Walras-Bowley Lecture 2003

Sergiu Hart

This version: September 2004

ADAPTIVE HEURISTICS

A Little Rationality Goes a Long Way

Sergiu Hart

Center for Rationality,
Dept. of Economics, Dept. of Mathematics
The Hebrew University of Jerusalem
`http://www.ma.huji.ac.il/~hart`

- Most of this talk is based on joint work with

Andreu Mas-Colell

(Universitat Pompeu Fabra, Barcelona)

- Most of this talk is based on joint work with

  Andreu Mas-Colell

  (Universitat Pompeu Fabra, Barcelona)

- All papers – and this presentation – are available on my home page

`http://www.ma.huji.ac.il/~hart`

Papers

http://www.ma.huji.ac.il/~hart

Papers

http://www.ma.huji.ac.il/~hart

- *Econometrica* (2000)
- *Journal of Economic Theory* (2001)
- *Economic Essays* (2001)
- *Games and Economic Behavior* (2003)
- *American Economic Review* (2003)

PART I

Introduction: Dynamics

Dynamics

Dynamics

“Learning”

Dynamics

"Learning"

- **START**: prior beliefs
- **STEP**:
  - observe
  - update (Bayes)
  - optimize (best-reply)
- **REPEAT**

Dynamics

“Evolution”

Dynamics

“Evolution”

- populations
- each individual $\longleftrightarrow$ fixed action (“gene”)
- frequencies of each action in the population
  $\longleftrightarrow$ mixed strategy

Dynamics

“Evolution”

- populations
- each individual $\leftrightarrow$ fixed action (“gene”)
- frequencies of each action in the population
  $\leftrightarrow$ mixed strategy
- Change:
  - **Selection**
    higher payoff $\Longrightarrow$ higher frequency
  - **Mutation**
    random and relatively rare

“Adaptive Heuristics”

Dynamics

“Adaptive Heuristics”

- “rules of thumb”
- myopic
- simple
- stimulus response, reinforcement
- behavioral, experiments
- non-Bayesian

Dynamics

"Adaptive Heuristics"

- "rules of thumb"
- myopic
- simple
- stimulus response, reinforcement
- behavioral, experiments
- non-Bayesian

Example: **Fictitious Play**

Dynamics

“Adaptive Heuristics”

- “rules of thumb”
- myopic
- simple
- stimulus response, reinforcement
- behavioral, experiments
- non-Bayesian

Example: **Fictitious Play**

(Play optimally against the empirical distribution of past play of the other player)

?

