On Reputation with Imperfect Monitoring

M. W. Cripps, G. Mailath, L. Samuelson

UCL, Northwestern, Pennsylvania, Yale

Theory Workshop

Reputation Effects or Equilibrium Robustness

Reputation Effects:

- Kreps, Wilson and Milgrom and Roberts: A small amount of uncertainty has a big effect on the set of equilibrium payoffs.
- This has come to be called a reputation effect.
- Usually this considers one long run player playing a sequence of short run players. (Sometimes these are very short run as in continuous time models Faingold and Sannikov 2007.)

Reputation Effects or Equilibrium Robustness

Equilibrium Robustness:

- Folk Theorem $\Rightarrow$ a repeated game has many equilibrium payoffs as $\delta \rightarrow 1$.
- Does introducing a small amount of uncertainty shrink this set significantly and sharpen predictive power?
- This is a continuity question: Can you find payoffs of limiting equilibria (as $\delta \rightarrow 1$) in games with incomplete information that are close to any folk-theorem payoff?

This is equivalent to thinking about the value of a reputation when playing against a long run opponent.

Weak Reputations Under Perfect Monitoring

Cripps and Thomas (1997,2003): When players are able to monitor each others actions perfectly and have equal discount factors, then adding a small amount of incomplete information will not change the set of equilibrium payoffs dramatically...

- Take a repeated strategic form game.
- Introduce uncertainty about the type of one of the players.
- Consider the set of equilibrium payoffs as $\delta \to 1$.
- Show that you can find equilibria in this set that give the informed player payoffs arbitrarily close to their minmax payoff.

Note: This approach is known to work in all but 3 special cases mentioned below. These conclusions were substantially generalized in Peski (2007).

Figure 1

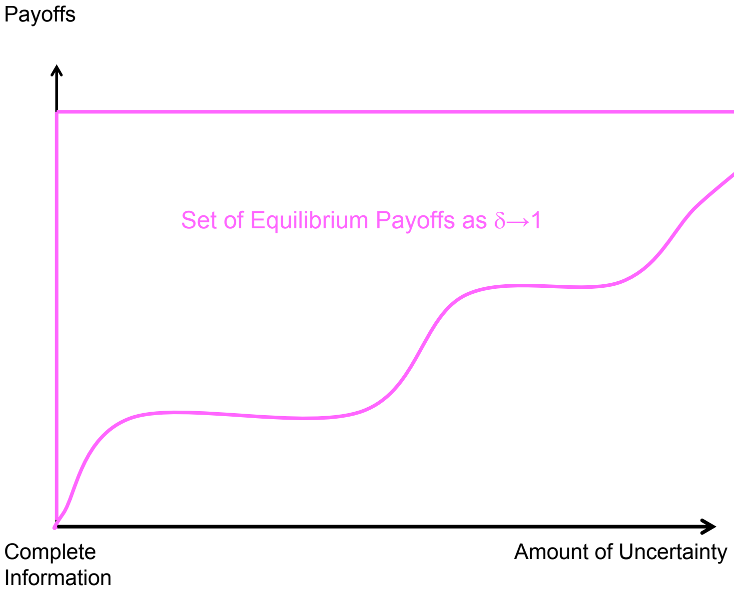
Payoffs
Set of Equilibrium Payoffs as δ→1
Complete
Information
Amount of Uncertainty

Our Example

Consider the game:

$$\begin{bmatrix} (1,1) & (0,0) \\ (0,0) & (0,0) \end{bmatrix}$$

In this game there are no reputation effects under perfect monitoring but full reputation effects with an arbitrary small amount of imperfect monitoring.

Notation

- Let $\delta < 1$ denote the discount factor for both players.
- There is uncertainty about the type of the row player.
- At time $t = -1$ the "type" of the row player is selected.
- With probability μ row is a commitment type.
- The commitment type always plays the top row.
- With probability $1 - \mu$ row is a normal type.
- The normal type has payoffs as in the above matrix.

