Insights into robust game theory and applications

Davide Radi

Dept. of Economics and Management, Università di Pisa

18 March, 2022, Milano, Italy.

Seminario di Dipartimento in dual mode

RISCHIO E INCERTEZZA

Overview

In game theory....

- Sometimes we assume a probabilistic description of the unknown values of the parameters of our model and agents maximize their expected payoffs (*Bayesian Games*);

- In other cases, it is not possible to have a unique probabilistic description of the unknown values of the parameters of our model;

- The values of the parameters are represented by a set of probability distributions and, following Gilboa and Schmeidler (1989), agents are maxmin expected utility maximizers (*incomplete information games with multiple priors*, see, e.g., Kajii and Ui (2005)).

- In the limit case where all possible probability distributions are possible, agents maximize their guaranteed payoffs, we obtain a *distribution-free model*: (*robust games, see Aghassi and Bertsimas (2006)*).

Robust-Game Framework:

Focusing on simultaneous move games with no private information, a **Robust Game** is defined as

$$G = \left\{ A_i, f_i, W_i^{\delta_i} : i \in \mathcal{N} \right\}$$

- $\mathcal{N} = \{1, \ldots, n\}$ is the finite set of players.
- Each player $i \in \mathcal{N}$ has a set of possible actions $A_i \subseteq \mathbb{R}_+$.
- $f_i\left(\boldsymbol{\alpha}_i; x_i, \mathbf{x}_{-i}\right) : W_i^{\delta_i} \times A \rightarrow \mathbb{R}$ records player i's payoff ($A := \times_{i \in N} A_i$ is the set of all possible strategies).
- f_i is known expect for the value of a parameter vector $\boldsymbol{\alpha}_i \in W_i^{\delta_i}$.
- $W_i^{\delta_i} = \delta_i U_i + (1 - \delta_i)\,\boldsymbol{\alpha}_i^0$, where $\delta_i \in [0, 1]$ measures the level of uncertainty ($\delta_i = \delta$ for simplicity) **(we introduce this representation)**.
- $U_i \subset \mathbb{R}^{\nu_i}$, $v_i \in \mathbb{Z}_+$, is the set of all possible realizations of the vector parameter $\boldsymbol{\alpha}_i$, while $\boldsymbol{\alpha}_i^0 \in U_i$ is a singleton **(nominal realization)**.

Remark: For easy understanding of what follows, player i's uncertainty set

$$W_i^{\delta_i} = \delta_i U_i + (1 - \delta_i)\, \boldsymbol{\alpha}_i^0 \tag{1}$$

has to be understood as follows

- $\boldsymbol{\alpha}_i^0$ must be interpreted as the vector of the real values (nominal realizations) of the parameters;
- A player knows that $\boldsymbol{\alpha}_i^0$ should be the real realization (nominal realization), but (s)he is not one hundred percent sure of what (s)he knows :

  ... I know it should be sunny tomorrow, but I take an umbrella because you never know ....

- Set U_i is the set of all possible realizations for the vector $\boldsymbol{\alpha}_i$ (all states of the world);
- δ measures player i's degree of confidence on the nominal realization.

Worst-case payoff functions and robust games

In a robust game, each player i is uncertainty averse and determines his/her action x_i by maximizing his/her **worst-case payoff function**:

$$\rho_i^{\delta}\left(x_i, \mathbf{x}_{-i}\right) \triangleq \min_{\boldsymbol{\alpha}_i \in W_i^{\delta}} f_i\left(\boldsymbol{\alpha}_i; x_i, \mathbf{x}_{-i}\right) \quad (2)$$

Assuming that the worst-case payoff functions are *well defined*:

Definition

Given the robust game $\left\{A_i, f_i, W_i^{\delta_i} : i \in \mathcal{N}\right\}$:

- The **nominal-game representation** is the complete information game $\left\{A_i, \rho_i^{\delta_i} : i \in \mathcal{N}\right\}$;
- The **nominal counterpart game** $(\delta_i = 0 \ \forall i)$ is the complete information game obtained setting to zero the level of uncertainty of each player, that is $\left\{A_i, f_i, W_i^0 : i \in \mathcal{N}\right\}$.

Robust-optimization equilibrium and Existence Th.

