Smart Transformations: or, The Evolution of Choice Principles

A Talk About Rationality

Paolo Galeazzi

ILLC, University of Amsterdam

Joint work with M. Franke

Outline

- model
- results
- comments
- conclusion

Outline

Introduction

- The model:
  - a population of players
  - a class of possible games
- Interpretation: different types in the population represent different reasoning/player types, or choice principles.

Game theory

Definition (Games)

A game $\mathbf{G}$ is a tuple $\mathbf{G} = \langle N, S_1, \ldots, S_n, \pi_1, \ldots, \pi_n \rangle$

As usual, $S := S_1 \times \ldots \times S_n$.

Definition (Solution concepts)

Given a game $\mathbf{G}$, a solution concept is a function F s.t. $F(\mathbf{G}) \subseteq S$

Nash equilibrium

Definition (Nash equilibrium)

A Nash equilibrium (NE) of a game **G** is a strategy profile $s \in S$ s.t. $\forall i \in N, \forall s_i' \in S_i,\ \ \pi_i(s) \geq \pi_i(s_i', s_{-i})$.

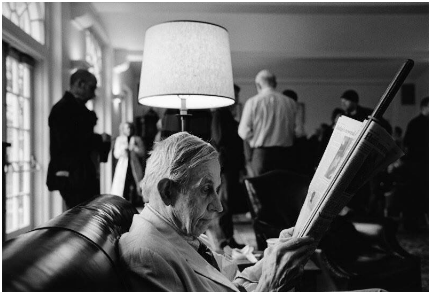

Regret

Definition (Regret minimization [4])

Given a game **G**, the regret of a strategy s_i for player i is defined as follows: $reg_i(s_i) := \max_{s_{-i}} \pi_i(s_i^*, s_{-i}) - \pi_i(s_i, s_{-i})$, where s_i^* is i's best reply to s_{-i}.
A strategy s_i' is a regret minimization strategy ($regmin_i$) of player i if $reg_i(s_i') = \min_{s_i} \max_{s_{-i}} \pi_i(s_i^*, s_{-i}) - \pi_i(s_i, s_{-i})$.

It will be useful to define a pairwise notion of regret.

Definition (Pairwise regret)

Given a game **G**, the regret of player i at profile (s_i, s_{-i}) is:
$reg_i(s_i, s_{-i}) := \pi_i(s_i^*, s_{-i}) - \pi_i(s_i, s_{-i})$.

the Travellers' dilemma

the Travellers' dilemma (TD)

- $N = Ann, Bob$
- $S_{Ann} = S_{Bob} = \{2, \ldots, 100\}$
- $\pi_i = \begin{cases} s_i \text{ if } s_i = s_{-i} \\ s_i + 2 \text{ if } s_i < s_{-i} \\ s_{-i} - 2 \text{ if } s_{-i} < s_i \end{cases}$

the Travellers' dilemma case

In the Travellers' dilemma there is only one NE: (2,2).

All the other solution concepts grounded on expected utility that solve the game (e.g., rationalizability, iterated elimination of weakly dominated strategies, ...) return the outcome (2,2), given perfect rationality of the players.

In the TD, iterated regret minimization gives the outcome (97,97).

However, in what follows we do not use the iterated version of regret minimization (epistemically not well-founded).

The pairwise notion defines a transformation of the game **G** into another game $\mathbf{G}^{reg}$.

G

		Bob *L*	Bob *R*
Ann	*U*	1,5	3,2
Ann	*D*	4,2	1,4

$\mathbf{G}^{reg}$

		Bob *L*	Bob *R*
Ann	*U*	3,0	0,3
Ann	*D*	0,2	2,0

The pairwise notion defines a transformation of the game **G** into another game $\mathbf{G}^{reg}$.

G

		Bob *L*	Bob *R*
Ann	*U*	1,5	3,2
Ann	*D*	4,2	1,4

$\mathbf{G}^{reg}$

		Bob *L*	Bob *R*
Ann	*U*	3,0	0,3
Ann	*D*	0,2	2,0

Inside regret minimization

Regret minimization may be seen as a security strategy.

Definition (Maximin)

Let us define $min_i(s_i) = \min_{s_{-i}} \pi_i(s_i, s_{-i})$. Then we can say that strategy s_i' is a maximin strategy of player i ($maximin_i$) if $min_i(s_i') = \max_{s_i} \min_{s_{-i}} \pi_i(s_i, s_{-i})$.

