History Invariance in Framed Repeated Games

(working paper #04/AC/02)

André Casajus[1]

(November 2004, this version: March 8, 2005, 12:41)

Abstract

In this paper, we extend the framing of games and language invariance (Casajus, 2001) to repeated games. Our concept of history invariance in framed repeated games unifies and generalizes the Crawford and Haller (1990) and Blume (2000) approaches for learning in repeated games. As a result, dynamic focal points in repeated coordination games can be supported in a general fashion.

Journal of Economic Literature Classification Number: C72.

Key Words: Dynamic focal point, symmetry, history, matching, coordination, learning.

[1]Universität Leipzig, Wirtschaftswissenschaftliche Fakultät, Professur für Mikroökonomik, PF 920, D-04009 Leipzig, Germany. Email: casajus@wifa.uni-leipzig.de

I thank three anonymous referees for helpful comments. Of course, the usual disclaimer applies. I also thank the School of Economics at the University of East Anglia for their hospitality during the final draft of this paper.

1. Introduction

Three persons are playing the following game: They are each given a basket containing one white ball, one red ball, and one black ball. Independently, they have to take out one object. If all balls drawn are red they all get the payoff 1; if exactly two black balls and one white ball were drawn then all player get a payoff 2. In all other cases the players get the payoff 0. While a player's strategies are payoff distinguished, the players themselves are not so. Hence, the Harsanyi and Selten (1988) symmetry invariance requirement prescribes them to take the "same" strategy. Payoff dominance then selects the always-red rule which gives the payoff 1 to all players. But the players could do better if they could single out one of them who takes the white ball while the other two take the black ones. This would be possible if the persons were not just anonymous players but for example Ada, Onno, and Max, i.e. the players were concrete persons commonly "labelled" in some or another way. Then, Ada as the single female stands out and should cast the white-ball player.

Schelling (1960) introduced the term of a focal point—uniqueness in some conspicuous respect—for this phenomenon. In the example where there is no repetition of the game or—more general—there are no sequential moves, the description or labelling is given from the beginning of the game. In a sense, these focal points are static. Besides Schelling's early sample, real-life observations and recent experiments (e.g. Mehta, Starmer and Sugden 1994a, 1994b; Bacharach and Bernasconi 1997, Blume and Gneezy 2000) support the focal point effect. Moreover, several attempts have been made to formalize these static focal points (Bacharach 1991, 1993; Sugden 1995; Bacharach and Stahl 2000; Casajus 2001; Janssen 2001). Janssen (1998) provides a survey of focal point theories.

Yet, Schelling already suggests that focal points may rest upon precedence, e.g. the players' choices at earlier occurrences of the same situation (as in repeated coordination games) or their choices in other situations (as in cheap talk games). In particular, the players may take some inferior action in early stages of the game in order to generate a description of actions or players that for example enables coordination.

Reconsider our example as an infinitely repeated game with uniform discounting and where the players commonly observe their choices at the end of each stage. On the one hand, the players can coordinate for sure by following the always-red rule. However, this would not distinguish the players in order to enable coordination at the later stages. On the other hand, the players could randomize on the white and the black ball which renders them with an expected payoff of $\frac{3}{4}$ at that stage which is lower than the payoff 1 in the case of coordination on the red ball. But doing so the players could be distinguished at the following stages: If they did not take the same ball at the first stage, one of them stands out. Hence in the repeated game, the players may consider the sophisticated rule: At the first stage do (a) "Randomize on the withe and the black ball." At the other stages do (b) "If all players took the same ball at the previous stage then do (a)" or (c) "If one player stands out with his choice at the previous stage, this player takes the white ball and the other two players take the black one." With the discount parameter δ, the expected payoff of the always-all-red rule is $\frac{1}{1-\delta}$. The expected payoff of the sophisticated rule can be calculated as follows: With probability $\frac{1}{4}$ all players choose the same ball in the first round. With this outcome, the players are still indistinguishable and therefore after this stage in the same position as before. With probability $\frac{3}{8}$, two players took a black ball and one player the with one. This yields a payoff of 2 at this and at all following stages. Again, with probability $\frac{3}{8}$, two players took a white ball and one player the black one. This yields a payoff of 0 at this stage but a payoff of 2 at all following stages. Since this rule is stationary, the expected payoff u of this rule satisfies

$$u = \frac{1}{4}\delta u + \frac{3}{8}\frac{2}{1-\delta} + \frac{3}{8}\frac{2\delta}{1-\delta},$$

i.e. we have

$$u = \frac{3(1+\delta)}{(1-\delta)(4-\delta)}.$$

It is easy to check, that the sophisticated rule is better then the always-red rule if $\delta > \frac{1}{4}$, i.e. if the players are sufficiently patient. Note that this rule involves a dominated stage-game equilibrium at the first stage.

Besides Kramarz (1996), who seems to have coined the term, Crawford and Haller (1990) (henceforth CH), Casajus (2001), and, in a sense, Blume (2000) made attempts to formalize these *dynamic focal points.* Blume and Gneezy (2000) find some empirical support for the CH rules for optimal play in repeated pure-coordination games.

In this paper, we combine framed games (Casajus, 2001) with learning in repeated games (CH, Blume 2000) This way we are able to unify and to generalize these approaches. In addition, we explicate some of the details CH mention in the Appendix. In particular, we obtain a learning rule similar to that of Blume. Departing from the symmetries of framed strategic games we derive a concept of history invariance for infinitely repeated framed games as general way to formalize the idea of history distinguishedness. Applying history invariance to pure coordination games, we are able to support the CH optimal rules. In addition, we generalize the CH results for n-player coordination games.

The paper is organized as follows: Basic definitions and notation as well as the framed strategic game approach are presented in the next section. The third section develops the idea of history invariance and presents its essential implications. In the fourth one, history invariance is applied to repeated pure coordination games. In addition, we establish the relation between history invariance in framed repeated games and the CH approach. The fifth section relates history invariance to Blume's learning with partial a language. An Appendix contains the omitted proofs.

2. Language invariance in the stage game

Casajus (2001) proposes the framed game approach to focal points as a generalization, unification, and extension of the ideas of Bacharach (1991, 1993), Sugden (1995), and Janssen (2001). Core of this approach are framed strategic games (FSG) and the language invariance requirement. While frames of strategic games represent Bacharach style multi-dimensional labellings of the strategies, the language invariance requirement just extends the Harsanyi and Selten (1988) invariance under isomorphism to these framed strategic games. As we will focus on implications of invariance within one game, we just provide the definitions of symmetries—isomorphisms into a the same game. Moreover, we slightly simplify the notion of a FSG symmetry in order to keep the exposition more coherent.

2.1. Basic notation.

For two sets X and X' let $S(X, X')$ denote the set of bijections from X to X'; we also write $S(X)$ for the symmetric group $S(X, X)$ which contains the identity mapping denoted by id_X . The set of natural numbers including the 0 is denoted by $\mathbb{N}$. For $n \in \mathbb{N}$, $\mathbb{N}_n$ denotes the set of the first n natural numbers excluding the 0, $S_n = S(\mathbb{N}_n)$.

2.2. Payoff invariance.

A finite strategic form is a tuple $\gamma = (I, (A_i)_{i\in I})$ where I is a non-empty and finite set of players, A_i the non-empty and finite set of player i's pure (stage) actions a_i; $A = \prod_{i\in I} A_i$. B_i denotes player i's set of mixed actions b_i where $b_i(a_i)$ is the probability of player i's action a_i in b_i. The set $\prod_{i\in I} B_i$ of mixed-action profiles b is denoted B. In particular, we consider the class $\mathcal{C}$ of general coordination games, i.e. of strategic games where all players have the same payoff function u, and we consider the two-player pure-coordination games C_n, $n > 1$: $I = \{1, -1\}$. There is a set Ω of n objects $\omega, \omega', \ldots$, $A_i = \{i\} \times \Omega$, $u((1,\omega), (-1,\omega))$ $= 1$, and $u((1,\omega), (-1,\omega')) = 0$ for $\omega \neq \omega'$.

A pre-symmetry of γ is a system of bijective mappings $\mathbf{r} = (\pi, (r_i)_{i\in I})$, $\pi \in S(I)$ and $r_i \in S(A_i, A_{\pi(i)})$. Abusing notation, $\mathbf{r}$ induces a bijection $\mathbf{r} = (\mathbf{r}_i)_{i\in I} : A \to \bar{A}$, $a \mapsto \mathbf{r}(a)$ via

$$\mathbf{r}_{\pi(i)}(a) := r_i(a_i), \quad i \in I,\ a \in A. \tag{2.1}$$

$\mathbf{r}$ is extended to B by $\mathbf{r}_{\pi(i)}(b)(r_i(a_i)) = b_i(a_i)$ for all $i \in I$, $a_i \in A_i$, and $b \in B$.