Strategies and Beliefs

- μ_t denotes the column player's posterior at the start of time t that row is the commitment type.
- $(p_t, 1-p_t)$ is the row player's time t behavior strategy.
- π_t is the probability the uninformed player attaches to the commitment action being played at time t.

$$\begin{pmatrix} \pi_t \\ 1-\pi_t \end{pmatrix} \equiv \mu_t \begin{pmatrix} 1 \\ 0 \end{pmatrix} + (1-\mu_t) \begin{pmatrix} p_t \\ 1-p_t \end{pmatrix}$$

and

$$\mu_{t+1} = \frac{\mu_t}{\pi_t}, \qquad or \qquad \mu_{t+1} = 0$$

How to build bad equilibria:

$$\begin{bmatrix} (1,1) & (0,0) \\ (0,0) & (0,0) \end{bmatrix}$$

We will now construct an equilibrium where: the column player plays Right for N periods and then $(1,1)$ is played forever.
The players get the payoffs

$$(1-\delta^N)0 + \delta^N 1 = \delta^N$$

where $\delta^N \to 0$ as $\delta \to 1$ and $\mu \to 0$.

Bad Equilibria

(1,1)	(0,0)
(0,0)	(0,0)

Events in the first period of Play

Bad Equilibria

Column player plays
Right

(1,1)	(0,0)
(0,0)	(0,0)

Events in the first period of Play

Bad Equilibria

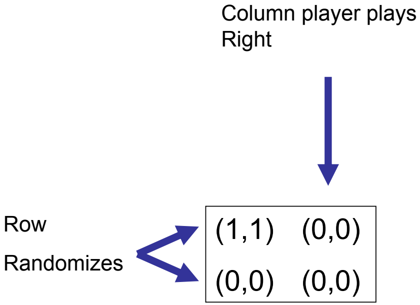

Events in the first period of Play: Then what?

Bad Equilibria

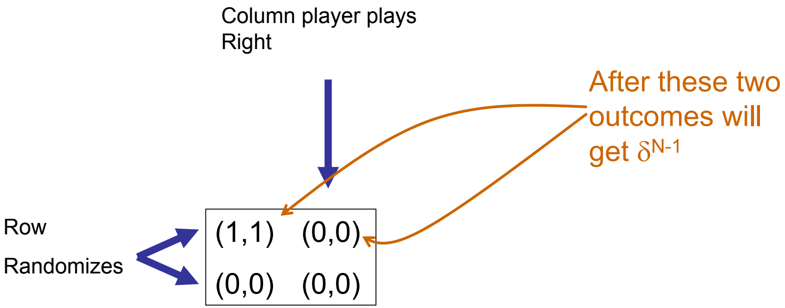

IF REPUTATION IS PRESERVED PLAY EQUILIBRIUM WITH N-1 PERIODS OF RIGHT

Bad Equilibria

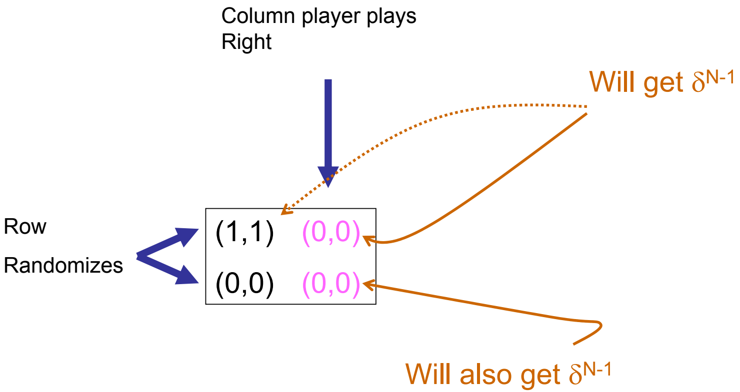

IF REPUTATION BROKEN THROUGH A CHOICE RANDOMIZATION TELLS US WHAT THE ROW PLAYER WILL GET

Bad Equilibria

(1,1)	(0,0)
(0,0)	(0,0)

1st period

(δ^N, δ^{N-1})	$(\delta^{N-1}, \delta^{N-1})$
(?,?)	$(\delta^{N-1}, \delta^{N-1})$

Afterwards

What do we put if Bottom Left is played?

Bad Equilibria

(1,1)	(0,0)
(0,0)	(0,0)

1^{st} period

(δ^N, δ^{N-1})	$(\delta^{N-1}, \delta^{N-1})$
(0,0)	$(\delta^{N-1}, \delta^{N-1})$

Afterwards

What do we put if Bottom Left is played?

Make this as Low as possible to punish column player for playing right.

