Available online at www.sciencedirect.com
SciVerse ScienceDirect
Journal of Economic Theory 148 (2013) 226–258
JOURNAL OF Economic Theory
www.elsevier.com/locate/jet

Hidden symmetries and focal points ☆

Carlos Alós-Ferrer [a,*], Christoph Kuzmics [b,1]

[a] Department of Economics, University of Cologne, 50923 Cologne, Germany
[b] Institute of Mathematical Economics, Bielefeld University (Germany), 33501 Bielefeld, Germany

Received 16 February 2011; final version received 2 July 2012; accepted 5 July 2012
Available online 13 December 2012

Abstract

We provide a general formal framework to define and analyze the concepts of focal points and frames for normal form games. The information provided by a frame is captured by a symmetry structure which is consistent with the payoff structure of the game. The set of symmetry structures has itself a clear structure (a lattice). Focal points are strategy profiles which respect the symmetry structure and are chosen according to some meta-norm, which is not particular to the framed game at hand.

JEL classification: C72; D83

Keywords: Symmetry; Focal points; Frames; Nash equilibria

1. Introduction

The aim of this paper is to provide a general framework for the analysis of focal points in the tradition of Schelling [24]. The main tool for the analysis is the concept of **symmetry structure**, which captures the various ways in which strategies and players can be viewed as symmetric in

☆ We would like to thank Andreas Blume, Vincent Crawford, Péter Esö, Michael Greinecker, Hans Haller, Maarten Janssen, Andrew McLennan, Roger Myerson, Larry Samuelson, Bill Sandholm, Itai Sher, Jonathan Weinstein, two anonymous referees, and seminar participants in Berlin, Cologne, Karlsruhe, Maastricht, Northwestern University, Paris, Vancouver, and Vienna for helpful comments. Financial support from the Austrian Science Fund (FWF) under Project P18141-G09 is gratefully acknowledged.
* Corresponding author.
E-mail addresses: carlos.alos-ferrer@uni-koeln.de (C. Alós-Ferrer), christoph.kuzmics@uni-bielefeld.de (C. Kuzmics).
1 Fax: +49 521 106 2997.

0022-0531/$ – see front matter
http://dx.doi.org/10.1016/j.jet.2012.12.003

a given normal-form game. In a first approximation, focal points are understood as those (and all those) Nash equilibria which respect the symmetries in a one-shot framed game, where a frame[2] endows strategies and players with objectively observable characteristics.

Our analysis shows that symmetry structures of games have a rich and useful mathematical structure. We establish two main results. First, the set of symmetry structures form a lattice. For each symmetry structure in the lattice, standard arguments imply the existence of a Nash equilibrium respecting the associated symmetries (i.e. treating symmetric strategies equally, and analogously for symmetric players), in such a way that as one moves "down" the lattice, i.e. as more symmetries are broken by increasingly more detailed frames, the set of associated equilibria grows. Second, symmetry structures are in a natural, essentially one-to-one, correspondence with the different possible frames of a given game. Thus symmetry structures become a useful tool for the study of frames.

The objects of study, in this paper, are truly one-shot finite normal form games. That is, they are not played recurrently, such as the game of which side of the road to drive on, for which conventions have been established through recurrent interaction. Rather, we assume that players are unfamiliar with the particular game at hand (and have no expectation of ever playing it again either). The game might be of a form that is recognized, but the game itself is new to the players.

While the game is thus assumed unfamiliar to the players, it might come with a setting or context, a frame, which could well be familiar to the players. Thus, this frame might provide players with more information, a "clue", as Schelling [24] puts it, as to how to play the game at hand. In this paper we will consider all possible frames a game might be accompanied with and analyze focal points for each frame.

Our approach is a normative one. For the sake of concreteness, imagine the following situation. A player is about to play a one-shot game which is completely new to her. Lacking experience and specialized knowledge, she decides to obtain advice from a game theory consultant. The consultant will first write down a description of the game, for instance using a matrix form. However, neither player positions nor strategy names in this description have any intrinsic meaning. Recommending to play whatever strategy has received a particular label, say A or "top-right", is completely arbitrary, as this label only makes sense in her depiction of the game. The consultant should thus realize that there is a great deal of arbitrariness in her representation of the game. She will conclude that a strategy must be solely identified by its associated payoff consequences. Further, if two opponents, engaged in the same game, seek advice from two different consultants, both consultants will most likely refer to their respective player as player 1. There might be information in the game which allows for a unique identification of the player (e.g. he or she might be the only player who could lose money), but sometimes this is not possible. Loosely speaking, if there are two ways of writing down the game leading to the same payoff tables although the ordering of two strategies or two players is different, then we should declare the two strategies or players symmetric. In the absence of further information, the labeling of those strategies and players by the different actors is effectively random. What one consultant calls A will be B for another consultant or for the player when she actually plays. Whether two strategies are actually considered symmetric can depend on additional (non-payoff) information provided by the players, i.e. on the frame of the game. A complete description of which strategies and players are considered symmetric will be called a **symmetry structure**.

[2] The concept of frame goes back to Tversky and Kahneman [26]. It has been formalized and used in different ways in the game-theoretic literature on focal points.

The consultant's analysis boils down to the identification of the appropriate symmetry structure given all available information entailed in the framed game. The recommendations provided by consultants are required to satisfy three axioms. One, it has to constitute a Nash equilibrium. We shall call this the **axiom of rationality**. The idea is that every consultant delivers both advice on how a particular player should play, and a prediction of her opponents' play, so that players can indeed check that the recommendation "makes sense". Second, the recommendation shall treat symmetric strategies equally, i.e. they must receive the same probability. We shall call this the **axiom of equal treatment of symmetric strategies**. This axiom can be and has been motivated by Laplace's Principle of Insufficient Reason. However, a simpler justification is that symmetric strategies are those which cannot reliably be distinguished from each other given the payoff structure and the available frame. Hence, a consultant is not able to distinguish among them in a recommendation and must treat them equally. Third, a recommendation should be such that two symmetric players receive the same advice (from one consultant). We shall call this the **axiom of equal treatment of symmetric players**. Any strategy profile that satisfies these three axioms shall be called a **rational symmetric recommendation**.

If we were to restrict ourselves to, say, coordination games, the question of which strategies and which players can be declared symmetric for a given normal-form game would be a simple one. For instance, in pure coordination games (where all off-diagonal payoffs are zeros), two strategies are symmetric if the corresponding equilibria yield the same payoffs. In matching games (with unity payoff matrices), all strategies are symmetric unless the symmetries are broken by the frame. A normative approach, however, cannot be restricted to a class of games because of its simplicity. Hence the framework we develop, and especially the theory of symmetry structures we provide, encompasses arbitrary normal-form games. One contribution of the paper is thus to show how one can find, in general games, the symmetries that are obvious in coordination games. This gives rise to surprisingly subtle points, and, indeed, in many games symmetries and asymmetries might be initially "hidden" from casual players (following with our analogy, it is the consultant's job to uncover them). It becomes necessary to generalize and clarify the relation between the different symmetry concepts which have been explicitly or implicitly used in game theory.

A notable approximation to the problem is the definition of symmetric strategies in two-player games of Crawford and Haller [10, Appendix], which we apply to general n-player games.[3] In many cases, in particular for the two-player matching games studied in (the main body of) Crawford and Haller [10], this concept is sufficient for the analysis. We show, however, that it is in general not enough. First, Crawford and Haller's [10] concept is based on pairwise strategy comparisons, but a global concept is needed once one moves away from pure coordination games. Second, once we move away from two-player games, strategy symmetries cannot be established independently from player symmetries. We are led to a more subtle definition of symmetries within a game, which leads to different predictions. For this global definition, we build on Nash's [20] concept of symmetries and Harsanyi and Selten's [15] game automorphisms, as also used e.g. in Casajus [8]. The resulting concept of global symmetry structures is closely related to the concept of frames itself. Every frame induces a global symmetry structure and for every global symmetry structure there is a frame that justifies it. We are thus led to study the structure of global symmetry structures and find that together with the partial order of "coarser than" they

[3] Crawford and Haller [10] use their framework to study how players, in a repeated two-player coordination game, can use history to coordinate in a pure equilibrium. See also Blume [6].

form a lattice with non-trivial joins and meets, where the meet of two symmetry structures is the symmetry structure resulting from the combination of (the information contained in) two appropriate frames.