Remark: the maximin solution to the TD is also (2,2).

Regret minimization as maximin

$-\mathbf{G}^{reg}$

		Bob L	Bob R
Ann	U	-3,0	0,-3
Ann	D	0,-2	-2,0

Finding the regret minimizing strategies in $\mathbf{G}$ corresponds to finding the maximin strategies in $-\mathbf{G}^{reg}$.

Regret minimization as maximin

$-\mathbf{G}^{reg}$ Bob

		L	R
Ann	U	-3,0	0,-3
	D	0,-2	-2,0

Finding the regret minimizing strategies in **G** corresponds to finding the maximin strategies in $-\mathbf{G}^{reg}$.

Evolutionary analysis of choice principles: the intuition

So far we have seen that regret defines a specific transformation of the game. This gives an intuition on how we might understand different player types (or choice principles).

Evolution of preferences

- An *objective fitness game* **G** is played in the population and each type θ is a preference type, i.e., $\theta : S_1 \times ... \times S_n \to \mathbb{R}$.
- The regret type has the subjective preference/utility function $\theta^{reg} : S_1 \times ... \times S_n \to \mathbb{R}$ s.t. $\theta^{reg}(s_i, s_{-i}) = -reg_i(s_i, s_{-i})$.
- The π-type has the subjective utility function $\theta^{\pi} : S_1 \times ... \times S_n \to \mathbb{R}$ s.t. $\theta^{\pi}(s_i, s_{-i}) = \pi_i(s_i, s_{-i})$.
- Θ is the set of all preference types.
- Indirect evolutionary approach.

Meta-game

Tabella : A coordination game

G	I	II
I	1,1	0,0
II	0,0	2,2

Tabella : Chicken game

G'	I	II
I	0,0	-1,1
II	1,-1	-2,-2

Meta-game

Tabella : A coordination game

G	I	II
I	1,1	0,0
II	0,0	2,2

Tabella : Chicken game

G'	I	II
I	0,0	-1,1
II	1,-1	-2,-2

A bit more formal: player types or choice principles

- Given a class of fitness games $\mathcal{G}$, we call *player type* a function $\tau : \mathcal{G} \to \Theta$.
- Intuition: a player type is interpreted as a way of thinking across games, it is a thread, a red line that relates different preference types across different games.
- The regret type τ^{reg} is defined s.t. for any $\mathbf{G} \in \mathcal{G}$, $\tau^{reg}(\mathbf{G}) = \theta^{reg}$.
- Similarly, for any $\mathbf{G} \in \mathcal{G}$, $\tau^{\pi}(\mathbf{G}) = \theta^{\pi}$.
- Let $\mathcal{T}$ denote the set of player types.

Evolutionary analysis of choice principles: the model

- $\mathcal{G}$: class of symmetric two-by-two games
- Players in the population are player types (i.e., τ-functions)
- Radical uncertainty, i.e., no probabilistic beliefs (because of lack of information, lack of cognitive capabilities, ...)
- Because of their radical uncertainty we assume that players play a security strategy wrt their player type (*secure players*).
- At each time a game **G** is randomly selected from $\mathcal{G}$ and players are randomly matched to play **G**.
- The evolutionary fitness of each type is determined by the fitness games **G** via the objective fitness function π_G.
- The players' action choices across games depend on their player type τ.

Evolutionary analysis of choice principles: the model

- $\mathcal{G}$: class of symmetric two-by-two games
- Players in the population are player types (i.e., τ-functions)
- Radical uncertainty, i.e., no probabilistic beliefs (because of lack of information, lack of cognitive capabilities, ...)
- Because of their radical uncertainty we assume that players play a security strategy wrt their player type (*secure players*).
- At each time a game **G** is randomly selected from $\mathcal{G}$ and players are randomly matched to play **G**.
- The evolutionary fitness of each type is determined by the fitness games **G** via the objective fitness function π_G.
- The players' action choices across games depend on their player type τ.

Evolutionary analysis of choice principles: the model

- $\mathcal{G}$: class of symmetric two-by-two games
- Players in the population are player types (i.e., τ-functions)
- Radical uncertainty, i.e., no probabilistic beliefs (because of lack of information, lack of cognitive capabilities, ...)
- Because of their radical uncertainty we assume that players play a security strategy wrt their player type (*secure players*).
- At each time a game **G** is randomly selected from $\mathcal{G}$ and players are randomly matched to play **G**.
- The evolutionary fitness of each type is determined by the fitness games **G** via the objective fitness function π_G.
- The players' action choices across games depend on their player type τ.

