Relational contracts in repeated interactions

Watson §22-23, pages 257-282

Bruno Salcedo

The Pennsylvania State University

Econ 402

Summer 2012

Repeated interactions

- When agents interact repeatedly they can use publicly observed history as a coordination devise
- Each agent can condition his/her choices on the observations of the past
- When agent's make their choices they do not consider only their direct impact on payoffs but also the way that other agents will react to them
- By reacting to past behavior, agents can enforce "relational" contracs that generate incentives to implement desirable outcomes
- For example, Anna might use the following strategy: "*I'll be nice to you as long as you are nice to me*"
- Other players might choose to "be nice" to Anna in the present because they want Anna to be nice to them in the future

Example: finitely repeated prisoner's dilemma

- Anna and Bob play the following prisoner's dilemma twice

	C	D
C	4 , 4	0 , 5
D	5 , 0	1 , 1

- They play the game once, and then they play it again *after observing the outcome of the first period*
- The total payoff of each player is the sum of the payoff that he/she gets on each period

Example: finitely repeated prisoner's dilemma

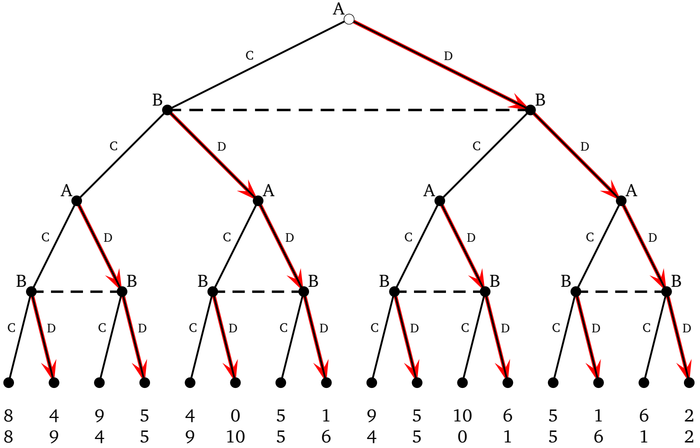

Example: finitely repeated prisoner's dilemma

- It doesn't matter how many times the game is repeated
- Subgame perfection requires that both players defect on the last period (because (D,D) is the only NE of the stage game)
- Since the payoff of the last period is independent of what happens one period before, on the previous period they are also playing a Prisoner's dilemma and the only SPNE implies that they will once again play (D,D)
- This argument can be extended towards the beginning of the game to conclude that in the only SPNE both players will always choose to defect

Theorem

Suppose that players play a simultaneous move game repeatedly (the stage game) a ***finite*** *number of times. If the stage game has a* ***unique*** *NE equilibrium, then the only SPNE of the repeated game has players playing this NE* ***on all periods***

Example: implementation in a two stage game

- Anna and Bob play the following stage game twice:

	C	D	P
C	4 , 4	0 , 5	−3 , −1
D	5 , 0	1 , 1	−2 , −1
P	−1 , −3	−1 , −2	−1 , −1

Example: implementation in a two stage game

- On the second period, Anna and Bob must play a NE, either (D,D), (P,P) or the mixed equilibrium
- However, they can choose which equilibrium to play depending on what happened on the first period
- For instance they could choose to play (D,D) if they played (C,C) on the first period and (P,P) otherwise
- The first period game then looks like:

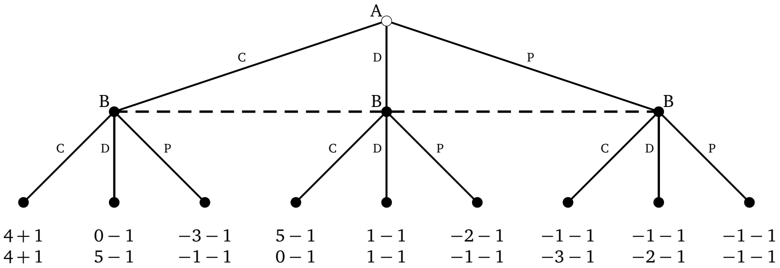

Example: implementation in a two stage game

- On the second period, Anna and Bob must play a NE, either (D,D), (P,P) or the mixed equilibrium
- However, they can choose which equilibrium to play depending on what happened on the first period
- For instance they could choose to play (D,D) if they played (C,C) on the first period and (P,P) otherwise
- The first period game then looks like:

