

Working Papers


Institute of
Mathematical
Economics

402

October 2008

Persistent Ideologies in an Evolutionary Setting

Lars P. Koch



IMW · Bielefeld University
Postfach 100131
33501 Bielefeld · Germany



email: imw@wiwi.uni-bielefeld.de
<http://www.wiwi.uni-bielefeld.de/~imw/Papers/showpaper.php?402>
ISSN: 0931-6558

Persistent Ideologies in an Evolutionary Setting

Lars P. Koch*

Institute for Mathematical Economics
Bielefeld University

October 14, 2008

Abstract

We analyse finite two player games in which agents maximize given arbitrary private payoffs which we call ideologies. We define an equilibrium concept and prove existence. Based on this setup, a monotone evolutionary dynamic governs the distribution of ideologies within the population. For any finite 2 player normal form game we show that there is an open set of ideologies being not equivalent to the objective payoffs that is not selected against by evolutionary monotonic dynamics. If the game has a strict equilibrium set, we show stability of non-equivalent ideologies. We illustrate these results for generic 2×2 -games.

JEL codes: C72, C73, D82

Keywords: *Dynamic Stability, Evolution of Preferences, Imitation Dynamics*

1 Introduction and Related Literature

Suppose rational agents meet to interact in a strategic situation and none of them has verifiable information about the payoffs of the game. Instead, each agent has been socialized by some ideology that specifies a private payoff matrix. Agents then play the game as to maximize the private payoffs. To give an example, consider a measure taken to prevent terrorist attacks in airplanes: passengers are not allowed to bring flasks with their hand luggage that can

*correspondence: lkoch@wiwi.uni-bielefeld.de

contain more than 100 ml of liquid. Since this measure was taken until now, no catastrophe was caused by a large amount of liquid in the hand luggage. The reason could be that the measure is indeed successful. Another reasoning could be that 100 ml of liquid in a single bottle is no threat to an airplane, in other words that the measure is useless. The passenger is not capable to verify the utility gained by the measure but is obliged to *believe* in a payoff. In this paper the view is taken that there is no objective knowledge on this matter, because the real payoffs are not observable, maybe not even available, at least not verifiable. Instead, interacting agents have private payoffs. To stay in the example, given the measure some agents might feel more secure and therefore more comfortable while airborne; other agents might be annoyed due to the security restrictions and additional waiting times. We do not claim that a particular position is *true* or *false* (actually, as the modeller of the problem we *assume* that certain payoffs are true). We aim to show that even if some perception of reality is wrong, there might be no selection against this perception based on the true payoffs. The key result from this paper is that for any (Nash-) equilibrium in situations of strategic interaction, there exist non-equivalent views regarding the payoffs – ideologies – that cannot be driven out by evolutionary dynamics.

In this paper we model an ideology as a bias that shifts the subjective payoffs. We allow for an unbounded continuum of biases, that is for any two player game with finite strategy space there may exist several groups of agents at the same time, each group believing to face a different payoff matrix. Once in a while, agents change their ideology, however we assume they do so boundedly rational and model the adaptation of ideologies with an evolutionary monotonic process: ideologies that result in behavior that produces relatively high *objective* payoffs spread faster within the population. Hence we assume that the “currency” of evolution are the objective payoffs. The gist of this paper is that *even if evolution* cares for an exclusive specification of payoffs, interacting agents may permanently maximize alternative private payoffs even though these payoffs are subject to evolution. This paper relies heavily on

Sandholm (2001) who constructs a similar model for symmetric 2×2 -games in which the payoffs of some *action* are biased, we build upon this model by allowing for biases for *outcomes* of the game. We focus less on the dynamic that yields equilibrium behavior, but assume equilibrium behavior explicitly. Hereby tractability is gained and one can study general finite $m \times n$ -normal form games. Symmetric 2-player games with a unique pure strategy equilibrium were also studied by Heifetz, Shannon, and Spiegel (2007) within a similar framework. Ok and Vega-Redondo (2001) construct a model in which agents either know the true payoffs of a symmetric game or maximize some alternative symmetric utility function, while they allow for the presence of only one such alternative utility function. For symmetric two player games they show global stability of states in which all agents maximize the true payoffs. We show that for any finite game there is an open set of ideologies that are not equivalent to the objective payoffs and that survive in an evolutionary scenario. This includes the special case of symmetric 2×2 -games. The main difference of our model to Ok and Vega-Redondo (2001) is that we allow for asymmetric payoffs and where we study symmetric games we do not assume behavioral distinguishability. Ok and Vega-Redondo (2001) state that behavioral distinguishability is necessary for the fact that agents who maximize the evolutionary relevant payoffs have an evolutionary advantage. However, this property is not sufficient: for some asymmetric games we show survival of non-individualistic preferences, even if they are behaviorally distinguishable. We do so right away for a simple example, to provide a flavor of the model.

Consider the well known matching pennies game. Suppose that the row population has agents of three types. One third maximizes the true payoffs. Another third has a bias towards “(head, head)”, the last third is biased towards “(tail, tail)”. The biases are depicted below:

	H	T
H	(-1,1)	(1,-1)
T	(1,-1)	(-1,1)

objective payoffs

	H	T
H	(-1+ 1 ,1)	(1,-1)
T	(1,-1)	(-1,1)

bias for $\frac{1}{3}$ of row

	H	T
H	(-1,1)	(1,-1)
T	(1,-1)	(-1+ 1 ,1)

bias for $\frac{1}{3}$ of row

To keep the example simple, the column population is completely unbiased. Each agent knows the sum of the true payoffs and her bias. Suppose the column population mixes equally between heads and tails. Then, head-biased agents optimally choose head, tail-biased agents optimally choose tail while unbiased row agents are just indifferent (assume they mix equally between heads and tails). In this situation each agent maximizes subjective payoffs and no agent has an incentive to deviate. Any agent gets the same real payoffs on average. Moreover, any type has different equilibrium behavior.

Ely and Yilankaya (2001) study which set of outcomes is supported by stable preferences in normal form games. As they rely on static concepts to infer stability properties, we define a dynamic process explicitly to analyze stability issues. In their model, any set of alternative payoff specifications that is robust to exogenous shocks implies equilibrium behavior that produces a probability distribution over the set of Nash equilibria. Ely and Yilankaya (2001) do not have results for the case in which there is zero mass on ideologies that represent the true preferences. Our results concern exactly the case that excludes positive mass on true preferences. The static equilibrium concept offered in this paper does not imply Nash behavior, while those equilibria that satisfy a stability property –dynamically stable equilibrium– do induce a Nash equilibrium.

2 Model

2.1 The Stage Game

We consider two infinite populations of agents from which at each point of continuous time, a pair (one agent from each population) is uniformly and independently drawn to play a finite two player normal form game. An agent of

population i chooses an element s^i of the finite strategy set S^i with cardinality n^i . For convenience, define $n = n_i \cdot n_j$. The objective payoffs are represented by the $n^i \times n^j$ matrix U^i , where $U^i(s^i, s^j)$ denotes i 's objective utility from the outcome generated by (s^i, s^j) . We assume that agents are unable to verify the objective payoffs. For each population i , each agent is characterized by a $n_i \times n_j$ -matrix of parameters $\theta^i = \{\theta^i(s)\}_{s \in S}$ that induces the subjective payoffs. An agent θ^i believes to play a game that specifies her payoffs as $u^i(s, \theta^i) = U^i(s^i, s^j) + \theta^i(s^i, s^j)$, $s^i \in S^i, s^j \in S^j$. Let u be extended to the space of mixed strategies and define $u^i((s^i, \sigma^j), \theta^i) = \sum_{s^j \in S^j} \sigma^j(s^j)(U^i(s^i, s^j) + \theta^i(s^i, s^j))$.