Strategic games $G = (\gamma, u)$ are strategic forms with a payoff function $u = (u_i)_{i\in I} : A \to \mathbb{R}^I$ which is extended to B as usual. Since players, actions, and payoffs are just the game theorists names for real-life agents, their options and preferences, a solution concept should not rest upon these to a certain extent arbitrary assignments. In order to formalize this idea, Harsanyi and Selten (1988) introduce isomorphisms of strategic games and require solutions to be invariant under isomorphism. For one-point solution concepts, this implies the selection of a symmetry invariant action profile. A payoff-symmetry of (γ, u) is a pre-symmetry $\mathbf{r} = (\pi, (r_i)_{i\in I})$ of γ such that for all $i \in I$ there are $\alpha_i, \beta_i \in \mathbb{R}$, $\alpha_i > 0$ such that

$$\bar{u}_{\pi(i)}(\mathbf{r}(a)) = \alpha_i u_i(a) + \beta_i, \quad a \in A. \tag{2.2}$$

Note that there is at least the payoff-symmetry $\mathrm{id}_\gamma = (\mathrm{id}_I, (\mathrm{id}_{A_i})_{i\in I})$.

Together with the composition of mappings the pre-symmetries of γ form the pre-symmetry group $\mathfrak{S}(\gamma)$, and the payoff-symmetries of $G = \gamma(u)$ form the group $\mathfrak{S}(G)$ which is a subgroup of $\mathfrak{S}(\gamma)$. The unit in both groups is id_γ . An action profile b is called $\mathfrak{S}'$-invariant for some subgroup $\mathfrak{S}' \subset \mathfrak{S}(\gamma)$ iff $\mathbf{r}(b) = b$ for all $\mathbf{r} \in \mathfrak{S}'$. Players i, i' (actions $a_i, a_{i'}$) are called $\mathfrak{S}'$-symmetric if there is some pre-symmetry $(\pi, (r_i)_{i\in I}) \in \mathfrak{S}'$ such that $\pi(i) = i'$ ($r_i(a_i) = a_{i'}$), otherwise they are called $\mathfrak{S}'$-indistinguishable. Finally, we define the stabilizer $\mathfrak{S}'_a$ of some action profile a with respect to $\mathfrak{S}'$ as the subgroup those of pre-symmetries $\mathbf{r}$ from $\mathfrak{S}'$ that leave a fixed, i.e. which satisfy $\mathbf{r}(a) = a$.

While pre-symmetries just preserve the assignment of actions to players and are restricted only by the cardinality of the action sets, the group $\mathfrak{S}(G)$ describes the extent to which players and actions can be distinguished in terms of payoffs. Players i, i' (actions $a_i, a_{i'}$) are called payoff-symmetric if they are $\mathfrak{S}(G)$-symmetric; they are called payoff distinguished if they are not so. An action profile b is called payoff invariant iff it is $\mathfrak{S}(G)$-invariant. Payoff invariant action profiles then can be characterized by assigning the same probabilities to symmetric actions.

Example 2.1. Reconsider our leading example. The following table presents a strategic game $G = \gamma(u)$ for this situation:

$(3, r)$

	$(2, r)$	$(2, b)$	$(2, w)$
$(1, r)$	1,1,1	0,0,0	0,0,0
$(1, b)$	0,0,0	0,0,0	0,0,0
$(1, w)$	0,0,0	0,0,0	0,0,0

$(3, b)$

	$(2, r)$	$(2, b)$	$(2, w)$
$(1, r)$	0,0,0	0,0,0	0,0,0
$(1, b)$	0,0,0	0,0,0	2,2,2
$(1, w)$	0,0,0	2,2,2	0,0,0

$(3, w)$

	$(2, r)$	$(2, b)$	$(2, b)$
$(1, r)$	0,0,0	0,0,0	0,0,0
$(1, b)$	0,0,0	2,2,2	0,0,0
$(1, w)$	0,0,0	0,0,0	0,0,0

The players 1, 2, and 3 choose among the objects in $\Omega = \{r, b, w\}$ where r (b, w) stands for the red (black,white) ball. This gives the actions $(i, \omega) \in I \times \Omega$ with the obvious interpretation. Player 1 $(2, 3)$ chooses the row (column, matrix). The first (second, third) number in each matrix entry is the payoff of player 1 $(2, 3)$. In this game, all of a player's strategies are payoff distinguished but the players itself are not. In particular, we have

$$\mathfrak{S}(\gamma) = \left\{ \left(\pi, (r_i)_{i \in I}\right) | \pi \in S_3, r_i : (i, \omega) \mapsto (\pi(i), \rho_i(\omega)), \rho_i \in S(\Omega) \right\}$$

and

$$\mathfrak{S}(G) = \left\{ \left(\pi, (r_i)_{i \in I}\right) | \pi \in S_3, r_i : (i, \omega) \mapsto (\pi(i), \omega) \right\}. \tag{2.3}$$

Since the players have the same number of actions, the pre-symmetries are not restricted. The payoff symmetries reflect the payoff distinguishedness of the actions by fixing the object part of the r_i to id_Ω. Hence, in a payoff symmetric action profile, the players, in fact, chose the same probability distributions on Ω. This gives two payoff invariant equilibria: All players take the red ball or all players choose the black and the white one each with probability $\frac{1}{2}$. The first equilibrium with payoffs $(1, 1, 1)$ dominates the second one with payoffs $\left(\frac{3}{4}, \frac{3}{4}, \frac{3}{4}\right)$.

2.3. **Language invariance.** A frame of a strategic form is a triple $F = (C, \Lambda, (\ell_i)_{i \in I})$ where C denotes the non-empty and finite set of properties (e.g. color or shape), Λ denotes the non-empty and finite set of labels (e.g. red or round), and $\ell_i : A_i \times C \to \Lambda$ denotes player i's label function which assigns multi-dimensional labels to his actions. A pair (G, F) is called a framed strategic game. Even though a frame represents the way agents view their options it is made up by the game theorist's names for properties and labels. Again, a solution concept should not rest upon these to a certain extent arbitrary assignments. This is formalized by language symmetries and language invariance. For a more detailed motivation see Casajus (2001). A language symmetry of (G, F) is a payoff symmetry $\mathbf{r} = (\pi, (r_i)_{i \in I})$ of G such that there are bijections $\mu \in S(P)$ and $\varphi \in S(\Lambda)$ that satisfy

$$\ell_{\pi(i)}(r_i(a_i), \varphi(c)) = \varphi(\ell_i(a_i, c)), \quad i \in I, a_i \in A_{i,}, c \in C. \tag{2.4}$$

The language symmetries of (G, F) form the language-symmetry group $\mathfrak{S}(G, F)$ which is a subgroup of $\mathfrak{S}(G)$. An action profile b is called language invariant if it is $\mathfrak{S}(G, F)$-invariant. Players i, i' (actions $a_i, a_{i'}$) are called language symmetric if they are $\mathfrak{S}(G, F)$-symmetric; they are called language distinguished if they are not so. Language invariant action profiles then can be characterized by assigning the same probabilities to language-symmetric actions. Hence, the language symmetry group $\mathfrak{S}(G, F)$ describes the extent to which players and actions can be distinguished by the frame and by the payoffs. In a sense, a frame represents the (common) language of the agents to describe or to distinguish the agents and their options beyond the payoffs. The absence of a common language can be modelled by a trivial frame with just one property and one label, for example. Obviously, we then have $\mathfrak{S}(G, F) = \mathfrak{S}(G)$.

Example 2.2. Reconsider our leading example. Suppose player 1 is Ada. The players' perception of their sexes in the game from Example 2.1 can be modelled by the following frame F: There is just one property, `sex`, and two labels, `female` and `male` which are assigned to the actions as in the following table:

(k, i), $k \in \Omega$	$(k, 1)$	$(k, 2)$	$(k, 3)$
$\ell_i((k, i), \texttt{sex})$	`female`	`male`	`male`

I.e., all of a player's actions are labelled with his/her sex. This way the players' actions remain frame symmetric, but player 1, Ada, is frame distinguished from the other two players who are frame symmetric. Therefore, the payoff symmetry group $\mathfrak{S}(F)$ in (2.3) shrinks to

$$\mathfrak{S}(G, F) = \left\{ \left(\pi, (r_i)_{i \in I}\right) | \pi \in S_3, \pi(1) = 1, r_i : (i, \omega) \mapsto (\pi(i), \omega) \right\}$$

i.e. language symmetries fix player 1 and do not allow for permutations of the objects. Hence, language invariance just says that player 2 and 3 have to choose the same probability distributions on the objects while player 1 is free to deviate.

In contrast to payoff symmetry, the equilibrium $((1,w),(2,b),(3,b))$ (payoff 2) is language invariant and dominates the all-red equilibrium (payoff 1) which also is language invariant.