Although this paper concentrates on the concept of symmetry, the main motivation is to provide the tools necessary to formalize a theory of focal points in finite, arbitrary games. Our position is as follows. We argue that a definition of focal points must necessarily come in two parts. The first part requires a focal point to be an equilibrium and to respect the symmetries as specified in the symmetry structure. Since the number of predictions grows as frames become increasingly detailed, the second part of the definition of focal points must involve the selection of equilibria through an appropriate (and maybe very intricate) "norm" which can be defined independently of the game at hand, i.e. a **meta-norm**. Our purpose is to provide a full formal theory of the first part of the definition (symmetries), and argue that it has to be complemented by a meta-norm. Accordingly, in our examples we will often rely on some particular meta-norm for clarity. In our view, however, the choice of a meta-norm is not part of what a formal theory of fully rational behavior should provide, but rather a behavioral question.

Meta-norms can range from fairly simple to very intricate. As a first example, one can rely on the fact that games have payoffs, in, say, monetary terms. Within a pure coordination game, familiarity with money is probably enough to ensure that players coordinate on a unique Pareto-efficient equilibrium, if one exists. That is, the (partial) meta-norm of always picking the strategy which gives rise to the Pareto-optimal outcome would enable players to coordinate even if the particular game has never been encountered before. We shall call this the meta-norm of Pareto-efficiency.[4]

In this paper, we thus define focal points as equilibria that **should** be played by highly rational players (consultants) who understand the symmetry structure induced by the framed game and have common knowledge of a given meta-norm. While we thus always assume that a commonly held meta-norm is in place, we do not investigate what this meta-norm could or even should be. In accordance with this view, we will also abstract from possible conflicts between alternative meta-norms. Thus our paper is not a descriptive one of how players behave in a given framed game in the lab (although, we believe our analysis could prove helpful even in these cases), but rather how players should and perhaps eventually will behave, after generations of teaching and learning.

We remark that our ultimately normative motivation differs from that of most of the existing literature. We are interested in what hypothetical consultants could advise players to play in a framed game, while the modern literature on focal points, at least in our view, aims at identifying how people actually do play, very specially in coordination games. In simple games, players can be safely assumed to have a set of thought routines and reasoning mechanisms which can be equivalent to having a hypothetical consultant; hence, the two approaches are very similar. In more general games, one can no longer expect every real-life player's thought routines to have worked out all the intricate symmetries inherent in the framed game and to be able to identify the unique best way to play the game. Hence, in a sense our normative goal is less ambitious than a descriptive theory of actual, boundedly-rational play. In exchange, by relying on the concept of symmetry structures, we are able to tackle all games and not just coordination games.

[4] Variations of this meta-norm are called the "Principle of Coordination" in Gauthier [12], Bacharach [2,3], Sugden [25], and Casajus [8], "Rationality in the Extended Sense" in Goyal and Janssen [13], and the "Principle of Individual Team Member Rationality" in Janssen [17].

Our motivation, however, also differs from other normative approaches as e.g. Harsanyi and Selten's [15] aim to provide a unique solution for every game. For a two-strategy matching game, their solution must be the unique completely mixed (and Pareto-inferior) Nash equilibrium. In our case, this remains the only solution for a particular frame, but more detailed frames (or symmetry structures) lead to Pareto-improved consultant's recommendations. Since actual players sometimes do manage to coordinate in such simple games, our approach is still compatible with such behavior.

The paper is structured as follows. Section 2 concentrates on strategy symmetry. This section presents the concepts of pairwise and global symmetry and analyzes the structure of symmetry structures, their correspondence to frames, and the existence of rational-symmetric recommendations. Section 3 further concentrates on global symmetry structures and extends the analysis to include the more involved concept of player symmetry. Subsection 3.5 discusses our concept of focal points based on the notion of a meta-norm. Section 4 briefly discusses the role of payoff transformations. Section 5 concludes. Proofs (which rely on some elementary group and lattice theory) are relegated to Appendix A.

2. Strategy symmetry

In this section, we analyze the concept of symmetry structures in games and their relevance for focal points. In order to simplify the presentation, we rely on one simplification: we will ignore symmetry among players and concentrate on symmetry among strategies (of a given player) only.

In Subsection 2.2 below, we will further restrict ourselves to concepts of **pairwise** symmetry, where strategies are compared in pairs in order to decide whether they can be declared symmetric or not. For two-player games, this first concept reduces to the one defined by Crawford and Haller [10].

Pairwise symmetry structures allow us to discuss most of the intuitions behind our approach while greatly reducing the necessary conceptual and analytical complexity. Furthermore, the concept is of interest in itself, since it already captures many of the examples that have been discussed in the literature. It is, however, not entirely satisfactory, as we will discuss further below. In Subsection 2.3, and building upon the intuitions developed in this section, we will discuss a **global** notion of symmetry. For some special games, such as matching games (the object of study in the main body of Crawford and Haller [10]), our global notion of symmetry is equivalent to pairwise symmetry. In Section 3, we will further develop the notion of global symmetry to allow for player symmetry.

Once the symmetry concepts are in place, Subsection 2.4 below analyzes the correspondence between pairwise symmetry structures and frames. As a first approximation to the idea of focal points, Subsection 2.5 then shows that it is always possible to spell out a recommendation in the form of a Nash equilibrium respecting the symmetries of the (framed) game, be they pairwise or global.

2.1. Games and labels

Consider a finite game $\Gamma = [I, (S_i, u_i)_{i\in I}]$, where I is a finite set of players, S_i is the finite set of pure strategies for player i, and $u_i : S \mapsto \mathbb{R}$ is the payoff function of player i, defined on the set of strategy profiles $S = \times_{i\in I} S_i$. The vector payoff function $u : S \to \mathbb{R}^{|I|}$ is the function whose i-th coordinate is u_i.

Following game-theoretic conventions, for all $s \in S$ we write $u_i(s) = u_i(s_i|s_{-i})$,[5] where $s_{-i} \in S_{-i} = \bigtimes_{j \neq i} S_j$. Abusing notation, we also write $u(s) = u(s_i|s_{-i})$ for the vector of payoffs whenever we want to single out player i's strategy but refer to the whole vector of payoffs. Further, denote by $\Theta_i = \Delta(S_i)$ the set of mixed strategies of player i, and let $\Theta = \bigtimes_{i \in I} \Theta_i$ be the set of mixed strategy profiles. We extend the payoff functions u_i to mixed strategies in the usual way, i.e. taking expectations over all mixed strategies.

Much of the literature on focal points and salience deals almost exclusively with two-player games of pure coordination, or even matching games. A two-player game $\Gamma = [\{1, 2\}, S_1, u_1, S_2, u_2]$ is a **game of pure coordination** if $S_1 = S_2$ and $u_i(s_i|s_{-i}) = 0$ if $s_i \neq s_{-i}$ and $u_i(s_i|s_{-i}) > 0$ if $s_i = s_{-i}$, for $i = 1, 2$. A game of pure coordination is a **matching game** if additionally $u_i(s_i|s_{-i}) = 1$ if $s_i = s_{-i}$, $i = 1, 2$. Let M_k denote the matching game with k strategies. For instance, the main results of Casajus [8, Theorem 5.6] and Janssen [17, Propositions 1 and 2] are restricted to matching games.