We will show that depending on the game, some ideologies that are non-equivalent to the objective payoffs won't have a long-run evolutionary disadvantage.

We assume that for each population i types are distributed among the agents by some atomless density $f_i : \mathbb{R}^n \rightarrow \mathbb{R}_+$ with cumulative distribution function $F_i : \mathbb{R}^n \rightarrow [0, 1]$. To be precise, "atomless" here means $\int_{\mathcal{H}} f_i(\theta^i) d\theta^i = 0$ for any hyperplane \mathcal{H} in \mathbb{R}^n . The types are independently distributed across populations, that is $F(\Theta^i, \Theta^j) = F_i(\Theta^i) \cdot F_j(\Theta^j) \forall \Theta^i, \Theta^j \subset \mathbb{R}^n$.

2.2 Equivalent Ideologies

An agent with parameters $\theta^i(s) = 0 \forall s \in S$ plays the game given the objective payoffs. These payoffs represent the evolutionary relevant payoffs.

An ideology θ^i is *equivalent* to the objective payoffs U^i , if $\text{sign}\{U^i(s) - U^i(s')\} = \text{sign}\{u^i(s, \theta^i) - u^i(s', \theta^i)\} \forall s, s' \in S$. Whenever θ^i is equivalent to the zero bias for $i = 1, 2$, we write $\theta \Leftrightarrow 0$.

The definition of equivalence is relatively weak. One could also demand a stronger version, for example that there must exist a positive affine function $h : \mathbb{R} \rightarrow \mathbb{R}$ such that $u^i(s) = h(U^i(s)) \forall s$. This, however, would only weaken the main result of this paper. The main result concerns a situation in which

no agent holds an ideology that is equivalent to the zero bias. Of course, this includes situations in which agents do not hold “strongly equivalent” ideologies.

2.3 Optimal Choices

Suppose agents of population i believe σ^j to be the probability distribution over the strategies of population $j \neq i$. An agent of type θ^i chooses s^i iff $s^i \in \arg \max u^i((\tilde{s}^i, \sigma^j), \theta^i)$. We assume that in equilibrium, for all i , all agents of population i hold the same belief σ^j . This assumption is critical to the concept we develop in the following and therefore also to the results of this paper. Nevertheless, note that it is not as ad hoc as it might seem at first sight. The beliefs σ^j are not about an object that is unobservable – like preferences, ideologies or tastes. The belief σ^j is defined as a probability measure on actual actions taken by opponents. Since all agents observe this information, homogeneous beliefs seem natural.

Define the set of types $\Theta_{s^i}(\sigma^j)$ that choose strategy s^i given σ^j as

$$\Theta_{s^i}(\sigma^j) := \{\theta^i \in \mathbb{R}^n \mid s^i \in \arg \max u^i((\tilde{s}^i, \sigma^j), \theta^i)\}$$

Any belief σ^j defines such a set $\Theta_{s^i}(\sigma^j)$ uniquely. We implicitly assume that agents maximize myopically.

Lemma 1

For $i, j = 1, 2, i \neq j$, given any $\sigma^j \in \Delta(S^j)$ there is no $s^i \in S^i$ such that $\Theta_{s^i}(\sigma^j) = \emptyset$.

Proof

Fix a constant $K > \max_{s^i \in S^i} |U^i(s^i, \sigma^j)|$. The max is well defined because S is finite. Then the bias θ^i with $\theta^i(s^i, \sigma^j) = K$ and $\theta^i(\tilde{s}^i, \sigma^j) = -K$ for $\tilde{s}^i \neq s^i$ imposes the choice of s^i , since $u^i((s^i, \sigma^j), \theta^i) = U^i(s^i, \sigma^j) + K > U^i(\tilde{s}^i, \sigma^j) - K = u^i((\tilde{s}^i, \sigma^j), \theta^i) \forall \tilde{s}^i \neq s^i$. There is an open set $\tilde{\Theta}^i$ around θ^i such that s^i is optimal for all $\theta^i \in \tilde{\Theta}^i$. Therefore, $\emptyset \neq \tilde{\Theta}^i \subset \Theta_{s^i}(\sigma^j)$ and $\Theta_{s^i}(\sigma^j)$ is non-empty. \square

The lemma states that given σ^j , for any s^i , the objective payoffs $U_i(\cdot, \cdot)$ can be biased so strongly towards s^i such that the agent maximizes her subjective payoffs by choosing s^i .

If all agents of population j share the same belief σ^i on the behavior of population i , the true probability of an agent from population j choosing s^j is:

$$F_j(\Theta_{s^j}(\sigma^i)) := \int_{\Theta_{s^j}(\sigma^i)} F_j(d\theta^j)$$

2.4 Equilibrium

We assume that in equilibrium j 's belief that i plays s^i , $\sigma^i(s^i)$, and the true probability that i plays s^i , $F_i(\Theta_{s^i}(\sigma^j))$, coincide: $\sigma^i(s^i) = F_i(\Theta_{s^i}(\sigma^j)) \forall s^i \in S^i, i = 1, 2$.

Definition Equilibrium Set

A collection $\{\Theta_{s^i}\}_{s^i}, i = 1, 2$ is an *equilibrium set* given cdf F , if

- E_1 $\Theta_{s^i} \subset \Theta_{s^i}(\sigma^j) \forall s^i \in S^i, i = 1, 2$ and
- E_2 $\sigma^i(s^i) = F_i(\Theta_{s^i}) \forall s^i \in S^i, i = 1, 2$ and
- E_3 $F_i(\cup_{s^i \in S^i} \Theta_{s^i}) = 1, i = 1, 2$.

In short we denote an equilibrium set by $\{\Theta_s\}$ given F .

E_1 demands that agents optimally choose their strategy given their beliefs and subjective payoffs. E_2 demands that beliefs on behavior equal actual probabilities. E_3 is a technical condition which becomes necessary as we assumed non-atomic probabilities. In equilibrium all agents in the same population hold the same beliefs over the behavior of the other population and maximize given their subjective (and heterogeneous) payoffs. Further beliefs equal true probabilities. The definition states that equilibrium set depends on a specific distribution F . Since F is arbitrary, the induced equilibrium distribution on the set of outcomes S does not at all need to be a Nash equilibrium distribution. However, any Nash equilibrium can be represented by an Equilibrium

set by an appropriate choice of F . This is implied by Theorem 1 below.

Lemma 2

The function $G_{sj}(\sigma^i) = F_j(\Theta_{sj}(\sigma^i))$ is continuous in $\sigma^i \forall s^j, i$.

Proof: see Appendix.

Proposition Existence

For any F , here exists at least one pair of beliefs $\sigma = (\sigma^i, \sigma^j)$ such that

$$\sigma^i(s^i) = F_i(\Theta_{si}(\sigma^j)) \forall s^i \in S^i, i \neq j = 1, 2 .$$

Proof:

To prove the result we apply Brouwer’s Fixed Point Theorem. Any continuous function $G : \Sigma \rightarrow \Sigma$ has a fixed point $b^* \in \Sigma$ such that $G(b^*) = b^*$, where $\Sigma = \Delta(S^1) \times \Delta(S^2)$. $\sigma = (\sigma^1, \sigma^2) \in \Sigma$. The function $G(\sigma) = (G^1(\sigma^2), G^2(\sigma^1))$ maps from Σ to Σ . Lemma 2 above states that $G_{si}^i(\sigma^j)$ is continuous in σ^j for all s^i . Therefore, all requirements are met and we can apply Brouwer’s Fixed Point Theorem to proof existence. \square

Note that the set $\{\hat{\Theta}_s\}_s$, $\hat{\Theta}_s = \Theta_s(\hat{\sigma})$ is an equilibrium set given F . Note further that $\Theta_s(\sigma)$ is not uhc, see the appendix for details. There may be many $\sigma = (\sigma^1, \sigma^2)$ satisfying the equilibrium conditions given the same distribution F . The assumption that all i -agents hold the same belief σ^j implies that they manage to coordinate on one of potentially many equilibria. However, we do not select any equilibrium; the results require only that the agents coordinate on some, but not on which equilibrium.¹

2.5 Evolutionary Dynamics

Given some cdf F with density f and some equilibrium set $\{\Theta_s\}$, agents of population i with type $\theta^i \in \Theta_{si}$ receive objective payoffs

¹Implicitly, we assume that behavior adjusts with infinite speed. Sandholm (2001) shows in a symmetric 2x2-games setting, that the infinite speed dynamics can be expressed as the limit of finite speed dynamics.