Definition (of robust-optimization equilibrium)

A robust-optimization equilibrium (ROE in short) of a robust game $\left\{A_i, f_i, W_i^{\delta} : i \in \mathcal{N}\right\}$ is a Nash equilibrium of its **nominal-game representation** $\left\{A_i, \rho_i^{\delta} : i \in \mathcal{N}\right\}$.

Assumptions

Let A_i, with $i \in \mathcal{N}$, be subsets of Euclidean spaces. We assume that:

1) *A_i is a non-empty, closed, bounded, and convex set, for all $i \in \mathcal{N}$;*
2) *$U_i \subset \mathbb{R}^{\nu_i}$ is a non-empty, closed, bounded, and convex set, for all $i \in \mathcal{N}$, where ν_i is the number of entries in vector $\boldsymbol{\alpha}_i$;*
3) *$f_i\left(\boldsymbol{\alpha}_i; x_i, \mathbf{x}_{-i}\right)$, $i = 1, \ldots, n$, are continuous on $W_i^{\delta_i} \times A$;*
4) *$f_i\left(\boldsymbol{\alpha}_i; x_i, \mathbf{x}_{-i}\right)$, $i = 1, \ldots, n$, are quasi-concave in x_i, for every $\left(\boldsymbol{\alpha}_i, \mathbf{x}_{-i}\right) \in W_i^{\delta_i} \times A_{-i}$.*

Theorem (existence of robust-optimization equilibrium)

The robust game $\left\{A_i, f_i, W_i^{\delta} : i \in \mathcal{N}\right\}$ has at least a robust-optimization equilibrium.

Theoretical insights into the link between ROE and ϵ-Nash equilibria of the nominal counterpart game

Definition (of ϵ-Nash equilibrium)

The action profile $(x_1^*, \ldots, x_n^*) \in A$ is an ϵ-Nash equilibrium of game $\{A_i, f_i : i \in \mathcal{N}\}$, when for each $i \in \mathcal{N}$

$$f_i\left(\boldsymbol{\alpha}_i^0; x_i^*, \mathbf{x}_{-i}^*\right) \geq f_i\left(\boldsymbol{\alpha}_i^0; x_i, \mathbf{x}_{-i}^*\right) - \epsilon \quad \forall x_i \in A_i \tag{3}$$

G. P. Crespi., D. Radi and M. Rocca (2020), *Insights into the Theory of Robust Games*, working paper arXiv:2002.00225.

Robust players play an ϵ-Nash equilibrium where ϵ is not lower than the cost of uncertainty aversion.

Theoretical insight (1)

We define the **cost of uncertainty aversion** (game-theory analogous of the **loss function proposed in Savage (1951)**) for player i

$$C_i^{\delta_i}(\mathbf{x}_{-i}) := \max_{x_i \in A_i} \rho_i^0(x_i, \mathbf{x}_{-i}) - \rho_i^0\left(x_i^+(\delta_i), \mathbf{x}_{-i}\right) \tag{4}$$

where

$$x_i^+(\delta_i) = \arg\max_{x_i} \rho_i^{\delta_i}(x_i, \mathbf{x}_{-i}) \tag{5}$$

Then, a robust-optimization equilibrium is an ϵ-Nash equilibrium of the nominal counterpart game, **where ϵ is the maximum of players' costs of uncertainty aversion**.

Theorem (ROE is an ϵ-Nash equilibrium of the nominal counterpart game)

Set $\epsilon = \max\left\{C_1^{\delta_1}(\mathbf{x}_{-1}^*), \ldots, C_n^{\delta_n}(\mathbf{x}_{-n}^*)\right\}$. If $(x_1^*, \ldots, x_n^*) \in A$ is a ROE of robust game $\left\{A_i, f_i, W_i^{\delta} : i \in \mathcal{N}\right\}$, then $(x_1^*, \ldots, x_n^*)$ is an ϵ-Nash equilibrium of the nominal counterpart game. Moreover, for $\hat{\epsilon} < \epsilon$, $(x_1^*, \ldots, x_n^*)$ is not an $\hat{\epsilon}$-Nash equilibrium of the nominal counterpart game.

In a robust game, reaching a better approximation is not possible in the sense that ϵ cannot be smaller than the higher of the players' costs of uncertainty aversion.

Theoretical insight (2)

Remarks: It is not straightforward to determine if and at which speed the cost of uncertainty aversion shrinks when uncertainty reduces.