A **label** is any observable characteristic that can be objectively established and reliably communicated, and that consultants can attach to strategies when analyzing the game without relying on factors out of their control. For instance, if a player is asked to choose from four identical objects arranged in a table in front of him in a square, "first-row, left-hand" might not qualify as a label unless the consultant know for sure that there is a pre-established position for the objects (and that the player will sit at a specific position around the table). However, if one of the objects is red, while the others are blue, the colors will qualify as labels (unless the player or the consultant are color-blind). We will focus on neutral adjectives like "red" or "shiny" for our examples. However, a label is anything which can be used to provide a strategy with a universally recognizable meaning, and hence other examples can range from "hire your opponent" to "the set of prime numbers larger than 42" or "go to Grand Central Station". In repeated games, labels can also be derived from the history of play, as argued by Crawford and Haller [10]. For instance, in the second stage a strategy (of the stage game) may be labeled "was played by two players in the first stage".

2.2. *Pairwise strategy symmetry*

To motivate our definition of a pairwise symmetry structure consider the following simple example, the matching game M_2.

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

Game 1.

In this section, players are distinct by assumption, i.e. the two players can, in principle, be given different recommendations (perhaps one is a man, the other a woman). There are two different situations one can envision. In the first, the two strategies are symmetric (as in Crawford and Haller [10]), because their names do not come with a salience ranking (the labels A and B are

[5] We prefer the notation $u_i(\cdot|\cdot)$ to the more extended $u_i(\cdot,\cdot)$, since the latter is not always unequivocal. For instance, in a symmetric two-player game, $u_2(z, z')$ might refer to the payoff when player 2 plays z against z' or vice versa, depending on interpretation, while $u_2(z|z')$ and $u_2(z'|z)$ cannot be confused with each other if the symbol "|" is read as "given".

meaningless, arbitrary choices by the consultant; a different consultant might have labeled them otherwise). In the second situation, the two strategies are clearly labeled. For example, A is heads and B is tails, as in Schelling's [24] example, with heads commonly known to be a more salient label. Alternatively, a commonly known label might refer to a result in a previous round of play of the game as in Crawford and Haller [10].

Since both situations are conceivable, we want this game to have two strategy symmetry structures, one in which both strategies of both players are declared symmetric (to cover the first case) and another in which the two strategies are kept distinct (to cover the second case). The following definition captures this idea.

Definition 2.1. A **pairwise strategy symmetry structure** of game Γ is a collection $\mathcal{T} = \{\mathcal{T}_i\}_{i\in I}$, where each $\mathcal{T}_i$ is a partition of S_i such that, for each $i \in I$, each $T_i \in \mathcal{T}_i$, and each pair of distinct strategies $s_i, s_i' \in T_i$, there exist renamings ρ_j of S_j (for all $j \neq i$) such that $\rho_j(T_j) = T_j$ for all $T_j \in \mathcal{T}_j$ and $u(s_i|s_{-i}) = u(s_i'|\rho_{-i}(s_{-i}))$ for all $s_{-i} \in S_{-i}$. The sets $T_i \in \mathcal{T}_i$ are called **pairwise symmetry classes** for player i. Two strategies s_i, s_i' are said to be pairwise strategy-symmetric (relative to $\mathcal{T}$) if they belong to the same symmetry class.[6]

This definition is a natural generalization to n-player games of the notion of strategy symmetry introduced by Crawford and Haller [10, Appendix] for two-player games. The existence of a pairwise strategy symmetry structure of any game is guaranteed by the observation that the partition which consists of all singleton sets is a pairwise strategy symmetry structure. We will refer to this as the **trivial symmetry structure**.

For the case in which strategies have no (commonly understood) labels whatsoever we would like to find the symmetry structure with the largest possible symmetry classes. It is not immediately obvious whether there is a unique such 'largest' symmetry structure.

First, we need to clarify what 'largest' means. The set of partitions of S_i is partially ordered as follows. A partition $\mathcal{T}_i'$ is **coarser** than another partition $\mathcal{T}_i$, if for each $T_i \in \mathcal{T}_i$ there exists $T_i' \in \mathcal{T}_i'$ with $T_i \subseteq T_i'$. If $\mathcal{T}_i'$ is **coarser** than another partition $\mathcal{T}_i$, the latter is **finer** than the former. We say that one symmetry structure $\mathcal{T}'$ is **coarser** than another symmetry structure $\mathcal{T}$, if $\mathcal{T}_i'$ is coarser than $\mathcal{T}_i$ for every $i \in I$. A **coarsest symmetry structure** is a maximal element of the set of symmetry structures according to the partial order of "coarser than". Note that the trivial symmetry structure is the unique **finest** symmetry structure.

Given two partitions $\mathcal{T}_i$ and $\mathcal{T}_i'$ of S_i, the **join** $\mathcal{T}_i \vee \mathcal{T}_i'$ is the finest partition which is coarser than both $\mathcal{T}_i$ and $\mathcal{T}_i'$. Dually, the **meet** $\mathcal{T}_i \wedge \mathcal{T}_i'$ is the coarsest partition which is finer than both partitions. Lemma A1 in Appendix A gives a useful characterization of the join of two partitions.

The join (least upper bound) $\mathcal{T} \vee \mathcal{T}'$ of two pairwise strategy symmetry structures $\mathcal{T}$ and $\mathcal{T}'$ can be defined as the finest pairwise strategy symmetry structure which is coarser than the two given ones. Analogously, the meet (greatest lower bound) $\mathcal{T} \wedge \mathcal{T}'$ is the coarsest pairwise symmetry structure which is finer than the two given ones. The following result shows that any two pairwise symmetry structures have a join and a meet, i.e. symmetry structures form a lattice. Since the set is finite, it follows that any arbitrary set of symmetry structures has both a join and a meet, i.e. they form a **complete lattice**.

[6] When it is clear from the context we will refer to a pairwise strategy symmetry structure as simply a symmetry structure.

Theorem 1. *For every finite game Γ the set of pairwise strategy symmetry structures endowed with the partial order of "coarser than" forms a complete lattice. The join of two pairwise strategy symmetry structures $\mathcal{T}'$ and $\mathcal{T}''$ is given by $\mathcal{T}' \vee \mathcal{T}'' = \{\mathcal{T}_i' \vee \mathcal{T}_i''\}_{i \in I}$.*

As a consequence of this result, we obtain that, for any finite normal-form game, there exists a coarsest symmetry structure.

Corollary 1. *Every finite game has a unique coarsest pairwise symmetry structure $\mathcal{T}^*$.*

The coarsest symmetry structure is important because it captures as much symmetry as actually exists in the payoff matrix of the game alone, i.e. without the addition of any external labels. In a sense, it provides a useful "symmetry benchmark". In any other symmetry structure, external labels have been added and symmetries have been consequently broken, creating a finer structure.

We conclude this subsection with a remark. As observed in Theorem 1, the join of two symmetry structures has a particularly simple form. This is not true for the meet. Although the meet of any two symmetry structures exists, it is in general not given by the collection of meets of the individual player partitions. To see this, consider the following two-player game, Game 2.

	E	F	G	H
A	1, 1	0, 0	1, 1	0, 0
B	0, 0	1, 1	0, 0	1, 1
C	2, 2	0, 0	2, 2	0, 0
D	0, 0	2, 2	0, 0	2, 2

Game 2.