Evolutionary analysis of choice principles: the model

- $\mathcal{G}$: class of symmetric two-by-two games
- Players in the population are player types (i.e., τ-functions)
- Radical uncertainty, i.e., no probabilistic beliefs (because of lack of information, lack of cognitive capabilities, ...)
- Because of their radical uncertainty we assume that players play a security strategy wrt their player type (*secure players*).
- At each time a game **G** is randomly selected from $\mathcal{G}$ and players are randomly matched to play **G**.
- The evolutionary fitness of each type is determined by the fitness games **G** via the objective fitness function π_G.
- The players' action choices across games depend on their player type τ.

Evolutionary analysis of choice principles: the model

- $\mathcal{G}$: class of symmetric two-by-two games
- Players in the population are player types (i.e., τ-functions)
- Radical uncertainty, i.e., no probabilistic beliefs (because of lack of information, lack of cognitive capabilities, ...)
- Because of their radical uncertainty we assume that players play a security strategy wrt their player type (*secure players*).
- At each time a game **G** is randomly selected from $\mathcal{G}$ and players are randomly matched to play **G**.
- The evolutionary fitness of each type is determined by the fitness games **G** via the objective fitness function π_G.
- The players' action choices across games depend on their player type τ.

Evolutionary analysis of choice principles: the model

- $\mathcal{G}$: class of symmetric two-by-two games
- Players in the population are player types (i.e., τ-functions)
- Radical uncertainty, i.e., no probabilistic beliefs (because of lack of information, lack of cognitive capabilities, ...)
- Because of their radical uncertainty we assume that players play a security strategy wrt their player type (*secure players*).
- At each time a game **G** is randomly selected from $\mathcal{G}$ and players are randomly matched to play **G**.
- The evolutionary fitness of each type is determined by the fitness games **G** via the objective fitness function π_G.
- The players' action choices across games depend on their player type τ.

Evolutionary analysis of choice principles: the model

- $\mathcal{G}$: class of symmetric two-by-two games
- Players in the population are player types (i.e., τ-functions)
- Radical uncertainty, i.e., no probabilistic beliefs (because of lack of information, lack of cognitive capabilities, ...)
- Because of their radical uncertainty we assume that players play a security strategy wrt their player type (*secure players*).
- At each time a game **G** is randomly selected from $\mathcal{G}$ and players are randomly matched to play **G**.
- The evolutionary fitness of each type is determined by the fitness games **G** via the objective fitness function π_G.
- The players' action choices across games depend on their player type τ.

Outline

Evolutionarily and neutrally stable strategies

Definition (ESS)

In evolutionary game theory a strategy s is said to be evolutionary stable (ESS) if for all the other strategies t:

1. $(s \text{ vs } s) > (t \text{ vs } s)$ or
2. $(s \text{ vs } s) = (t \text{ vs } s)$ and $(s \text{ vs } t) > (t \text{ vs } t)$

Definition (NSS)

A strategy s is said to be neutrally stable (NSS) if for all the other strategies t:

1. $(s \text{ vs } s) > (t \text{ vs } s)$ or
2. $(s \text{ vs } s) = (t \text{ vs } s)$ and $(s \text{ vs } t) \geq (t \text{ vs } t)$

Remark 1: the assumption of radical uncertainty is necessary to distinguish τ^{π} from τ^{reg}. I.e., with probabilistic beliefs maximizing expected utility would be the same as minimizing expected regret, for any **G**.

Formally,

$MaxExp(\text{-}\mathbf{G}^{reg}, \mu) = argmax_i \sum_j (\text{-}\mathbf{G}_{ij}^{reg} \cdot \mu_j)$ (by def)
$= argmax_i \sum_j (\mathbf{G}_{ij} - max_k(\mathbf{G}_{kj})) \cdot \mu_j$
$= argmax_i \sum_j (\mathbf{G}_{ij} \cdot \mu_j - max_k(\mathbf{G}_{kj}) \cdot \mu_j)$
$= argmax_i \sum_j (\mathbf{G}_{ij} \cdot \mu_j)$
$= MaxExp(\mathbf{G}, \mu)$.