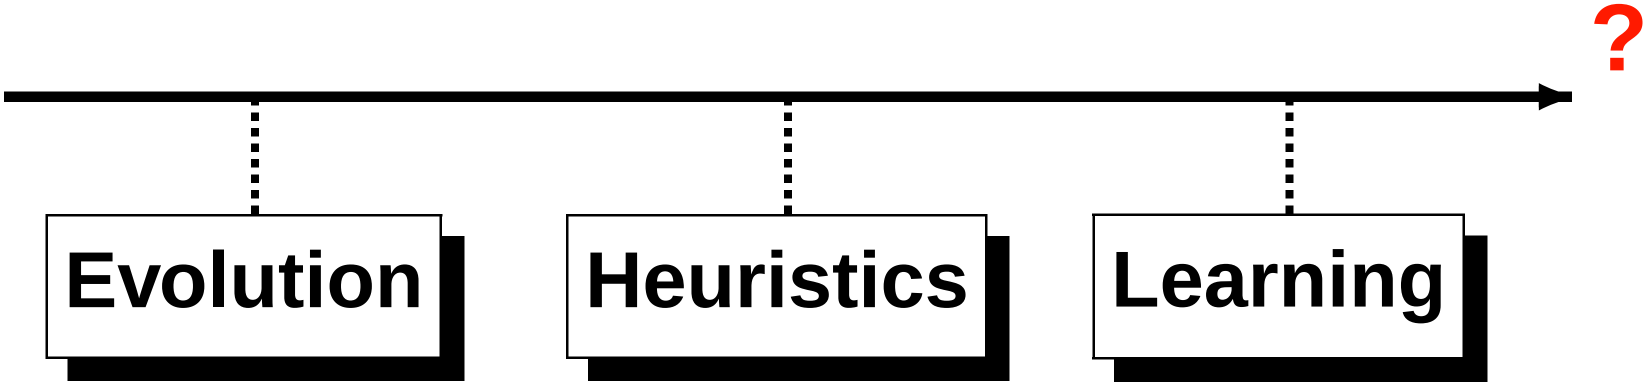
?
Evolution
Heuristics
Learning

Rationality

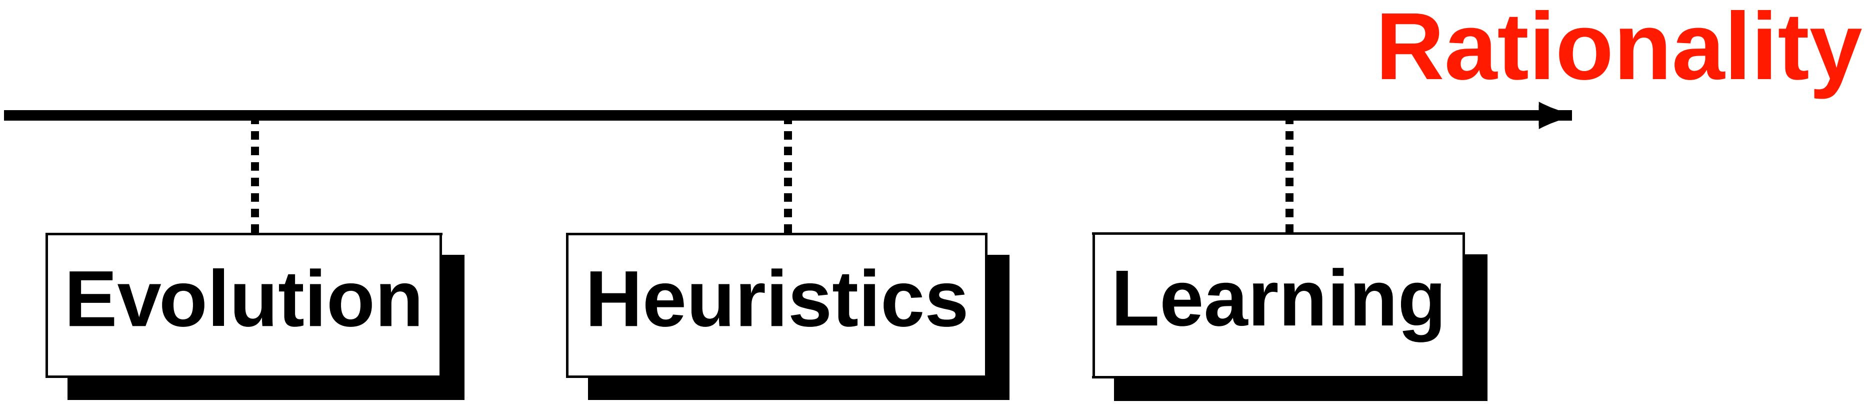

Rationality

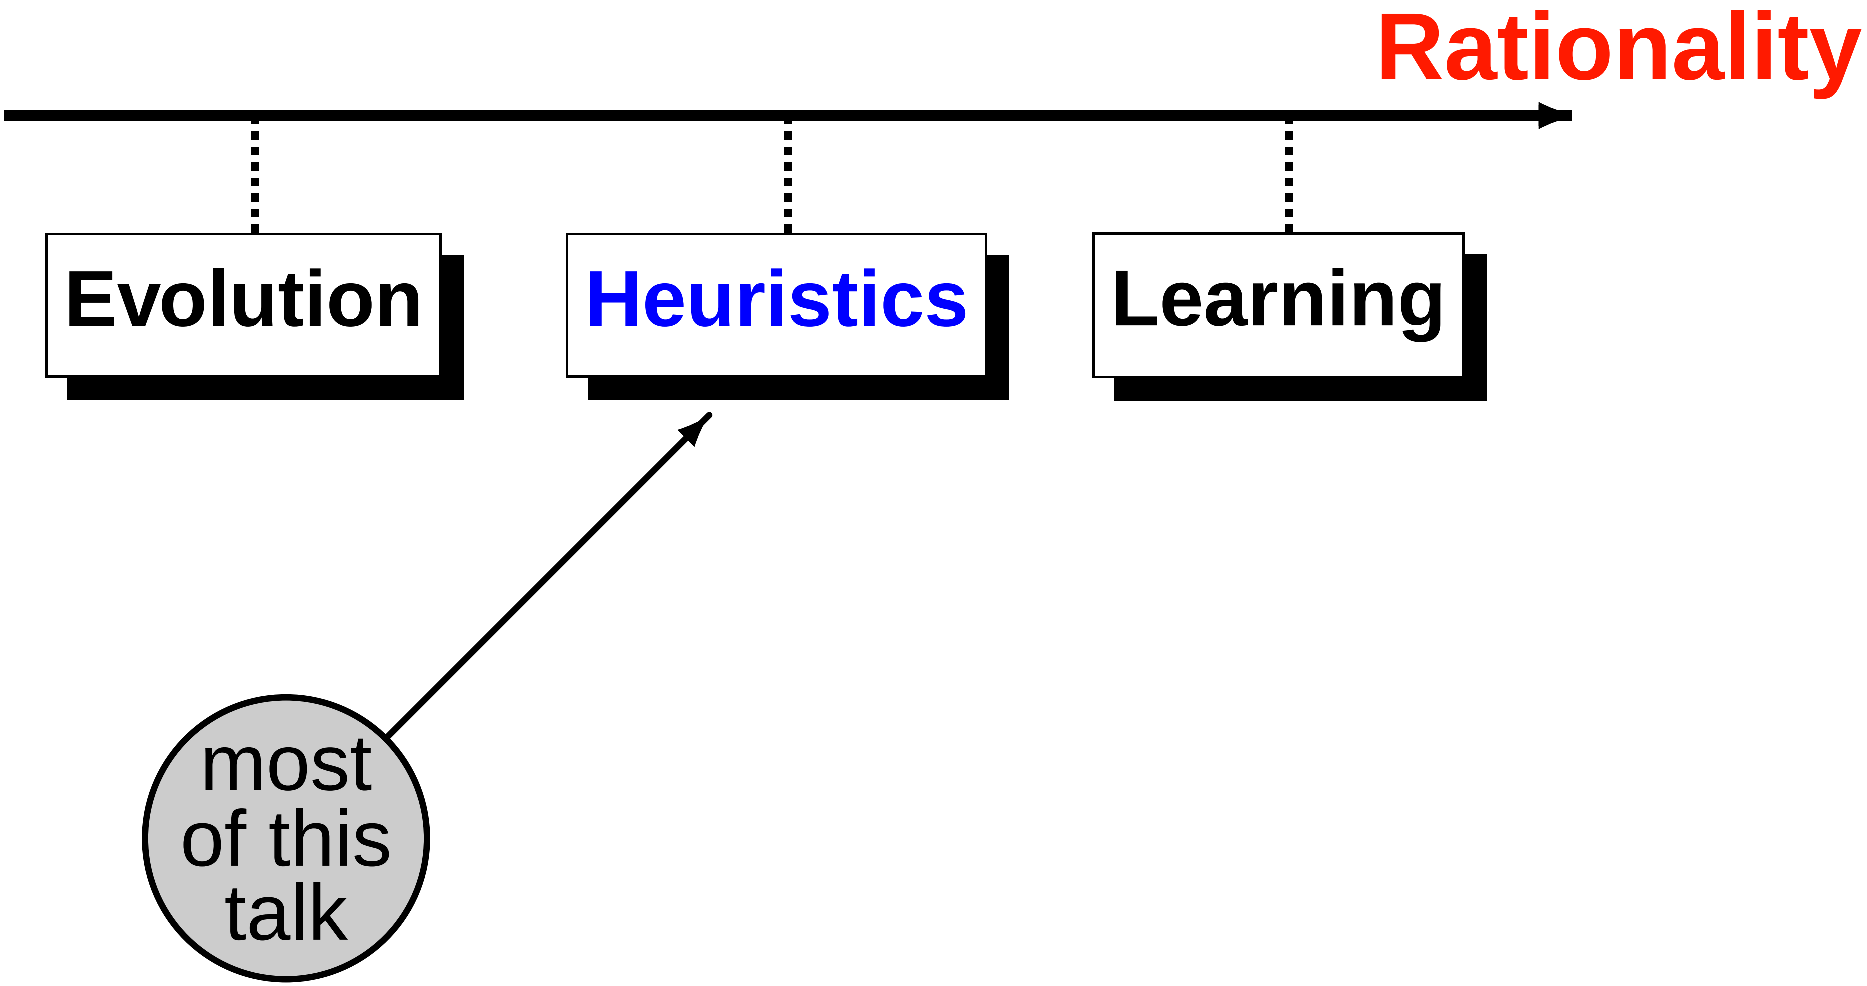

Can simple
adaptive heuristics
lead to
sophisticated
rational behavior ?

Game

N-person game in strategic (normal) form

- **Players**

$$i = 1, 2, ..., N$$

Game

N-person game in strategic (normal) form

- **Players**

$$i = 1, 2, ..., N$$

- For each player i: **Actions**

$$s^i \text{ in } S^i$$

Game

N-person game in strategic (normal) form

- **Players**

$$i = 1, 2, \ldots, N$$

- For each player i: **Actions**

$$s^i \text{ in } S^i$$

- For each player i: **Payoffs** (**utilities**)

$$u^i(s) \equiv u^i(s^1, s^2, \ldots, s^N)$$

Dynamics

- **Time**

$$t = 1, 2, \ldots$$

Dynamics

Dynamics

- **Time**

$$t = 1, 2, \ldots$$

- At time t each player i chooses an **action**

$$s_t^i \text{ in } S^i$$

Regret Matching

Advertisement

DON'T YOU FEEL A PANG OF REGRET?

47.15%
YIELD

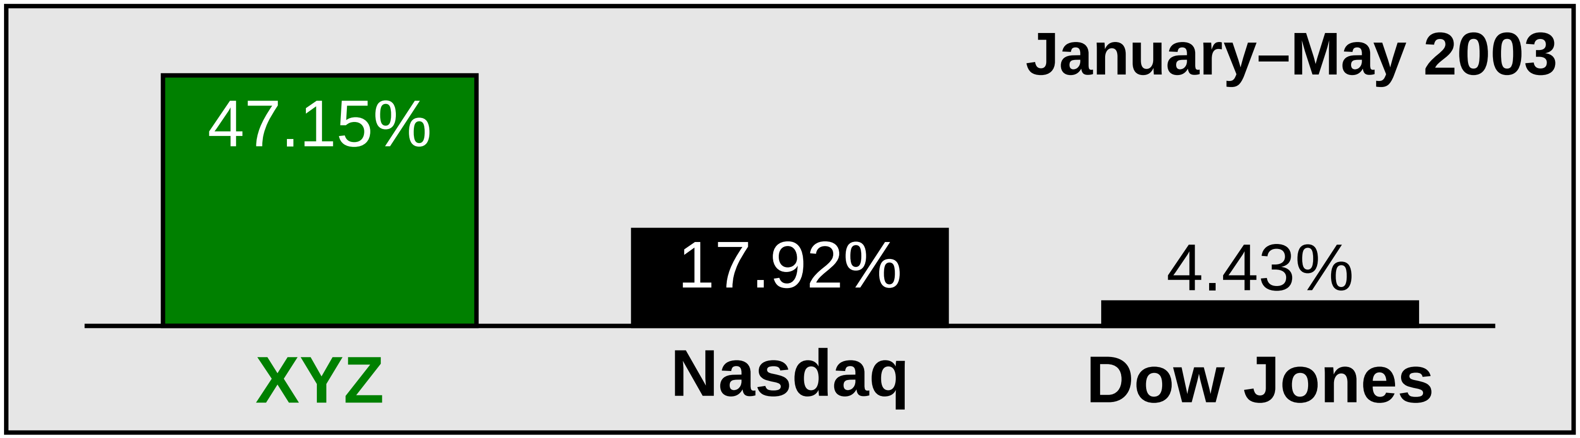

Don't wait! Ask your broker today

Haaretz – June 3, 2003

Regret Matching

REGRET MATCHING =

Switch next period to a different action with a probability that is proportional to the regret for that action

Regret Matching

REGRET MATCHING =
Switch next period to a different action
with a probability that is proportional
to the regret for that action

REGRET = increase in payoff
had such a change
always been made in the past

Regret

- U = average payoff up to now

Regret

- U = average payoff up to now
- $V(k)$ = average payoff if action k had been played instead of **the current action** j every time in the past that j was played

Regret

- U = average payoff up to now
- $V(k)$ = average payoff if action k had been played instead of **the current action** j every time in the past that j was played
- $R(k) = [V(k) - U]_+$ = **regret** for action k

Regret

- U = average payoff up to now
- $V(k)$ = average payoff if action k had been played instead of **the current action** j every time in the past that j was played
- $R(k) = [V(k) - U]_+$ = **regret** for action k

$$R(k) \equiv R_T^i(j \to k) =$$

$$\left[\frac{1}{T} \sum_{t \le T:\ s_t^i = j} \left(u^i(k, s_t^{-i}) - u^i(s_t)\right)\right]_+$$

Regret

- U = average payoff up to now
- $V(k)$ = average payoff if action k had been played instead of **the current action** j every time in the past that j was played
- $R(k) = [V(k) - U]_+$ = **regret** for action k

$$R(k) \equiv R_T^i(j \to k) =$$

$$\left[\frac{1}{T} \sum_{t \le T:\, s_t^i = j} \left(u^i(k, s_t^{-i}) - u^i(s_t) \right) \right]_+$$

Regret

- U = average payoff up to now
- $V(k)$ = average payoff if action k had been played instead of **the current action j** every time in the past that j was played
- $R(k) = [V(k) - U]_+$ = **regret** for action k

$$R(k) \equiv R_T^i(j \to k) =$$

$$\left[\frac{1}{T}\sum_{t \le T:\, s_t^i = j} \left(u^i(k, s_t^{-i}) - u^i(s_t)\right)\right]_+$$

Regret

- U = average payoff up to now
- $V(k)$ = average payoff if action k had been played instead of **the current action** j every time in the past that j was played
- $R(k) = [V(k) - U]_+$ = **regret** for action k

$$R(k) \equiv R^i_T(j \to k) =$$

$$\left[\frac{1}{T}\sum_{t \le T:\, s^i_t = j}\Big(u^i(k, s^{-i}_t) - u^i(s_t)\Big)\right]_+$$