$U^i(s^i, F_j) = \sum_{s^j} U^i(s^i, s^j) F_j(\Theta_{s^j})$. Once in a while agents adopt different ideologies. We assume that this dynamic process can be captured by the deterministic differential equation

$$\dot{F}_i(\Theta^i) = \int_{\Theta^i} g_i(\theta^i, F, \{\Theta_s\}) F_i(d\theta^i), \quad \forall \Theta^i \subset \mathbb{R}^n, \quad i = 1, 2 \quad (1)$$

where $g_i(\theta^i, F, \{\Theta_s\})$ is the growth rate of the marginal density f_i at θ^i given cdf F and some equilibrium set $\{\Theta_s\}$.² We require $g_i(\theta^i, F, \{\Theta_s\})$ to maintain the probability property of (F_1, F_2) , that is $\int_{\mathbb{R}^n} g_i(\theta^i, F, \{\Theta_s\}) f_i(\theta^i) d\theta^i = 0$, $i = 1, 2$ and we require $g_i(\theta^i, F, \{\Theta_s\})$ to be Lipschitz continuous in F . Then, as Oechssler and Riedel (2001) show, for any initial distribution $F(0)$ a solution $F(t)$ exists. If $g_i(\theta^i, F, \{\Theta_s\}) = U^i(s^i, F_j) - \sum_{\hat{s}^i} F_i(\Theta_{\hat{s}^i}) \cdot U^i(\hat{s}^i, F_j)$, the dynamic is the well known replicator dynamic. We assume that $g(\theta, F, \{\Theta_s\})$ fulfills the less demanding requirement of monotonicity³, that is

$$\mathbf{M} \quad U^i(s^i, F_j) > (=) U^i(\hat{s}^i, F_j) \Rightarrow g_i(\theta^i, F, \{\Theta_s\}) > (=) g_i(\hat{\theta}^i, F, \{\Theta_s\})$$

for $\theta^i \in \Theta_{s^i}$ and $\hat{\theta}^i \in \Theta_{\hat{s}^i}$. Suppose strategy s^i yields higher objective payoffs than strategy \hat{s}^i in a particular equilibrium. Payoff monotonicity implies that the mass of agents that follow an ideology that recommends the choice of s^i grows. In the context of Björnerstedt and Weibull (1996), this means that “unsuccessful” agents who choose \hat{s}^i revise their ideology more often than agents who choose s^i .

Two agents that have biases $\theta^i, \tilde{\theta}^i \in \Theta_{s^i} \subset \Theta_{s^i}(\sigma^j)$ both optimally choose s^i and gain the same objective payoffs $U_i(s^i, \sigma^j)$. For payoff monotonic dynamics this implies that $g_i(\theta^i, F, \{\Theta_s\}) = g_i(\tilde{\theta}^i, F, \{\Theta_s\})$. We will sometimes denote this growthrate by $g_{s^i}(F, \{\Theta_s\})$.

In this model, agents face a twofold decision problem. Firstly, agents adapt an ideology. We assume that this is done boundedly rational, for example

²Björnerstedt and Weibull (1996) give an interpretation of imitating agents that produces such a dynamic on the space of strategies.

³Samuelson and Zhang (1992)

by imitation, maybe even subconsciously. In particular, agents do not apply “backwards induction” by anticipating a particular equilibrium and thereby induced payoffs. Secondly, agents act rationally by maximizing their subjective payoffs induced by their ideology.

3 Results

In this section we collect the results of this paper. The first two statements concern the unperturbed model, while the second two consider stability and therefore a dynamic environment with mutations.

We begin with Lemma 3 that states that if an equilibrium set represents a Nash equilibrium, then it is a restpoint of the dynamics.

Lemma 3

Let $\{\Theta_{s^i}\}_{s^i \in S^i, i = 1, 2}$ be an equilibrium set given $F = (F_1, F_2)$. Let $\sigma^i(s^i) = F_i(\Theta_{s^i}) \forall s^i \in S^i, i = 1, 2$. If σ is a Nash equilibrium given objective payoffs of the game, then F is a restpoint of (1).

Proof

(σ^i, σ^j) being a Nash equilibrium of the unbiased game $\langle N, S, U \rangle$ implies for all $s^i, \tilde{s}^i \in S^i : F_i(\Theta_{s^i}), F_i(\Theta_{\tilde{s}^i}) > 0$ that $U^i(s^i, \sigma^j) = U^i(\tilde{s}^i, \sigma^j)$. From monotonicity of (1) follows that $g_i(\theta^i, F, \{\Theta_s\}) = g_i(\tilde{\theta}^i, F, \{\Theta_s\}) = \bar{g}_i(F, \{\Theta_s\}) \forall \theta^i \in \Theta_{s^i}, \tilde{\theta}^i \in \Theta_{\tilde{s}^i}$. We can rewrite (1) as

$$\begin{aligned} \dot{F}_i(\Theta^i) &= \sum_{s^i \in S^i: F(\Theta_{s^i}) > 0} \int_{\Theta^i \cap \Theta_{s^i}} g_i(\theta^i, F, \{\Theta_s\}) f_i(\theta^i) d\theta^i \\ &= \bar{g}_i(F, \{\Theta_s\}) \sum_{s^i \in S^i: F(\Theta_{s^i}) > 0} \int_{\Theta^i \cap \Theta_{s^i}} f_i(\theta^i) d\theta^i \\ &= \bar{g}_i(F, \{\Theta_s\}) \cdot F(\Theta^i) \end{aligned}$$

Since $\dot{F}(\mathbb{R}^n) = 0$ and $F(\mathbb{R}^n) = 1$ we have $\bar{g}_i(F, \{\Theta_s\}) = 0$. Therefore $\dot{F}_i(\Theta) = \bar{g}_i(F, \{\Theta_s\}) \cdot F_i(\Theta) = 0 \forall \Theta \subset \mathbb{R}^n$. \square

We would like to mention Ely and Yilankaya (2001) who show also “stability” in a related framework. We suspect that this is a very courageous interpretation of a system being in a steady state. In the model presented here, dynamic stability does not apply to all Nash equilibria. We give an example in the next section.

The next theorem states that for any Nash equilibrium we can find an equilibrium set such that in fact no agent maximizes payoffs that are equivalent to the unbiased payoffs of the game. The proof is constructive.

Theorem 1

Given a finite game, any Nash equilibrium σ of that game can be represented by an equilibrium set $\{\Theta_{s^i}\}_{s^i \in S^i}, i = 1, 2$ with respect to a cdf F such that

- i) $F_i(\Theta_{s^i}) = \sigma^i(s^i)$ $s^i \in S^i, i = 1, 2$ and
- ii) if $\Theta^i \subset \mathbb{R}^n$ is equivalent to the zero-bias, then $F_i(\Theta^i) = 0$.