The coarsest symmetry structure of this game is the one where $\mathcal{T}_1^* = \{\{A, B\}, \{C, D\}\}$ and $\mathcal{T}_2^* = \{\{E, F, G, H\}\}$. Consider two alternative symmetry structures, $\mathcal{T}'$ and $\mathcal{T}''$ with $\mathcal{T}_1' = \mathcal{T}_1'' = \mathcal{T}_1^* = \{\{A, B\}, \{C, D\}\}$ and $\mathcal{T}_2' = \{\{E, F\}, \{G, H\}\}$ and $\mathcal{T}_2'' = \{\{E, H\}, \{F, G\}\}$. The join of these two structures is the coarsest one, $\mathcal{T}^*$. If we consider the greatest lower bounds for the individual player partitions, we obtain a "meet candidate" $\widetilde{\mathcal{T}}$ given by $\widetilde{\mathcal{T}}_1 = \mathcal{T}_1' = \mathcal{T}_1'' = \mathcal{T}_1^* = \{\{A, B\}, \{C, D\}\}$ and $\widetilde{\mathcal{T}}_2 = \{\{E\}, \{F\}, \{G\}, \{H\}\}$. However, this is not a symmetry structure. Note that player 2's symmetry partition is the finest possible, consisting only of singletons. Given this, two strategies of player 1 can only be symmetric if they always deliver identical payoffs, which is not the case for any pair of strategies. Hence, this is not a pairwise symmetry structure. In this example, the meet $\mathcal{T}' \wedge \mathcal{T}''$ is the trivial symmetry structure.

2.3. Global strategy symmetry

The concept of pairwise strategy symmetry, however, is not entirely satisfactory for general games. In a sense, in some games it declares too many strategies symmetric.[7] To illustrate the problem, consider Game 3.

[7] This is not a problem for matching games, for which Crawford and Haller [10] developed the concept.

	b_1	b_2	b_3	b_4	b_5	b_6
a_1	1, 1	0, 0	0, 0	3, 3	0, 0	0, 0
a_2	0, 0	1, 1	0, 0	0, 0	3, 3	0, 0
a_3	0, 0	0, 0	1, 1	0, 0	0, 0	3, 3
a_4	−1, −1	−1, −1	4, 4	2, 2	0, 0	0, 0
a_5	−1, −1	4, 4	−1, −1	0, 0	2, 2	0, 0
a_6	4, 4	−1, −1	−1, −1	0, 0	0, 0	2, 2

Game 3.

The coarsest pairwise symmetry structure, as commented above, captures the symmetries present in the payoff matrix alone in the absence of any labels (including strategy names). In this game, this structure has $\widetilde{\mathcal{T}}_1 = \{\{a_1, a_2, a_3\}, \{a_4, a_5, a_6\}\}$ and $\widetilde{\mathcal{T}}_2 = \{\{b_1, b_2, b_3\}, \{b_4, b_5, b_6\}\}$. To see this, notice for instance that a_1 and a_2 are seen to be symmetric if we permute the strategies of player 2 in such a way that $b_1 \to b_2$ and $b_4 \to b_5$. Note also for reference that **only** such permutations allow us to declare a_1 and a_2 symmetric. For analogous reasons any pair of strategies within one of the classes above can be declared pairwise symmetric. Hence, $\widetilde{\mathcal{T}}$ is a pairwise symmetry structure.

Yet, there is a sense in which a_1 and a_2 are different from each other. Suppose a consultant wrote the game down in matrix form by starting with a_2 as the first strategy of player 1 and a_1 as the second. By pairwise symmetry of a_1 and a_2 we know that we can rename (or reorder in the matrix) the strategies in such a way that the complete row of payoff vectors for strategy a_2 in this new depiction of the game is exactly the same as that of a_1 in the depiction above and vice versa. In order to do so we need to swap player 2's strategies b_1 and b_2 as well as b_4 and b_5. We thus obtain the following partial depiction.

	b_2	b_1	b_3	b_5	b_4	b_6
a_2	1, 1	0, 0	0, 0	3, 3	0, 0	0, 0
a_1	0, 0	1, 1	0, 0	0, 0	3, 3	0, 0

This confirms that given $\widetilde{\mathcal{T}}$, player 1's strategies a_1 and a_2 are indeed pairwise symmetric. Intuitively, this should mean that these strategies are indistinguishable. That is, a consultant who starts with this partial depiction should be able to complete her depiction of the game and, except for the names of the strategies, still obtain exactly the same payoff matrix as above. It turns out, however, that having started the depiction of the game as just explained, it is not possible to obtain a fully identical depiction to the previous one. To see this, note that in order to obtain the two upper blocks in the original payoff matrix, the next strategy to be listed for player 1 is a_3. Now we should continue the depiction of the game in such a way that the 3×3 matrix in the lower right hand corner is the same as its counterpart in the original depiction. Since the ordering on the column player strategies is already fixed, this leaves only one possibility for the ordering of the row player's strategies and leads to the following complete depiction.

	b_2	b_1	b_3	b_5	b_4	b_6
a_2	1, 1	0, 0	0, 0	3, 3	0, 0	0, 0
a_1	0, 0	1, 1	0, 0	0, 0	3, 3	0, 0
a_3	0, 0	0, 0	1, 1	0, 0	0, 0	3, 3
a_5	4, 4	−1, −1	−1, −1	2, 2	0, 0	0, 0
a_4	−1, −1	−1, −1	4, 4	0, 0	2, 2	0, 0
a_6	−1, −1	4, 4	−1, −1	0, 0	0, 0	2, 2

This depiction differs from the original one. Given our original choice of a_2 rather than a_1 as the first of the row players' strategies, it is not possible to match both lower 3×3 corner matrices simultaneously to the original depiction. Hence, a clever consultant can distinguish a_1 and a_2 on the basis of this difference. While strategy a_1, if used as the first strategy, can give rise to the original depiction, strategy a_2 cannot. Thus, we need to weaken our definition of strategy symmetry. The key difference with the original concept is that declaring two strategies symmetric should allow to perform a global renaming of strategies of the players in such a way that the whole payoff matrix (and not just two rows or columns) remains unchanged.

Definition 2.2. A **strategy symmetry** of a normal form game Γ is a collection $\tau = (\tau_i)_{i \in I}$, where, for each $i \in I$, $\tau_i : S_i \to S_i$ is a bijection, fulfilling that, for all $i \in I$ and all $s = (s_i | s_{-i}) \in S$,

$$u_i(s_i | s_{-i}) = u_i\big(\tau_i(s_i) \big| \tau_{-i}(s_{-i})\big). \quad (1)$$

This concept captures the key difference between the global and the pairwise strategy symmetry concepts. While the renamings used in Definition 2.1 are only required to keep the payoff matrix unchanged when comparing two specific strategies of a fixed player, strategy symmetries are required to leave the payoff matrix unchanged **globally** after reordering strategies of all players simultaneously. This concept of strategy symmetry is essentially the one used in Nash [20] or Harsanyi and Selten [15]. The definition of strategy symmetry structure can now be adapted by using this more stringent concept.

Definition 2.3. A **global strategy symmetry structure** of game Γ is a collection $\mathcal{T} = \{\mathcal{T}_i\}_{i \in I}$ with each $\mathcal{T}_i$ a partition of S_i such that for each $i \in I$, each $\hat{T}_i \in \mathcal{T}_i$, and each $s_i, s_i' \in \hat{T}_i$, there exists a strategy symmetry τ such that $\tau_i(s_i) = s_i'$ and $\tau_j(T_j) = T_j$ for all $T_j \in \mathcal{T}_j$ and all $j \in I$. The sets $T_i \in \mathcal{T}_i$ are called **strategy symmetry classes** for player i. Two strategies s_i, s_i' are said to be (globally) strategy-symmetric (relative to $\mathcal{T}$) if they belong to the same symmetry class.