Regret

- U = average payoff up to now
- $V(k)$ = average payoff if action k had been played instead of **the current action** j every time in the past that j was played
- $R(k) = [V(k) - U]_+$ = **regret** for action k

$$R(k) \equiv R_T^i(j \to k) =$$

$$\left[\frac{1}{T} \sum_{t \le T:\, s_t^i = j} \left(u^i(k, s_t^{-i}) - u^i(s_t) \right) \right]_+$$

Regret Matching

Next period play:

- **Switch** to action k with a probability that is proportional to the regret $R(k)$ (for $k \neq j$)

Regret Matching

Next period play:

- **Switch** to action k with a probability that is proportional to the regret $R(k)$ (for $k \neq j$)
- Play **the same** action j of last period with the remaining probability

Regret Matching

Next period play:

- **Switch** to action k with a probability that is proportional to the regret $R(k)$ (for $k \neq j$)
- Play **the same** action j of last period with the remaining probability

---

$$\sigma(k) \equiv \sigma^i_{T+1}(k) = cR(k), \quad \text{for } k \neq j$$

$$\sigma(j) \equiv \sigma^i_{T+1}(j) = 1 - \sum_{k \neq j} cR(k)$$

Regret Matching

Next period play:

- **Switch** to action k with a probability that is proportional to the regret $R(k)$ (for $k \neq j$)
- Play **the same** action j of last period with the remaining probability

---

$$\sigma(k) \equiv \sigma^i_{T+1}(k) = cR(k), \quad \text{for } k \neq j$$

$$\sigma(j) \equiv \sigma^i_{T+1}(j) = 1 - \sum_{k \neq j} cR(k)$$

- c = a fixed positive constant (so that the probability of not switching is > 0)

Regret Matching Theorem

Theorem

If all players play Regret Matching then the joint distribution of play converges to the set of CORRELATED EQUILIBRIA of the game

Joint Distribution of Play

Joint distribution of play z_T $=$
The relative frequencies that the N-tuples of actions have been played up to time T

Joint Distribution of Play

Joint distribution of play $z_T =$
The relative frequencies that the N-tuples of actions have been played up to time T

Joint Distribution of Play

Joint distribution of play z_T $=$
The relative frequencies that the N-tuples of actions have been played up to time T

*		

$$T = 1$$

Joint Distribution of Play

Joint distribution of play z_T $=$
The relative frequencies that the N-tuples of actions have been played up to time T

*		

$T = 1$

0	0	0
1	0	0

z_1

Joint Distribution of Play

Joint distribution of play z_T $=$
The relative frequencies that the N-tuples of actions have been played up to time T

		*
*		

$T = 2$

0	0	1/2
1/2	0	0

z_2

Joint Distribution of Play

Joint distribution of play z_T =
The relative frequencies that the N-tuples of actions have been played up to time T

		* *
*		

$T = 3$

0	0	$2/3$
$1/3$	0	0

z_3

Joint Distribution of Play

Joint distribution of play $z_T =$
The relative frequencies that the N-tuples of actions have been played up to time T

* * *		* *
*	* * *	*

$T = 10$

$3/10$	0	$2/10$
$1/10$	$3/10$	$1/10$

z_{10}

Joint Distribution of Play

Note 1: The fact that the players randomize **independently at each period does not imply** that the **joint distribution is independent** !

$* \quad * \quad *$		$* \quad *$
$*$	$* \quad * \quad *$	$*$

$T = 10$

$3/10$	0	$2/10$
$1/10$	$3/10$	$1/10$

z_{10}

Joint Distribution of Play

Note 1: The fact that the players randomize **independently at each period** **does not imply** that the **joint distribution is independent** !

Note 2: Players **observe** the **joint distribution** (the history of play)

Joint Distribution of Play

Note 1: The fact that the players randomize **independently at each period** **does not imply** that the **joint distribution is independent** !

Note 2: Players **observe** the **joint distribution** (the history of play)

Note 3: Players **react to** the **joint distribution** (patterns, “coincidences”, communication, signals, ...)

Correlated Equilibrium

A **Correlated Equilibrium** is a Nash equilibrium when the players receive payoff-irrelevant signals before playing the game (Aumann 1974)

Correlated Equilibrium

A **Correlated Equilibrium** is a Nash equilibrium when the players receive payoff-irrelevant signals before playing the game (Aumann 1974)

- **Examples**:

Correlated Equilibrium

A **Correlated Equilibrium** is a Nash equilibrium when the players receive payoff-irrelevant signals before playing the game (Aumann 1974)

- **Examples**:
  - Independent signals

Correlated Equilibrium

A **Correlated Equilibrium** is a Nash equilibrium when the players receive payoff-irrelevant signals before playing the game (Aumann 1974)

- **Examples**:
  - Independent signals $\Leftrightarrow$ Nash equilibrium

Correlated Equilibrium

A **Correlated Equilibrium** is a Nash equilibrium when the players receive payoff-irrelevant signals before playing the game (Aumann 1974)

- **Examples**:
  - Independent signals $\Leftrightarrow$ Nash equilibrium
  - Public signals ("sunspots")

Correlated Equilibrium

A **Correlated Equilibrium** is a Nash equilibrium when the players receive payoff-irrelevant signals before playing the game (Aumann 1974)

- **Examples**:
  - Independent signals $\Leftrightarrow$ Nash equilibrium
  - Public signals ("sunspots") $\Leftrightarrow$ convex combinations of Nash equilibria

Correlated Equilibrium

A **Correlated Equilibrium** is a Nash equilibrium when the players receive payoff-irrelevant signals before playing the game (Aumann 1974)

- **Examples**:
  - Independent signals $\Leftrightarrow$ Nash equilibrium
  - Public signals ("sunspots") $\Leftrightarrow$ convex combinations of Nash equilibria
  - Butterflies play the Chicken Game ("Speckled Wood" *Pararge aegeria*)

"Chicken" game

	LEAVE	STAY
LEAVE	$5, 5$	$3, 6$
STAY	$6, 3$	$0, 0$

Correlated Equilibria

"Chicken" game

	LEAVE	STAY
LEAVE	$5, 5$	$3, 6$
STAY	$6, 3$	$0, 0$

a **Nash equilibrium**

Correlated Equilibria

"Chicken" game

	LEAVE	STAY
LEAVE	$5, 5$	$3, 6$
STAY	$6, 3$	$0, 0$

another **Nash equilibrium**

Correlated Equilibria

"Chicken" game

	LEAVE	STAY
LEAVE	$5, 5$	$3, 6$
STAY	$6, 3$	$0, 0$

0	$1/2$
$1/2$	0

a **(publicly) correlated equilibrium**

Correlated Equilibria

"Chicken" game

	LEAVE	STAY
LEAVE	$5,5$	$3,6$
STAY	$6,3$	$0,0$

	L	S
L	$1/3$	$1/3$
S	$1/3$	0

another **correlated equilibrium**

- after signal L play LEAVE
- after signal S play STAY

Correlated Equilibrium

A **Correlated Equilibrium** is a Nash equilibrium when the players receive payoff-irrelevant signals before playing the game (Aumann 1974)

- **Examples**:
  - Independent signals $\Leftrightarrow$ Nash equilibrium
  - Public signals ("sunspots") $\Leftrightarrow$ convex combinations of Nash equilibria
  - Butterflies play the Chicken Game ("Speckled Wood" *Pararge aegeria*)

Correlated Equilibrium

A **Correlated Equilibrium** is a Nash equilibrium when the players receive payoff-irrelevant signals before playing the game (Aumann 1974)

- **Examples**:
  - Independent signals $\Leftrightarrow$ Nash equilibrium
  - Public signals ("sunspots") $\Leftrightarrow$ convex combinations of Nash equilibria
  - Butterflies play the Chicken Game ("Speckled Wood" *Pararge aegeria*)
  - Boston Celtics' front line

Correlated Equilibrium

- Signals (public, correlated) are **unavoidable**

Correlated Equilibrium

- Signals (public, correlated) are **unavoidable**
- **Bayesian Rationality** $\Leftrightarrow$ **Correlated Equilibrium** (Aumann 1987)

Correlated Equilibrium

- Signals (public, correlated) are **unavoidable**
- **Bayesian Rationality** $\Leftrightarrow$ **Correlated Equilibrium** (Aumann 1987)

A joint distribution z is a **correlated equilibrium**

$$\Leftrightarrow$$

$$\sum_{s^{-i}} u(j, s^{-i}) z(j, s^{-i}) \geq \sum_{s^{-i}} u(k, s^{-i}) z(j, s^{-i})$$

for all $i \in N$ and all $j, k \in S^i$

Regret Matching Theorem [recall]

Theorem

If all players play Regret Matching then the joint distribution of play converges to the set of CORRELATED EQUILIBRIA of the game

Regret Matching Theorem

- CE = set of **correlated equilibria**
- z_T = **joint distribution of play** up to time T

$$\text{distance}(z_T, \mathrm{CE}) \to 0 \quad \text{as } T \to \infty \quad \text{(a.s.)}$$

Regret Matching Theorem

- CE = set of **correlated equilibria**
- z_T = **joint distribution of play** up to time T

$$\boxed{\text{distance}(z_T, \mathrm{CE}) \longrightarrow 0 \quad \text{as} \quad T \rightarrow \infty \quad \text{(a.s.)}}$$

$$\Longleftrightarrow$$

z_T is **approximately** a **correlated equilibrium**
(or z_T is a **correlated approximate equilibrium**)
from some time on (for all large enough T)