The incentive to play Right

- When plays right gets $(1-\delta)0 + \delta(\delta^{N-1})$.
- If play left and up is played will get 1 today and δ^{N-1} tomorrow.

$$(1-\delta) + \delta^N$$

- If play left and down is played will get 0 today and 0 tomorrow.
- Right is optimal iff

$$\delta^N \geq \pi\left((1-\delta) + \delta^N\right) + (1-\pi)0$$

- Equivalently

$$1-\pi \geq \frac{1-\delta}{1-\delta+\delta^N}$$

Summary: This is a potential equilibrium as long as the probability the row player plays down, $1-\pi$, isn't too small.

The incentive to play Right for N periods

We have 3 conditions that need to be satisfied

- $\pi_t = \mu_t + (1 - \mu_t)p_t$ is the probability that the row player plays top.
- $\mu_{t+1} = \mu_t / \pi_t$ Bayesian updating.
- $(1 - \mu_t)(1 - p_t) = 1 - \pi_t \geq \frac{1-\delta}{1-\delta+\delta^{N-t}}$ gives incentive to play right
- Solving iteratively give

$$\mu_0 \leq \prod_{n=1}^{N} \frac{\delta^n}{1 - \delta + \delta^n}$$

The set of equilibria

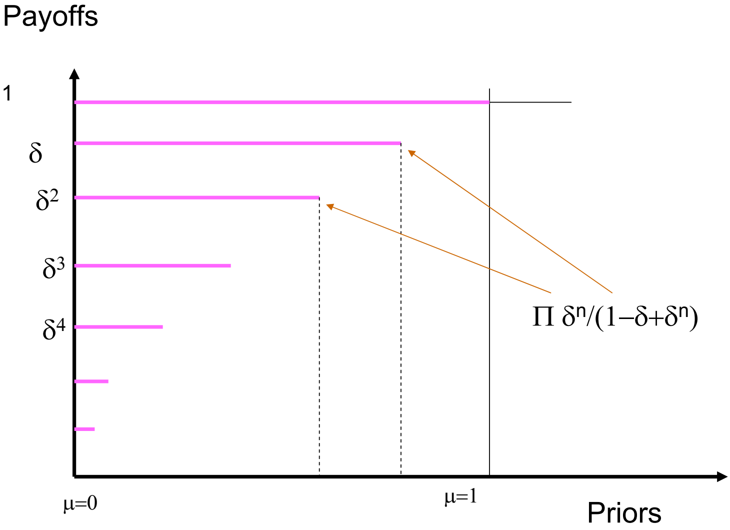

Behavior as $\delta \to 1$.

Taking logarithms

$$\log \mu_0 \leq \sum_{n=1}^{N} \log \frac{\delta^n}{1-\delta+\delta^n}$$

Now use $\log x \geq 1-(1/x)$ to get the sufficient condition

$$\log \mu_0 \leq \sum_{n=1}^{N} \frac{\delta-1}{\delta^n} = 1-\delta^{-N}$$

This implies we can choose

$$\delta^N = \frac{1}{1-\log \mu}$$

Which tends to zero as $\mu \to 0$.

Why do we fail to get reputation effects?

- Key feature is that the uninformed player does not want to play a best response to the reputation.
- He is punished if he plays right and the row player plays down.
- The punishment cannot occur too frequently because otherwise there is a big loss of reputation. So the punishment is is a vanishingly (as $\delta \to 1$) chance of a big (He gets $(0, 0)$) loss.

In 3 Known Cases this Breaks down:

- Chan: The commitment action is strictly dominant in the stage game — in this case can never provide incenties for the row player to randomize.
- Cripps, Dekel, Pesendorfer: Games of conflicting interests — in this case playing a best response to the reputation action minmaxes the uninformed player and nothing worse than this can be done to him!
- Atakan and Ekmekci: Repeated Extensive form games — The punishment has to occur after the deviation has occurred and therefore cannot be too bad.

Imperfect Public Monitoring

- We deal with the simplest possible case: The column player's action is perfectly observable.
- The row player's action is imperfectly monitored.
- With probability $1 - \epsilon$ the column player sees the true action.
- With probability ϵ the column player sees the reverse action.
- Payoffs are unobservable.