With this definition in hand,[8] let us return to Game 3. The previously identified pairwise strategy symmetry structure $\widehat{\mathcal{T}}$ is **not** a global strategy symmetry structure. To see this, let us try to find a symmetry allowing us to declare a_1 and a_2 symmetric. As commented above, that symmetry must necessarily permute the strategies of player 2 in such a way that $b_1 \to b_2$ and $b_4 \to b_5$. But permuting $b_1 \to b_2$ implies that the strategies of player 1 must be permuted in such a way that $a_6 \to a_5$. On the other hand, permuting $b_4 \to b_5$ implies that the strategies of player 1 must be permuted in such a way that $a_4 \to a_5$, a contradiction (τ_1 would fail to be a bijection). Hence, there exists no symmetry allowing us to declare a_1 and a_2 symmetric.

Careful examination of this game shows that the coarsest global symmetry structure without player symmetry is given by $\widehat{\mathcal{T}_1} = \{\{a_1, a_3\}, \{a_2\}, \{a_4, a_6\}, \{a_5\}\}$ and $\widehat{\mathcal{T}_2} = \{\{b_1, b_3\}, \{b_2\}, \{b_4, b_6\}, \{b_5\}\}$. To see that this is indeed a symmetry structure, just consider the symmetry which simultaneously swaps the two strategies in every non-singleton class and leaves the remaining

8 It is possible to provide an alternative, equivalent definition which does not use the partitions of strategies as a primitive, but rather appropriate restrictions on the sets of symmetries which are allowed. Those restrictions boil down to the requirement that the symmetries form a subgroup of the group of all symmetries. However, an added difficulty with that approach is that different subgroups of symmetries might generate the same partition of strategies. See Appendix A.4 for details.

unchanged. The fact that it is the coarsest one follows from the observation that a_1 and a_2 cannot be declared symmetric, and the analogous reasoning for all other strategies in singleton classes.

Thus, Game 3 provides an answer to the question of whether every pairwise strategy symmetry structure is a global strategy symmetry structure. For the particular case of the coarsest symmetry structure, one can rephrase the question as follows. Suppose two strategies can be declared pairwise symmetric (following Crawford and Haller [10]). Can they always be declared globally symmetric (i.e. symmetric in the sense implicit in Nash [20])? This question was already posed (as an open question) by Casajus [9, p. 20] (except for differences with respect to whether symmetries are allowed to transform payoffs or not; see Section 4). Game 3 shows that the answer is negative.

Even though settling the question of the (non-)equivalence between pairwise symmetry concepts as in Crawford and Haller [10] and global ones as in Nash [20] or Harsanyi and Selten [15] is interesting in itself, this is more than a technical point. Game 3 also illustrates that the recommendations might differ qualitatively under both approaches. Getting ahead of ourselves for a second, suppose a consultant is allowed to recommend any Nash equilibrium where strategies which are declared symmetric are played with equal probability. Suppose also that this particular consultant aims to recommend Pareto-efficient equilibria in this class. Adopting the pairwise approach, the coarsest symmetry structure is given by $\widetilde{\mathcal{T}}$ above and hence we are left with three equilibrium candidates: the first randomizes uniformly among a_1, a_2, a_3 and among b_4, b_5, b_6; the second randomizes uniformly among a_4, a_5, a_6 and among b_1, b_2, b_3; while the third randomizes uniformly among a_4, a_5, a_6 and among the whole set $b_1, \ldots, b_6$. The expected payoffs of the first are larger than the payoffs in the other two recommendations, and hence the pairwise approach delivers a unique recommendation resulting in an expected payoff of 1. On the other hand, the global approach delivers the coarsest symmetry structure $\widehat{\mathcal{T}}$, which enables coordination in the equilibrium (a_5, b_2), with a payoff of 4. This is then the unique recommendation, which is qualitatively different from the one arrived at under the pairwise approach.

While Game 3, thus, demonstrates that not every pairwise strategy symmetry structure is also a global strategy symmetry structure, the converse is true. The proof is immediate.

Proposition 1. *If $\mathcal{T}$ is a global strategy symmetry structure, then $\mathcal{T}$ is a pairwise strategy symmetry structure.*

In view of this result, and although we take the position that the global version is more appropriate, for the remainder of the section we will spell out the results for pairwise strategy symmetry structures.

2.4. Pairwise strategy symmetry structures and basic frames

We now turn to frames. Let $\mathcal{Z}_i$ be a universal set of **labels** for the strategy set of each player i. A **basic frame** for the game Γ is a collection $L = (L_i)_{i \in I}$ where $L_i : S_i \to \mathcal{Z}_i$ for each $i \in I$. It is important to focus on the interpretation of a frame as reporting on universally observable, objective characteristics. In particular, each consultant will be able to observe the labels $L_i(s_i)$ of all strategies of all players. As in the rest of this section, here we assume that players are readily identifiable, i.e. there is no question of player symmetry.

The coarsest pairwise symmetry structure $\mathcal{T}^*$ delivers the strongest (coarsest) reclassification of strategies that a consultant can obtain from the game, based on payoffs alone. In this sense, $\mathcal{T}^*$

is associated to the game without frames. It is useful to consider how other symmetry structures might arise.

Suppose the consultant analyzes the game in two steps. First, she extracts as much information as she can from the payoff structure alone. Thus she will arrive at the symmetry structure $\mathcal{T}^*$. Second, she considers the basic frame L. Consider two strategies which are not symmetric in $\mathcal{T}^*$. Since they can already be distinguished on the basis of payoffs, whether they receive the same or different labels adds no further information. Labels are important, however, to distinguish among symmetric strategies. That is, a basic frame induces a refinement of $\mathcal{T}^*$ by further partitioning the symmetry classes.[9] Given a basic frame component L_i for player i, the L_i-**partition** of S_i is the partition given by the sets $T_i \cap L_i^{-1}(a)$ for all $T_i \in \mathcal{T}_i^*$ symmetry classes of the coarsest symmetry structure and all $a \in \mathcal{Z}_i$.

It is, however, not true that the refined partitions will automatically form a symmetry structure. In other words, the consultant is in general left with some work to do to integrate the new information into a new symmetry structure. A simple example of how additional information provided by a frame can change the symmetry structure is given by Game 4 below. Another example in the next subsection, Game 5, illustrates that a frame can lead to subtle changes in the symmetry structure with radical consequences for the ultimate recommendation consultants can provide.

Definition 2.4. Let L be a basic frame for game Γ. The pairwise strategy symmetry structure induced by L, $\mathcal{T}(L)$, is the coarsest pairwise strategy symmetry structure $\mathcal{T}$ such that, for each player i, $\mathcal{T}_i$ is finer than the L_i-partition of S_i.

Note that $\mathcal{T}(L)$ is always well defined by Theorem 1. The argument is as follows. Consider the set of all symmetry structures whose players' partitions are finer than the L_i-partitions. This set is nonempty (since it contains the trivial one). It is easy to see that if two symmetry structures are in this set, the join also fulfills the characterizing property. It follows that the join of all symmetry structures in the set delivers the coarsest one.

To see that $\mathcal{T}(L)$ is in general not just given by the repartitioning of symmetry classes according to the labels, consider the following two-player framed game, where $\mathcal{Z}_1 = \{\bullet, \circ\}$ and $\mathcal{Z}_2 = \{\blacksquare, \square\}$.

		$\blacksquare$	$\blacksquare$	$\square$
		D	E	F
$\bullet$	A	1, 2	0, 0	0, 0
$\circ$	B	0, 0	1, 2	0, 0
$\circ$	C	0, 0	0, 0	1, 2

Game 4.

In the coarsest (frame-free) symmetry structure, all strategies are symmetric, for both players. If we instead repartition the symmetry classes according to the observed labels, we obtain $\{\{A\}, \{B, C\}\}$ for player 1 and $\{\{D, E\}, \{F\}\}$ for player 2. These partitions do not form a symmetry structure. For, in order to declare B and C symmetric for player 1, it is necessary to permute

[9] We assume that frames summarize universally observable information. One might conceive of interesting situations with frames which are only observed partially by some of the players, leading to incomplete-information settings.