Proof

Case: strategy s^i is strictly dominant. Fix some $\check{s}^j : \sigma^j(\check{s}^j) < 1$ and some $\check{s}^i \neq s^i$. Consider the conditions A and B :

$$A \quad \begin{aligned} \theta^i(s^i, \check{s}^j) &< -U_i(s^i, \check{s}^j) \\ \theta^i(\check{s}^i, \check{s}^j) &> -U_i(\check{s}^i, \check{s}^j) \end{aligned}$$

$$B \quad \begin{aligned} \sum_{s^j \neq \check{s}^j} \sigma^j(s^j) [U_i(s^i, s^j) + \theta^i(s^i, s^j)] &> -\sigma^j(\check{s}^j) [U_i(s^i, \check{s}^j) + \theta^i(s^i, \check{s}^j)] \\ \sum_{s^j \neq \check{s}^j} \sigma^j(s^j) [U_i(\check{s}^i, s^j) + \theta^i(\check{s}^i, s^j)] &< -\sigma^j(\check{s}^j) [U_i(\check{s}^i, \check{s}^j) + \theta^i(\check{s}^i, \check{s}^j)] \end{aligned}$$

These conditions restrict $\{\theta^i(s^i, s^j)\}_{s^j \in S^j}$ and $\{\theta^i(\check{s}^i, s^j)\}_{s^j \in S^j}$ to be in a half-space of \mathbb{R}^{n_j} . Define $\hat{\Theta}_{s^i} = \{\theta^i \in \Theta_{s^i}(\sigma^j) \mid \theta^i \text{ satisfies } A, B\}$. Condition B implies that $\Theta_{s^i}(\sigma^j)$ strictly contains $\hat{\Theta}_{s^i}$. Condition A implies $U_i(s^i, \check{s}^j) + \theta^i(s^i, \check{s}^j) < U_i(\check{s}^i, \check{s}^j) + \theta^i(\check{s}^i, \check{s}^j)$ and hence that no $\theta^i \in \hat{\Theta}_{s^i}$ is equivalent to the zero bias.

Case: strategy s^i is not strictly dominant. Fix some $\check{s}^i \neq s^i$. For $K = \max_{\check{s}^i, s^j} U_i(\check{s}^i, s^j) - U_i(s^i, s^j)$ define

$$\hat{\Theta}_{s^i} = \{\theta^i \in \text{int}(\Theta_{s^i}(\sigma^j)) \mid \theta^i(s^i, s^j) > K, \theta^i(\check{s}^i, s^j) < -K \forall s^j \in S^j\} .$$

For any $\theta^i \in \hat{\Theta}_{s^i}$ the following holds for any s^j :

$$\begin{aligned} U_i(s^i, s^j) + \theta^i(s^i, s^j) &> U_i(s^i, s^j) + K \\ &\geq U_i(s^i, s^j) + U_i(\check{s}^i, s^j) - U_i(s^i, s^j) \geq \\ &U_i(\check{s}^i, s^j) - K > U_i(\check{s}^i, s^j) + \theta^i(\check{s}^i, s^j) \end{aligned}$$

Therefore, no $\theta^i \in \hat{\Theta}_{s^i}$ is equivalent to the zero bias.

Finally, choose a cdf $\hat{F} = (\hat{F}_i, \hat{F}_j)$ such that $\hat{F}_i(\hat{\Theta}_{s^i}) = \sigma^i(s^i) \forall s^i \in S^i$, $i = 1, 2$. $\{\hat{\Theta}_{s^i}\}_{s^i \in S^i, i = 1, 2}$ is an equilibrium set given \hat{F} and $\hat{F}_i(\Theta) > 0$ implies that $\theta^i \in \Theta$ is not equivalent to the zero bias. \square

The proof distinguishes two cases for each strategy of each player i , the one in which s^i is strictly dominant and the one in which it is not. For the first case, there is a set of ideologies $\hat{\Theta}_{s^i}$ that still induce the choice of s^i given σ^j , but also bias the player toward some alternative strategy \check{s}^i , if player j would play that strategy \check{s}^j which is not used with certainty in equilibrium. Strategy s^i is therefore not strictly dominant given the biases in $\hat{\Theta}_{s^i}$.

For the second case, in which s^i is not strictly dominant, a fraction $\sigma^i(s^i)$ of population i has an ideology that induces the choice of strategy s^i given any strategy of population j – agents with such an ideology believe strategy s^i to be strictly dominant.

The distribution \hat{F}_i has mass one on $\hat{\Theta}_{s^i}$. In both cases there is no agent who maximizes payoffs that are equivalent to the objective payoffs. Secondly, we know from Lemma 3 that no agent would profit in terms of objective payoffs by deviating from his strategy choice, because the distribution \hat{F} is a restpoint of the monotonic dynamic (1).

The next results concern stability with respect to small perturbations. The stability notion considered here is Lyapunov stability. A state is Lyapunov stable, if once it is perturbed very slightly, it does not move further away. ‘‘Slightliness’’ is measured by the variational norm, that –from our point of

view— represents best the model of mutations.

Let $\|F - F'\|$ denote the variational norm:

$\|F - F'\| = \sup_h \left| \int_{\mathbb{R}^n} h(\theta)(f(\theta) - f'(\theta))d\theta \right|$, where the supremum is taken over all measurable functions $h : \mathbb{R}^n \rightarrow \mathbb{R}$ satisfying $\sup_{\theta} |h(\theta)| \leq \frac{1}{2}$.⁴

As in Oechssler and Riedel (2001)'s remark, we discuss the relation of the variational norm to the concept of mutations in evolutionary game theory. Consider a cdf F that has measure zero on some set Θ . Now suppose a small group (of size $\frac{\epsilon}{1-\epsilon} > 0$) enters the population, that has a cdf F' with measure one on Θ . The new normalized distribution is $F^\epsilon = (1 - \epsilon) \cdot F + \epsilon \cdot F'$ and has distance $\sup_h \left| \int_{\mathbb{R} \setminus \Theta} h(\theta) \underbrace{(f(\theta) - (1 - \epsilon)f(\theta))}_{>0} d\theta + \int_{\Theta} h(\theta) \underbrace{(0 - \epsilon f'(\theta))}_{<0} d\theta \right| = \epsilon \frac{1}{2} \underbrace{(1 - F(\Theta))}_{=1} + \epsilon \frac{1}{2} \underbrace{F'(\Theta)}_{=1} = \epsilon$ from F . Therefore, if the group of entering mutants is relatively small, the distribution changes only little. Consider instead that the original distribution F has mass one on a ball Θ with radius δ around the point θ and that each agent changes his belief of the payoffs only a little, say shifts it by $\epsilon \cdot (1, 1, \dots, 1)$, $\epsilon > 2\delta$ such that the new distribution F^ϵ has measure one on a ball $\Theta^\epsilon = \Theta + \epsilon \cdot (1, 1, \dots, 1)$. Then, F and F^ϵ have distance $\sup_h \left| \int_{\Theta} h(\theta) \underbrace{(f(\theta) - f^\epsilon(\theta))}_{>0} d\theta + \int_{\Theta^\epsilon} h(\theta) \underbrace{(f(\theta) - f^\epsilon(\theta))}_{<0} d\theta \right| = \frac{1}{2} \underbrace{F(\Theta)}_{=1} + \frac{1}{2} \underbrace{F^\epsilon(\Theta^\epsilon)}_{=1} = 1$, which is the maximal distance. That means, if all agents mutate, even very little, the measure of change of the distribution is maximal; if only a very small fraction of the population mutates, even very starkly, the measure of change is very small. From my point of view, the norm reflects the logic of mutations adequately: mutations occur independently and with probability of equal order (in ϵ) across agents. Multiple instantaneous mutations are therefore less probable than a single mutation. On the other hand, since mutations are arbitrary per se, a mutation “far away” has probability of the same order (in ϵ) as mutation “close by”. Other norms such as

⁴We take the definition of the variational norm from Oechssler and Riedel (2001).

for example the weak topology (here: \mathcal{L}^2 -norm) also captures changes that concern all agents, which could be a global change of the payoffs or statistical environment. Such changes are not under consideration in the model studied in this paper.