There is, however, a class of robust games for which the cost of uncertainty aversion can be approximated (by excess) by a function which is linear w.r.t. to the level of uncertainty:

As suggested in Crespi *et al.* (2017), $\forall i \in \mathcal{N}$ define

$$\bar{E}_i(\mathbf{x}_{-i}) := \max_{x_i \in A_i} E_i(x_i, \mathbf{x}_{-i}) \quad \text{where} \quad E_i(x_i, \mathbf{x}_{-i}) := \rho_i^0(x_i, \mathbf{x}_{-i}) - \rho_i^1(x_i, \mathbf{x}_{-i}) \tag{6}$$

Lemma

Let $f_i(\boldsymbol{\alpha}_i; x_i, \mathbf{x}_{-i})$ be concave on U_i, for all $(x_i, \mathbf{x}_{-i})$ and $i \in \mathcal{N}$. Then, $\delta \bar{E}_i(\mathbf{x}_{-i}) \geq C_i^{\delta}(\mathbf{x}_{-i})$ $\forall \mathbf{x}_{-i} \in A_{-i}$ and $\forall i \in \mathcal{N}$.

$\downarrow$

Theorem (in Crespi *et al.* (2017))

Let $f_i(\boldsymbol{\alpha}_i; x_i, \mathbf{x}_{-i})$ be concave on U_i, for all $(x_i, \mathbf{x}_{-i})$ and $i \in \mathcal{N}$. If $(x_1^*, \ldots, x_n^*) \in A$ is a ROE of a robust game, then it is an ϵ-Nash equilibrium of the nominal counterpart game, where $\epsilon = \delta \max\left\{\bar{E}_1(\mathbf{x}_{-1}^*), \ldots, \bar{E}_n(\mathbf{x}_{-n}^*)\right\}$.

An ϵ-Nash equilibrium can be motivated in terms of uncertainty aversion.

Theoretical insight (3)

- We know that

... a Nash equilibrium is an ϵ-Nash equilibrium but it is not true the vice versa ...

- We proved that

... a robust-optimization equilibria is an ϵ-Nash equilibrium ...

- But, there is something more:

... each ϵ-Nash equilibrium admits a representation in terms of a robust-optimization equilibrium.

Theorem (Each ϵ-Nash equilibrium can be considered a robust-optimization equilibrium)

Consider a nominal game $\{A_i, f_i : i \in \mathcal{N}\}$. Then, for each ϵ-Nash equilibrium $(x_1^*, \ldots, x_n^*)$ of this game it exists at least a robust game such that $(x_1^*, \ldots, x_n^*)$ is a robust-optimization equilibrium of the robust game and $\{A_i, f_i : i \in \mathcal{N}\}$ is its nominal counterpart.

Implication of theoretical insight (3): The ϵ profit that players give up when they play an ϵ-Nash equilibrium can be interpreted as the **cost of uncertainty aversion**.

Theoretical insight (4)

The proposed interpretation of ϵ-approximation provides a selection criterion for ϵ-Nash equilibria.

Theorem (Discriminating among ϵ-Nash equilibria)

Consider an ϵ-Nash equilibrium of the nominal counterpart game which is not an ϵ^1-Nash equilibrium, with $\epsilon^1 < \epsilon$.

1) *If the cost of uncertainty aversion is lower than ϵ for all players, then for at least one player the unilateral deviation obtained by playing a robust strategy is desirable in terms of both robust game and nominal counterpart game (Do not consider this ϵ-Nash equilibrium!);*

2) *If the cost of uncertainty aversion of one player is higher than ϵ, then for at least one player the unilateral deviation obtained by playing a robust strategy is desirable in terms of robust game but it is not in terms of nominal counterpart game;*

3) *If the cost of uncertainty aversion is higher than ϵ for all players, then for all players the unilateral deviation obtained by playing a robust strategy is desirable in terms of robust game but it is not in terms of nominal counterpart game.*

In addition to the theoretical insights, we can say that the worst-case approach to uncertainty introduces nonlinearities in the payoff functions of players....... and this has a relevant effect when .....

UNCERTAINTY AVERSION MEETS CONSTANT ESPECTATIONS

$\downarrow$

D. Radi and L. Gardini (2021), *Ambiguity aversion as a route to randomness in a duopoly game*, under review.

Remark: In the following $\delta = 1$ and $W_i^\delta = U_i$.