Notation Strategies and Beliefs

- μ_t denotes the column player's posterior at the start of time t that row is the commitment type.
- $(p_t, 1 - p_t)$ is the row player's time t behavior strategy.
- π_t is the probability the uninformed player attaches to the commitment action being played at time t.
- $\tilde{\pi}_t = \epsilon + (1 - 2\epsilon)\pi_t$ is the probability with which the uninformed player observes a signal that says the commitment action was played at time t.

Bayes' Theorem

$$\mu_{t+1} = \frac{\mu_t(1-\epsilon)}{\tilde{\pi}_t} \equiv \mu' \qquad \text{or} \qquad \mu_{t+1} = \frac{\mu_t \epsilon}{1 - \tilde{\pi}_t} \equiv \mu''$$

Intuition for the result

- Recall our earlier construction..
- Playing optimally against the reputation type is punished.
- Punishment = a very small probability of a very large loss.
- A large loss is possible because when the row player plays down they reveal their type and play an equilibrium of the complete information game.
- The very small probability is necessary because this has to occur in many periods.

Intuition for the result

- *Under Imperfect Monitoring playing down with very small probability does not reveal your type*!
- It results in an arbitrarily small revision of beliefs and consequently arbitrarily small punishment.
- The noise in the signals means very small actions by the row player are very hard for the column player to detect.

Bayes' Theorem after down

$$\mu_{t+1} = \frac{\mu_t \epsilon}{1 - \tilde{\pi}_t} = \frac{\mu_t \epsilon}{1 - \epsilon - (1 - 2\epsilon)\pi_t} \rightarrow \mu_t$$

AS $\pi_t \rightarrow 1$.

The Result

Let $b_\delta(\mu)$ be the worst public equilibrium payoff to the column/row player in the game with prior μ and discount factor $\delta < 1$.

Proposition

For any $\mu > 0$ we have that $\lim_{\delta \to 1} b_\delta(\mu) = 1$.

Strategy of Proof

Step 1 Find a set that includes the equilibrium payoff correspondence.

Step 2 Show that this set can be described as the unique fixed point of a simple operator.

Step 3 Show that this fixed point converges to 1 as $\delta \to 1$.

The Equilibrium Correspondence 1

For any δ the set of public equilibrium payoffs is a closed graph correspondence (set-valued map) from $\mu \in [0,1]$ to equilibrium payoffs in $[0,1]$.

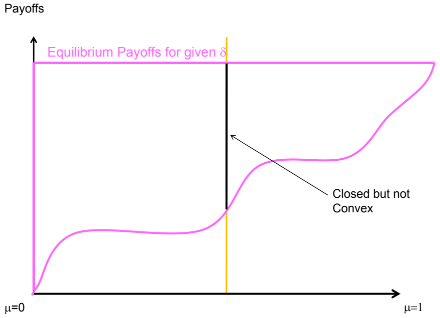

Imposing the restriction that the players get the same equilibrium payoffs!

The Equilibrium Correspondence 2

$\mathcal{E}_\delta : [0,1] \rightrightarrows [0,1]$ is closed but not necessarily convex, take its convex hull.

$$\mathcal{C}^0\mathcal{E}_\delta(\mu)$$

This may allow us to provide incentives for the players to do more so let's write down the set of payoffs that can be enforced using $\mathcal{C}^0\mathcal{E}_\delta(\mu)$ as continuations. Take the convex hull of this.

$$\mathcal{C}^1\mathcal{E}_\delta(\mu)$$

Iterate calculating $\mathcal{C}^n\mathcal{E}_\delta(\mu)$ in the same way.
This is an increasing sequence of closed sets so let

$$\mathcal{E}^*_\delta(\mu) \equiv \overline{\cup_{n=0}^{\infty}\mathcal{C}^n\mathcal{E}_\delta(\mu)} \supseteq \mathcal{E}_\delta(\mu)$$

Characterizing the Correspondence $\mathcal{E}^*_\delta(\mu)$

Let $b^*_\delta(\mu)$ be the minimum value of this correspondence.