	C	D	P
C	5 , 5	−1 , 4	−4 , −2
D	4 , −1	0 , 0	−3 , −2
P	−2 , −4	−2 , −3	−2 , −2

Example: implementation in a two stage game

- In summary, the following strategies constitute a SPNE:
  - Play C on the 1st period
  - Play D on the 2nd period if the outcome of the first period was (C,C)
  - Play P on the 2nd period otherwise

- This equilibrium achieves the outcome (C,C) on the first period *even though C is a dominated strategy in the stage game*
- Players can generate incentives to play (C,C) on the first period because the stage game has two NE, a good one and a bad one
- They can thus agree on different (self-enforceable) outcomes for the second period that are contingent on the outcome of the first

- Let $v(s)$ be continuation value of outcome s, ie the payoff that a player expects to get on the second period *if the outcome of the first period is* s
- **On the first period players don't maximize just the stage payoff u, they also care about the continuation value so they maximize $u + v$**

Infinitely repeated games

- From now on we are interested in infinitely repeated games with uniform discounting
- Players live for an infinite sequence of period indexed by $t = 0, 1, 2, 3, \dots$
- On each period, player play a simultaneous move game (the stage game) and observe the outcome of the game before proceeding to the next period
- Players discount their payoffs with a common and constant discount factor $\delta \in (0, 1)$

$$v_i(\{s_t\}) = \sum_{t=0}^{\infty} \delta^t u_i(s_t) = u_i(s_0) + \delta u_i(s_1) + \delta^2 u_i(s_2) + \delta^3 u_i(s_3) + \dots$$

- Interpretations:
    - Firms paying interest $r \geq 0$ with $\delta = 1/(1+r)$
    - Uncertainty about the end of the game with hazard rate δ
    - Overlapping generations with concern about the future

Present value

- To compute the present value $v = v(\{u_t\})$ of a constant stream of payoffs $u_t = \bar{u}$ notice that:

$$
\begin{aligned}
& v = \bar{u} + \delta\bar{u} + \delta^2\bar{u} + \delta^3\bar{u} + \dots \\
& \delta v = \delta\bar{u} + \delta^2\bar{u} + \delta^3\bar{u} + \delta^4\bar{u} + \dots \\
\Rightarrow \quad & (1-\delta)v = \bar{u} \\
\Rightarrow \quad & v = \left(\frac{1}{1-\delta}\right)\bar{u}
\end{aligned}
$$

Example: The present value of an investment

- Suppose that an investment generates the stream of payoffs $(-50, 2, 20, 5, 5, 5, 5, \ldots)$ and $\delta = 0.9$
- The present value of the investment is:

$$
\begin{aligned}
v &= -50 + \delta 2 + \delta^2 20 + \delta^3 5 + \delta^4 5 + \delta^5 5 + \ldots \\
&= -50 + \delta 2 + \delta^2 20 + \delta^3 \left(5 + \delta 5 + \delta^2 5 + \ldots \right) \\
&= -50 + \delta 2 + \delta^2 20 + \delta^3 \sum_{t=0}^{\infty} \delta^t 5 \\
&= -50 + \delta 2 + \delta^2 20 + \delta^3 \left(\frac{1}{1-\delta} \right) 5 \\
&= -50 + \frac{9}{10} 2 + \frac{81}{100} 20 + \frac{729}{1000} \frac{10}{1} 5 \\
&= -50 + \frac{180 + 1620 + 3645}{100} = -50 + 54.45 = 4.45
\end{aligned}
$$