The starting point

- In this work, we consider an example of robust game, specifically a robust duopoly game, proposed in:

G. P. Crespi., D. Radi and M. Rocca (2020), *Insights into the Theory of Robust Games*, working paper arXiv:2002.00225.

- In this framework, we introduce players' expectations (of type naive) on the quantity produced by the competitor.

- The aim is to show that the worst-case approach to uncertainty leads to complicated dynamics (chaos) in an otherwise stable system.

- Moreover, we want to show that this depends on the shape of the uncertainty set rather than on the level of the uncertainty set.

Model Setup

- Cournot-duopoly game where firms produce differentiated goods and each firm produces one type of output only:
  - x is the quantity of commodity 1 produced by firm 1;
  - y is the quantity of commodity 2 produced by firm 2.

- Focusing on a symmetric setting, the inverse demand functions for commodity 1 and commodity 2 are, respectively:

$$P_1(x, y; b, \gamma) = \max\{a - bx - \gamma y, 0\} \quad \text{and} \quad P_2(x, y; b, \gamma) = P_1(y, x; b, \gamma) \tag{7}$$

  where $a > 0$, $b > 0$ and $\gamma = \sigma b(n-1)$, with $\sigma \in [0, b]$ that measures the degree of substitutability between the two products and $n - 1$ is the number of competitors.

- Costs of production are set equal to zero.

- Players' payoff functions are

$$\Pi_1(x, y; b, \gamma) = P_1(x, y; b, \gamma)\, x \quad \text{and} \quad \Pi_2(x, y; b, \gamma) = \Pi_1(y, x; b, \gamma) \tag{8}$$

Uncertainty set

- Uncertainty set $U \subset \mathbb{R}^2$ contains the possible realizations of b and γ and is given by the segment of points joining the two realizations $(\bar{b}, \underline{\gamma})$ and $(\underline{b}, \bar{\gamma})$:

$$U = \left\{ (b, \gamma) \mid \bar{b} \geq b \geq \underline{b} \text{ and } \gamma = -\frac{\bar{\gamma} - \underline{\gamma}}{\bar{b} - \underline{b}} b + \frac{\bar{\gamma}\bar{b} - \underline{b}\underline{\gamma}}{\bar{b} - \underline{b}} \right\} \tag{9}$$

- To have U that includes the homogeneous setting $b = \gamma$, we impose $\bar{b} \geq \bar{\gamma} \geq \underline{b} \geq \underline{\gamma}$:
  - Realization $(\bar{b}, \underline{\gamma})$ is then consistent with a duopoly game, which requires $b \geq \gamma = \sigma b$;
  - For $\bar{\gamma} \geq \underline{b}$, realization $(\underline{b}, \bar{\gamma})$ implies that a firm does not know if she plays a duopoly game or a more general oligopoly game.
- Setting $r = \frac{\bar{\gamma} - \underline{\gamma}}{\bar{b} - \underline{b}}$, we have more uncertainty on γ when $r > 1$, more uncertainty on b when $r < 1$, and the same level of uncertainty on both parameters when $r = 1$.

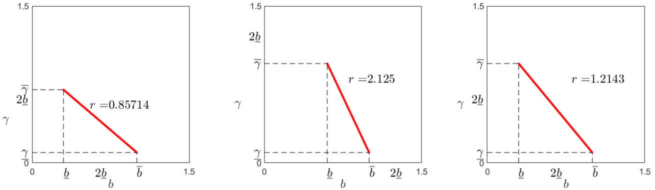

Robust optimization

- Players maximize their guaranteed payoff: Given x the level of production of the competitor, the worst-case best-reply function of firm 2 is given by

$$
\begin{aligned}
f(x) &= \arg\max_{y\geq 0}\left[\arg\min_{(b,\gamma)\in U}\Pi_2(x,y;b,\gamma)\right] \\
&= \begin{cases}
f_l(x) = \frac{1-\underline{\gamma}x}{2\overline{b}} & \text{if} \quad 0\leq x\leq x_l \\
f_m(x) = rx & \text{if} \quad x_l\leq x\leq x_u \\
f_r(x) = \frac{1-\overline{\gamma}x}{2\underline{b}} & \text{if} \quad x_u\leq x\leq x_m \\
0 & \text{if} \quad x\geq x_m
\end{cases}
\end{aligned} \tag{10}
$$

where f_l is also denoted the left branch of f, f_m the middle branch, f_r the right branch, and

$$
x_l = \frac{1}{\underline{\gamma}+2\overline{b}r};\quad x_u = \frac{1}{2\underline{b}+\overline{\gamma}r};\quad x_m = \frac{1}{\overline{\gamma}};\quad r = \frac{\overline{\gamma}-\underline{\gamma}}{\overline{b}-\underline{b}} \tag{11}
$$

- By symmetry, the worst-case best-reply function of firm 1 is $f(y)$.