Regret Matching Theorem

Proof

- z_T **is a correlated equilibrium**
  $\Leftrightarrow$ **there is no regret**:
  $R^i_T(j \to k) = 0$ for all players and all actions

Regret Matching Theorem

Proof

- z_T **is a correlated equilibrium**
  $\Leftrightarrow$ **there is no regret**:
  $R^i_T(j \to k) = 0$ for all players and all actions
- **Regret Matching**
  $\Rightarrow$ **all regrets converge to** 0
  (Proof: **Blackwell Approachability** for the vector of regrets + approximate eigenvector probabilities by transition probabilities)

Regret Matching Theorem

Proof

- z_T **is a correlated equilibrium**
  $\Leftrightarrow$ **there is no regret**:
  $R^i_T(j \to k) = 0$ for all players and all actions
- **Regret Matching**
  $\Rightarrow$ **all regrets converge to** 0
  (Proof: **Blackwell Approachability** for the vector of regrets + approximate eigenvector probabilities by transition probabilities)

Note: z_T converges to the *set* $\mathbf{CE}$, *not* to a point

Remarks

- Correlating device: the **history** of play

Remarks

- Correlating device: the **history** of play
- Other procedures leading to correlated equilibria:
  - Foster–Vohra 1997
    Calibrated Learning: best-reply to calibrated forecasts
  - Fudenberg–Levine 1999
    Conditional Smooth Fictitious Play
    Eigenvector strategy

Remarks

- Correlating device: the **history** of play
- Other procedures leading to correlated equilibria:
  - Foster–Vohra 1997
    Calibrated Learning: best-reply to calibrated forecasts
  - Fudenberg–Levine 1999
    Conditional Smooth Fictitious Play
    Eigenvector strategy

Not heuristics!

Behavioral Aspects

Behavioral aspects of **Regret Matching**:

Behavioral Aspects

Behavioral aspects of **Regret Matching**:

- **Commonly used** rules of behavior

Behavioral Aspects

Behavioral aspects of **Regret Matching**:

- **Commonly used** rules of behavior
  - Never change a winning team

Behavioral Aspects

Behavioral aspects of **Regret Matching**:

- **Commonly used** rules of behavior
  - Never change a winning team
  - The higher would have been the payoff from another action – the higher the tendency to switch to it

Behavioral Aspects

Behavioral aspects of **Regret Matching**:

- **Commonly used** rules of behavior
  - Never change a winning team
  - The higher would have been the payoff from another action – the higher the tendency to switch to it
  - Small probability of switching (the “status quo bias")

Behavioral Aspects

Behavioral aspects of **Regret Matching**:

- **Commonly used** rules of behavior
  - Never change a winning team
  - The higher would have been the payoff from another action – the higher the tendency to switch to it
  - Small probability of switching (the "status quo bias")
- **Stimulus-response, reinforcement**

Behavioral Aspects

Behavioral aspects of **Regret Matching**:

- **Commonly used** rules of behavior
  - Never change a winning team
  - The higher would have been the payoff from another action – the higher the tendency to switch to it
  - Small probability of switching (the “status quo bias")
- **Stimulus-response, reinforcement**
- **No beliefs** (defined directly on actions)
  No best-reply (better-reply ?)

Behavioral Aspects

- Similar to models of **learning, experimental and behavioral** economics:

Behavioral Aspects

- Similar to models of **learning, experimental and behavioral** economics:
  - Bush–Mosteller 1955
  - Erev–Roth 1995, 1998
  - Camerer–Ho 1997, 1998, 1999
  - ...

Behavioral Aspects

- Similar to models of **learning, experimental and behavioral** economics:
  - Bush–Mosteller 1955
  - Erev–Roth 1995, 1998
  - Camerer–Ho 1997, 1998, 1999
  - ...

Behavioral Aspects

- Similar to models of **learning, experimental and behavioral** economics:
  - Bush–Mosteller 1955
  - Erev–Roth 1995, 1998
  - Camerer–Ho 1997, 1998, 1999
  - ...
- N. Camille *et al*,
  "The Involvement of the **Orbitofrontal Cortex** in the Experience of **Regret**"
  Science May 2004 (304: 1167–1170)

PART III

Generalized Regret Matching

Questions

- How special is Regret Matching?

Questions

- How special is Regret Matching?
- Why does conditional smooth fictitious play work?

Questions

- How special is Regret Matching?
- Why does conditional smooth fictitious play work?
- Any connections?

Generalized Regret Matching

Regret Matching = Switching probabilities are **proportional** to the **regrets**: $\sigma(k) = cR(k)$

Generalized Regret Matching

Regret Matching = Switching probabilities are **proportional** to the **regrets**: $\boldsymbol{\sigma(k) = cR(k)}$

Generalized Regret Matching = Switching probabilities are a **function** of the **regrets**:

$$\boldsymbol{\sigma(k) = f(R(k))}$$

Generalized Regret Matching

Regret Matching = Switching probabilities are **proportional** to the **regrets**: $\sigma(k) = cR(k)$

Generalized Regret Matching = Switching probabilities are a **function** of the **regrets**:

$$\sigma(k) = f(R(k))$$

- f is a **sign-preserving** function:
  $f(0) = 0$, and $x > 0 \Rightarrow f(x) > 0$
- f is a Lipschitz **continuous** function

(in fact, much more general: $f_{k,j}$, potential)

Generalized Regret Matching

Theorem

If all players play Generalized Regret Matching then the joint distribution of play converges to the set of correlated equilibria of the game

Generalized Regret Matching

Theorem

If all players play
Generalized Regret Matching
then the joint distribution of play
converges to the set of
correlated equilibria of the game

Proof: “Universal” approachability strategies
+ Amotz Cahn, M.Sc. thesis, 2000

Special Cases

Play probabilities proportional to the **m-th power of the regrets**
($f(x) = cx^m$, for $m \geq 1$)

Special Cases

Play probabilities proportional to the
m-th power of the regrets
($f(x) = cx^m$, for $m \geq 1$)

- $m = 1$: **Regret Matching**

Special Cases

Play probabilities proportional to the
m-th power of the regrets
$(f(x) = cx^m$, for $m \geq 1)$

- $m = 1$: **Regret Matching**
- $m = \infty$: Positive probability only to actions with maximal regret

Special Cases

Play probabilities proportional to the
m**-th power of the regrets**
($f(x) = cx^m$, for $m \geq 1$)

- $m = 1$: **Regret Matching**
- $m = \infty$: Positive probability only to actions with maximal regret $\Leftrightarrow$ **Conditional Fictitious Play**

Special Cases

Play probabilities proportional to the
m-th power of the regrets
($f(x) = cx^m$, for $m \geq 1$)

- $m = 1$: **Regret Matching**
- $m = \infty$: Positive probability only to actions with maximal regret $\Leftrightarrow$ **Conditional Fictitious Play**
  - But: **Not continuous**

Special Cases

Play probabilities proportional to the
m-th power of the regrets
($f(x) = cx^m$, for $m \geq 1$)

- $m = 1$: **Regret Matching**
- $m = \infty$: Positive probability only to actions with maximal regret $\Leftrightarrow$ **Conditional Fictitious Play**
  - But: **Not continuous**
  - Therefore: **Smooth Conditional Fictitious Play**

PART IV

Unknown Game

Unknown Game

The case of the **"Unknown game"**:

- The player **knows only**
  - Its own set of actions
  - Its own past actions and payoffs

Unknown Game

The case of the **"Unknown game"**:

- The player **knows only**
  - Its own set of actions
  - Its own past actions and payoffs
- The player **does not know the game**
  (other players, actions, payoff functions, history of other players' actions and payoffs)

Proxy Regret

- **Unknown game** $\Rightarrow$ **Unknown regret**
  (The player does not know what the payoff would have been if he had played a different action k)

Proxy Regret

- **Unknown game** $\Rightarrow$ **Unknown regret** (The player does not know what the payoff would have been if he had played a different action k)
- **"Proxy Regret"** for k: Use the payoffs received when k has been actually played in the past

Proxy Regret

- **Unknown game $\Rightarrow$ Unknown regret**
  (The player does not know what the payoff would have been if he had played a different action k)
- **"Proxy Regret"** for k: Use the payoffs received when k has been actually played in the past

Theorem. If all players play strategies based on proxy regret, then the joint distribution of play converges to the set of correlated equilibria of the game

PART V

Uncoupled Dynamics

Nash Equilibrium

Question:

Adaptive heuristics $\longrightarrow$ **Nash equilibria**?

Nash Equilibrium

Question:

Adaptive heuristics $\longrightarrow$ **Nash equilibria**?

- In **SPECIAL** classes of games: **YES**

Nash Equilibrium

Question:

Adaptive heuristics $\longrightarrow$ **Nash equilibria**?

- In **SPECIAL** classes of games: **YES**

  Fictitious play, Regret-based, ...
  - Two-person zero-sum games
  - Two-person potential games
  - Supermodular games
  - ...

Nash Equilibrium

Question:

Adaptive heuristics $\longrightarrow$ **Nash equilibria**?

- In **SPECIAL** classes of games: **YES**

  Fictitious play, Regret-based, ...
  - Two-person zero-sum games
  - Two-person potential games
  - Supermodular games
  - ...
- In **GENERAL** games: **NO**

Nash Equilibrium

Question:

Adaptive heuristics $\longrightarrow$ **Nash equilibria**?

- In **SPECIAL** classes of games: **YES**

  Fictitious play, Regret-based, ...
  - Two-person zero-sum games
  - Two-person potential games
  - Supermodular games
  - ...
- In **GENERAL** games: **NO**

WHY ?