Definition

Let F^* be a restpoint satisfying $\dot{F}_i^*(\Theta) = 0 \forall \Theta \in \mathbb{R}^n, i = 1, 2$. Then F^* is Lyapunov stable if $\forall \epsilon > 0 \exists \eta > 0$:

$$\|F_i(0) - F_i^*\| < \eta \forall i \Rightarrow \|F_i(t) - F_i^*\| < \epsilon \forall t > 0, \forall i$$

This stability notion is less demanding than for example being weakly- or strongly- attracting, which would mean that a perturbation would be forced to converge back to the original distribution. We cannot demand this stronger property because there always exist many distributions with the same support that are Lyapunov stable.

Definition

An equilibrium set $\{\Theta_s\}$ is *dynamically stable* with respect to some cdf F , if F is Lyapunov stable given $\{\Theta_s\}$ and $\{\Theta_s\}$ is an equilibrium set given F .

Dynamic stability is the natural requirement of an equilibrium set being robust to small perturbations in the distribution of biases. The next lemma states that Lyapunov stability implies Nash equilibrium. This property is satisfied in most models of evolutionary game theory.

Lemma 4

Let $\{\Theta_s\}$ be dynamically stable with respect to some cdf F . Then σ with $\sigma_i(s^i) = F_i(\Theta_{s^i}) \forall s^i \in S^i$ is a Nash equilibrium.

Proof

F is Lyapunov stable. Hence F is a rest point, $\dot{F}(\Theta) = 0 \forall \Theta \in \mathbb{R}^n$ and

$$\text{i) } F_i(\tilde{\Theta}) > 0 \Rightarrow g^i(\tilde{\theta}^i, F\{\Theta_s\}) = 0 \forall \tilde{\theta}^i \in \tilde{\Theta}$$

$$\text{ii) } F_i(\hat{\Theta}) = 0 \Rightarrow g^i(\hat{\theta}^i, F\{\Theta_s\}) \leq 0 \quad \forall \hat{\theta}^i \in \hat{\Theta}$$

For a strategy s^i with $\sigma^i(s^i) = F_i(\Theta_{s^i}) > 0$ it follows $\forall \hat{s}^i$ that $g_{\hat{s}^i}^i(F, \{\Theta_s\}) \geq g_{s^i}^i(F, \{\Theta_s\})$. Monotonicity then implies $U_i(s^i, \sigma^j) \geq U_i(\hat{s}^i, \sigma^j) \quad \forall s^i \in S^i : \sigma^i(s^i) > 0$ and $\hat{s}^i \in S^i$, therefore σ is a Nash equilibrium. \square

We note that the reverse statement, namely that a Nash equilibrium implies dynamic stability is generally not true. Consider, for example, a Nash equilibrium involving strategies that are weakly dominated. It is clear, that for an open environment of distributions with full support around the Nash equilibrium distribution there are ideologies that induce the choice of strategies that gain strictly higher (real) payoffs than the Nash equilibrium strategies.

We state our main theorem for a refinement of the Nash equilibrium concept firstly defined by Balkenborg (1994). Unfortunately, its name consists also of the words “equilibrium set”, but is distinct from the concept equilibrium set defined in this paper.

Definition Strict Equilibrium Set (SEset) (Balkenborg (1994))

A set $\Delta \subset \Delta(S^i) \times \Delta(S^j)$ is a *Strict Equilibrium Set (SEset)*, if

each $\sigma \in \Delta$ is a Nash equilibrium and if

for $(\sigma^i, \sigma^j) \in \Delta$, $\tilde{\sigma}^i$ is an alternative best reply to σ^j , then $(\tilde{\sigma}^i, \sigma^j) \in \Delta$.

As Lemma 3 states, all elements of a SEset can be represented by restpoints of the dynamic that exclude ideologies that are equivalent to the unbiased payoffs. The following theorem states that these states are also dynamically stable.

Main Theorem

Let the game $\langle N, S, U \rangle$ have a strict equilibrium set Δ . Then any $\sigma \in \Delta$ can be represented by a dynamically stable equilibrium set $\{\Theta_s\}_s$ with respect to a cdf F such that

$$\text{i) } F_i(\Theta_{s^i}) = \sigma^i(s^i) \quad \forall s^i \in S^i, i = 1, 2 \text{ and}$$

ii) if all ideologies in Θ^i are equivalent to the zero bias, then $F_i(\Theta^i) = 0$.

Proof

Fix $\sigma \in \Delta$. For each $s^i \in S^i$ and $i = 1, 2$, choose the set of ideologies $\hat{\Theta}_{s^i}$ as defined in Theorem 1. Define cdf F such that $F_i(\hat{\Theta}_{s^i}) = \sigma^i(s^i)$. Since $\hat{\Theta}_{s^i} \subset \Theta_{s^i}(\sigma^j)$, $\{\hat{\Theta}_s\}_s$ is an equilibrium set with $F(\Theta) = 0$ whenever $\theta \Leftrightarrow 0 \forall \theta \in \Theta$. Since Δ is a strict equilibrium set, for any $s^i \in \text{supp}(\Delta^i)$ it holds that $U_i(s^i, \sigma^j) = U_i(\tilde{s}^i, \sigma^j) \forall \tilde{s}^i \in \text{supp}(\Delta^i)$ and $U_i(s^i, \sigma^j) > U_i(\tilde{s}^i, \sigma^j) \forall \tilde{s}^i \notin \text{supp}(\Delta^i)$. Therefore, for all $s^i : \sigma^i(s^i) > 0$, $i = 1, 2$ monotonicity implies $g_{s^i}^i(F, \{\hat{\Theta}_s\}) \geq g^i(\theta^i, F, \{\hat{\Theta}_s\}) \forall \theta^i \in \mathbb{R}^n$. As a consequence of regularity $\dot{F}(\Theta) = 0 \forall \Theta \subset \mathbb{R}^n$ and F is Lyapunov stable. Therefore $\{\hat{\Theta}\}$ is dynamically stable given F . \square

All results stated in this section hold for games in strategic form with finite strategy spaces and two players and for monotonic dynamics in the sense of Samuelson and Zhang (1992). Since we did not exploit the restriction to two players in any stage of the proofs, we do believe that these results extend to n -player games. In the next section, we focus on 2×2 games, give examples and a counter example – the matching pennies game mentioned in the introduction of this paper. The class of games including the matching pennies game does not have a SEset.

4 Dynamically Stable Sets in 2×2 Games

A 2×2 game has four outcomes, for each outcome there is one bias for each player. Hence we have to account for eight parameters. To reduce the parameter space, we normalize the asymmetric 2×2 -game to a game with off-diagonal zero payoffs:

Payoffs for player i

$$\begin{array}{cc|cc} & A & B & \\ \hline A & a_1^i + \theta_{a_1}^i & a_2^i + \theta_{a_2}^i & \\ \hline B & b_2^i + \theta_{b_2}^i & b_1^i + \theta_{b_1}^i & \\ \hline \end{array} \rightarrow \begin{array}{cc|cc} & A & B & \\ \hline A & a^i + \theta_a^i & 0 & \\ \hline B & 0 & b^i + \theta_b^i & \\ \hline \end{array}$$

The figure depicts the payoffs for player i , where $a^i = a_1^i - b_2^i$, $b^i = b_1^i - a_2^i$, $\theta_a^i = \theta_{a_1}^i - \theta_{b_2}^i$ and $\theta_b^i = \theta_{b_1}^i - \theta_{a_2}^i$. The payoffs for player j are constructed accordingly. Assume without loss of generality $a^1 \geq b^1$. We assume $a^i \cdot b^i \neq 0$ for $i = 1, 2$ to have less case distinctions. A bias θ^i is equivalent to the true payoffs, if $\frac{\theta_a^i}{a^i} > -1$, $\frac{\theta_b^i}{b^i} > -1$, and $\frac{\theta_a^i - \theta_b^i}{a^i - b^i} > -1$ for $a^i \neq b^i$. For any $\sigma^j \in \Delta(\{A, B\})$, i 's belief over j 's actions, the set of i -biases that induce the choice of A is

$$\Theta_A^i(\sigma^j) = \{\theta^i \in \mathbb{R}^2 \mid \sigma^j(A)(\theta_a^i + a^i) \geq \sigma^j(B)(b^i + \theta_b^i)\}.$$


