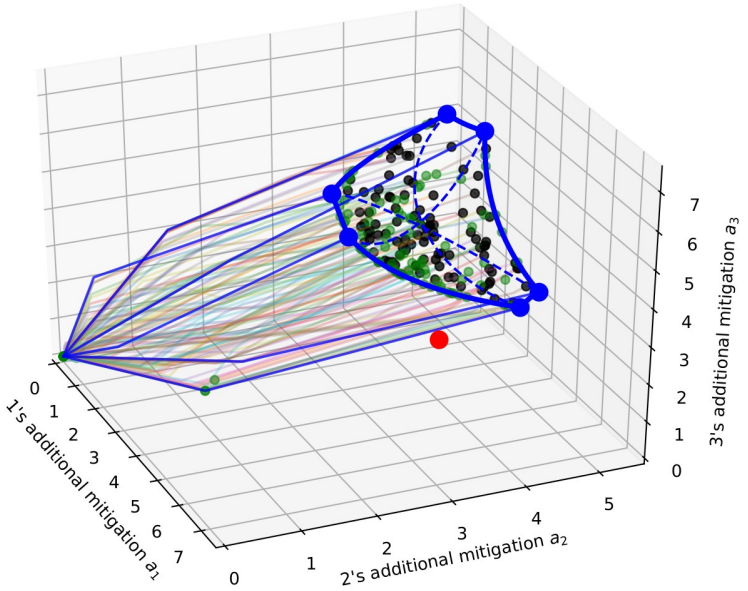

Assume all players use in some CPEP the above collective CCF-learning rule and some fairly weak conditions apply.

Then the outcomes almost surely converge to a "core" outcome.

Thank you! – Questions?

(Heitzig 2019,
seven times desk-rejected,
ssrn.com/abstract=3449004)