Uncoupled Dynamics

General dynamic for 2-person games:

$$\dot{x}(t) = F\ (\ x(t)\ ,\ y(t)\ ;\ u^1\ ,\ u^2\ )$$

$$\dot{y}(t) = G\ (\ x(t)\ ,\ y(t)\ ;\ u^1\ ,\ u^2\ )$$

$$x(t) \in \Delta(S^1), \qquad y(t) \in \Delta(S^2)$$

Uncoupled Dynamics

General dynamic for 2-person games:

$$\dot{x}(t) = F\,(\,x(t)\,,\ y(t)\,;\ u^1\,,\ u^2\,)$$

$$\dot{y}(t) = G\,(\,x(t)\,,\ y(t)\,;\ u^1,\ u^2\,)$$

Uncoupled dynamic:

$$\dot{x}(t) = F\,(\,x(t)\,,\ y(t)\,;\ u^1\qquad\quad)$$

$$\dot{y}(t) = G\,(\,x(t)\,,\ y(t)\,;\qquad\quad u^2\,)$$

$$x(t) \in \Delta(S^1), \qquad y(t) \in \Delta(S^2)$$

Uncoupled Dynamics

- **"Adaptive"** ("rational") dynamics

  (best-reply, better-reply, payoff-improving, monotonic, fictitious play, regret-based, replicator dynamics, ...)

  - are **uncoupled**

Uncoupled Dynamics

- **"Adaptive"** ("rational") dynamics

  (best-reply, better-reply, payoff-improving, monotonic, fictitious play, regret-based, replicator dynamics, ...)

  - are **uncoupled**

- **Uncoupledness** is a natural **informational** condition

Nash-Convergent Dynamics

- Consider a **family of games**, each having a **unique Nash equilibrium**
  (no “coordination problems”)

Nash-Convergent Dynamics

- Consider a **family of games**, each having a **unique Nash equilibrium** (no "coordination problems")
- A dynamic is **Nash-convergent** if it always converges to the unique Nash equilibrium
  - Regularity conditions: The unique Nash equilibrium is a stable rest-point of the dynamic

Impossibility

There exist no uncoupled dynamics which guarantee Nash convergence

Impossibility

There exist no uncoupled dynamics which guarantee Nash convergence

There are simple families of games
whose **unique Nash equilibrium**
is **unstable**
for every **uncoupled** dynamic

Impossibility

- **“Adaptive”** (“rational”) dynamics

  (best-reply, better-reply, payoff-improving, monotonic, fictitious play, regret-based, replicator dynamics, ...)

  - are **uncoupled**

Impossibility

- **“Adaptive”** (“rational”) dynamics

  (best-reply, better-reply, payoff-improving, monotonic, fictitious play, regret-based, replicator dynamics, ...)

  - are **uncoupled**
  - $\Longrightarrow$ cannot always converge to Nash equilibria

Nash vs Correlated

Correlated equilibria $\longleftrightarrow$ **Uncoupled** dynamics

Nash vs Correlated

Correlated equilibria $\longleftrightarrow$ **Uncoupled** dynamics

Nash equilibria $\longleftrightarrow$ **Coupled** dynamics

Nash vs Correlated

Correlated equilibria $\longleftrightarrow$ **Uncoupled** dynamics

Nash equilibria $\longleftrightarrow$ **Coupled** dynamics

"Law of Conservation of Coordination"

There must be coordination
either in the equilibrium concept
or in the dynamic

Summary

Where Do We Go From Here?

Where Do We Go From Here?

- Dynamics and equilibria
  - Which equilibria?
  - Which dynamics?

Where Do We Go From Here?

- Dynamics and equilibria
  - Which equilibria?
  - Which dynamics?
- Correlated equilibria: theory and practice
  - Coordination
  - Communication
  - Bounded complexity

Where Do We Go From Here?

- Dynamics and equilibria
  - Which equilibria?
  - Which dynamics?
- Correlated equilibria: theory and practice
  - Coordination
  - Communication
  - Bounded complexity
- Experiments, empirics $\longleftrightarrow$ Theory

Where Do We Go From Here?

- Dynamics and equilibria
  - Which equilibria?
  - Which dynamics?
- Correlated equilibria: theory and practice
  - Coordination
  - Communication
  - Bounded complexity
- Experiments, empirics $\longleftrightarrow$ Theory
- **Joint distribution** of play
  (instead of just the marginals)

Summary

Summary

- **There is** a simple **adaptive heuristic** always leading to **correlated equilibria**

Summary

- **There is** a simple **adaptive heuristic** always leading to **correlated equilibria**

(Regret Matching)

Summary

- **There is** a simple **adaptive heuristic** always leading to **correlated equilibria**

  (Regret Matching)

- There are **many** **adaptive heuristics** always leading to **correlated equilibria**

Summary

- **There is** a simple **adaptive heuristic** always leading to **correlated equilibria**

  (Regret Matching)

- There are **many** **adaptive heuristics** always leading to **correlated equilibria**

  (Generalized Regret Matching)

Summary

- **There is** a simple **adaptive heuristic** always leading to **correlated equilibria**
  (Regret Matching)

- There are **many** **adaptive heuristics** always leading to **correlated equilibria**
  (Generalized Regret Matching)

- There can be **no** **adaptive heuristics** always leading to **Nash equilibria**

Summary

- **There is** a simple **adaptive heuristic** always leading to **correlated equilibria**

  (Regret Matching)

- There are **many** **adaptive heuristics** always leading to **correlated equilibria**

  (Generalized Regret Matching)

- There can be **no** **adaptive heuristics** always leading to **Nash equilibria**

  (Uncoupledness)