The worst-case best-reply function (dynamics in Δ)

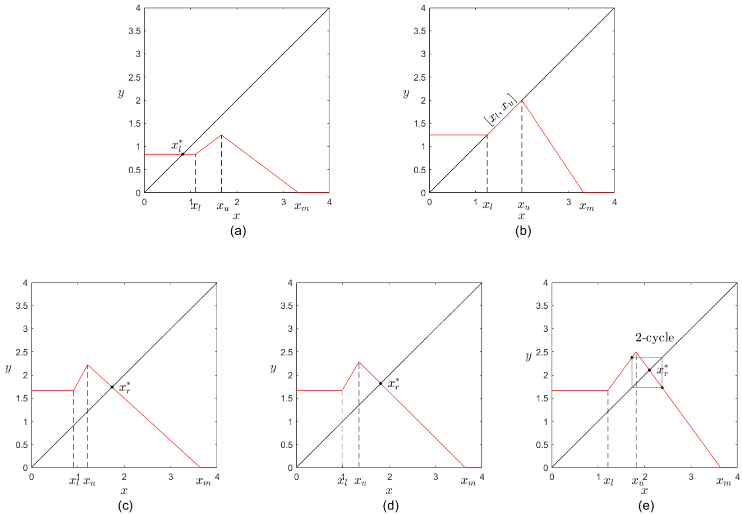

Figure: In (a) $r < 1$. In (b), $r = 0$. Second row $r > 0$. In (c), $\overline{\gamma} > 2\underline{b}$. In (d), $\overline{\gamma} = 2\underline{b}$. In (e), $\overline{\gamma} < 2\underline{b}$.

Constant expectations

- Assuming that firms have constant expectations about the level of production of the competitor, the quantity-dynamics of the robust duopoly game is

$$(x(t+1), y(t+1)) = T(x(t), y(t)) \tag{12}$$

where

$$T(x(t), y(t)) = (f(y(t)), f(x(t))) \tag{13}$$

- In case of no uncertainty, the best-reply function is downward slopping and we have the following *well-known* result:

Proposition (No uncertainty)

Assume no uncertainty, i.e. $\overline{b} = \underline{b} = b$ and $\overline{\gamma} = \underline{\gamma} = \gamma$. The oligopoly game has a unique Cournot-Nash equilibrium that is globally stable under T:

$$\left(\frac{1}{\gamma + 2b}; \frac{1}{\gamma + 2b}\right) \tag{14}$$

- T is a so-called decoupled square system, that is $(x, y) \to \left(f^2(x), f^2(y)\right)$, see Bischi *et al.* (2000) and Tramontana *et al.* (2010).

General properties of the map T

Proposition (Results on decoupled square systems)

The map $T(x, y)$ is such that:

(a) *The diagonal Δ (the straight line $x = y$), is a trapping set (i.e. $T(\Delta) \subseteq \Delta$), the restriction is represented by the one-dimensional map $x(t+1) = f(x(t))$;*

(b) *Any invariant set (stable/unstable sets and basins of attraction) $\mathcal{A}$ of the phase plane (i.e. $T(\mathcal{A}) = \mathcal{A}$), either is symmetric w.r.t. Δ, or the symmetric one is also invariant.*

(c) *Let $\{p_1, \ldots, p_N\}$ be the set of all the periodic points of the one-dimensional map $f(x)$, then the points of the Cartesian product $\{p_1, \ldots, p_N\} \times \{p_1, \ldots, p_N\}$ give all the periodic points of the two-dimensional map T.*

Theorem

Consider map $f(x)$:

(i) *If x^* is a repelling (resp. attracting) fixed point of $f(x)$, then (x^*, x^*) is a repelling (resp. attracting) fixed point of T belonging to the diagonal Δ;*

(ii) *For $r > 1$ and $\overline{\gamma} > 2\underline{b}$ a 2-cycle $\{x_1, x_2\}$ of $f(x)$ leads to a pair of fixed points repelling nodes of T outside the diagonal Δ given by (x_1, x_2) and (x_2, x_1), and to a repelling 2-cycle of T on Δ given by $\{(x_1, x_1), (x_2, x_2)\}$.*