3. History invariance in repeated games

3.1. **Repeated games.** We first define an infinitely repeated framed game $\left(G^{\delta},F\right)$ in some detail: Let $\delta\in\left]0,1\right[$ be the discount parameter and (G,F) a FSG. The infinitely repeated strategic game $G^{\delta}=(I,(\Sigma_i)_{i\in I},(u_i^{\delta})_{i\in I})$ then is defined as usual: We set $H^0=\{\eta\}$, $H^t=\prod_{\tau=0}^{t-1}A$ for $t\in\mathbb{N}$, $t>0$, and $H=\{H^t|t\in\mathbb{N}\}$. The elements h of H^t are called t-histories, η is the unique 0-history—the empty one, $t\left(h\right)$ denotes the order of h, i.e. $h\in H^{t(h)}$. For $h\in H$ and $\tau<t\left(h\right)$, $h_{\tau}\in A$ denotes the component of h at stage τ. For $t'\leq t\left(h\right)$, $h^{t'}$ denotes the t'-subhistory of h, i.e. the t'-history that coincides with h till stage $t'-1$, i.e. $h_{\tau}=h_{\tau}^{t'}$ for all $\tau<t'$.

The set of player i's strategies $\sigma_i:H\rightarrow B_i$ is denoted Σ_i, $\Sigma=\prod_{i\in I}\Sigma_i$ denotes the set of strategy profiles. The payoff functions u_i^{δ} are defined inductively: $\operatorname{pr}\left(\eta|\sigma\right)=1$ and

$$\operatorname{pr}\left(\left(h,a\right)|\sigma\right)=\operatorname{pr}\left(h|\sigma\right)\cdot\prod_{i\in I}\sigma_i\left(h\right)\left(a_i\right)\tag{3.1}$$

for $h\neq\eta$, where $\operatorname{pr}\left(h,\sigma\right)$ denotes the probability that the h is generated by σ. The ex-ante-probability of the action profile a at stage t under σ then is given by

$$\operatorname{pr}\left(a|t,\sigma\right)=\sum_{h\in H^t}\operatorname{pr}\left(h|\sigma\right)\cdot\prod_{i\in I}\sigma_i\left(h\right)\left(a_i\right),\tag{3.2}$$

the expected payoff of player i at stage t and in the whole game by

$$u_i^t\left(\sigma\right):=\sum_{a\in A}\operatorname{pr}\left(a|t,\sigma\right)u_i\left(a\right)\quad\text{and}\quad u_i^{\delta}\left(\sigma\right)=\left(1-\delta\right)\sum_{t=0}^{\infty}\delta^t u_i^t\left(\sigma\right).\tag{3.3}$$

The continuation payoff of σ at h is denoted $u_i^{\delta}\left(\sigma|h\right)$. Note that $u_i^{\delta}\left(\sigma\right)=u_i^{\delta}\left(\sigma|\eta\right)$. A strategy profile σ^* is called subgame optimal (SOSP) iff $u_i^{\delta}\left(\sigma^*|h\right)\geq u_i^{\delta}\left(\sigma|h\right)$ for all $i\in I$; $h\in H$, $\sigma\in\Sigma$, it is just called optimal iff the former holds for $h=\eta$. A strategy combination σ^* is a subgame perfect equilibrium (SPE) iff $u_i^{\delta}\left(\sigma^*|h\right)\geq u_i^{\delta}\left(\sigma_i\sigma_{-i}^*|h\right)$ for all $i\in I$, $h\in H$, $\sigma_i\in\Sigma_i$, it is just called an equilibrium iff the former holds for $h=\eta$.

3.2. **Symmetry invariance.** Language invariant action profiles give a precise meaning to our intuition that action profiles to be chosen should reflect the symmetry of actions and players in the stage game with respect to both the payoffs and the descriptions of the agents. For strategy profiles in the repeated game one may want more, namely that the symmetry of histories induced by the language symmetry of stage game actions is respected. This is can be formalized by an extension of the language invariance of action profiles in the stage game to strategy profiles in the repeated game.

In this section, we first extend strategic game payoff invariance to repeated games. Our definition is inspired by the weak isomorphism of extensive games (Casajus, 2001). Basically, invariance under weak isomorphism means that a solution does not depend on the particular way in which the game theorist draws the game tree and scales the payoffs. A detailed motivation is given in the Appendix. A symmetry of a repeated game G^δ is a system of payoff symmetries of G: $\ddot{\mathbf{r}} = \left(\mathbf{r}^h\right)_{h\in H}$, $\mathbf{r}^h = \left(\pi, \left(r_i^h\right)_{i\in I}\right)$. I.e., we have symmetries for all histories that involve the same bijection on players. By the latter, repeated game symmetry respects the identity of players during the course of the game. The former indicates that repeated game symmetry abstracts from the fact that a fixed game is repeated. This is justified if one takes the viewpoint that the names for the agents' options are the game theorist's.

Any symmetry $\ddot{\mathbf{r}}$ induces a bijection on the histories, abusing notation, $\ddot{\mathbf{r}} \in S(H)$ by applying its components $\mathbf{r}^{h^\tau}$ to the components h_τ of a history h,

$$\ddot{\mathbf{r}}(h) = \ddot{\mathbf{r}}\left(h_0, h_1, \ldots, h_{t(h)-1}\right) = \left(\mathbf{r}^{h^0}(h_0), \mathbf{r}^{h^1}(h_1), \ldots, \mathbf{r}^{h^{t(h)-1}}\left(h_{t(h)-1}\right)\right)$$

i.e. by applying $\ddot{\mathbf{r}}$ stagewise by $\ddot{\mathbf{r}}(\eta) = \eta$ and

$$\ddot{\mathbf{r}}_\tau(h) := \mathbf{r}^{h^\tau}(h_\tau), \quad h \in H, \tau < t(h). \tag{3.4}$$

I.e. how the action profile h_τ is mapped depends on the τ-subhistory h^t of h. $\ddot{\mathbf{r}}$ also induces a notion of symmetry of histories: h and $\bar{h}$ are symmetric if there is a history symmetry $\ddot{\mathbf{r}}$ such that $\ddot{\mathbf{r}}(h) = \bar{h}$. It is clear that histories of the same order only can be symmetric. Again, abusing notation, $\ddot{\mathbf{r}}$ is extended to a bijection $\ddot{\mathbf{r}} \in S(\Sigma)$ by

$$\ddot{\mathbf{r}}_{\pi(i)}(\sigma)(\ddot{\mathbf{r}}(h))\left(r_i^h(a_i)\right) = \sigma_i(h)(a_i), \quad h \in H, i \in I, \sigma \in \Sigma. \tag{3.5}$$

This way the players' actions are mapped in accordance with the histories: While the history h is mapped to $\ddot{\mathbf{r}}(h)$ the action profile $\sigma(h)$ at h is mapped to the action profile $\ddot{\mathbf{r}}(\sigma(h))$ at $\ddot{\mathbf{r}}(h)$. The following lemma shows that these definitions are compatible with the repeated game payoffs.

Lemma 3.1. *For all symmetries* $\ddot{\mathbf{r}}$ *of* G^δ *and all* $i \in I$, *there are* $\alpha_i, \beta_i \in \mathbb{R}$, $\alpha_i > 0$ *such that* $u^\delta_{\pi(i)}(\ddot{\mathbf{r}}(\sigma)) = \alpha_i u^\delta_i(\sigma) + \beta_i$ *for all* $\sigma \in \Sigma$.

Now, we can extend the concept of language invariance to framed repeated games in a straightforward way. In essence, it means that symmetric actions are played at symmetric histories.

Definition 3.2. *A strategy profile* σ *is symmetry invariant iff we have* $\ddot{\mathbf{r}}(\sigma) = \sigma$ *for all symmetries* $\ddot{\mathbf{r}}$ *of* G^δ.

As one might have suspected, however, symmetry invariance in repeated games does not restrict the players behavior to much: At some fixed history, the players are just restricted to choose a payoff invariant action profile. This immediate from the fact that a repeated game symmetry $\ddot{\mathbf{r}}$ consists of unrelated stage game symmetries for all histories. This way it is possible to fix some history h by setting all $\mathbf{r}^{\bar{h}}$ to id_G while $\mathbf{r}^h$ can be arbitrarily chosen. The only power of repeated-game symmetry lies in its implication for the players behavior at different histories of the same order. At symmetric histories, the players have to choose the same symmetry invariant action profile.

Theorem 3.3. *In* G^δ, σ *is symmetry invariant iff* $\sigma(h)$ *is payoff invariant for all* h *and* $\sigma(h) = \sigma(\bar{h})$ *for all symmetric histories* h *and* $\bar{h}$.

The impotence of repeated-game symmetry-invariance is especially apparent in highly symmetric games. For the two-player pure-coordination games we have

$$\mathfrak{S}(C_n) = \left\{ \left(\pi, (r_i)_{i \in I}\right) | \pi \in S_2, \rho \in S(\Omega),\ r_i : (i, \omega) \mapsto (\pi(i), \rho(\omega)) \right\}, \tag{3.6}$$

i.e. both players and all actions are payoff-symmetric in C_n. Therefore, randomization is the unique payoff-invariant action profile in C_n which gives the payoff $\frac{1}{n}$ instead of 1 in the case of coordination. In view of our remarks above, repeating C_n does not help: Randomization at all stages still is the unique symmetry invariant strategy profile in C^δ_n. This again emphasizes that players and actions are the game theorist's representatives for the real-world agents and their options. Hence, the players can not exploit the fact that the *same* game is repeated. I.e. symmetry invariance implies that the players have no common language to describe actions across the stages.