SPNE of repeated games

- An outcome *of the stage game* is just strategy profile s of the stage game
- The history up to period T is the truncated history of past outcomes $(s^0, s^1, \ldots, s^T)$
- A strategy *of the repeated game* for player i is a function that specifies a strategy *of the stage game* to be played on each period as a function of the observed history
- A SPNE *of the repeated game* is just a strategy profile *of the repeated game* that induces a NE on every subgame

- Our understanding of the set of SPNE of repeated games is due to two observations from Abreu-Pearce-Stacchetti
  1. The recursive nature of repeated games allows to decompose payoffs as the sum of a stage payoff plus a discounted continuation value
  2. A strategy profile is a SPNE if and only if no player can deviate at a single history and be strictly better off

Example: Repeated Prisoner's dilemma

Grim trigger

- Suppose that Anna and Bob play our prisoner's dilemma repeatedly and $\delta = 0.5$
- Consider the following "Grimm trigger" strategy:
  - As long as everybody has played C in the past, play C
  - If at least one person has played D in the past, play D
- Notice that the present value of the payoffs if both players play C forever of if both players play D forever are:

$$\frac{1}{1-\delta}4 = 8 \qquad \frac{1}{1-\delta}1 = 2$$

Example: Repeated Prisoner's dilemma

Grim trigger

- Suppose that no-one has deviated then the continuation values and stage payoffs are:

Stage payoffs

	C	D
C	4 , 4	0 , 5
D	5 , 0	1 , 1

Continuation values

	C	D
C	8 , 8	2 , 2
D	2 , 2	2 , 2

- Suppose that someone has already deviated then the continuation values and stage payoffs are:

Stage payoffs

	C	D
C	4 , 4	0 , 5
D	5 , 0	1 , 1

Continuation values

	C	D
C	2 , 2	2 , 2
D	2 , 2	2 , 2

Example: Repeated Prisoner's dilemma

Grim trigger

- Suppose that no-one has deviated then the total payoffs are:

	C	D
C	8 , 8	1 , 6
D	6 , 1	2 , 2

- Suppose that someone has already deviated then the total payoffs are:

	C	D
C	5 , 5	1 , 6
D	6 , 1	2 , 2

Example: Repeated Prisoner's dilemma

Tit for tat

- Suppose that Anna and Bob play our prisoner's dilemma repeatedly and $\delta = 0.5$
- Consider the following "tit for tat" strategy:
  - Play C on the first period
  - On every period other than the first play whatever action your opponent played on the last period
- The continuation values and stage payoffs after any history are:

	C	D
C	4 , 4	0 , 5
D	5 , 0	1 , 1

Stage payoffs

	C	D
C	8 , 8	$6.\bar{6}$, $3.\bar{3}$
D	$3.\bar{3}$, $6.\bar{6}$	2 , 2

Continuation values

Example: Repeated Prisoner's dilemma

Tit for tat

- The total payoffs after any history are thus the stage payoffs plus δ times the continuation values:

	C	D
C	8 , 8	$3.\bar{3}$, $6.\bar{6}$
D	$6.\bar{6}$, $3.\bar{3}$	2 , 2

Single deviation principle

Theorem

A strategy profile for the repeated game is a SPNE if and only if no player can unilaterally deviate ***at a single history*** *and be strictly better off*

- Fix a history and the continuation values, players can enforce an agreement that results in the outcome s^* today if and only if for every player i:

$$u_i(s_i^*, s_{-i}^*) + \delta v_i(s_i^*, s_{-i}^*) \geq u_i(s_i', s_{-i}^*) + \delta v_i(s_i', s_{-i}^*)$$

  for every other strategy s_i'
- The single deviation principle means that we can easily verify whether a strategy profile is a SPNE, we don't have to consider the entire game at the same time we can consider individual histories one at a time
- Repeated games are recursive in nature: after each history a new subgame begins and this subgame is identical to the entire game
- Hence continuation values must be SPNE payoffs for the entire game