The worst-case best-reply function (dynamics in Δ)

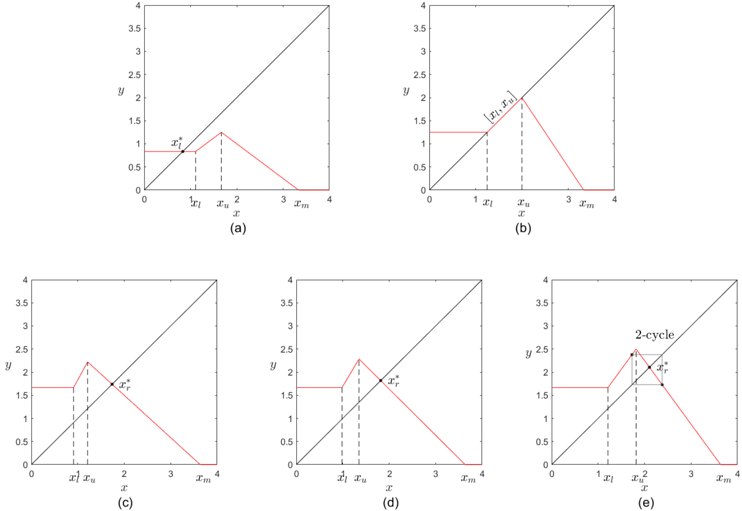

Figure: In (a) $r < 1$. In (b), $r = 0$. Second row $r > 0$. In (c), $\overline{\gamma} > 2\underline{b}$. In (d), $\overline{\gamma} = 2\underline{b}$. In (e), $\overline{\gamma} < 2\underline{b}$.

The worst-case best-reply function (dynamics in Δ)

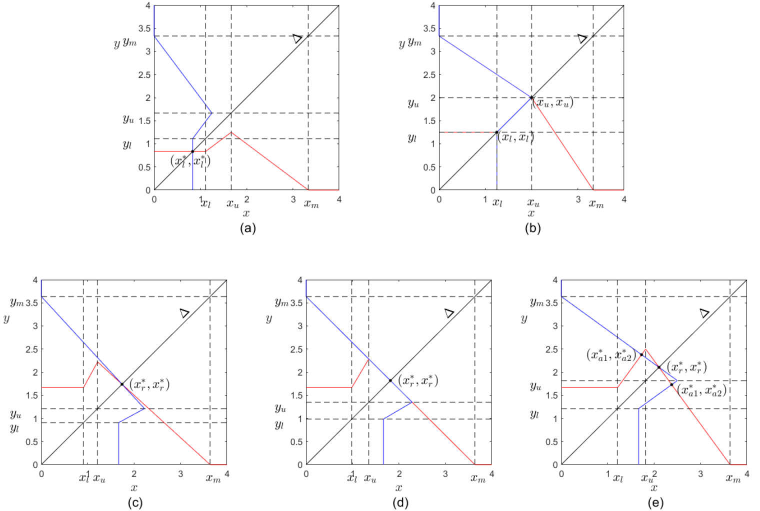

Figure: In (a) $r < 1$. In (b), $r = 0$. Second row $r > 0$. In (c), $\overline{\gamma} > 2\underline{b}$. In (d), $\overline{\gamma} = 2\underline{b}$. In (e), $\overline{\gamma} < 2\underline{b}$.

The worst-case best-reply function

Theorem (The global dynamics of f)

Consider the map f.

(i) *For $(\overline{\gamma} - \underline{\gamma}) < (\overline{b} - \underline{b})$, i.e. $r < 1$, the unique fixed point x_l^* is globally attracting.*

(ii) *For $(\overline{\gamma} - \underline{\gamma}) = (\overline{b} - \underline{b})$, i.e. $r = 1$, the segment $[x_l, x_u]$ is filled with fixed points, stable but not attracting.*

(iii) *For $(\overline{\gamma} - \underline{\gamma}) > (\overline{b} - \underline{b})$, i.e. $r > 1$, and we can have the following dynamics:*

  (a) *$\overline{\gamma} < 2\underline{b}$, then the unique fixed point x_r^* is globally attracting;*