3.3. **History invariance.** In this section, we combine the ideas of language invariance and repeated game symmetry invariance. A framed repeated game (FRG) simply is a pair (G^δ, F) consisting of repeated game G^δ and a frame $F = (C, \Lambda, (\ell_i)_{i\in I})$ for the stage game G. A history symmetry for a FRG is a symmetry $\ddot{\mathbf{r}} = \left(\mathbf{r}^h\right)_{h\in H}$, $\mathbf{r}^h = (\pi, \left(r_i^h\right)_{i\in I})$ of G^δ such that there are bijections $\mu \in S(C)$ and $\varphi \in S(\Lambda)$ such that $\ell_{\pi(i)} \circ \left(r_i^h \times \mu\right) = \varphi \circ \ell_i$, i.e.

$$(3.7) \qquad \ell_{\pi(i)}\left(r_i^h(a_i), \mu(c)\right) = \varphi\left(\ell_i(a_i, c)\right), \quad h \in H, c \in C, i \in I, a_i \in A_i.$$

A strategy profile σ in G^δ is called history invariant (HI) if $\ddot{\mathbf{r}}(\sigma) = \sigma$ for all history symmetries $\ddot{\mathbf{r}}$ of G^δ.

From (2.4) it is clear that the $\ddot{\mathbf{r}}$ constitutes a system of language symmetries $\mathbf{r}^h$ of the framed stage game (G, F). The crucial point in this definition is that the bijections μ and φ are the same for all components of $\ddot{\mathbf{r}}$. Via μ and φ, the payoff symmetries $\ddot{\mathbf{r}}^h$ are related across the stages. This feature of history symmetry models that players perceive a repeated game indeed as a repeated game and therefore may be able to exploit past behavior.

Absence of a common language can be modelled by the trivial frame with just one property and one label. In this case, (3.7) has no effect—history invariance and symmetry invariance then coincide.

In real life, agents have a description of the game that allows to distinguish all their actions. Otherwise, one could argue, it would not make sense to speak of different actions. Consider three 'identical' balls on a tray among which the players have to choose. Despite being identical as such they have different positions on the tray which could be modelled by a frame. Suppose now that these balls are in a bag and the players first have to take them out, one by one, and then have to make their choice. In this case, the balls are just labelled by the order they have been taken from the bag. But as all orders can be regarded as equally probable, choosing the first one, for example, amounts to randomizing on all balls. In fact, this is the players' only real option in this case.

In the following, we therefore focus on atomic frames, i.e. frames where the actions have pairwise different label tuples, i.e. for all actions a_i and $a_{i'}$ there is some property c such that $\ell_i(a_i, c) \neq \ell_{i'}(a_{i'}, c)$. The frame in Example 2.2 is not atomic. Yet, any frame can be transformed into an atomic frame without affecting the language symmetry group $\mathfrak{S}(G, F)$. We just have to add a new property `dist` and the labels `dist`a_i for all $i \in I$ and $a_i \in A_i$ and to extend ℓ_i by $\ell_i(a_i, \texttt{dist}) = \texttt{dist}a_i$. From (2.4), it is immediate that atomic frames have the following property: If $\mathbf{r}$ and $\bar{\mathbf{r}}$ are language symmetries via the same bijections $\mu \in S(P)$ and $\varphi \in S(\Lambda)$ then

$\mathbf{r} = \bar{\mathbf{r}}$. By (3.7), the language symmetries $\mathbf{r}^h$ that make up a history symmetry $\ddot{\mathbf{r}}$ are identical. This makes it easier to determine history invariant strategy profiles.

By (3.5), the local strategies at different t-histories are related by HI in the obvious way. In the following, we explore the implications of HI for the local strategies, i.e. the actions at some given history $h \in H$. In order to do this, we have to consider those language symmetries that leave h fixed. Let $\mathfrak{S}_h(G, F)$ denote the stabilizer of h in $\mathfrak{S}(G, F)$,

$$\mathfrak{S}_h(G, F) = \{\mathbf{r} | \mathbf{r} \in \mathfrak{S}(G, F) : \mathbf{r}(h_\tau) = h_\tau, \tau < t(h)\}.$$

The relation between the group $\mathfrak{S}_h(G, F)$ and HI then is immediate from (3.4) and Definition 3.2.

Lemma 3.4. *(i) If σ is HI in G^δ then $\sigma(h)$ is $\mathfrak{S}_h(G, F)$-invariant for all $h \in H$. (ii) If $b \in B$ is $\mathfrak{S}_h(G, F)$-invariant, then there is a history invariant strategy profile σ such that $\sigma(h) = b$.*

Hence, the task of determining HI strategy profiles basically consists in the determination of the groups $\mathfrak{S}_h(G, F)$. In view of (3.4), the stabilizer of h can be derived from the symmetry groups of the underlying stage game: For $G = \gamma(u)$ and an atomic frame F, we have

$$\mathfrak{S}_h(\gamma(u, F)) = \bigcap_{\tau=0}^{t(h)-1} \mathfrak{S}_{h_\tau}(\gamma(u, F)) = \mathfrak{S}(G, F) \cap \bigcap_{\tau=0}^{t(h)-1} \mathfrak{S}_{h_\tau}(\gamma) \tag{3.8}$$

where $\mathfrak{S}_{h_\tau}(\gamma)$ denotes the stabilizer of h_τ in $\mathfrak{S}(\gamma)$.

Hence, the ideas behind HI for atomic FRG are the following: At the beginning of the game, players and actions are distinguished only by the payoffs and by the initial frame, $\mathfrak{S}_\eta(G, F) = \mathfrak{S}(G, F)$. Hence, at stage 0, the players have to choose a language invariant action profile.

In the course of the repeated game, the description of players and actions may change, e.g. by the fact that an action has been chosen before or that some player had chosen a certain action. Of course, this is possible only if the players can identify actions across stages. Note that this is not already implied by the fact that the stage game does not change. In the leading example, the fixed object set enables this identification. But changing the objects from stage to stage leaves the stage game unaffected; yet playing this repeated game would not produce exploitable histories. Formally, the possibility of identification across stages is modelled by considering the same atomic frame for all stages together with the application of a single pre-symmetry to all stages—implied by atomicness—in (3.4) and (3.5). Practically, this can be interpreted as a that the actions are virtually the same, e.g. the same objects to be chosen.

Further, it is implicit in this definition that the players can distinguish the stages and that they remember all previous actions. The former is embodied in that (3.4) applies the pre-symmetry stagewise, the latter in that (3.4) does so for all stages.

Example 3.5. Reconsider our leading example. We have demonstrated in Example 2.2 that framing the players' sex destroys the symmetry of players and enables the "better" equilibrium. Suppose the game is accompanied by an atomic frame that does not distinguish the players, i.e. we have $\mathfrak{S}(G,F) = \mathfrak{S}(G)$, but the game is repeated infinitely. Since the objects (and the related actions) are payoff distinguished, at the 1-history $h = ((1,b),(2,b),(3,b))$, this gives the stabilizer $\mathfrak{S}_h(G,F) = \mathfrak{S}(G)$, i.e. the players remain indistinguishable. At the 1-history $\bar{h} = ((1,r),(2,w),(3,w))$, in contrast, HI allows the interchange of player 2 and player 3 but not the interchange of player 1 and player 2 or 3, i.e.

$$\mathfrak{S}(G) = \left\{\left(\pi,(r_i)_{i\in I}\right) | \pi \in S_3, \pi(1) = 1, r_i = \pi \times \mathrm{id}_\Omega, i \in I\right\}.$$

Hence, for example, the sophisticated rule is attainable under history invariance.

There always is some history invariant strategy profile—uniform randomization at all histories. Since framed strategic games have language invariant equilibria (Casajus, 2001), repeated games do also have HI SPE: Playing some fixed language invariant stage game equilibrium irrespective of the history is a SPE, and by (3.8) it is history invariant.

Theorem 3.6. *Any repeated game G^δ has a HI SPE.*

Hence, history invariant SPE of FRG may serve as candidates for a (unique) solutions of repeated games that account for the history. Note that HI does not undo the folk theorem: Since in the repeated prisoners' dilemma a player's stage actions are not symmetric (while the players itself are), e.g. the GRIM strategy (both players begin cooperative and punish for ever in case of deviation) is a HI SPE.

4. Optimal learning in repeated coordination games

Crawford and Haller (1990) consider learning in infinitely repeated two-person coordination games, mainly without a common language. In order to capture the symmetries of the stage game, they introduce the concept of attainability similar to the Harsanyi and Selten (1988) symmetry invariance. For the class of two-person coordination games, Casajus (2001, Theorem 2.3.2) shows that symmetry invariance implies attainability. But it is not yet clear whether the converse holds in general. Nevertheless, the following Lemma allows to extend Theorem 1 of CH to our approach for general coordination games.