In equilibrium $F_i(\Theta_{s^i}(\sigma^j)) = \sigma^i(s^i)$ and $U^i(A, \sigma^j) = a^i \cdot \sigma^j(A)$ and $U^i(B, \sigma^j) = b^i \cdot \sigma^j(B)$. Therefore $g^i(\theta^i, F, \{\Theta_s\})$ can take only two values for any given F and equilibrium $\{\Theta_s\}$; $g^i(\theta^i, F, \{\Theta_s\}) = g_A^i(F, \{\Theta_s\})$ if $\theta^i \in \Theta_A^i$ and $g^i(\tilde{\theta}^i, F, \{\Theta_s\}) = g_B^i(F, \{\Theta_s\})$ if $\tilde{\theta}^i \in \Theta_B^i$.

This implies then since

$$\dot{F}_i(\mathbb{R}^2) = g_A^i(F, \{\Theta_s\}) \cdot F_i(\Theta_A^i) + g_B^i(F, \{\Theta_s\}) \cdot (1 - F_i(\Theta_A^i)) = 0$$

that $g_A^i(F, \{\Theta_s\}) = 0$ if $F_i(\Theta_A^i) = 1$ and that $g_B^i(F, \{\Theta_s\}) = 0$ if $F_i(\Theta_A^i) = 0$. Hence $F_i(\Theta_A^i) = 1$ and $F_i(\Theta_A^i) = 0$ are restpoints of the dynamic of population i . An interior⁵ steady state ($0 < F_i(\Theta_A^i) < 1$) has the following characteristics: $g_A^i(F, \{\Theta_s\}) = g_B^i(F, \{\Theta_s\}) = 0$ which implies that $U_A^i(F, \{\Theta_s\}) = U_B^i(F, \{\Theta_s\}) \Leftrightarrow F_j(\Theta_A^j) = \frac{b^i}{a^i + b^i}$, which can only be in $(0, 1)$ if $a^i(a^i + b^i) > 0$ and $b^i(a^i + b^i) > 0$.

⁵If $F_i(\Theta_A^i) \in (0, 1)$, one can substitute $g_B^i(F, \{\Theta_s\}) = -g_A^i(F, \{\Theta_s\}) \frac{F_i(\Theta_A^i)}{F_i(\Theta_B^i)}$ and the dynamics simplify for all $\Theta^i \subset \mathbb{R}^n$ to $\dot{F}_i(\Theta^i) = g_A^i(F, \{\Theta_s\}) (F_i(\Theta^i | \Theta_A^i) - F_i(\Theta^i | \Theta_B^i)) F_i(\Theta_A^i)$.



Restpoint classes of 2×2 -games

To summarize, depending on $\frac{b^i}{a^i+b^i} \in (0, 1)$ $i \in \{1, 2\}$, there are at most nine restpoints: $(F_i(\Theta_A^i), F_j(\Theta_A^j)) \in \left(\left\{ 0, \frac{b^j}{a^j+b^j}, 1 \right\} \times \left\{ 0, \frac{b^i}{a^i+b^i}, 1 \right\} \right) \cap [0, 1]^2$

5 Two examples

We illustrate the results for two examples. One example is a game with a unique Nash equilibrium that is strict. In this case the Main Theorem is applicable and we construct a dynamically stable equilibrium set of ideologies that are non-equivalent to the zero bias in the spirit of Theorem 1. The second example is the matching pennies game, which does not have an SEset.

An example (unique strict Nash equilibrium $\hat{\sigma} = (A, A)$):

	A	B
A	(1,1)	(0,0)
B	(0,0)	(-1,2)

For population 1, action A is not strictly dominant in terms of objective payoffs.

The set of types in population 1 that prefer strategy A given σ^2 is


$$\Theta_A^1(\sigma^2) = \{\theta^1 \in \mathbb{R}^2 \mid \sigma^2(A)(1 + \theta_a^1) \geq \sigma^2(B)(-1 + \theta_b^1)\} .$$

The set of ideologies that are equivalent to the zero bias is

$$\tilde{\Theta}^1 = \{\theta^1 \in \mathbb{R}^2 \mid 1 + \theta_a^1 > 0 > -1 + \theta_b^1\} .$$

Suppose population 1 consists only of agents who have types in the set

$$\hat{\Theta}_A^1 = \{\hat{\theta}^1 \in \mathbb{R}^2 \mid \hat{\theta}_a^1 \in (-3, -2), \hat{\theta}_b^1 < 0\} .$$



Clearly, $\hat{\Theta}_A^1 \subset \Theta_A^1(\sigma^2)$ for all σ^2 close enough to equilibrium $\hat{\sigma}^2$, but $\hat{\Theta}_A^1 \cap \tilde{\Theta}^1 = \emptyset$.

For population 2 the set of biases that induce the choice of A is


$$\Theta_A^2(\sigma^1) = \{\theta^2 \in \mathbb{R}^2 \mid (1 + \theta_a^2)\sigma^1(A) \geq (2 + \theta_b^2)\sigma^1(B)\} .$$

The set of biases the induce the same ranking over the set of outcomes as the zero bias is

$$\tilde{\Theta}^2 = \{\theta^2 \in \mathbb{R}^2 \mid 2 + \theta_b^2 > 1 + \theta_a^2 > 0\} .$$

Suppose population 2 consist only of agents who believe that action A is dominant. For example the set $\hat{\Theta}_A^2$ consist of such types:

$$\hat{\Theta}_A^2 = \{\theta^2 \in \mathbb{R}^2 \mid \theta_a^2 > 0, \theta_b^2 < -3\} .$$



Again, $\hat{\Theta}_A^2 \subset \Theta_A^2(\sigma^1) \forall \sigma^1$ close enough to $\hat{\sigma}^1$ (in fact for all σ^1) and $\hat{\Theta}_A^2 \cap \tilde{\Theta} = \emptyset$. Define the equilibrium set as


$$\Theta_A^i = \hat{\Theta}_A^i, \Theta_B^i = \Theta_B^i(\hat{\sigma}^j) \text{ and } F_i(\Theta_A^i) = 1$$

For players $i = 1, 2$. In the following we argue that this equilibrium set is dynamically stable.


For this game, $U^1(A, \sigma^2) = \sigma^2(A) > -(1 - \sigma^2(A)) = U^1(B, \sigma^2) \forall \sigma^2(A) \in [0, 1]$. From monotonicity and regularity of g it follows that $g_A^1(F, \{\Theta_s\}) > 0 > g_B^1(F, \{\Theta_s\})$ for all F_2 and $F_1(\Theta_A^1) < 1$. Therefore $\dot{F}_1(\Theta_A^1) > 0 \forall F_1(\Theta_A) < 1$ and $\hat{F}_1(\hat{\Theta}_A^1) = 1$ is (Lyapunov-) stable. For population 2, the unique best reply to A^1 is A^2 , hence $g_A^2((\hat{F}_1, F_2), \{\Theta_s\}) > 0$ and $\dot{F}_2(\Theta_A^2) > 0 \forall F_2(\Theta_A^2) < 1$. Therefore $\hat{F}_2(\hat{\Theta}_A^2) = 1$ is also (Lyapunov-) stable.