  (b) *$\overline{\gamma} = 2\underline{b}$, then x_r^* undergoes a degenerate flip-bifurcation;*

  (c) *$\overline{\gamma} > 2\underline{b}$, then x_r^* is repelling, a repelling 2-cycle $\{x_1, x_2\}$ exists, and close to the bifurcation the unique attracting set consists in 2^k-cyclic chaotic intervals, where $k \geq 0$ depends on the value r.*

  (d) *For $\overline{\gamma} > 2\underline{b}$ and $f^2(x_u) > x_l$, at $f_m \circ f_r \circ f_m(x_u) = x_r^*$ the first homoclinic bifurcation of the repelling fixed point in the right branch causes the transition from two to one chaotic interval. For $f_m \circ f_r \circ f_m(x_u) < x_r^*$ there is one unique chaotic interval (globally attracting).*

4-cyclical chaotic attractors

From the dynamics of f to the dynamics of T....

Theorem (from chaos in f to chaos in T)

Consider map f in (10)-(11) for $r > 1$, $\overline{\gamma} > 2\underline{b}$ and $f(x_u) < x_m$ (i.e. $\frac{1}{\overline{\gamma}/r+2\underline{b}} < \frac{1}{\overline{\gamma}}$), assuming that map f has 2^k-cyclic chaotic intervals $I_1, ..., I_{2^k}$ for $k \geq 0$. Then map T has one 2^k-cyclic chaotic set crossing the diagonal Δ, given by the squares $I_1 \times I_1, \ldots, I_{2^k} \times I_{2^k}$. Moreover, for $k = 1$, there are two 1-cyclical chaotic attractors outside Δ, that is, $I_1 \times I_2$ and $I_2 \times I_1$, while, for $k > 1$, map T has $(2^k - 1)$ distinct 2^k-cyclic chaotic rectangles external to Δ, consisting of chaotic rectangles $I_i \times I_j$ for $i \neq j$ and $i, j \in \{1, \ldots, 2^k\}$.

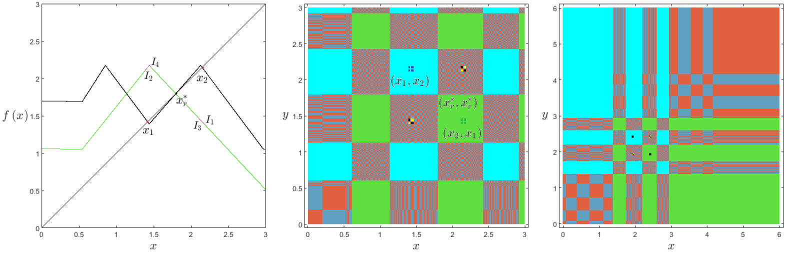

(a) (b) (c)

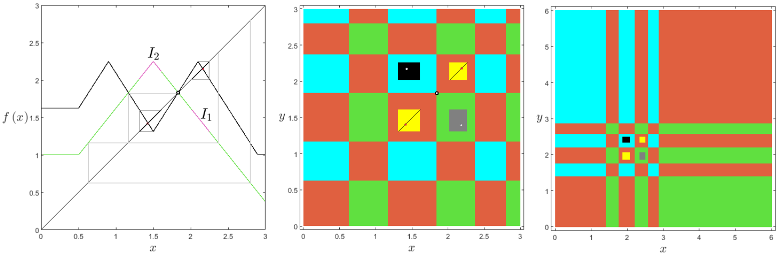

(a) (b) (c)

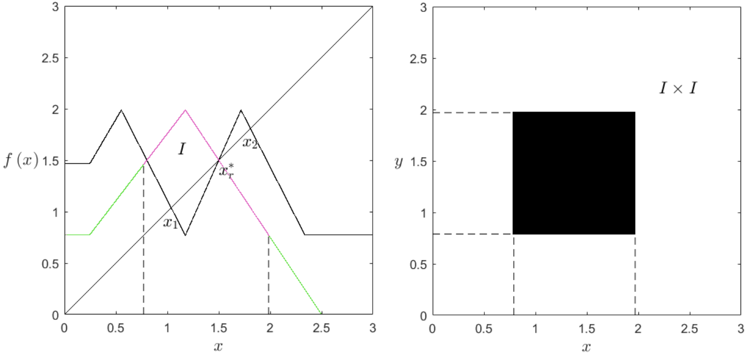

(a) (b)

The parameter space of T

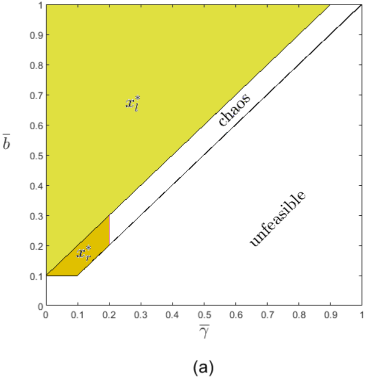

Less uncertainty on b implies a larger chaotic region!!!