E and F for player 2. But the latter strategies are in a different element of the L_2-partition. The symmetry structure induced by the frame in this example is the trivial one.[10]

The mapping $L \to \mathcal{T}(L)$ gives us a natural translation of frames into symmetry structures. This mapping is actually onto, that is, for every symmetry structure a consultant might come up with, there exists a frame which rationalizes it.

Theorem 2. *For any pairwise strategy symmetry structure there exists a basic frame L such that* $\mathcal{T}(L) = \mathcal{T}$.

Proof. Fix $\mathcal{T}$ and let $\mathcal{Z}_i = \mathcal{T}_i$. Define $L_i(s_i) = T_i$ where $T_i \in \mathcal{T}_i$ is such that $s_i \in T_i$. The L_i-partitions just reproduce $\mathcal{T}_i$ and thus $\mathcal{T}(L) = \mathcal{T}$. □

Although this result is straightforward, we find its interpretation interesting. We can rephrase it through the usual appeal to the canonical decomposition of a mapping as follows. Call two basic frames L and L' **similar** if they generate the same symmetry structure, i.e. $\mathcal{T}(L) = \mathcal{T}(L')$. If we consider the mapping $\mathcal{T}$ to be defined on the quotient set, i.e. the set of similarity classes of basic frames, then it becomes bijective. Thus, at an abstract level, we could identify basic frames (up to similarity) with symmetry structures. In other words, given a basic frame, the corresponding symmetry structure becomes a complete model of all symmetries implied by the frame and the payoff structure of the game. In that sense, symmetry structures help us understand the implications of frames.

We remark that the equivalence between symmetry structures and (classes of) frames respects the lattice structure in the natural way. As an illustration, consider a situation where, as in Casajus [8], Janssen [17], or Binmore and Samuelson [4], players might observe the realizations of several sets of attributes, e.g. color $L_i^C(s_i)$ out of certain sets $\mathcal{Z}_i^C$ and shape $L_i^H(s_i)$ out of certain sets $\mathcal{Z}_i^H$. The problem can be easily reformulated by defining the composite labels $L_i(s_i) = (L_i^C(s_i), L_i^H(s_i)) \in \mathcal{Z}_i = \mathcal{Z}_i^H \times \mathcal{Z}_i^C$. The corresponding symmetry structure is then just the meet of the color and shape symmetry structures, $\mathcal{T}(L) = \mathcal{T}(L^C) \wedge \mathcal{T}(L^H)$, which always exists by Theorem 1.

2.5. *Equal treatment of symmetric strategies*

We are now ready to spell out the first two of the axioms we require a consultant's recommendation to satisfy for any given game and basic frame (Γ, L), i.e. for any given (pairwise or global) strategy symmetry structure $\mathcal{T}$. A recommendation is simply a mixed strategy profile $x \in \Theta$.

Axiom 1. A recommendation $x \in \Theta$ is **rational** if it is a Nash equilibrium of the game.

That x_i should be a best response to x_{-i} is a minimal rationality requirement. When confronted with a specific recommendation, which includes a prediction for the play of the opponents, players should be able to recognize whether they have an individual incentive to deviate; likewise, they should be able to check whether the prediction for the opponents' play is reasonable, in the same sense.

[10] One could alternatively define frames directly as the partitions they induce. This example shows that the induced partition is not always obvious, and hence we prefer to define frames through the labeling of strategies.

Axiom 2. A recommendation $x \in \Theta$ satisfies the axiom of **equal treatment of symmetric strategies** for pairwise strategy symmetry structure $\mathcal{T}$ if, whenever $s_i, s_i' \in T_i$ for some $T_i \in \mathcal{T}_i$ then $x_i(s_i) = x_i(s_i')$.

This axiom says that if there is a meaningful sense in which two (pure) strategies can be considered equivalent or symmetric, then the consultant must treat those strategies symmetrically. Of course, as discussed previously (and often argued in the literature), we could link this requirement to Laplace's Principle of Insufficient Reason; that is, in the absence of information distinguishing two options, they should be ascribed equal probability, or, in our terms, treated equally in the recommendation. Our interpretation, however, is different. Two strategies are declared symmetric because there is no conceivable way to distinguish them, or, more specifically, a way for a consultant to communicate a distinction to either a player or another consultant. While a consultant can use available language (the payoff structure of the game and the commonly known external labels) to identify an element of the partition, if she further calls two strategies within the same element "A" and "B", it might well be that the player, or another consultant, have called them "B" and "A", respectively. Under this interpretation, it is simply not possible to formulate a recommendation treating symmetric strategies differently. Equal treatment of symmetric strategies simply reflects the fact that, when two consultants write down the game, the labels they use for symmetric strategies are necessarily random, without any possible bias in that randomness.[11]

Definition 2.5. A recommendation $x \in \Theta$ is a **rational strategy-symmetric recommendation** with respect to a pairwise strategy symmetry structure $\mathcal{T}$ if it satisfies the axioms of rationality and equal treatment of symmetric strategies.

The following two-player game illustrates how a small change in the symmetry structure (or the frame) might enlarge the set of recommendations.

	D_1	D_2	E_1	E_2	F_1	F_2
A_1	1, 1	0, 0	2, 2	0, 1	0, 0	0, 0
A_2	0, 0	1, 1	0, 1	2, 2	0, 0	0, 0
B_1	2, 2	1, 0	3, 3	1, 1	1, 1	1, 1
B_2	1, 0	2, 2	1, 1	3, 3	1, 1	1, 1
C_1	0, 0	0, 0	1, 1	1, 1	4, 4	1, 1
C_2	0, 0	0, 0	1, 1	1, 1	1, 1	4, 4

Game 5.

In the absence of a frame, the coarsest (pairwise) symmetry structure is given by $\mathcal{T}_1^* = \{\{A_1, A_2\}, \{B_1, B_2\}, \{C_1, C_2\}\}$ and $\mathcal{T}_2^* = \{\{D_1, D_2\}, \{E_1, E_2\}, \{F_1, F_2\}\}$. Notice that strategies A_1, A_2 for player 1 and D_1, D_2 for player 2 are strictly dominated. There are three Nash equilibria where symmetric strategies are played with identical probabilities. The first, given by $x_1^{BE} = x_2^{BE} = (0, 0, \frac{1}{2}, \frac{1}{2}, 0, 0)$, involves randomization over the symmetry classes "B" and "E". The second, given by $x_1^{CF} = x_2^{CF} = (0, 0, 0, 0, \frac{1}{2}, \frac{1}{2})$, prescribes to randomize over the symmetry classes "C" and "F". The third, given by $x_1^m = x_2^m = (0, 0, \frac{3}{10}, \frac{3}{10}, \frac{1}{5}, \frac{1}{5})$, recommends a more complex randomization among all non-dominated strategies. Among these rational strategy-symmetric recommendations, the second one involves the largest expected payoff. Hence, the

[11] We thank the associate editor for help clarifying this point.

payoff-dominant recommendation is x^{CF}. In a world where consultants focus on payoffs, this would be a candidate focal point for this game.

Suppose now that one of the strictly dominated strategies, A_1, receives a distinctive label ("shiny"). Examination of the payoff table reveals that four symmetry classes break down, and we are led to the symmetry structure given by $\mathcal{T}_1' = \{\{A_1\}, \{A_2\}, \{B_1\}, \{B_2\}, \{C_1, C_2\}\}$ and $\mathcal{T}_2' = \{\{D_1\}, \{D_2\}, \{E_1\}, \{E_2\}, \{F_1, F_2\}\}$. Obviously, x^{CF} is also a rational strategy-symmetric recommendation for this symmetry structure. However, there are now two alternative rational symmetric recommendations leading to strictly higher payoffs, namely those represented by the Nash equilibria (B_1, E_1) and (B_2, E_2). Hence we see that labeling a single, strictly dominated strategy might lead to a change in what constitutes a reasonable recommendation for a game.