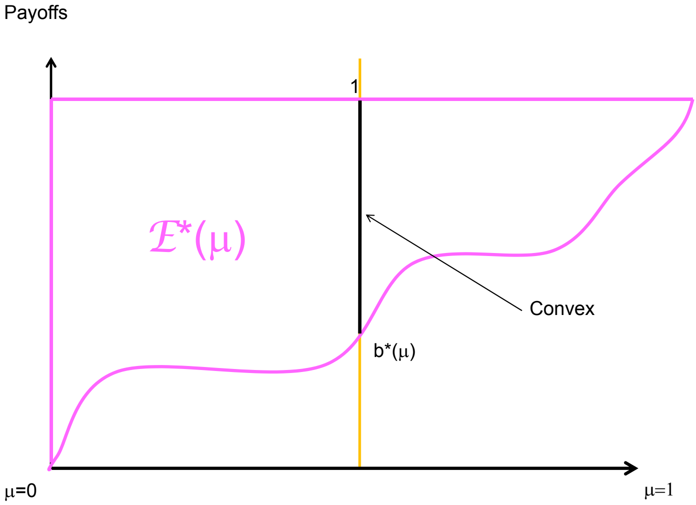

Properties of $\mathcal{E}^*_\delta(\mu)$: 1

Column: Can always play *Left* and get

$$(1-\delta)\pi + \delta\left(\tilde{\pi} b^*_\delta(\mu') + (1-\tilde{\pi}) b^*_\delta(\mu'')\right)$$

This is true at the worst payoff so there exists $\tilde{\pi}$, μ', μ'' such that

$$b^*_\delta(\mu) \geq (1-\delta)\pi + \delta\left(\tilde{\pi} b^*_\delta(\mu') + (1-\tilde{\pi}) b^*_\delta(\mu'')\right)$$

Row: At a worst equilibrium must randomize and be indifferent btwn *Top* and *Bottom*. The continuations to playing bottom cannot be less than those from playing top. (As top payoffs better than the bottom.) This implies

$$b^*_\delta(\mu) \geq \delta \text{B}ottom\ Cont \geq \delta \text{T}op\ Cont \geq \delta b^*_\delta(\mu')$$

Properties of $\mathcal{E}_\delta^*(\mu)$: 2

Combining these

$$b_\delta^*(\mu) \geq \min_{\mu' \mu''} \max \left\{ \begin{array}{c} (1-\delta)\pi + \delta\left(\tilde{\pi} b_\delta^*(\mu') + (1-\tilde{\pi}) b_\delta^*(\mu'')\right) \\ \delta b_\delta^*(\mu') \end{array} \right\}$$

Here the minimum is taken over all pairs μ', μ'' that are consistent with some value of $\tilde{\pi} \in [\epsilon, 1-\epsilon]$.

Properties of $\mathcal{E}_\delta^*(\mu)$: 4

Define the operator

$$\mathcal{T}_\delta \circ f(\mu) \equiv \min_{\mu' \mu''} \max \left\{ \begin{array}{c} (1-\delta)\pi + \delta\left(\tilde{\pi} f(\mu') + (1-\tilde{\pi}) f(\mu'')\right) \\ \delta f(\mu') \end{array} \right\}$$

We have that $b_\delta^*(\mu)$ satisfies.

$$b_\delta^*(.) \geq \mathcal{T}_\delta \circ b_\delta^*(.)$$

Properties of $\mathcal{E}_\delta^*(\mu)$: 5

We can study the properties of $\mathcal{T}_\delta$ and its fixed points:

- **Uniqueness:** $\mathcal{T}_\delta$ is a contraction by Blackwell's Theorem so it has a unique fixed point for all $\delta < 1$.
- **Increasing:** $\mathcal{T}_\delta$ maps increasing functions to increasing functions so the fixed point is increasing.
- **Continuous and Increasing:** $\mathcal{T}_\delta$ maps continuous increasing functions to continuous increasing functions so the fixed point is continuous increasing.
- **$f(\mu) \geq \mu$:** Iterating $\mathcal{T}_\delta$ we can deduce that $f_\delta^*(\mu) \geq \mu$ for any fixed point.
- **Equality:** If f_δ^* is cont. and increasing then the solution to the min max problem has a simple outcome...

Properties of $\mathcal{E}^*_\delta(\mu)$: 5 1/2

Given the above we can conclude that if f^*_δ is the unique fixed point of $\mathcal{T}_\delta$ then $f^*_\delta \leq b^*_\delta$:

Step 1: By definition $\mathcal{T}_\delta b^*_\delta \leq b^*_\delta$.

Step 2: $\mathcal{T}_\delta$ is an increasing map $f \leq g \Rightarrow \mathcal{T}_\delta f \leq \mathcal{T}_\delta g$.