Lemma 4.1. *If b^* is payoff invariant in $G \in \mathcal{C}$ and $u(b^*) \geq u(b)$ for all payoff invariant $b \in B$ then b^* is an equilibrium.*

In contrast to CH, we provide both an initial language (the frame together with frame invariance) and an algorithm that determines the “attainable” action profiles at a given history (HI). Even though CH do not specify such an algorithm, they require that it must not use future information to derive implications for a given history—a requirement joint by Blume (2000). In view of (3.8), HI meets this requirement.

In our approach, HI takes the place of admissability. For repeated games based on coordination games, we define an optimal HI strategy profile as a HI strategy profile σ^* such that for all HI strategy profiles σ we have $u^\delta(\sigma^*) \geq u^\delta(\sigma)$; σ^* is called subgame-optimal history-invariant (SOHI) iff $u^\delta(\sigma^*|h) \geq u^\delta(\sigma|h)$ for all t and $h \in H^t$. Together with Lemma 4.1 arguments similar to those of CH's proof, we have

Theorem 4.2. *Any SOHI strategy profile in a framed repeated coordination game is subgame perfect.*

It should be clear that CH's Theorem 2 about the convergence of optimal learning also holds within our setup.

Theorem 4.3. *With probability 1, an optimal HI strategy profile converges to a set of action profiles that all yield players the same stage-game payoffs.*

As an example, CH consider repeated pure coordination games C_n. They derive rules for optimal learning in absence of a common language: At stage 0, uniformly randomize on all objects. Take the same object as before if coordination occurred at the previous stage. In case of discoordination at the previous stage, take the action that had not been taken at that stage if there are three actions or randomize on the two objects taken at the previous stage if there are two or more than five objects. In the following, we demonstrate that these rules can be supported within our approach and are—to some extent—unique.

We explore the implications of HI and subgame optimality for C_n, $n = 2, 3$ and $n \geq 6$. Absence of a common language implies $\mathfrak{S}(C_n, F) = \mathfrak{S}(C_n)$. In addition, we assume that F is atomic. By (3.8), we just have to derive $\mathfrak{S}(C_n)$ and $\mathfrak{S}_a(C_n)$ for all $a \in A$. In C_n, both players and all actions are payoff symmetric, but the objects must be permuted in the same way for both players by a symmetry, i.e.

$$\mathfrak{S}(C_n) = \left\{ (\pi, (r_i)_{i\in I}) | \pi \in S(I), r_i : (i, \omega) \mapsto (i, \rho(\omega)), \rho \in S(\Omega) \right\}. \tag{4.1}$$

For $a = ((1, \bar{\omega}), (-1, \bar{\omega}))$, i.e. both players took the same object, $\mathfrak{S}_a(\gamma_n)$ contains those pre-symmetries from $\mathfrak{S}(C_n)$ that fix the object $\bar{\omega}$. I.e. the players are still symmetric as well as all objects except of $\bar{\omega}$ which now can be distinguished from the other objects, i.e. with the notation from (4.1)

$$\mathfrak{S}_n(\bar{\omega}) := \mathfrak{S}_a(C_n) = \left\{ (\pi, (r_i)_{i\in I}) \in \mathfrak{S}(C_n) | \rho(\bar{\omega}) = \bar{\omega} \right\}. \tag{4.2}$$

By (3.8), this implies that if a appears in a history then coordination at ω is HI at all continuation histories. For $a' = ((1, \bar{\omega}_1), (-1, \bar{\omega}_{-1}))$, $\bar{\omega}_1 \neq \bar{\omega}_{-1}$, i.e. the players took different objects, things are a bit more complicated: The players are still symmetric as well as all objects except of $\bar{\omega}_1$ and $\bar{\omega}_{-1}$. Interestingly, but not too astonishingly, $(i, \bar{\omega}_i)$ and $(i, \bar{\omega}_{-i})$ are not symmetric. The reason for this, of course, is that a player can distinguish between the object he took himself and the object the other player took. As the same is true for the other player, the players are symmetric—$(i, \bar{\omega}_i)$ and $(-i, \bar{\omega}_{-i})$ as well as $(i, \bar{\omega}_{-i})$ and $(-i, \bar{\omega}_i)$ are symmetric. Therefore, $\mathfrak{S}_{a'}(C_n)$ does not change if the players interchange the objects. Again, with the notation from (4.1), we have

$$\mathfrak{S}_{a'}(C_n) = \left\{ (\pi, (r_i)_{i\in I}) \in \mathfrak{S}(C_n) | \pi = \mathrm{id}, \rho(\bar{\omega}_i) = \bar{\omega}_i, \text{ or } \pi(i) = -i, \rho(\bar{\omega}_i) = \bar{\omega}_i \right\}; \tag{4.3}$$

set $\mathfrak{S}_n(\bar{\omega}_1, \bar{\omega}_{-1}) := \mathfrak{S}_{a'}(C_n)$. Hence, if a' appears in a history then uniform randomization on ω_i and ω_{-i} is HI at all continuation histories. Together with (4.2), this already reveals that the CH rules are HI and by Theorem 4.2 are subgame perfect.

Now, we explore to which extent these rules are unique: It is clear that uniform randomization is the unique local strategy at stage 0.

Two objects. If the players did not coordinate at the first t stages, i.e. at the t-history $h = (a^0, a^1, \ldots, a^{t-1})$ where $a^\tau = ((1, \omega), (-1, \omega'))$ or $a^\tau = ((1, \omega'), (-1, \omega))$, we have $\mathfrak{S}_h = \mathfrak{S}_n(\omega, \omega')$. Again, by (4.3), uniform randomization on the two objects is HI at this stage and also is the HI local strategy profile that maximizes the stage payoff, i.e. that maximizes the probability of coordination. Since coordination at the next stage makes onward coordination HI, maximizing the stage payoff also maximizes the probability of creating of the most favorable history for the next stage. Hence, there is a unique SOHI behavior at h. Uniqueness fails if the players got coordinated at some stage of a history: By (3.8) and (4.2), the objects are history distinguished in this case and repeated coordination on any one of them is a SOHI continuation. The latter is fully in line with Goyal and Janssen (1996).

t stages of coordination, $n > 2$. For $a = ((1, \omega), (-1, \omega))$ and $t > 0$, consider the t-history $h[t]$ in which a just appears t times, i.e. the players coordinate on some object ω from stage 0 to stage $t-1$. By (3.8), it is clear that we have $\mathfrak{S}_{h[t]}(C_n) = \mathfrak{S}_n(\omega)$. By (4.2), the players have to take the "same" local strategy which gives all objects except of ω the same probability. Hence, maintaining coordination once achieved is HI. It is clear that this also is the unique SOHI behavior at these histories.

Three objects. The case of coordination at stage 0 has already been done. Consider now discoordination at stage 0, e.g. the 1-history $h = ((1, \omega_1), (-1, \omega_{-1}))$, $\omega_1 \neq \omega_{-1}$. By (4.3), the local strategies $(1, \omega_1)$ and $(-1, \omega_{-1})$ have to be chosen with the same probabilities as well as $(1, \omega_{-1})$ and $(-1, \omega_1)$. Hence, coordination at the third object (ω_3), i.e. the local strategy profile $((1, \omega_3), (-1, \omega_3))$ is the unique stage-payoff maximizing HI local strategy. In addition, it induces the 2-history $\bar{h} = (((1, \omega_1), (-1, \omega_{-1})), ((1, \omega_3), (-1, \omega_3)))$ with $\mathfrak{S}_{\bar{h}}(C_n) = \mathfrak{S}_n(\omega_3) \cap \mathfrak{S}_n(\omega_1, \omega_{-1}) = \mathfrak{S}_n(\omega_1, \omega_{-1})$ which enables coordination on object ω_3 at stage 2 as the unique stage-payoff maximizing HI local strategy at h_1. Thus, coordination on object ω_3 is the unique SOHI behavior at h and at all continuation histories induced by coordinating on ω_3 from then on.

More than five objects. The case of coordination on the same object from the beginning has already been done. If the players did not coordinate at the first t stages but always took the same two objects, i.e. at the t-history $h = (a^0, a^1, \ldots, a^{t-1})$ where $a^\tau = ((1, \omega), (-1, \omega'))$ or $a^\tau = ((1, \omega'), (-1, \omega))$, we have $\mathfrak{S}_h = \mathfrak{S}_n(\omega, \omega')$. Since $n > 5$, randomizing on ω and ω' is the unique stage-payoff maximizing HI local strategy. In addition, CH show that it does not pay deviate from this local strategy in order to create a more favorable history for future play. Hence, it is the unique SOIH behavior at these histories. The same arguments as for $n = 2$ imply that there is no unique SOIH behavior when the players coordinate on ω at stage $t+1$ after discoordinating on ω and ω' until stage t.