Summary

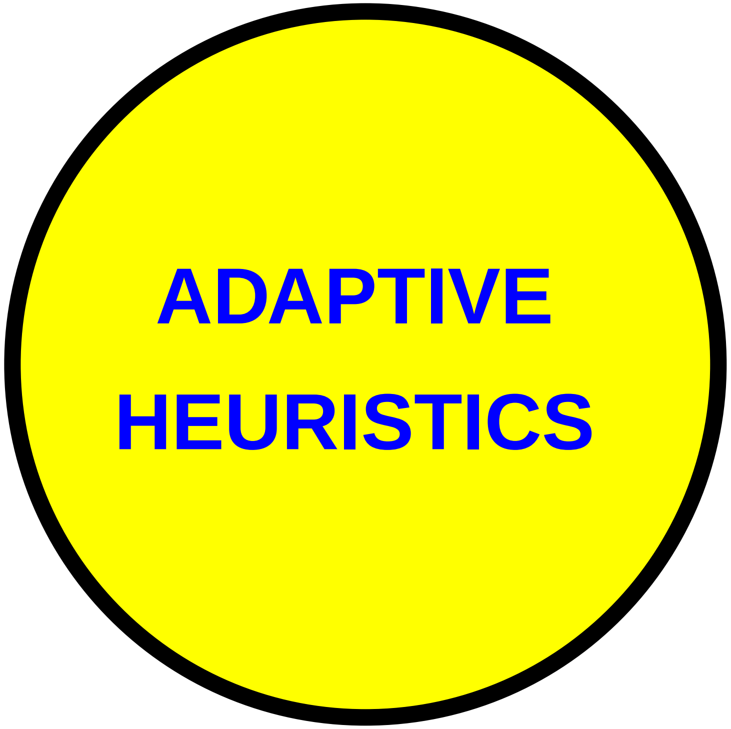

Summary

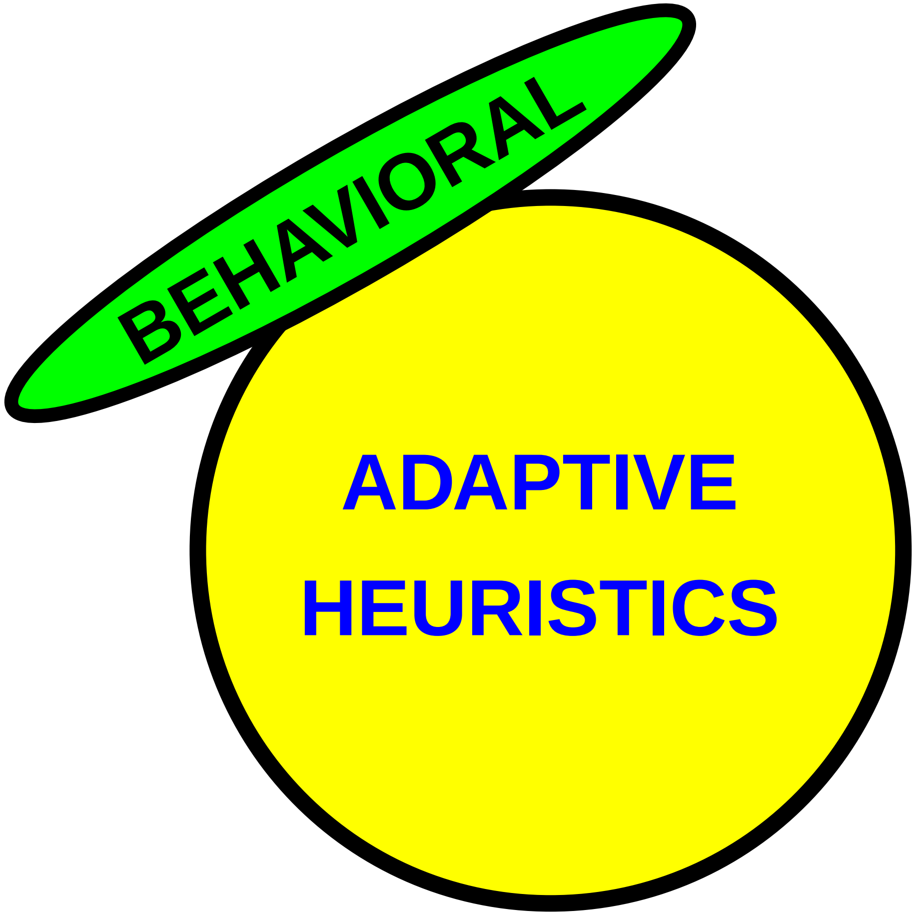

Summary

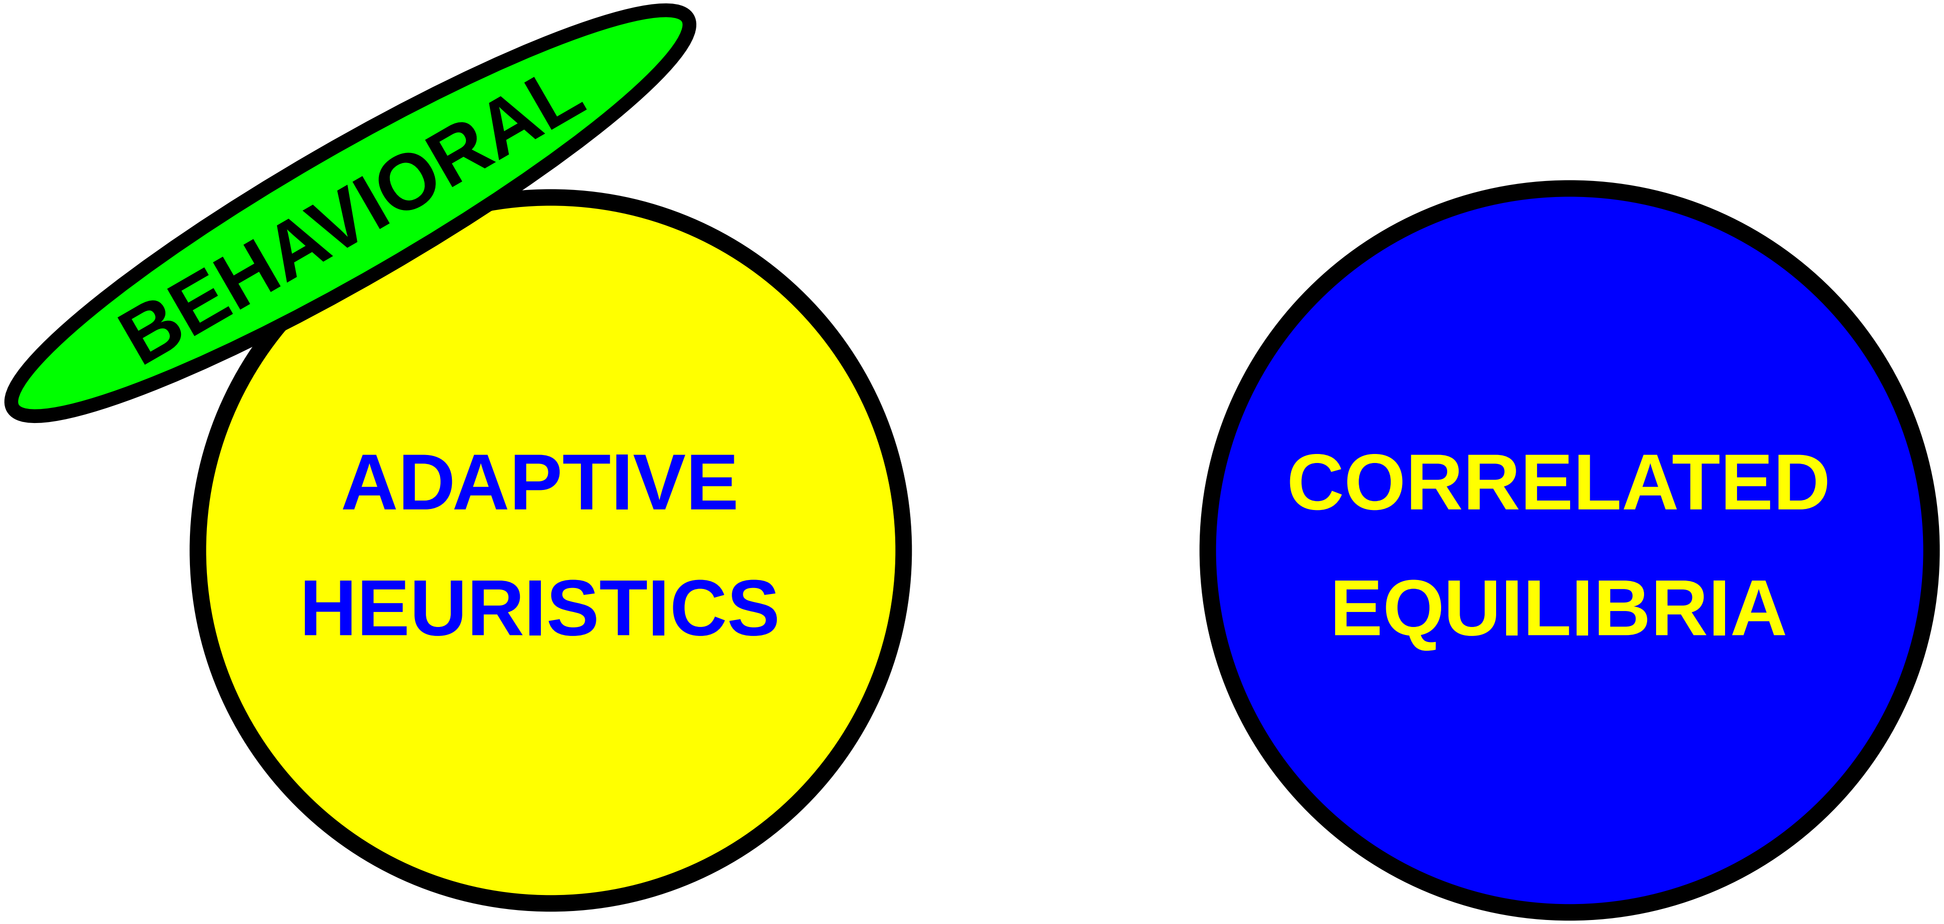

Summary

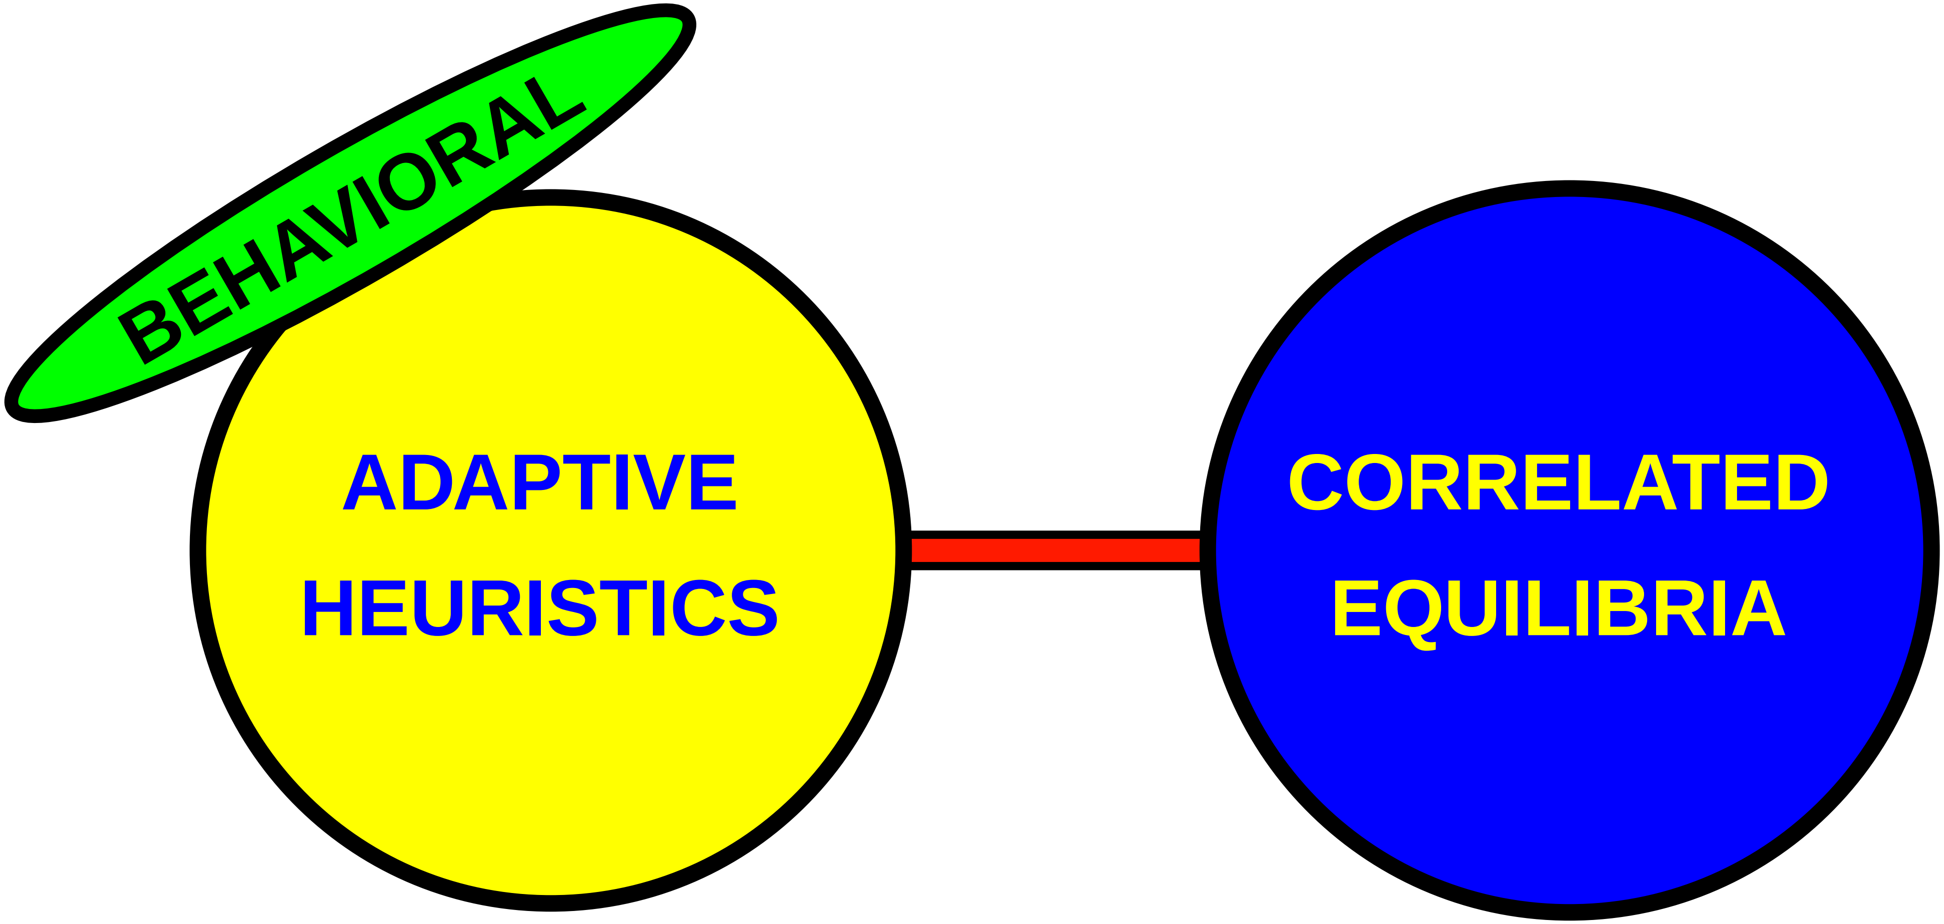

Regret Matching

Summary

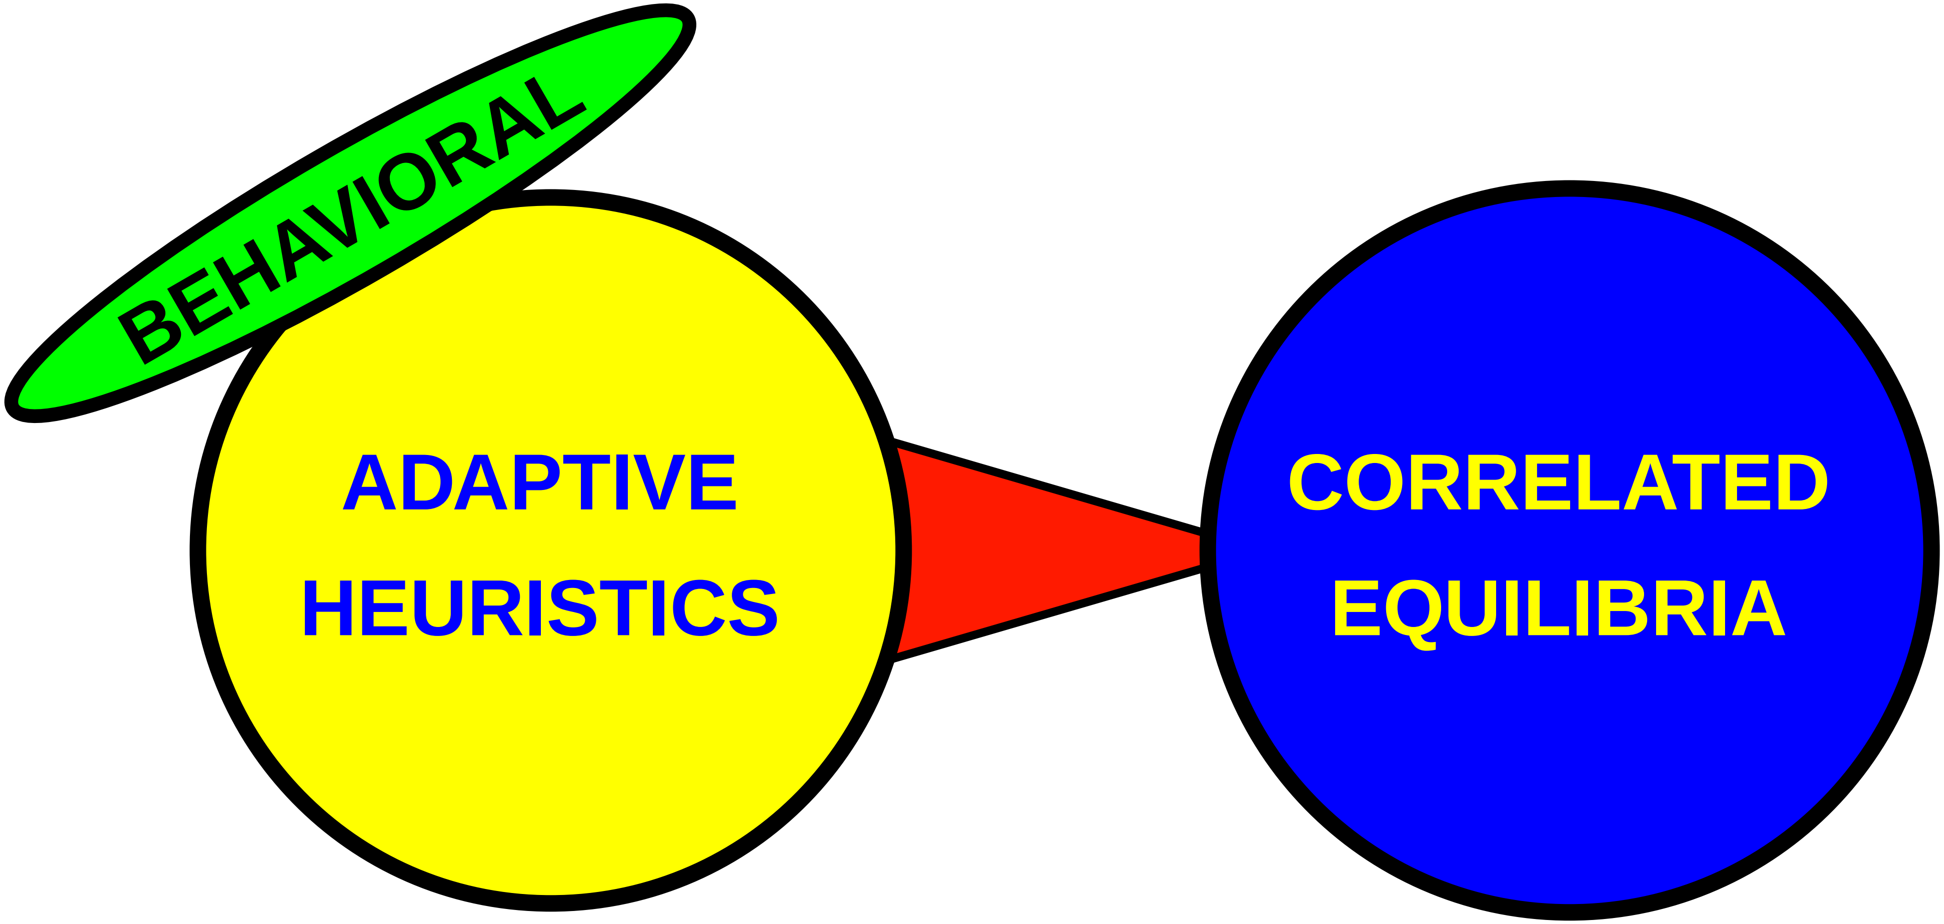

Generalized Regret Matching

Summary

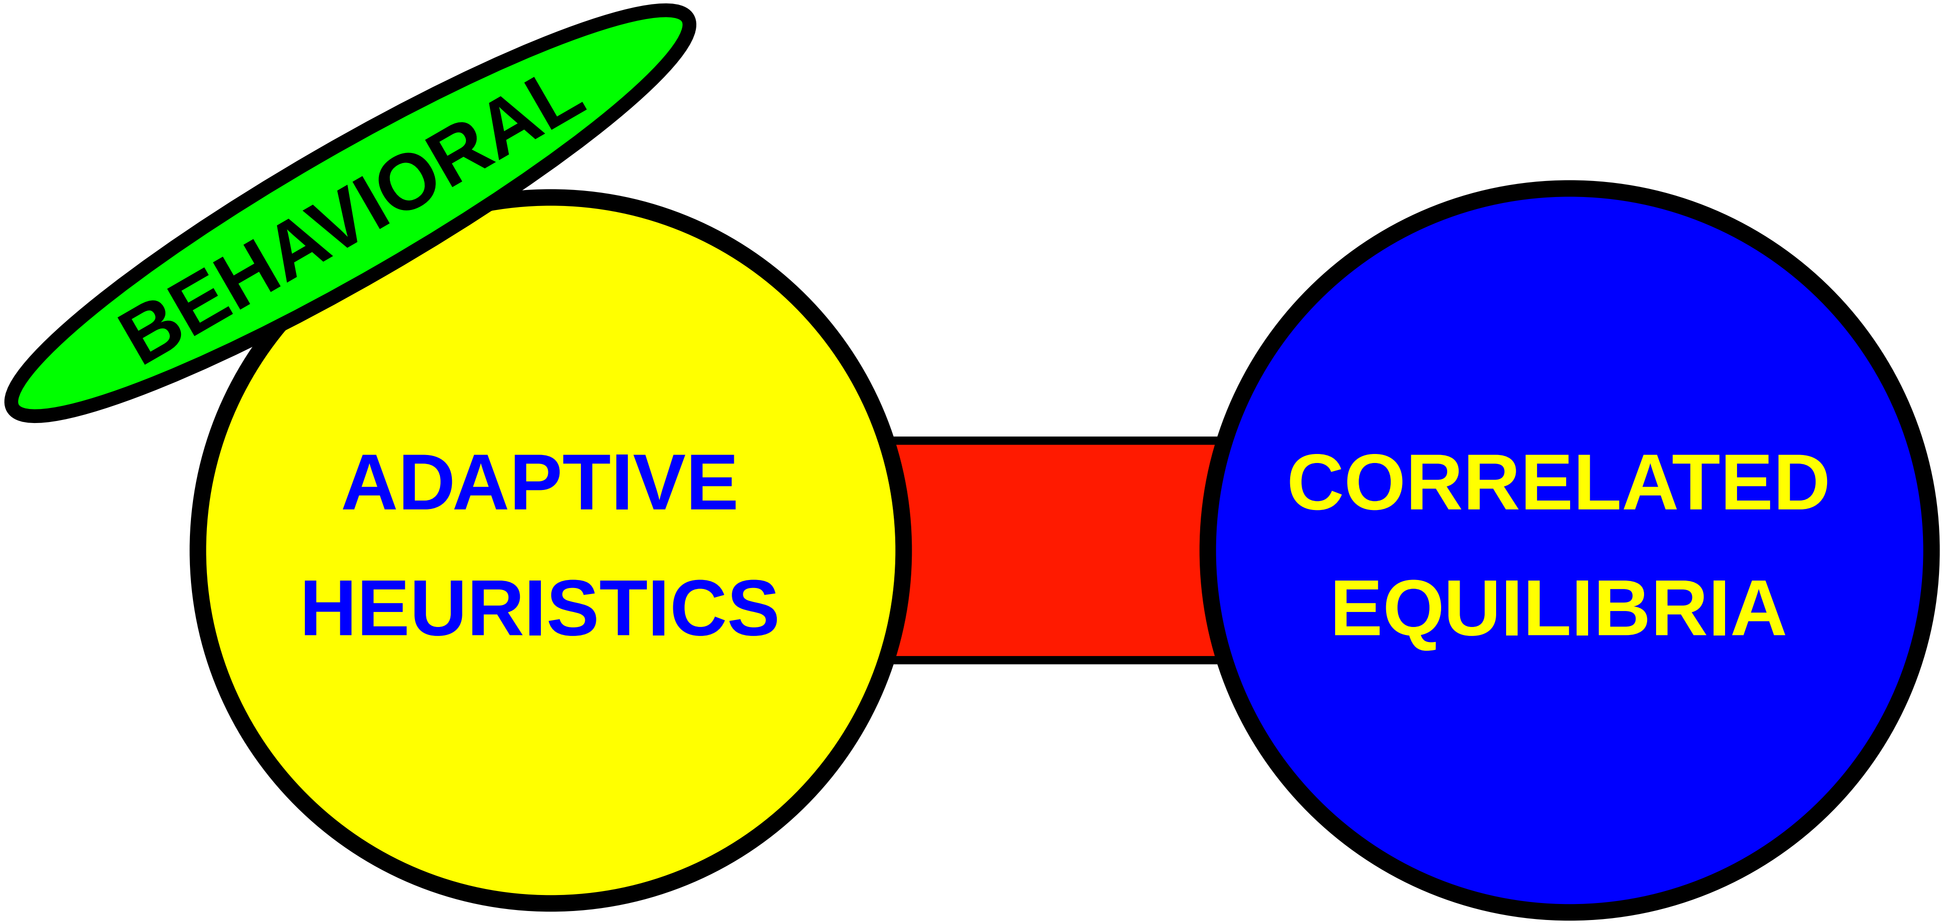

Uncoupledness