Step 3: The sequence $(\mathcal{T}_\delta)^n \circ b^*_\delta$ is decreasing and converges to f^*_δ.

Properties of $\mathcal{E}^*_\delta(\mu)$: 6

There is a unique increasing, continuous solution to the operator equation satisfying:

$$
\begin{aligned}
f^*_\delta(\mu) &= (1-\delta)\pi + \delta\left(\tilde{\pi} f^*_\delta(\mu') + (1-\tilde{\pi}) f^*_\delta(\mu'')\right) \\
f^*_\delta(\mu) &= \delta f^*_\delta(\mu')
\end{aligned}
$$

Letting $\delta \to 1$: 1

- Consider a sequence of $\delta \to 1$
- This generates a sequence of increasing continuous functions $f_\delta^* : [0,1] \to [0,1]$.
- This has a convergent subsequence (Helly).
- Let us study the properties of this convergent subsequence.

Letting $\delta \to 1$: 2

The limit is continuous on the interior of $[0,1]$. Along this subsequence:

$$\delta f^*_\delta(\mu') = (1-\delta)\pi + \delta\left(\tilde{\pi} f^*_\delta(\mu') + (1-\tilde{\pi}) f^*_\delta(\mu'')\right)$$

This implies

$$(1-\delta)\pi/\delta = (1-\tilde{\pi})(f^*_\delta(\mu') - f^*_\delta(\mu'')) \geq 0$$

But $\tilde{\pi} \leq 1-\epsilon$, so as $\delta \to 1$ we have

$$f^*_\delta(\mu') - f^*_\delta(\mu'') \to 0$$

when $\mu' \geq \mu \geq \mu''$. So the limiting function must be continuous for interior μ.

Letting $\delta \to 1$: 3

Along this subsequence:

$$\begin{aligned} f_\delta^*(\mu) &= (1-\delta)\pi + \delta\left(\tilde{\pi} f_\delta^*(\mu') + (1-\tilde{\pi}) f_\delta^*(\mu'')\right) \\ 0 &= \frac{1-\delta}{\mu\delta(1-\epsilon-\tilde{\pi})}(\pi - f_\delta^*(\mu)) + \Delta_\mu^+ - \Delta_\mu^- \end{aligned}$$

Where the incentives are given by slopes:

$$\begin{aligned} \Delta_\mu^+ &\equiv \frac{f_\delta^*(\mu') - f_\delta^*(\mu)}{\mu' - \mu} \\ \Delta_\mu^- &\equiv \frac{f_\delta^*(\mu) - f_\delta^*(\mu'')}{\mu - \mu'} \end{aligned}$$

Letting $\delta \to 1$: 4

Along this subsequence:

$$\begin{aligned} f_\delta^*(\mu) &= \delta f_\delta^*(\mu') \\ \frac{(1-\delta) f_\delta^*(\mu)}{\delta(\mu'-\mu)} &= \Delta_\mu^+ \end{aligned}$$

Combining this with what came before:

$$0 = \Delta_\mu^+ \left(1 + \frac{\pi - b_\delta^*(\mu)}{\tilde{\pi}}\right) - \Delta_\mu^-$$

Letting $\delta \to 1$: 5

The limit of $b^*_\delta(.)$ is increasing $\Rightarrow$ it is differentiable almost everywhere.
We now show that it is constant on the interior of $[0,1]$.
At a point of differentiability the up-slope and the down-slope converge to the same thing:

$$0 = \Delta^+_\mu \left(1 + \frac{\pi - b^*_\delta(\mu)}{\tilde{\pi}}\right) - \Delta^-_\mu$$

Becomes

$$0 = Db^*_1 \left(\frac{\pi - b^*_1(\mu)}{\tilde{\pi}}\right)$$

Almost everywhere the continuous limit is constant $(Db^*_1 = 0)$ or $\pi = b^*_1(\mu)$. If $\pi < 1$ this implies $\mu' >> \mu''$ and that the slope is constant here too.

And Finally

The limit $b_1^*(\mu)$ is constant on the interior of $[0,1]$. The limiting function also satisfies $b_\delta^*(\mu) \geq \mu$.

The limit is

$$b_1^*(\mu) = \begin{cases} 1, & \mu > 0; \\ 0, & \mu = 0. \end{cases}$$

This is the limit of all convergent subsequences.■