Example: Repeated Prisoner's dilemma

Grim trigger

- With $\delta = 0.5$, the grim trigger strategies are a SPNE
- Suppose that no-one has deviated then the total payoffs are:

	C	D
C	**8 , 8**	1 , 6
D	6 , 1	2 , 2

- Suppose that someone has already deviated then the total payoffs are:

	C	D
C	5 , 5	1 , 6
D	6 , 1	**2 , 2**

Example: Repeated Prisoner's dilemma

Grim trigger

- We can ask for which values of δ are the grim trigger strategies a SPNE
- Suppose that no-one has deviated then the total payoffs are:

	C	D
C	$4+\frac{\delta}{1-\delta}8\,,\ 4+\frac{\delta}{1-\delta}8$	$\frac{\delta}{1-\delta}2\,,\ 5+\frac{\delta}{1-\delta}2$
D	$5+\frac{\delta}{1-\delta}2\,,\ \frac{\delta}{1-\delta}2$	$1+\frac{\delta}{1-\delta}2\,,\ 1+\frac{\delta}{1-\delta}2$

- For players to be willing to choose C at this point we need:

$$4+\frac{\delta}{1-\delta}4 \geq 5+\frac{\delta}{1-\delta}1 \quad \Leftrightarrow \quad 4(1-\delta)+4\delta \geq 5(1-\delta)+1\delta$$

$$\Leftrightarrow \quad 4 \geq 5-4\delta \quad \Leftrightarrow \quad \delta \geq \frac{1}{4}$$

Example: Repeated Prisoner's dilemma

Grim trigger

- Suppose that someone has already deviated then the total payoffs are:

	C	D
C	$4+\frac{\delta}{1-\delta}2\,,\,4+\frac{\delta}{1-\delta}2$	$\frac{\delta}{1-\delta}2\,,\,5+\frac{\delta}{1-\delta}2$
D	$5+\frac{\delta}{1-\delta}2\,,\,\frac{\delta}{1-\delta}2$	$1+\frac{\delta}{1-\delta}2\,,\,1+\frac{\delta}{1-\delta}2$

- For players to be willing to choose D at this point we need:

$$1+\frac{\delta}{1-\delta}2 \geq \frac{\delta}{1-\delta}2 \quad \Leftrightarrow \quad 1 \geq 0$$

- Which is always satisfied
- Hence the grim trigger strategies are a SPNE if and only if $\delta \geq 1/7$

Example: Repeated Prisoner's dilemma

Tit for tat

- The total payoffs after any history are:

	C	D
C	8 , 8	$3.\bar{3}$, $6.\bar{6}$
D	$6.\bar{6}$, $3.\bar{3}$	2 , 2

- Notice that C is a dominant strategy with these payoffs, hence Tit for Tat is a NE when $\delta = 0.5$
- However it is not a SPNE because players are not willing to play D after a deviation
- There is a modification of Tit for Tat that results in a SPNE (cf Boyd)

Example: Repeated Prisoner's dilemma

Modified Tit for tat

- Given a given history we say that a player, say Anna, is in good standing if:
  1. It is the first period (everybody begins in good standing)
  2. On the last period she played C and Bob was in good standing
  3. On the last period she played C and she was in bad standing
  4. On the last period she played D, she was in good standing and Bob was in bad standing

- The modified Tit for Tat strategies are as follows: *"Play C unless you are in good standing and your opponent is in bad standing in which case you should play D"*

- Notice that:
  1. All players remain in good standing as long as they don't deviate
  2. If a player deviates only at a single period he/she goes to bad standing *for a single period*
- The strategy thus only punishes unilateral deviations, and only punishes them for one period