Completely distinguishing histories. So far, we have considered the implications of HI on the possible equilibrium paths of the CH rule. Consider now the history h off the equilibrium paths. Let the object set be $\Omega_n = \{\omega_0, \ldots, \omega_{n-1}\}$ and

$$h = (((1, \omega_0), (-1, \omega_0)), ((1, \omega_0), (-1, \omega_1)), \ldots, ((1, \omega_0), (-1, \omega_{n-1}))),$$

i.e. player 1 chose object ω_0 at the first n stages and player -1 chose all n objects (in their natural order). We then have

$$\mathfrak{S}_h(C_n) = \bigcap_{t=0}^{n-1} \mathfrak{S}_n(\omega_0, \omega_t) = \left\{\mathrm{id}_{\gamma_n}\right\},$$

i.e. all actions (objects) and the players are history distinguished at this history and therefore HI does not restrict players' behavior at the continuation histories.

5. Learning with a partial language

Blume (2000) defines a language for a set of objects Ω as a group action on Ω or—equivalently—as a subgroup $\mathfrak{L}$ of the symmetric group $S(\Omega)$. The group $\mathfrak{L}$ acts on Ω in the natural way giving the orbits $\mathfrak{L}\cdot\omega$ which partition Ω. The interpretation of these orbits is the following: The language cannot distinguish between the objects in the same orbit, but objects in different orbits are distinguishable.

Blume languages are special cases of frames in the following sense: Any object set Ω can be identified with the strategic game $\Omega \equiv (\{1\}, \Omega, u_1)$ where u_1 is constant with $\mathfrak{S}(\Omega) \cong S(\Omega)$. Hence, any frame for Ω gives rise to the group $\mathfrak{S}(\Omega, F) \subset S(\Omega)$. Frame invariance requires to assign the same probabilities to the objects within the same orbit of $\mathfrak{S}(\Omega, F)$, i.e. the frame F does not distinguish between them. This is just the interpretation of $\mathfrak{S}(\Omega, F)$ as a Blume language on Ω. The other way round, any Blume language for an object set can be realized by a frame.

Theorem 5.1. *For any subgroup $\mathfrak{L}$ of $S(\Omega)$ there is a frame $F(\mathfrak{L})$ such that $\mathfrak{S}(\Omega, F(\mathfrak{L})) = \mathfrak{L}$.*

While the focus of CH and our approach is, in a sense, on learning enough in order to coordinate effectively, Blume is concerned with learning a complete description of the objects based on observations of objects. As it turns out, the Blume learning rule can be supported by HI.

Blume starts with a partial language on Ω, i.e. a subgroup $\mathfrak{L}$ of $S(\Omega)$, $\mathfrak{L} \neq S(\Omega)$. Given an observation, i.e. a subset O of Ω, he refines the language $\mathfrak{L}$ by considering those elements of $\mathfrak{L}$ that leave the observed subset O fixed. These elements form a subgroup of $\mathfrak{L}$, the stabilizer $\mathfrak{L}_O$ of O, i.e. a language Blume language on Ω. Note that as $\mathfrak{L}_O$ is a subgroup of $\mathfrak{L}$, the orbit partition of $\mathfrak{L}_O$ is not coarser than the orbit partition of $\mathfrak{L}$. Given enough observations, the stabilizer $\mathfrak{L}_O$ becomes the trivial subgroup $\{\mathrm{id}_\Omega\}$, i.e. all objects are distinguished. Of course, the latter holds for $|O| = |\Omega| - 1$. Blume shows that for a partial language less observations may suffice.

In the repeated strategic form $(\Omega, F(\mathfrak{L}))$, the $|\Omega|$-histories can be interpreted as subsets of Ω: The $|\Omega|$-history $h = \left(\omega^0, \ldots, \omega^{|\Omega|-1}\right)$ stands for the set $\Omega(h) := \left\{\omega^0, \ldots, \omega^{|\Omega|-1}\right\}$, i.e. $\Omega(h)$ abstracts from the order of objects and their possibly multiple appearances in h. By (3.8), it is clear that $\mathfrak{S}_h(\Omega, F(\mathfrak{L}))$ is the stabilizer of $\Omega(h)$ within $\mathfrak{S}(\Omega, F) = \mathfrak{L}$. Hence, HI provides the same learning rule as Blume.

Blume languages may be not general enough to explain subtle kinds of focal points. Consider the following framed option set (Ω, F): There are four objects, $\omega_1, \omega_2, \omega_3$, and ω_4 that are labelled with respect to one property c according to the following table:

ω_i	ω_1	ω_2	ω_3	ω_4
$\ell(\omega_i, c)$	A	B	C	C

The group $\mathfrak{S}(\Omega, F)$ is generated by the transposition of ω_1 and ω_2, τ_{12}, and the transposition of ω_3 and ω_4, τ_{34} which gives the orbits $\{\omega_1, \omega_2\}$ and $\{\omega_3, \omega_4\}$. In addition, the stabilizer of ω_1 and of ω_2, i.e. the subgroup of $\mathfrak{S}(\Omega, F)$ that leaves $\omega_1$or ω_2 fixed, is generated by τ_{34} and the stabilizer of ω_3 and of ω_4, the group generated by τ_{12}, are isomorphic. Hence, from $\mathfrak{S}(\Omega, F)$ alone, the orbits $\{\omega_1, \omega_2\}$ and $\{\omega_3, \omega_4\}$ are indistinguishable. Yet, the frame allows a distinction between the two orbits—$\{\omega_1, \omega_2\}$ has two labels whereas $\{\omega_3, \omega_4\}$ has just one.

Yet, the latter may be of importance if the objects are to be chosen within a pure coordination game: Since the objects within one orbit are indistinguishable by the payoffs and by the frame, the players have to choose them with the same probabilities. Since both orbits contain the same number of objects, the frame's structure may be exploited in order to break the tie between the orbits. One could argue that $\{\omega_3, \omega_4\}$ is easier to describe (by one label) than $\{\omega_1, \omega_2\}$ and therefore should be chosen. Of course, other arguments may suggest that the players should randomize on $\{\omega_1, \omega_2\}$. Anyway, $\mathfrak{S}(\Omega, F)$ does not distinguish between both orbits.

6. Appendix

Motivation of repeated game symmetry in Section 3.2. Repeated games can be viewed as infinite extensive games: Any history $h \in H$ corresponds to one information set (h, i) for each player i and at which player i has a pure local strategies (a_i, h, i) for all of his (pure) actions $a_i \in A_i$; the corresponding set is denoted $\mathcal{A}(h, i)$. Let $\mathcal{H}_i$ ($\mathcal{H}_i^t$) denote the set of player i's information sets (at stage t), and let $\mathcal{H}$ ($\mathcal{H}^t$) denote the set of all information sets (at stage t). Let $\mathcal{A}_i$ ($\mathcal{A}_i^t$) denote the set of player i's local strategies (at stage t), and let $\mathcal{A}$ ($\mathcal{A}^t$) denote the set of all local strategies (at stage t). The role of terminal nodes play the infinite histories, i.e. the sequences $\ddot{h} \in \ddot{H} := \prod_{t=0}^{\infty} A$. The set of local actions $\ddot{h}$ runs through is denoted by $\mathcal{A}(\ddot{h}) = \left\{(\ddot{h}_{t,i}, \ddot{h}^t, i) | t \in \mathbb{N}, i \in I\right\}$.

One could think of transferring notions of symmetry of a finite extensive game to the infinite case. Since weak isomorphism (Casajus, 2001) fits the traditional extensive representation of strategic games, it may serve as point of departure for repeated game symmetry. For games without chance mechanism, weak isomorphism basically is isomorphism of the standard form (Harsanyi and Selten, 1988) plus preservation of unordered histories. In particular, a weak isomorphism involves a bijections between the sets of players, between the sets of information sets, and between the sets of actions which are compatible with each other. Translated into the repeated game formalism we have the following: Since G^δ and G have the same set of players, we have a bijection $\pi \in S(I)$. Compatibility of $\ddot{\nu} \in S(\mathcal{H})$ with π then means $\nu(\mathcal{H}_i) = \mathcal{H}_{\pi(i)}$ for all i, i.e. $\ddot{\nu}(h, i) = \left(\bar{h}, \pi(i)\right)$—information sets of one player are mapped onto information sets of the same player. Similarly, compatibility of $\ddot{r} \in S(\mathcal{A})$ with $\ddot{\nu}$ and π means $\ddot{r}(\mathcal{A}(h, i)) = \mathcal{A}(\nu(h, i))$ for all (h, i), i.e. $\ddot{r}(a_i, h, i) = \left(\bar{a}_{\pi(i)}, \nu(h, i)\right)$—local strategies of one information set (player)are mapped onto local strategies of the same information set (player).