Question

Can simple
adaptive heuristics
lead to
sophisticated
rational behavior ?

Answer

Can simple
adaptive heuristics
lead to
sophisticated
rational behavior ?

YES !

Answer

Can simple
adaptive heuristics
lead to
sophisticated
rational behavior ?

YES !

in time ...

Summary – Macro

BEHAVIORAL

RATIONAL

Summary – Macro

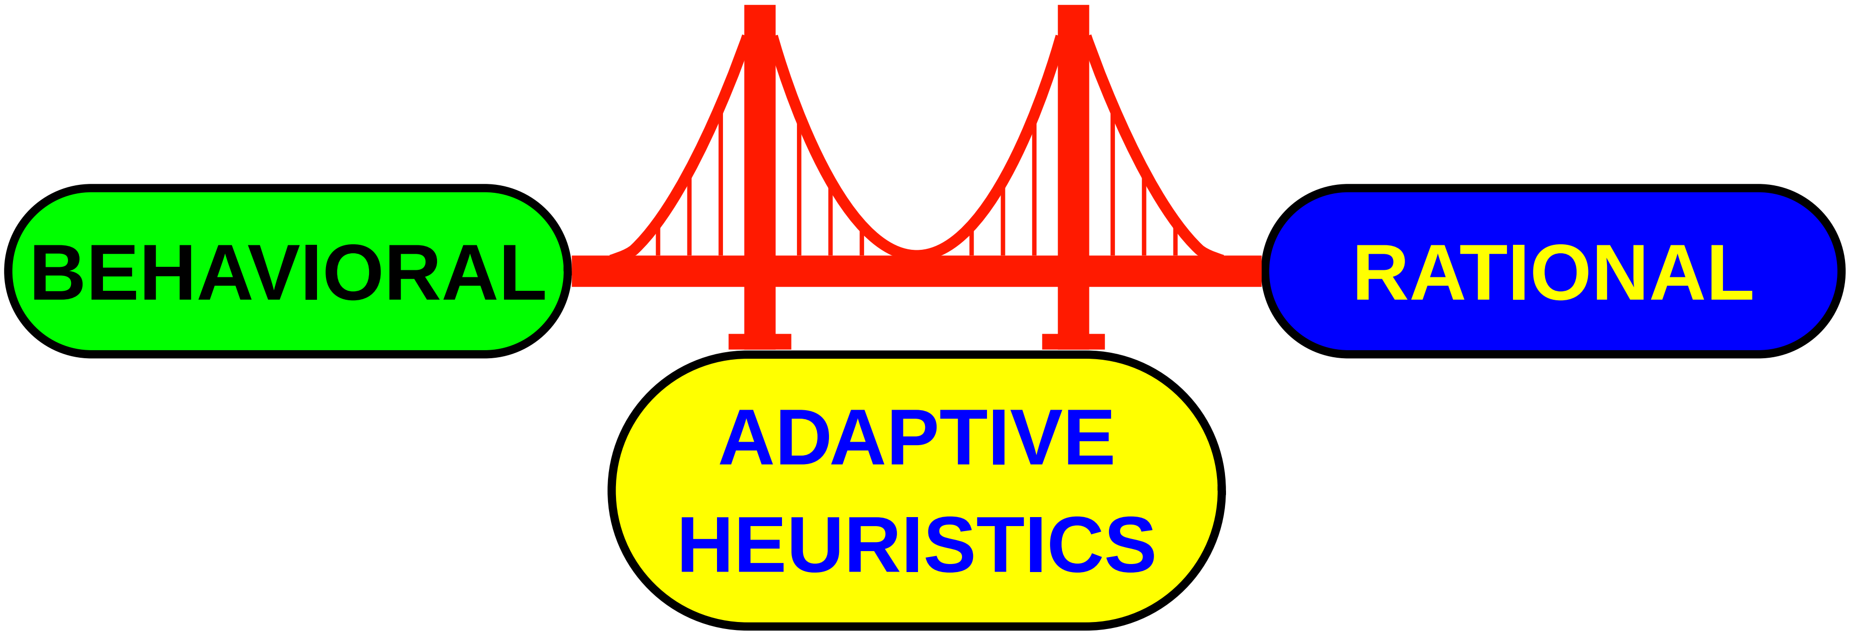

ADAPTIVE HEURISTICS

A Little Rationality Goes a Long Way

Title

ADAPTIVE HEURISTICS

(A Little Rationality Goes a Long Way)

Rationality Takes Time

Title

ADAPTIVE HEURISTICS

(A Little Rationality Goes a Long Way)

Rationality Takes Time

Regret ? ...

ADAPTIVE HEURISTICS

(A Little Rationality Goes a Long Way)

Rationality Takes Time

Regret ? ... No Regret !