Example: Repeated Prisoner's dilemma

Modified Tit for tat

- To verify whether the modified tit for tat strategies are a SPNE one must check that there are no unilateral *single history* deviations in four cases: (B,B), (G,G), (B,G) and (G,B)
- We will only verify this for (G,G) the remaining cases will be part of HW3
- In this case, players will play C forever if both players choose C or both players choose D
- If the outcome is (C,D), the outcome on the next period will be (D,C) and, after that, players will play (C,C) forever the continuation values are:

	C	D
C	$4+\frac{\delta}{1-\delta}4\,,\,4+\frac{\delta}{1-\delta}4$	$5+\frac{\delta}{1-\delta}4\,,\,\frac{\delta}{1-\delta}4$
D	$\frac{\delta}{1-\delta}4\,,\,5+\frac{\delta}{1-\delta}4$	$4+\frac{\delta}{1-\delta}4\,,\,4+\frac{\delta}{1-\delta}4$

Example: Repeated Prisoner's dilemma

Modified Tit for tat

- When both players are in good standings, the total payoffs are:

	C	D
C	$4+4\delta+\frac{\delta^2}{1-\delta}4\,,\ 4+4\delta+\frac{\delta^2}{1-\delta}4$	$5\delta+\frac{\delta^2}{1-\delta}4\,,\ 5+\frac{\delta^2}{1-\delta}4$
D	$5+\frac{\delta^2}{1-\delta}4\,,\ 5\delta+\frac{\delta^2}{1-\delta}4$	$1+4\delta+\frac{\delta^2}{1-\delta}4\,,\ +4\delta+\frac{\delta^2}{1-\delta}4$

- Both players are supposed to play C, which is a NE if and only if:

$$4+4\delta+\frac{\delta^2}{1-\delta}4 \geq 5+\frac{\delta^2}{1-\delta}4 \quad \Leftrightarrow \quad 4+4\delta \geq 5 \quad \Leftrightarrow \quad \delta \geq \frac{1}{4}$$

Simple punishments

- Abreu showed that the structure of the modified Tit for Tat strategies are sufficient to implement any outcome that can be implemented
- There is a sequence of outcomes to be implemented, and a sequence of outcomes to punish *each player*
- All players begin in good standing and remain to be in good standing as long as they don't deviate
- If someone deviates unilaterally he/she goes to bad standing and everybody switches to the strategies that punish him/her
- If someone deviates during a punishment phase, then he/she goes to bad standing and everybody switches to punish him or her

The folk theorem

- In repeated situations players can implement (as SPNE) outcomes that are not NE of the stage game
- The way to do so is to "punish" players who deviate by playing "against" them in the future
- The more patient that players are, the more they value the future and thus the more willing they are to comply today in order to avoid future punishments
- The folk theorems loosely speaking states that when players are patient enough, the coordination possibilities arising from repeated interactions are almost a perfect substitute for complete enforceable contracts

Theorem

The set of SPNE payoffs converges to the set of individually rational outcomes

- The folk theorem is robust to the perfect monitoring assumption
- See slides 9 on Moral hazard for a definition of individually rational outcomes

Example: A 5×5 game

Stage game

	A	B	C	D	E
A	$2, -1$	$0, 0$	$\underline{2}, -1$	$-1, 2$	$0, \underline{5}$
B	$0, 0$	$\boxed{\underline{2}, \underline{1}}$	$0, 0$	$0, 0$	$0, 0$
C	$-1, \underline{2}$	$\underline{2}, -1$	$\boxed{\underline{2}, \underline{2}}$	$0, 0$	$-1, \underline{2}$
D	$0, 0$	$0, 0$	$-1, \underline{2}$	$\boxed{\underline{1}, \underline{2}}$	$\underline{2}, -1$
E	$\underline{5}, 0$	$0, 0$	$\underline{2}, -1$	$0, 0$	$-1, \underline{2}$