For weak isomorphism, preservation of unordered histories is defined in terms of sets of actions that lie on a path from a terminal node to the root. Obviously, this is impossible for infinitely repeated games. Yet, we have the following analogue: Preservation of unordered histories implies a bijection $\theta \in S\left(\ddot{H}\right)$ such that $\ddot{r}$ maps the local strategies $\ddot{h} \in \ddot{H}$ intersects onto the local strategies $\theta(\ddot{h})$ runs through,

$$\ddot{r}(\mathcal{A}(\ddot{h})) = \mathcal{A}(\theta(\ddot{h})). \tag{6.1}$$

The set of all local strategies preceding (a_i, h, i) at previous stages is denoted

$$\mathcal{A}(a_i, h, i) = \{(h_{\tau,i'}, h^\tau, i) \,|\tau < t(h), i' \in I\}. \tag{6.2}$$

Obviously, this set just depends on h. As a first implication of (6.1) actions are mapped stagewise if the stage game has non-trivial players (all players have more than one action). This can be seen as follows: Suppose there is some $\ddot{h}$ such that $(a_i, h, i) \in \mathcal{A}(\ddot{h})$ is mapped onto some action $\ddot{r}(a_i, h, i) = (\bar{a}_{\pi(i)}, \bar{h}, \pi(i))$, $t(\bar{h}) > t(h)$. By (6.1), there is some $(a_i, h', i) \in \mathcal{A}(\ddot{h})$, $t(h') > t(h)$, such that $\ddot{r}(a_i, h', i) \in \mathcal{A}(\bar{a}_{\pi(i)}, \bar{h}, \pi(i))$. By non-triviality, there is some $(a_i', h', i) \in \mathcal{A}(h', i)$, $a_i' \neq a_i$ and some $\ddot{h}'$ such that $(a_i, h, i), (a_i', h', i') \in \mathcal{A}(\ddot{h}')$. Since $\ddot{r}(a_i, h', i)$ and $\ddot{r}(a_i', h', i)$ belong to the same information set and $\ddot{r}(a_i, h, i)$ is on a higher stage than $\ddot{r}(a_i, h', i)$ in $\theta(\ddot{h})$ and $\theta(\ddot{h}')$ runs through $\ddot{r}(a_i, h, i)$, we have $\ddot{r}(a_i, h, i) \notin \theta(\ddot{h}')$—contradicting (6.1). Therefore, $\ddot{\nu}$ and $\ddot{r}$ can be split into a sequences $(\ddot{\nu}^t)_{t\in\mathbb{N}}$, $\ddot{\nu}^t \in S(\mathcal{H}^t)$, $\ddot{\nu}|_{\mathcal{H}^t} = \ddot{\nu}^t$ and $(\ddot{r}^t)_{t\in\mathbb{N}}$, $\ddot{r}^t \in S(\mathcal{A}^t)$, $\ddot{r}|_{\mathcal{A}^t} = \ddot{r}^t$, respectively.

For $a \in A$ and $h \in H$, there is some $\ddot{h}$ such that $(a_i, h, i) \in \mathcal{A}(\ddot{h})$, hence $\mathcal{A}(a_i, h, i) \subset \mathcal{A}(\ddot{h})$ all $i \in I$. By (6.1), we have $\ddot{r}(a_i, h, i) \in \mathcal{A}(\theta(\ddot{h}))$ and in view of our previous findings, $\ddot{r}(\mathcal{A}(a_i, h, i)) \subset \mathcal{A}(\theta(\ddot{h}))$ for all $i \in I$. This implies that histories as part of information sets are mapped in the same way for all players, i.e. $\ddot{\nu}^t$ induces a bijection $\nu^t \in S(H^t)$ such that $\ddot{\nu}^t(h^t, i) = (\nu^t(h^t), \pi(i))$. In addition, we thus have systems $(r_i^h)_{i\in I}$, $r_i^h \in S(A_i, A_{\pi(i)})$ for all $h \in H$ such that $\ddot{r}(a_i, h, i) = (r^h(a_i), \nu^{t(h)}(h), \pi(i))$. I.e., we have a system of pre-symmetries $(\mathbf{r}^h)_{h\in H}$, $\mathbf{r}^h = (\pi, (r_i^h)_{i\in I})$. Thus by (6.2), (6.1) relates the mappings (r_i^h) and (ν^t) as follows: $\nu_\tau^{t(h)}(h) = \mathbf{r}^{h^\tau}(h_\tau)$ for all $h \in H$ and $\tau < t(h)$—histories are mapped stagewise by the appropriate pre-isomorphisms of G.

In finite extensive games, symmetries preserve the preferences on terminal nodes. Analogously, one could require that repeated game symmetry just respects the payoffs associated with infinite histories. As a repeated game is made of identical stage games, however, we feel that repeated game symmetries should respect the preferences at all stages, i.e. that the pre-symmetries of the stage game that constitute repeated game symmetry should be payoff symmetries. As we have seen (Lemma 3.1) this implies preservation of the preferences on infinite histories.

Proof of Lemma 3.1. Claim: $\mathrm{pr}(h|\sigma) = \mathrm{pr}(\ddot{\mathbf{r}}(h)|\ddot{\mathbf{r}}(\sigma))$ for all $h \in H$ and $\sigma \in \Sigma$. Of course, we have $\mathrm{pr}(\ddot{\mathbf{r}}(\eta)|\ddot{\mathbf{r}}(\sigma)) = \mathrm{pr}(\eta|\ddot{\mathbf{r}}(\sigma)) = 1 = \mathrm{pr}(\eta|\sigma)$. By induction on the order of histories

$$\begin{aligned}
\mathrm{pr}\left(\ddot{\mathbf{r}}\left(h,a\right)|\ddot{\mathbf{r}}\left(\sigma\right)\right) &= \mathrm{pr}\left(\left(\ddot{\mathbf{r}}\left(h\right),\mathbf{r}^{h}\left(a\right)\right)|\ddot{\mathbf{r}}\left(\sigma\right)\right) \\
&= \mathrm{pr}\left(\ddot{\mathbf{r}}\left(h\right)|\ddot{\mathbf{r}}\left(\sigma\right)\right)\cdot\prod_{i\in I}\ddot{\mathbf{r}}_{\pi(i)}\left(\sigma\right)\left(\ddot{\mathbf{r}}\left(h\right)\right)\left(r_{i}^{h}\left(a_{i}\right)\right) \\
&= \mathrm{pr}\left(h|\sigma\right)\cdot\prod_{i\in I}\sigma_{i}\left(h\right)\left(a_{i}\right) \\
&= \mathrm{pr}\left(\left(h,a\right)|\sigma\right)
\end{aligned}$$

from (3.4), (3.1) and (2.1), the induction hypothesis and (3.5), and (3.1). This proves the claim. Since all $\mathbf{r}^h$ involve the same bijection π, we are allowed to assume that the constants α_i and β_i from (2.2) do not depend on h. We then have

$$\begin{aligned}
u_{\pi(i)}^{t}\left(\ddot{\mathbf{r}}\left(\sigma\right)\right) &= \sum_{h\in H^{t}}\mathrm{pr}\left(\ddot{\mathbf{r}}\left(h\right)|\ddot{\mathbf{r}}\left(\sigma\right)\right)\sum_{a\in A}\left(\prod_{i\in I}\ddot{\mathbf{r}}_{i}\left(\sigma\right)\left(\ddot{\mathbf{r}}\left(h\right)\right)\left(a_{i}\right)\cdot u_{\pi(i)}\left(a\right)\right) \\
&= \sum_{h\in H^{t}}\mathrm{pr}\left(\ddot{\mathbf{r}}\left(h\right)|\ddot{\mathbf{r}}\left(\sigma\right)\right)\sum_{a\in A}\left(\prod_{i\in I}\ddot{\mathbf{r}}_{\pi(i)}\left(\sigma\right)\left(\ddot{\mathbf{r}}\left(h\right)\right)\left(r_{i}^{h}\left(a_{i}\right)\right)\cdot u_{\pi(i)}\left(\mathbf{r}^{h}\left(a\right)\right)\right) \\
&= \sum_{h\in H^{t}}\mathrm{pr}\left(h|\sigma\right)\sum_{a\in A}\left(\prod_{i\in I}\sigma_{i}\left(h\right)\left(a_{i}\right)\cdot u_{\pi(i)}\left(\mathbf{r}^{h}\left(a\right)\right)\right) \\
&= \sum_{h\in H^{t}}\mathrm{pr}\left(h|\sigma\right)\sum_{a\in A}\left(\prod_{i\in I}\sigma_{i}\left(h\right)\left(a_{i}\right)\cdot\left(\alpha_{i}u_{i}\left(a\right)+\beta_{i}\right)\right) \\
&= \alpha_{i}u_{i}^{t}\left(\sigma\right)+\beta_{i}
\end{aligned}$$

where the single equations follow from (3.2), (3.3), re-arranging sums and re-indexing h, re-indexing a and (2.1), the Claim, (3.5), and (2.2), and again (3.2) and (3.3). By (3.3), we then have

$$u_{i}^{\delta}\left(\ddot{\mathbf{r}}\left(\sigma\right)\right)=\left(1-\delta\right)\sum_{t=0}^{\infty}\delta^{t}\left(\alpha_{i}u_{i}^{t}\left(\sigma\right)+\beta_{i}\right)=\alpha_{i}u_{i}^{\delta}\left(\sigma\right)+\beta_{i}.$$