The next theorem states that a recommendation satisfying both of the axioms above always exists.

Theorem 3. *For any finite normal form game Γ and any pairwise strategy symmetry structure $\mathcal{T}$, a rational strategy-symmetric recommendation with respect to $\mathcal{T}$ exists.*

The proof (see Appendix A) relies on the appropriate appeal to Kakutani's fixed point theorem. The only difficulty is to show that the restriction of the best reply correspondence to rational strategy-symmetric recommendation is nonempty-valued; in other words, whenever the opponents of a player i give the same weight to their symmetric strategies, there exists a best response of player i which gives the same weight to any two of her symmetric strategies.

Theorem 3, in conjunction with Proposition 1, implies existence of a rational strategy-symmetric recommendation whenever strategy symmetry is defined through global strategy symmetry.

The lattice structure of symmetry structures has implications for the set of Nash equilibria, due to the following observation.

Proposition 2. *Let $\mathcal{T}$ and $\mathcal{T}'$ be pairwise strategy symmetry structures of a finite, normal-form game Γ. If $\mathcal{T}$ is coarser than $\mathcal{T}'$, then any rational recommendation which is strategy-symmetric with respect to $\mathcal{T}$ is also strategy-symmetric with respect to $\mathcal{T}'$.*

The proof is immediate. Note, in particular, that the set of rational recommendations which are strategy-symmetric with respect to the trivial structure is the set of all Nash equilibria, while rational recommendations which are strategy-symmetric with respect to the coarsest structure are also strategy-symmetric with respect to any structure.

This raises an interesting point. Suppose we have a framed game, and new information arrives in the form of further attributes, additional history, etc. The effect is to refine the frame and hence the symmetry structure. The set of strategy-symmetric Nash equilibria is consequently enlarged (not refined) to a (weakly) larger set.

For example, in Crawford and Haller [10], as the base game is repeated, the outcomes of past play form histories which incorporate more and more information, acting as more and more detailed frames, and thus enlarging the set of equilibria until coordination on a desired equilibrium is possible. Crawford and Haller [10] then rely on an additional principle, Pareto efficiency, in order to select an equilibrium (see also Goyal and Janssen [13]). The argument above shows that such a "meta-norm" is necessary for any full definition of focal points, because refining symmetry structures results in enlarged sets of equilibria.

3. Global strategy and player symmetry structures

In this section we add considerations of possible symmetries between players to our symmetry structures. In order to simplify the exposition, we now take a position and concentrate on global symmetry structures.

3.1. Global symmetry structures

To motivate our definition of global symmetry structures for a game consider the following game, Game 6, which is a version of the battle-of-the-sexes.

	A_2	B_2
A_1	4, 3	0, 0
B_1	0, 0	3, 4

Game 6.

For each player in this game, the two strategies have to be declared not symmetric. For player 1, for instance, strategy A_1 can provide a payoff of 4, which is not possible for strategy B_1. If the players are exogenously declared not symmetric, then the only symmetry structure is the trivial one. If the game is indeed the battle-of-the-sexes, which suggests that it is common knowledge that player 1 is a woman and player 2 a man (or vice versa), then this is indeed the appropriate symmetry structure. However, we argue that the game should also have a second symmetry structure, in which the two players are declared symmetric.

Consider the game being played by two randomly chosen students from some subject pool, who are not informed about their opponent's identity. They are only told that they have two strategies at their disposal, one that can provide a high payoff of 4, the other a lower payoff of 3, and payoffs are paid out only if one of the two players chooses the high payoff strategy and the other the lower payoff strategy. The game is symmetric, in the sense that the two players are in symmetric positions (they both face exactly the same situation). Formally, declaring the two players symmetric can be achieved by mapping player 1's high-payoff strategy A_1 to player 2's high-payoff strategy B_2, and B_1 to A_2. This mapping is a (player) symmetry, as defined below, again taken from Nash [20].

Definition 3.1. A **symmetry** of a normal form game Γ is a tuple (σ, τ) where $\sigma : I \to I$ is a permutation of the players' names and $\tau = (\tau_i)_{i \in I}$, where, for each $k \in I$, $\tau_k : S_k \to S_{\sigma(k)}$ is a bijection, fulfilling that, for all $k \in I$ and all $s = (s_k|s_{-k}) \in S$,

$$u_k(s_k|s_{-k}) = u_{\sigma(k)}\big(\tau_k(s_k)\big|\tau_{-k}(s_{-k})\big), \tag{2}$$

where the vector notation $\tau_{-k}(s_{-k}) \in S_{-k}$ involves the appropriate permutation of coordinates, i.e. $\tau_{-k}(s_{-k}) = (\tau_{\sigma^{-1}(k)}(s_{\sigma^{-1}(k)}))_{k \neq \sigma(i)}$.

The concept of symmetry of a game was introduced by Nash [20]. Harsanyi and Selten [15] reformulated and further generalized it by allowing for positive affine transformations of the payoffs. We discuss the difference in Section 4 below. Note that, if we constrain σ to be the identity, i.e. require players not to be considered symmetric, this definition reduces to Definition 2.2.

When player symmetry is taken into account, it turns out that partitioning player sets and strategy sets into symmetry classes is not sufficient. We need to add another component to the definition of symmetry structure, which explains in what way any two or more symmetric players are symmetric. The reason is that in some cases, declaring players symmetric can be done in two or more qualitatively different ways (with potential equilibrium-payoff consequences). Consider again the two-player, two-strategy matching game M_2 (Game 1). In the absence of further information (labels) the strategies of each player are symmetric. Of course, without any labels for the players' identities, players are also symmetric. One symmetry structure should clearly be that both strategies as well as players are declared symmetric. However, players might still be symmetric if strategies are not, e.g. if strategy A of each player is labeled "heads". In this case, players are symmetric with the mapping A for player 1 to A for player 2 (and the same for B). There is another case in which players are symmetric, where A is heads for player 1 and B is heads for player 2. Now the mapping carries A for player 1 to B for player 2 (and vice versa for the other strategies). These two ways of declaring players symmetric (when two actually symmetric strategies are exogenously kept not symmetric) are qualitatively different. In the first case there are three Nash equilibria that are feasible under the symmetry structure (coordinate on A, coordinate on B, and uniform mixing). In the second case there is only one Nash equilibrium consistent with the symmetry structure (uniform mixing). This second case is very similar to the player-symmetric Battle of the Sexes, Game 6.

This means that a "symmetry structure" which only specifies that the two players are symmetric but none of the strategies are, can have two different interpretations leading to different sets of feasible recommendations. As a consequence, when considering player symmetries, it is not enough to provide a partition of the set of players and partitions of the set of strategies for each player. We need to add a specification as to *how* the two players are symmetric. This is accomplished through the concept of **identification**.

Definition 3.2. Let $(\mathcal{I}, \mathcal{T})$ be a pair where $\mathcal{I}$ is a partition of I and $\mathcal{T} = \{\mathcal{T}_i\}_{i\in I}$ with $\mathcal{T}_i$ a partition of S_i. An **identification** of the players (relative to $(\mathcal{I}, \mathcal{T})$) is a vector of bijective mappings $\alpha = (\alpha_i)_{i\in I}$, $\alpha_i : \mathcal{T}_i \mapsto \Omega_i$, where the Ω_i are sets such that $\Omega_i = \Omega_j$ whenever $i, j \in J \in \mathcal{I}$.