Example: A 5×5 game

$\delta < 0.45$

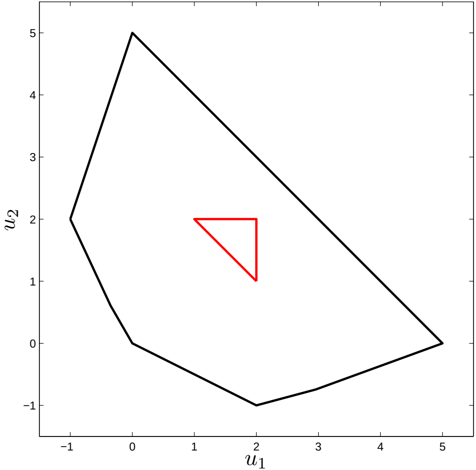

Example: A 5×5 game

$\delta = 0.45$

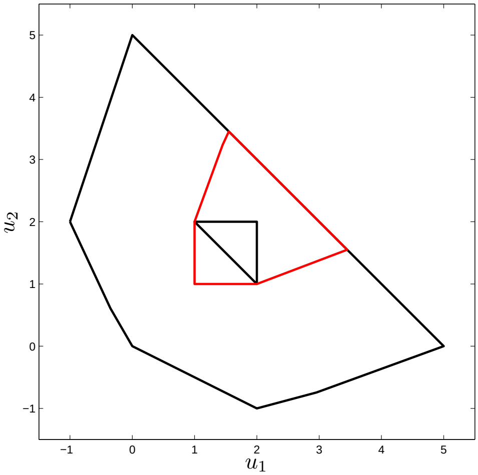

Example: A 5×5 game

$\delta = 0.80$

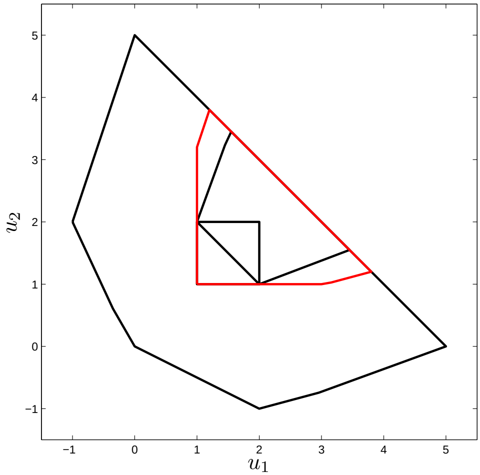

The Golden Rule

- One should treat others as one would like others to treat oneself
- Tit for tat
- *Hamurabi*.– An eye for an eye, a tooth for a tooth
- *Egypt*.– That which you hate to be done to you, do not do to another
- *Bible*.– Therefore all things whatsoever would that men should do to you, do ye even so to them
- *Hinduism*.– One should never do that to another which one regards as injurious to one's own self
- *Buddhism*.– Hurt not others in ways that you yourself would find hurtful
- *Confucius*.– Never impose on others what you would not choose for yourself
- *Qu'ran*.– That which you want for yourself, seek for mankind
- *Qu'ran*.– The most righteous person is the one who consents for other people what he consents for himself, and who dislikes for them what he dislikes for himself
- *Kant*.– Act only according to that maxim whereby you can, at the same time, will that it should become a universal law
- *Thales*.– Avoid doing what you would blame others for doing