Evolutionary analysis of choice principles

Just for fun, we define two other possible player types.

Definition (Competitive)

$\theta^{com} : S_1 \times S_2 \to \mathbb{R}$ s.t. $\theta^{com}(s_i, s_{-i}) = \pi_i(s_i, s_{-i}) - \pi_{-i}(s_i, s_{-i})$.

Definition (Altruistic)

$\theta^{alt} : S_1 \times S_2 \to \mathbb{R}$ s.t. $\theta^{alt}(s_i, s_{-i}) = \pi_i(s_i, s_{-i}) + \pi_{-i}(s_i, s_{-i})$.

Then we have: $\tau^{com}(\mathbf{G}) = \theta^{com}$ and $\tau^{alt}(\mathbf{G}) = \theta^{alt}$.

And the corresponding security strategies are
$\max_{s_i} \min_{s_{-i}} \theta^{com}(s_i, s_{-i})$ and $\max_{s_i} \min_{s_{-i}} \theta^{alt}(s_i, s_{-i})$.

Two-by-two symmetric games: matrix of average payoffs

	τ^{reg}	τ^{π}	τ^{com}	τ^{alt}
τ^{reg}	6.302	6.305	5.561	6.682
τ^{π}	6.128	6.131	5.750	6.331
τ^{com}	5.974	6.339	5.226	6.536
τ^{alt}	5.588	5.413	5.029	6.018

τ^{reg} turns out to be the only evolutionary stable (secure) type.

More in detail

More in detail:

- Coordination games: $a > c, d > b$
- Anti-coordination games: $a < c, d < b$
- Strong dominance games: $a > c, d < b$ or $a < c, d > b$
- Weak dominance games: aut $a - c = 0$ aut $d - b = 0$
- Indifferent games: $a - c = 0$ and $d - b = 0$

Average payoff matrix for coordination games

	τ^{reg}	τ^{π}	τ^{com}	τ^{alt}
τ^{reg}	7.676	5.914	5.794	6.436
τ^{π}	6.430	6.367	6.341	5.905
τ^{com}	6.462	6.434	6.346	5.920
τ^{alt}	5.441	4.281	4.237	5.833

Average payoff matrix for anti-coordination games

	τ^{reg}	τ^{π}	τ^{com}	τ^{alt}
τ^{reg}	4.751	6.202	4.713	6.141
τ^{π}	5.275	5.288	5.405	5.314
τ^{com}	5.342	7.074	3.426	7.016
τ^{alt}	5.207	5.303	5.387	5.270

Average payoff matrix for strong dominance games

	τ^{reg}	τ^{π}	τ^{com}	τ^{alt}
τ^{reg}	7.212	7.212	6.704	7.749
τ^{π}	7.212	7.212	6.704	7.749
τ^{com}	6.909	6.909	6.365	7.446
τ^{alt}	6.588	6.588	6.080	7.197

Average payoff matrix for weak dominance games

	τ^{reg}	τ^{π}	τ^{com}	τ^{alt}
τ^{reg}	6.402	6.342	5.187	7.168
τ^{π}	6.155	6.099	5.120	6.735
τ^{com}	5.959	6.157	4.706	6.787
τ^{alt}	5.710	5.492	4.532	6.486

Average payoff matrix for indifferent games

	τ^{reg}	τ^{π}	τ^{com}	τ^{alt}
τ^{reg}	5.833	5.75	3.541	7.166
τ^{π}	5.833	5.75	3.541	7.166
τ^{com}	5.833	5.75	3.541	7.166
τ^{alt}	5.833	5.75	3.541	7.166

These results from simulations can also be proven analytically.

Proposition 1

Fix $\mathcal{T} = \{\tau^{\pi}, \tau^{reg}, \tau^{com}, \tau^{alt}\}$ and $\mathcal{G}$ the class of symmetric 2×2 games with i.i.d. sampled payoffs from a finite or compact and convex set of values. Then τ^{reg} is the only *evolutionary stable (secure) type* in the population.

With probabilistic beliefs

Suppose now that players can have both probabilistic beliefs (in some situations) and radical uncertainty (in other situations). Moreover, suppose that the information (beliefs) available is symmetric. Then, from Proposition 1 and Remark 1 it follows that:

Corollary 1

Fix $\mathcal{T} = \{\tau^{\pi}, \tau^{reg}\}$ and $\mathcal{G}$ the class of symmetric 2×2 games with i.i.d. sampled payoffs from a finite or compact and convex set of values. Then τ^{reg} is the only *evolutionary stable type* in the population.