Example (matching pennies game):

	A	B
A	1,-1	0,0
B	0,0	1,-1



The matching pennies game does not have a SEset but only a unique Nash equilibrium in mixed strategies, $(\hat{\sigma}^1, \hat{\sigma}^2) = (\frac{1}{2}, \frac{1}{2})$. There are nine restpoints: $(F_1(\Theta_A^1), F_2(\Theta_A^2)) \in \{0, \frac{1}{2}, 1\} \times \{0, \frac{1}{2}, 1\}$. The phase diagram above indicates that none of the restpoints needs to be stable. Consider the replicator dynamic with $g_A^1(F, \{\Theta_s\}) = 2(1 - F_1(\Theta_A))(F_2(\Theta_A) - \hat{\sigma}^2)$ and $g_A^2(F, \{\Theta_s\}) = 2(1 - F_2(\Theta_A))(\hat{\sigma}^1 - F_1(\Theta_A))$. The collection $\hat{\Theta}_A^1 = \{\theta^1 \in \mathbb{R}^2 \mid \theta_a^1 > 0, \theta_b^1 < -1\}$, $\hat{\Theta}_B^1 = \{\theta^1 \in \mathbb{R}^2 \mid \theta_a^1 < -1, \theta_b^1 > 0\}$, $\hat{\Theta}_A^2 = \{\theta^2 \in \mathbb{R}^2 \mid \theta_a^2 > 1, \theta_b^2 < 0\}$ and $\hat{\Theta}_B^2 = \{\theta^2 \in \mathbb{R}^2 \mid \theta_a^2 < 0, \theta_b^2 > 1\}$ defines an equilibrium set given any F that fulfills $F_i(\hat{\Theta}_A^i \cup \Theta_B^i) = 1$. The trajectories cycle around the equilibrium $\hat{\sigma} = (\frac{1}{2}, \frac{1}{2})$, as the figure below illustrates. For the replicator dynamics it can be shown that $\hat{\sigma}$ is Lyapunov stable.



Trajectories for the replicator dynamics

6 Conclusions

We study a model in which interacting agents follow an ideology that specifies the unverifiable payoffs of a 2-player game in strategic form with finite strategy sets. When drawn from an infinite population, an agent knows his own ideology, but has incomplete information of the ideology that his opponent follows. Given a belief σ^j on j 's set of strategies, each agent optimizes given the payoffs specified by his or her ideology. In equilibrium all agents of the same population i have the same belief σ^j . We allow for a continuous variety of ideologies within each population. Given any distribution of ideologies, we define equilibrium sets and show the existence thereof. We assume equilibrium play in each instant of time. We assume that evolution selects against ideologies that induce behavior which yields relatively low evolutionary relevant objective payoffs. We characterize distributions that are stable with respect to small changes (Lyapunov stability). The main result states that for any two player game with a strict (Nash) equilibrium set, there exist dynamically stable equilibrium sets that consist only of non-equivalent ideologies.

Appendix

Lemma 2

$G_{s^i}(\sigma^j) = F_i(\Theta_{s^i}(\sigma^j))$ is continuous in $\sigma^j \forall i, j$.

Proof:

Outline:

Given a convergent sequence $\{\sigma_n\}_n$ in $\Delta(S^j)$ and associated halfspaces $\{H_n\}_n$ we show $H_n \cup H \downarrow H$ and $H_n \cap H \uparrow H$ and prove with Lemmata A1 & A2 continuity from below and above. We define H_n such that $\Theta_{s^i}(\sigma^j)$ is the intersection of finitely many halfspaces and therefore is also continuous from above and below.

For any $\sigma, \hat{\sigma} \in \Delta(S^j)$ define $\sigma_n = \sigma \frac{n-1}{n} + \hat{\sigma} \frac{1}{n}$, $n \in \mathbb{N}$. For any $s^i, \tilde{s}^i \in S^i$ and $\sigma_n \in \Delta(S^j)$ define $H_n = \{\theta^i \in \mathbb{R}^n \mid U_i(s^i, \sigma_n) + \theta^i(s^i, \sigma_n) \geq U_i(\tilde{s}^i, \sigma_n) + \theta^i(\tilde{s}^i, \sigma_n)\}$, the set of types that weakly prefer strategy s^i over strategy \tilde{s}^i given belief σ_n .⁶

$H_n \cap H \uparrow H$:

We show $\theta^i \in H_n \cap H \Rightarrow \theta^i \in H_{n+1}$. Fix some $\theta^i \in H_n \cap H$. Multiplying the inequality implied by $\theta^i \in H_n$ with $\frac{n}{n+1}$, the inequality implied by $\theta^i \in H$ with $\frac{1}{n+1}$ and summing up yields

$$\begin{aligned} & U_i(s^i, \sigma_n \frac{n}{n+1} + \sigma \frac{1}{n+1}) + \theta^i(s^i, \sigma_n \frac{n}{n+1} + \sigma \frac{1}{n+1}) \\ \geq & U_i(\tilde{s}^i, \sigma_n \frac{n}{n+1} + \sigma \frac{1}{n+1}) + \theta^i(\tilde{s}^i, \sigma_n \frac{n}{n+1} + \sigma \frac{1}{n+1}) \\ \Leftrightarrow & U_i(s^i, \sigma_{n+1}) + \theta^i(s^i, \sigma_{n+1}) \geq U_i(\tilde{s}^i, \sigma_{n+1}) + \theta^i(\tilde{s}^i, \sigma_{n+1}), \end{aligned}$$

which implies $\theta^i \in H_{n+1}$. $\theta^i \in H$ is trivially implied, we conclude $H_{n+1} \cap H \subset H_n \cap H \forall n \in \mathbb{N}$. Since $H \cap H = H$, we have shown $H_n \cap H \uparrow H$.

$H_n \cup H \downarrow H$:

We start by showing $\theta^i \in H_{n+1} \cap H_1 \Rightarrow \theta^i \in H_n$. Fix some $\theta^i \in H_{n+1} \cap H_1$. Multiplying the inequality implied by $\theta^i \in H_{n+1}$ with $\frac{n^2-1}{n^2}$, the inequality implied by $\theta^i \in H_1$ with $\frac{1}{n^2}$ and summing up yields

$$\begin{aligned} & U_i(s^i, \sigma_{n+1} \frac{n^2-1}{n^2} + \hat{\sigma} \frac{1}{n^2}) + \theta^i(s^i, \sigma_{n+1} \frac{n^2-1}{n^2} + \hat{\sigma} \frac{1}{n^2}) \\ \geq & U_i(\tilde{s}^i, \sigma_{n+1} \frac{n^2-1}{n^2} + \hat{\sigma} \frac{1}{n^2}) + \theta^i(\tilde{s}^i, \sigma_{n+1} \frac{n^2-1}{n^2} + \hat{\sigma} \frac{1}{n^2}) \\ \Leftrightarrow & \\ & U_i(s^i, \sigma \frac{n-1}{n} + \hat{\sigma} \frac{1}{n}) + \theta^i(s^i, \sigma \frac{n-1}{n} + \hat{\sigma} \frac{1}{n}) \\ \geq & U_i(\tilde{s}^i, \sigma \frac{n-1}{n} + \hat{\sigma} \frac{1}{n}) + \theta^i(\tilde{s}^i, \sigma \frac{n-1}{n} + \hat{\sigma} \frac{1}{n}), \end{aligned}$$

which implies $\theta^i \in H_n$. We conclude $\theta^i \in H_{n+1} \setminus H_n \Rightarrow \theta^i \notin H_1$.

We proceed by showing $\theta^i \in H_{n+1} \setminus H_n \Rightarrow \theta^i \in H$.

⁶The proofs for $\tilde{s}^i = s^i$ are trivially valid, however, only the cases $\tilde{s}^i \neq s^i$ are relevant.