Computing the set of SPNE payoffs

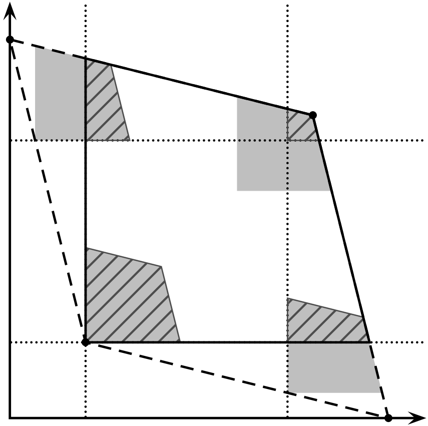

Computing the set of SPNE payoffs

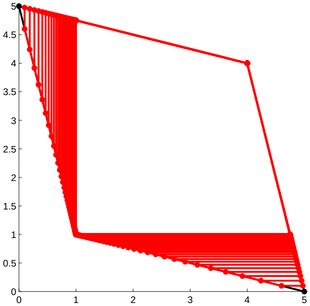

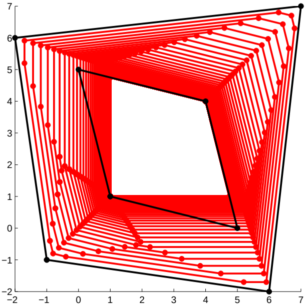

Example: Cournot competition

- Consider the Cournot duopoly with firms 1 and 2 producing the same good with constant marginal cost $c = 10$ and inverse demand function:

$$P(q_1, q_2) = 100 - q_1 - q_2$$

- Recall the the unique NE of this has both firms producing $q^C = 30$ and results in profits $u^C = 900$
- The symmetric Pareto efficient outcome has both firms producing $q^C = 22.5$ and results in profits $u^* = 1012.5$

- If the game was played repeatedly, firms could use the Grimm trigger strategy: *"Choose q^* as long as everyone has produced q^* in the past and produce q^C otherwise"*

Example: Cournot competition

- The continuation value associated to any deviation from q^* is:

$$v' = \frac{1}{1-\delta} u^C = \frac{1}{1-\delta} 900$$

- If one firm is producing q^*, the most profitable deviation for the other firm is q':

$$q' = BR(q^*) = 45 - \frac{1}{2} q^* = 33.75$$

which results in the stage payoff:

$$u' = (90 - q^* - q')q' = (33.75)(33.75) = 1139.0625$$

Example: Cournot competition

- Producing (q^C, q^C) is always incentive compatible because it is an equilibrium of the stage game
- Hence, the Grimm trigger strategies are a SPNE of the repeated game if and only if:

$$
\begin{aligned}
& \frac{1}{1-\delta} u^* \geq u' + \delta v' \\
\Leftrightarrow \quad & \frac{1}{1-\delta} 1012.5 \geq 1139.0625 + \frac{\delta}{1-\delta} 900 \\
\Leftrightarrow \quad & 1012.5 \geq (1-\delta) 1139.0625 + \delta 900 \\
\Leftrightarrow \quad & 1012.5 \geq 1139.0625 + \delta \left(900 - 1139.0625 \right) \\
\Leftrightarrow \quad & \delta \left(1139.0625 - 900 \right) \geq 1139.0625 - 1012.5 \\
\Leftrightarrow \quad & \delta \geq \frac{1139.0625 - 1012.5}{1139.0625 - 900} \approx 0.53
\end{aligned}
$$

Example: Bertrand Competition

SPNE payoffs

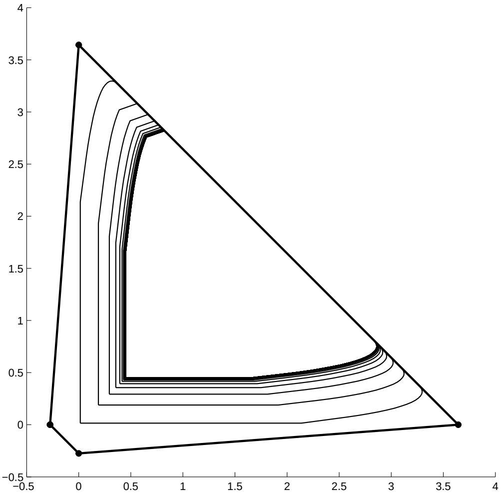

Imperfect monitoring