A bifurcation diagram of f

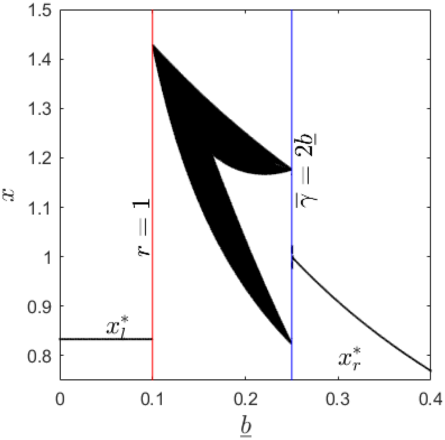

Less uncertainty on b implies a chaotic attractor.... a further reduction of uncertainty on b leads back to an equilibrium!!!

To conclude:

- Inspired by Rand (1978), there are several Cournot duopoly models that exhibit *complicated dynamics*.

- This is due to:
  - Non monotonic best reply functions related to sophisticated cost functions, see Kopel (1996), Bischi *et al.* (2000), Bischi and Kopel (2001);
  - Non monotonic best reply functions related to sophisticated demand functions as the isoelastic demand curve, see T. Puu and several other co-authors, see, e.g., Puu (1991) and Tramontana *et al.* (2010);
  - Sophisticated dynamic adjustment mechanisms, see, e.g., Bischi and Naimzada (2000), or evolutionary selection processes to discriminate between expectation schemes, see, e.g., Droste *et al.* (2002).

- We add to this literature by showing that *complicated* dynamics can emerge even in a very stylized duopoly model when a simple form of uncertainty is introduced.

- Counter-intuitive result: A conservative approach to uncertainty conditioned on constant expectations leads to chaotic dynamics, i.e., it generates further uncertainty instead of reducing it.

M. Aghassi, D. Bertsimas. Robust game theory. *Mahematical Programming Series B*, 107(1–2): 231–273, 2006.

G. I. Bischi, M. Kopel. Equilibrium selection in a nonlinear duopoly game with adaptive expectations. *Journal of Economic Behavior & Organization*, 46:73–100, 2001.

G. I. Bischi, A. Naimzada. *Advances in Dynamic Games and Applications*, chapter Global Analysis of a Duopoly Game with Bounded Rationality, pages 361–385. Birkhauser, 2000.

G. I. Bischi, C. Mammana, L. Gardini. Multistability and cyclic attractors in duopoly games. *Chaos, Solitons and Fractals*, 11:543–564, 2000.

G. P. Crespi, D. Radi, M. Rocca. Robust games: Theory and application to a Cournot duopoly model. *Decisions in Economics and Finance*, 40(1-2):177–198, 2017.

E. Droste, C. H. Hommes, J. Tuinstra. Endogenous fluctuations under evolutionary pressure in Cournot competition. *Games and Economic Behavior*, 40(2):232–269, 2002.

I. Gilboa, D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.

A. Kajii, T. Ui. Incomplete information games with multiple priors. *The Japanese Economic Review*, 56(3):332–351, 2005.

M. Kopel. Simple and complex adjustment dynamics in cournot duopoly models. *Chaos, Solitons & Fractal*, 7(12):2031–2048, 1996.

T. Puu. Chaos in duopoly pricing. *Chaos, Solitons & Fractals*, 1(6):573–581, 1991.

References II

D. Rand. Exotic phenomena in games and duopoly models. *Journal of Mathematical Economics*, 5 (2):173–184, 1978.

L. J. Savage. The theory of statistical decision. *Journal of the American Statistical Association*, 46 (253):55–67, 1951.

F. Tramontana, L. Gardini, T. Puu. Global bifurcations in a piecewise-smooth Cournot duopoly game. *Chaos, Solitons & Fractals*, 43(1-12):15–24, 2010.