Fix some $\theta^i \in H_{n+1} \setminus H_n$. Multiplying the inequality implied by $\theta^i \in H_{n+1}$ with $\frac{n^2-1}{n^2}$ and adding the inequality implied by $\theta^i \notin H_n$ yields

$$\begin{aligned} & U_i(s^i, \sigma) + \theta^i(s^i, \sigma) - U_i(\tilde{s}^i, \sigma) - \theta^i(\tilde{s}^i, \sigma) \\ & > \frac{2n}{n+1} [U_i(\tilde{s}^i, \hat{\sigma}) + \theta^i(\tilde{s}^i, \hat{\sigma}) - U_i(s^i, \hat{\sigma}) - \theta^i(s^i, \hat{\sigma})] \end{aligned}$$

Since $\theta^i \notin H_1$, the right hand side of the inequality is positive. Therefore $\theta^i \in H$.

Since $H_{n+1} = (H_{n+1} \setminus H_n) \cup (H_{n+1} \cap H_n)$ we conclude $H_{n+1} \subset H \cup H_n \forall n \in \mathbb{N}$. As $H \cup H = H$, we have shown $H_n \cup H \downarrow H$.

We apply Lemmata A1 & A2 below with $A_n = H_n \cap H$ and $B_n = H_n \cup H$ and conclude that F is continuous from above and below. From the definition of H , $\Theta_{s^i}(\sigma^j) = \cap_{\tilde{s}^i \in S^i} H_n^{\tilde{s}^i}$ is a finite intersection and the desired properties of F carry over to $\Theta_{s^i}(\sigma^j)$. \square

Lemmata A1 & A2 are taken from and proved in Bauer (1992).

Lemma A1 (continuity from below)

Consider a sequence $\{A_n\}_n$ of subsets of $\mathbb{R}^{n_i \cdot n_j}$ with $A_n \uparrow A \subset \mathbb{R}^{n_i \cdot n_j}$.

Then $\lim_{n \rightarrow \infty} F(A_n) = F(A)$.

Proof:

Define $A_0 := \emptyset$ and $a_n := A_n \setminus A_{n-1}, n \in \mathbb{N}$. We have $A_n = \cup_{i=1}^n a_i$ and $A = \cup_{n=1}^{\infty} a_n$. Since $a_n \cap a_m = \emptyset \forall n \neq m$ and σ -additivity of the measure F we have

$$F(A) = \sum_{n=1}^{\infty} F(a_n) = \lim_{n \rightarrow \infty} \sum_{i=1}^n F(a_i) = \lim_{n \rightarrow \infty} F(A_n) . \square$$

Lemma A2 (continuity from above)

Consider a sequence $\{B_n\}_n$ of subsets of $\mathbb{R}^{n_i \cdot n_j}$ with $B_n \downarrow B \subset \mathbb{R}^{n_i \cdot n_j}$.

Then $\lim_{n \rightarrow \infty} F(B_n) = F(B)$.

Proof:

Since $B_n \subset B_1 \forall n \in \mathbb{N}$ it holds that $F(B_1 \setminus B_n) = F(B_1) - F(B_n) \forall n \in \mathbb{N}$.

Clearly, $B_1 \setminus B_n \uparrow B_1 \setminus B$. From Lemma A1 we know that $F(B_1 \setminus B) = \lim_{n \rightarrow \infty} F(B_1 \setminus B_n)$ and therefore $F(B_1 \setminus B) = F(B_1) - \lim_{n \rightarrow \infty} F(B_n)$. We also have $B \subset B_1$ and therefore $F(B_1 \setminus B) = F(B_1) - F(B)$, which establishes $F(B) = \lim_{n \rightarrow \infty} F(B_n)$. \square

REMARKS

Symmetric difference $A \ominus B = A \setminus B \cup B \setminus A$:

Note since we do not require $F(\Theta) > 0 \forall \Theta \subset \mathbb{R}^n$, $F(A \ominus B) = 0$ does not imply $A = B$. Therefore $d_F(A, B) = F(A \ominus B)$ is only a pseudo metric (satisfying symmetry and the triangle inequality).

Upper hemi continuity:

Note further that $\Theta(\sigma)$ is not upper hemi continuous. To give an example,

consider the doubly symmetric 2×2 coordination game:

	s_1^j	s_2^j
s_1^i	(1,1)	(0,0)
s_2^i	(0,0)	(1,1)

Then each $\theta^i \in \Theta_{s_1^i}(\frac{1}{3})$ satisfies

$$[1 + \theta^i(s_1^i, s_1^j)]\frac{1}{3} + \theta^i(s_1^i, s_2^j)\frac{2}{3} \geq \theta^i(s_2^i, s_1^j)\frac{1}{3} + [1 + \theta^i(s_2^i, s_2^j)]\frac{2}{3}$$


Define E as the set of θ^i that satisfy

$$[1 + \theta^i(s_1^i, s_1^j)]\frac{1}{3} + \theta^i(s_1^i, s_2^j)\frac{2}{3} > \theta^i(s_2^i, s_1^j)\frac{1}{3} + [1 + \theta^i(s_2^i, s_2^j)]\frac{2}{3} - 1$$

Clearly, $\Theta_{s_1^i}(\frac{1}{3}) \subset E$. Consider some $\hat{\sigma}^j(s_1^j)$ arbitrarily close to $\frac{1}{3}$, for example $\hat{\sigma}^j(s_1^j) = \frac{1}{3} + \epsilon$ for $\epsilon \in (0, \frac{2}{3})$. Any $\theta^i \in \Theta_{s_1^i}(\hat{\sigma}^j)$ satisfies

$$\begin{aligned} & [1 + \theta^i(s_1^i, s_1^j)](\frac{1}{3} + \epsilon) + \theta^i(s_1^i, s_2^j)(\frac{2}{3} - \epsilon) \\ & \geq \\ & \theta^i(s_2^i, s_1^j)(\frac{1}{3} + \epsilon) + [1 + \theta^i(s_2^i, s_2^j)](\frac{2}{3} - \epsilon) . \end{aligned}$$

The reader can verify that $\theta^i = (\theta^i(s_1^i, s_1^j), \dots, \theta^i(s_2^i, s_2^j)) = (2\frac{1-3\epsilon}{3\epsilon}, 0, 0, \frac{1}{3\epsilon})$ does belong to $\Theta_{s_1^i}(\hat{\sigma}^j)$ but not to E . Therefore, $\Theta_{s_1^i}(\cdot)$ is not upper hemi continuous.



References

- BALKENBORG, D. (1994): “Strictness and Evolutionary Stability,” *Hebrew University of Jerusalem Center of Rationality and Interactive Decision Theory Discussion Paper*, 52.
- BAUER, H. (1992): *Maß- und Integrationstheorie*. De Gruyter, Berlin.
- BJÖRNERSTEDT, J., AND J. W. WEIBULL (1996): *Nash Equilibrium and Evolution by Imitation* pp. 155–71. New York: Macmillan.
- ELY, J. C., AND O. YILANKAYA (2001): “Nash Equilibrium and the Evolution of Preferences,” *Journal of Economic Theory*, 97, 255–72.
- HEIFETZ, A., C. SHANNON, AND Y. SPIEGEL (2007): “The dynamic evolution of preferences,” *Economic Theory*, 32, 251–286.
- OECHSSLER, J., AND F. RIEDEL (2001): “Evolutionary dynamics on infinite strategy spaces,” *Economic Theory*, 17, 141–62.
- OK, E. A., AND F. VEGA-REDONDO (2001): “On the Evolution of Individualistic Preferences: An Incomplete Information Scenario,” *Journal of Economic Theory*, 97, 231–54.
- SAMUELSON, L., AND J. ZHANG (1992): “Evolutionary Stability in Asymmetric Games,” *Journal of Economic Theory*, 57, 363–91.
- SANDHOLM, W. H. (2001): “Preference Evolution, Two-Speed Dynamics, and Rapid Social Change,” *Review of Economic Dynamics*, 4, 637–79.