Digression: Ellsberg's urn, ambiguity, and security strategies

	R	Y	B
f_R	100	0	0
f_Y	0	100	0
f_{RB}	100	0	100
f_{YB}	0	100	100

90 marbles in the urn, 30 red marbles.

Preferences that satisfy this pattern are called *uncertainty,* or *ambiguity, averse*.
The DM is not representable by a probability measure as in Savage's theorem.

Digression: Ellsberg's urn, ambiguity, and security strategies

	R	Y	B
f_R	100	0	0
f_Y	0	100	0
f_{RB}	100	0	100
f_{YB}	0	100	100

90 marbles in the urn, 30 red marbles.
The standard pattern is: $f_R > f_Y$ and $f_{RB} < f_{YB}$.
Preferences that satisfy this pattern are called *uncertainty, or ambiguity, averse.*
The DM is not representable by a probability measure as in Savage's theorem.

Digression: Ellsberg's urn, ambiguity, and security strategies

	R	Y	B
f_R	100	0	0
f_Y	0	100	0
f_{RB}	100	0	100
f_{YB}	0	100	100

90 marbles in the urn, 30 red marbles.
The standard pattern is: $f_R \succ f_Y$ and $f_{RB} \prec f_{YB}$.
Preferences that satisfy this pattern are called *uncertainty*, or *ambiguity, averse*.
The DM is not representable by a probability measure as in Savage's theorem.

Digression: Ellsberg's urn, ambiguity, and security strategies

	R	Y	B
f_R	100	0	0
f_Y	0	100	0
f_{RB}	100	0	100
f_{YB}	0	100	100

90 marbles in the urn, 30 red marbles.
The standard pattern is: $f_R \succ f_Y$ and $f_{RB} \prec f_{YB}$.
Preferences that satisfy this pattern are called *uncertainty*, or *ambiguity, averse*.
The DM is not representable by a probability measure as in Savage's theorem.

Digression: Ellsberg's urn, ambiguity, and security strategies

It has been shown [3] that ambiguity averse preferences are representable by a closed convex set P of probability measures together with the maximin rule: *Maxmin Expected Utility with Non-unique Prior*.

Working example: $P = [(\frac{1}{3}, \frac{2}{3}, 0), (\frac{1}{3}, 0, \frac{2}{3})]$.

$min_{\mu \in P} \sum_{s \in S} u(f_R(s)) \cdot \mu(s) = \frac{100}{3}$
$min_{\mu \in P} \sum_{s \in S} u(f_Y(s)) \cdot \mu(s) = 0$

$min_{\mu \in P} \sum_{s \in S} u(f_{RB}(s)) \cdot \mu(s) = \frac{100}{3}$
$min_{\mu \in P} \sum_{s \in S} u(f_{YB}(s)) \cdot \mu(s) = \frac{200}{3}$

Digression: Ellsberg's urn, ambiguity, and security strategies

It has been shown [3] that ambiguity averse preferences are representable by a closed convex set P of probability measures together with the maximin rule: *Maxmin Expected Utility with Non-unique Prior*.

Working example: $P = [(\frac{1}{3}, \frac{2}{3}, 0), (\frac{1}{3}, 0, \frac{2}{3})]$.

$$min_{\mu \in P} \sum_{s \in S} u(f_R(s)) \cdot \mu(s) = \frac{100}{3}$$
$$min_{\mu \in P} \sum_{s \in S} u(f_Y(s)) \cdot \mu(s) = 0$$
$$min_{\mu \in P} \sum_{s \in S} u(f_{RB}(s)) \cdot \mu(s) = \frac{100}{3}$$
$$min_{\mu \in P} \sum_{s \in S} u(f_{YB}(s)) \cdot \mu(s) = \frac{200}{3}$$

Digression: Ellsberg's urn, ambiguity, and security strategies

It has been shown [3] that ambiguity averse preferences are representable by a closed convex set P of probability measures together with the maximin rule: *Maxmin Expected Utility with Non-unique Prior*.