Proof of Lemma 4.1. We adapt the idea of the CH proof: Let b^* be payoff invariant in $G\in\mathcal{C}$ and $u\left(b^*\right)\geq u\left(b\right)$ all payoff invariant $b\in B$. Suppose, b^* is not an equilibrium. Then there is some player $\hat{\imath}$ and some pure action $a_{\hat{\imath}}\in A_{\hat{\imath}}$ such that $u\left(a_{\hat{\imath}}b_{-\hat{\imath}}^*\right)>u\left(b^*\right)$. Let $[\hat{\imath}]$ denote the set of all players symmetric to $\hat{\imath}$ and let $\left[\hat{a}_{\hat{\imath}}\right]_i$ denote the actions of player $i\in[\hat{\imath}]$ that are symmetric to $\hat{a}_{\hat{\imath}}$. For $i\in[\hat{\imath}]$ let c_i denote the action $b_i\left(a_i\right)=\left|\left[\hat{a}_{\hat{\imath}}\right]_i\right|^{-1}$ for $a_i\in\left[\hat{a}_{\hat{\imath}}\right]_i$ and $b_i\left(a_i\right)=0$ for $a_i\notin\left[\hat{a}_{\hat{\imath}}\right]_i$ and let $c_{[\hat{\imath}]}=\left(c_i\right)_{i\in[\hat{\imath}]}$. By construction, $(c_{[\hat{\imath}]},b_{-[\hat{\imath}]}^*)$ is symmetry invariant. Since it is also clear that $u(a_{\hat{\imath}}b_{-\hat{\imath}}^*)=u\left(c_{\hat{\imath}}b_{-\hat{\imath}}^*\right)$ and therefore $u\left(\mathbf{r}\left(c_{\hat{\imath}}b_{-\hat{\imath}}^*\right)\right)>u\left(b^*\right)$ for all $\mathbf{r}\in\mathfrak{S}\left(G\right)$, i.e.

$u\left(c_i b^*_{-\hat{\imath}}\right) > u\left(b^*\right)$ for all $i \in [\hat{\imath}]$. For $\delta \in (0,1)$ consider the action profile $\delta(c_{[\hat{\imath}]}, b^*_{-[\hat{\imath}]})$ $+ (1-\delta)\, b^*$ which is symmetry invariant too. We have

$$u\left(\delta(c_{[\hat{\imath}]}, b^*_{-[\hat{\imath}]}) + (1-\delta)\, b^*\right) = \sum_{J \subset [\hat{\imath}]} \delta^{|J|} (1-\delta)^{|[\hat{\imath}]|-|J|}\, b\left(c_J b^*_{-J}\right),$$

and as u is continuous in δ,

$$\lim_{\delta \to 0} \frac{u\left(\delta(c_{[\hat{\imath}]}, b^*_{-[\hat{\imath}]}) + (1-\delta)\, b^*\right) - u\left(b^*\right)}{1-(1-\delta)^{|[\hat{\imath}]|}} = \frac{1}{|[\hat{\imath}]|}\left(\sum_{i \in [\hat{\imath}]} u\left(c_i b^*_{-i}\right)\right) - u\left(b^*\right) > 0.$$

Hence, there is some δ^* such that $u\left(\delta(c_{[\hat{\imath}]}, b^*_{-[\hat{\imath}]}) + (1-\delta)\, b^*\right) > u\left(b^*\right)$. Contradiction.

Proof of Theorem 5.1. We construct such a frame: Let ξ be a bijection $\Omega \to \mathbb{N}_{|\Omega|}$, i.e. ξ is an enumeration of Ω. Consider the set $\bar{C} = \bigcup_{\omega \in \Omega} \{\omega\} \times \mathbb{N}_{\xi(\omega)}$ which contains for each object ω the elements $(\omega,1), (\omega,2), \ldots, (\omega, \xi(\omega))$, i.e. $\xi(\omega)$ enumerated copies of ω. Set $C = \mathfrak{G} \cup \bar{C}$, $\Lambda = \Omega$,

$$\ell(\omega, c) = \begin{cases} \mathfrak{g}(\omega), & c = \mathfrak{g} \\ \omega', & c = (\omega', k) \end{cases}$$

Fix $\overline{\mathfrak{g}} \in \mathfrak{G}$. Consider the bijections $r \in S(\Omega)$, $\omega \mapsto \overline{\mathfrak{g}}(\omega)$, $\mu \in S(C)$,

$$\mu(c) = \begin{cases} \mathfrak{g}\overline{\mathfrak{g}}^{-1}, & c = \mathfrak{g} \\ (\omega', k), & c = (\omega', k) \end{cases}$$

and $\tau = \mathrm{id}_\Omega$. This gives $\ell(r(\omega), \mu(\mathfrak{g})) = \left(\mathfrak{g}\overline{\mathfrak{g}}^{-1}\right)(\overline{\mathfrak{g}}(\omega)) = \mathfrak{g}(\omega) = \tau(\ell(\omega, \mathfrak{g}))$ and $\ell(r(\omega), \mu(\omega', k)) = \omega' = \tau(\ell(\omega, (\omega', k)))$, hence r is a frame symmetry of (Ω, F) that acts on Ω as $\overline{\mathfrak{g}}$.

Suppose, there were some frame symmetry r of (Ω, F) such that for all $\mathfrak{g} \in \mathfrak{G}$ there is some $\bar{\omega}_{\mathfrak{g}}$ such that $r(\bar{\omega}) \neq \mathfrak{g}(\bar{\omega})$. Then, there are $\mu \in S(C)$ and $\tau \in S(\Lambda)$ that satisfy (2.4). By Casajus (2001, Corollary 2.5.1), we have $\mu\left(\bar{C}\right) = \bar{C}$ since $\ell(\Omega, (\omega', k)) = \{\omega'\}$ but $\ell(\Omega, \mathfrak{g}) = \mathfrak{g}(\Omega) = \Omega$. In addition, we have $\tau = \mathrm{id}_\Omega$, since there are pairwise different numbers of properties (ω', k) with respect to which the objects ω' as labels are assigned. Note that to enforce $\tau = \mathrm{id}_\Omega$ for any frame symmetry was the reason for introducing the properties in $\bar{C}$. Thus, by (2.4), we have

$$\mu(\mathrm{id}_\Omega)(r(\omega)) = \ell(r(\omega), \mu(\mathrm{id}_\Omega)) = \tau(\ell(\omega, \mathrm{id}_\Omega)) = \mathrm{id}_\Omega(\mathrm{id}_\Omega(\omega)) = \omega,$$

i.e. $r(\omega) = \mu(\mathrm{id}_\Omega)^{-1}(\omega)$ for all ω where $\mu(\mathrm{id}_\Omega)^{-1} \in \mathfrak{G}$. A contradiction.

References

Bacharach, M. (1991). Games with concept-sensitive strategy spaces, Oxford University, mimeo.

Bacharach, M. (1993). Variable universe games, *in* K. Binmore, A. Kirman and P. Tani (eds), *Frontiers of Game Theory*, MIT Press, Cambridge, Mass., pp. 255–275.

Bacharach, M. and Bernasconi, M. (1997). The variable frame theory of focal points: An experimental study, *Games and Economic Behavior* **19**: 1–45.

Bacharach, M. and Stahl, D. O. (2000). Variable-frame level-n theory, *Games and Economic Behavior* **32**: 220–246.

Blume, A. (2000). Coordination and learning with a partial language, *Journal of Economic Theory* **95**: 1–36.

Blume, A. and Gneezy, U. (2000). An experimental investigation of optimal learning in coordination games, *Journal of Economic Theory* **90**: 161–172.

Casajus, A. (2001). *Focal Points in Framed Games: Breaking the Symmetry*, Vol. 499 of *Lecture Notes in Economics and Mathematical Systems*, Springer, Berlin et al.

Crawford, V. P. and Haller, H. (1990). Learning how to cooperate: Optimal play in repeated coordination games, *Econometrica* **58**(3): 571–595.

Goyal, S. and Janssen, M. C. W. (1996). Can we rationally learn to coordinate?, *Theory and Decision* **40**(1): 29–49.

Harsanyi, J. C. and Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*, MIT Press, Cambridge, Mass., London, England.

Janssen, M. C. W. (1998). Focal points, *in* P. Newman (ed.), *New Palgrave Dictionary of Economics and the Law*, MacMillan, London.

Janssen, M. C. W. (2001). Rationalizing focal points, *Theory and Decision* **50**: 119–148.

Kramarz, F. (1996). Dynamic focal-points in n-person coordination games, *Theory and Decision* **40**(3): 277–313.

Mehta, J., Starmer, C. and Sugden, R. (1994a). Focal points in pure coordination games: An experimental investigation, *Theory and Decision* **36**: 163–185.

Mehta, J., Starmer, C. and Sugden, R. (1994b). The nature of salience: An experimental investigation of pure coordination games, *American Economic Review* **84**(3): 658–673.

Schelling, T. C. (1960). *The Strategy of Conflict*, 8 edn, Harvard University Press, Cambridge, Mass., London, England.

Sugden, R. (1995). A theory of focal points, *Economic Journal* **105**: 533–550.