Working example: $P = [(\frac{1}{3}, \frac{2}{3}, 0), (\frac{1}{3}, 0, \frac{2}{3})]$.

$$min_{\mu \in P} \sum_{s \in S} u(f_R(s)) \cdot \mu(s) = \frac{100}{3}$$
$$min_{\mu \in P} \sum_{s \in S} u(f_Y(s)) \cdot \mu(s) = 0$$

$$min_{\mu \in P} \sum_{s \in S} u(f_{RB}(s)) \cdot \mu(s) = \frac{100}{3}$$
$$min_{\mu \in P} \sum_{s \in S} u(f_{YB}(s)) \cdot \mu(s) = \frac{200}{3}$$

Games with (some) uncertainty

In our framework, secure players are maxmin expected utility players (à la Gilboa-Schmeidler) with $P = \Delta(S_{-i})$.

Remark 2. τ^{reg} (secure) players outperform τ^{π} (secure) players for any non-singleton set P.

Outline

On the model: Theoretical biology and behavioral ecology

- Behavioral gambit in theoretical biology: the evolutionary models put all the focus on the expressed behavior, and neglect the underlying mechanisms that generate that behavior.

  > *[...] we should expect animals to have evolved a set of psychological mechanisms which enable them to perform well on average across a range of different circumstances. [2]*

- Psychological mechanisms and subjective conceptualizations as the phenotype under selection.
- Our model:
  - retains the indirect evolutionary approach
  - works on a class of possible games

On the result: What is rational?

This talk takes the cue from a paper by J.Halpern and R.Pass [4] introducing iterated regret minimization as a new solution concept.

The notion of regret had already been introduced in decision theory as an alternative to expected utility theory.

Regret theory

> *[...] we shall challenge the idea that the conventional axioms constitute the only acceptable basis for rational choice under uncertainty. We shall argue that it is no less rational to act in accordance with regret theory, and that conventional expected utility theory therefore represents an unnecessarily restrictive notion of rationality.* [5]

On the result: What is rational?

- Rationality is a descriptive notion: maxEU fails.
  (Btw, regret can explain and describe observed behavior better than maxEU [4],[5].)
- Rationality is a normative notion: apparently, EU-maximizers can be outperformed by regret minimizers.

On the result: What is rational?

- Rationality is a descriptive notion: maxEU fails.
  (Btw, regret can explain and describe observed behavior better than maxEU [4],[5].)
- Rationality is a normative notion: apparently, EU-maximizers can be outperformed by regret minimizers.

Outline

Conclusion

- We found *subjective* transformations of the games (”player types”) that can perform better than the maximization of the *objective* payoff.

Evolution of preferences

These papers [i.e., [6];[1]] highlight the dependence of indirect evolutionary models on observable preferences, posing a challenge to the indirect evolutionary approach that can be met only by allowing the question of preference observability to be endogenously determined within the model. [8]

The indirect evolutionary approach with unobservable preferences gives us an alternative description of the evolutionary process, one that is perhaps less reminiscent of biological determinism, but leads to no new results. [7]

Future work

- more general $\mathcal{G}$
- different/asymmetric beliefs and theory of mind
- other player types
- adding *learning/adaptation process*

Thanks for your attention.

References:

- J.C. Ely and O. Ylankaya.
  Nash Equilibrium and the Evolution of Preferences.
  Journal of Economic Theory, 97:255–272, 2001.
- Tim W. Fawcett, Steven Hamblin, and Luc-Alain Giraldeaub.
  Exposing the behavioral gambit: the evolution of learning and decision rules.
  Behavioral Ecology, 24(1):2–11, 2013.
- I. Gilboa and D. Schmeidler.
  Maxmin Expected Utility with Non-unique Prior.
  Journal of Mathematical Economics, 18:141–153, 1989.
- J. Halpern and R. Pass.
  Iterated Regret Minimization: a New Solution Concept.
  Games and Economic Behavior, 74:1:194–207, 2012.
- G. Loomes and R. Sugden.

Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty.
The Economic Journal, 92:368:805–824, 1982.

E. Ok and F. Vega Redondo.
On the Evolution of Individualistic Preferences: An Incomplete Information Scenario.
Journal of Economic Theory, 97:231–254, 2001.

A. Robson and L. Samuelson.
The Evolutionary Foundations of Preferences.
In J. Benhabib, M. Jackson, and A. Bisin, editors, *Handbook of Social Economics, 1A*. North-Holland, 2011.

L. Samuelson.
Introduction to the Evolution of Preferences.
Journal of Economic Theory, 97:225–230, 2001.