The sets Ω_i in the definition are inconsequential, and could be taken to be equal to $\mathcal{T}_i$ for some player $i \in J$, for each symmetry class $J \in \mathcal{I}$. A player identification merely couples together the symmetry classes of symmetric players by giving a common "label" or name to them. As a consequence, symmetric players will need to have the same number of symmetry classes.

The necessary consideration of identifications sets us apart from previous concepts. The reason is that this difficulty does not arise if the objective is to declare as many strategies and players symmetric as is possible according to payoffs only. Specifically, we will establish below (Proposition 3) that for the coarsest symmetry structure, the player identification is unique. Hence, the concept of identification is not needed in e.g. Nash [20] and Harsanyi and Selten [15], who, in our terms, deal with coarsest symmetry structures only.

The next definition summarizes how a symmetry should agree with a candidate symmetry structure and an identification thereof.

Definition 3.3. Let $(\mathcal{I}, \mathcal{T})$ be a pair where $\mathcal{I}$ is a partition of I and $\mathcal{T} = \{\mathcal{T}_i\}_{i\in I}$ with $\mathcal{T}_i$ a partition of S_i. Let α be a player identification relative to $(\mathcal{I}, \mathcal{T})$. A symmetry (σ, τ) agrees with $(\mathcal{I}, \mathcal{T}, \alpha)$ if (i) $\sigma(J) = J$ for every $J \in \mathcal{I}$, and (ii) for every $k \in I$ and $T_k \in \mathcal{T}_k$, there is a $T_{\sigma(k)} \in \mathcal{T}_{\sigma(k)}$ such that $\tau_k(T_k) = T_{\sigma(k)}$ and $\alpha(T_k) = \alpha(T_{\sigma(k)})$.

Note that, whenever $\sigma(i) = i$, the definition of identification and the second condition imply that, for every $T_i \in \mathcal{T}_i$, $\tau_i(T_i) = T_i$ (thus there is no need to spell out this condition in the definition separately).

We are now ready to present our definition.

Definition 3.4. A **global symmetry structure** of game Γ is a triple $(\mathcal{I}, \mathcal{T}, \alpha)$ where $\mathcal{I}$ is a partition of I, $\mathcal{T}$ is a collection $\mathcal{T} = \{\mathcal{T}_i\}_{i \in I}$ with each $\mathcal{T}_i$ a partition of S_i, and α is a player identification relative to $(\mathcal{I}, \mathcal{T})$, such that the following hold.

(i) For each $i \in I$, each $T_i \in \mathcal{T}_i$, and each $s_i, s_i' \in T_i$, there exists a symmetry (σ, τ) which agrees with $(\mathcal{I}, \mathcal{T}, \alpha)$ such that $\sigma(i) = i$ and $\tau_i(s_i) = s_i'$.

(ii) For each $J \in \mathcal{I}$ and each pair of (not necessarily different) players $i, j \in J$, there exists a symmetry (σ, τ) which agrees with $(\mathcal{I}, \mathcal{T}, \alpha)$ such that $\sigma(i) = j$.

The sets $T_i \in \mathcal{T}_i$ are called **strategy symmetry classes** for player i. Two strategies s_i, s_i' are said to be (globally) symmetric (relative to $\mathcal{T}$) if they belong to the same symmetry class. The sets of $\mathcal{I}$ are called **player symmetry classes**. Two players are symmetric if they belong to the same symmetry class.

When all players are identified, the definition just spelt out reduces to Definition 2.3. That is, a global strategy symmetry structure is just a global symmetry structure where all players are uniquely labeled.

Definition 3.4 identifies strategy symmetries and player symmetries simultaneously. It is important to note that it is not possible to tackle the issues of strategy symmetry and player symmetry separately. This is not evident from the examples presented so far. However, there are games in which strategies and players can only be declared symmetric when this is done simultaneously. Game 7 below (where player 1 chooses rows, player 2 chooses columns, and player 3 chooses matrices) demonstrates this point.

A

	A	B
A	2, 0, 1	0, 2, 0
B	0, 2, 0	2, 0, 1

B

	A	B
A	0, 2, 1	2, 0, 0
B	2, 0, 0	0, 2, 1

Game 7.

It is easy to see that in the coarsest symmetry structure, players 1 and 2 are symmetric, and both strategies of player 3 are symmetric. This is established using the symmetry which maps player 1 to player 2 and vice versa, strategies A and B of player 1 to strategies B and A of player 2, respectively, and at the same time swaps strategies A and B of player 3. Applying all these changes leaves the payoff matrix unchanged, which is the meaning of Eq. (2).

Suppose, however, that one adopts a different approach and identifies symmetric strategies first and only then looks at symmetries among players. Then, both strategies are symmetric for players 1 and 2 (just consider a symmetry swapping these strategies within each player's strategy set, while leaving everything else fixed), but we are not allowed to declare the two strategies of player 3 symmetric, because the symmetry which establishes this property needs to swap players 1 and 2 (who still have not been declared symmetric). But, if the two strategies of player 3 cannot be declared symmetric, when we look at player symmetry we will find that players 1 and 2 cannot be declared symmetric, because the symmetry establishing this needs to swap both strategies of player 3. This inconsistency arises because the determination of player symmetry and strategy symmetry needs to be simultaneous.

3.2. The structure of global symmetry structures

We say that one global symmetry structure $(\mathcal{I}', \mathcal{T}', \alpha')$ is **coarser** than another symmetry structure $(\mathcal{I}, \mathcal{T}, \alpha)$, if $\mathcal{I}'$ is coarser than $\mathcal{I}$, $\mathcal{T}_i'$ is coarser than $\mathcal{T}_i$ for every $i \in I$, and $\alpha(T_i) = \alpha(T_j)$ implies $\alpha'(T_i') = \alpha'(T_j')$ for all $i, j \in J \in \mathcal{I}$ and each $T_i \in \mathcal{T}_i, T_j \in \mathcal{T}_j$, $T_i' \in \mathcal{T}_i', T_j' \in \mathcal{T}_j'$ with $T_i \subseteq T_i'$ and $T_j \subseteq T_j'$. A **coarsest global symmetry structure** is a maximal element of the set of global symmetry structures according to the partial order of "coarser than". Of course, the trivial pairwise symmetry structure together with the finest partition of the set of players (each player is an element of the partition) form a trivial global symmetry structure which is finer than any other one.

The structure of global symmetry structures is more involved than the one of pairwise symmetry structures. However, analogously to Theorem 1, existence of meets and joins can also be established. This requires a group-theoretic detour (details are in Appendix A). Essentially, the set of all symmetries of a game, denoted $\mathrm{Sym}(\Gamma)$, forms a group with the composition of symmetries defined in the natural way. Each global symmetry structure defines one subgroup of this group in a natural way (the set of all symmetries which do not break symmetry classes in the global symmetry structure). The group of symmetries associated to the coarsest global symmetry structure is the grand group $\mathrm{Sym}(\Gamma)$. Global symmetry structures which coincide except for the player identification define different subgroups. The proof of the theorem below crucially uses the fact that the set of subgroups of a group has a lattice structure, which can be translated to global symmetry structures. A difficulty in the proof is that the mapping between symmetry structures and subgroups of symmetries is not onto, because a subgroup does not only specify which strategies are permuted, but also "in which way" (see Appendix A.4 for details).

Theorem 4. *For every finite game Γ the set of global symmetry structures endowed with the partial order of "coarser than" forms a (complete) lattice. There exists a coarsest global symmetry structure.*

Contrary to Theorem 1, it is not true that the join of two global symmetry structures $(\mathcal{I}', \mathcal{T}', M')$ and $(\mathcal{I}'', \mathcal{T}'', M'')$ can be constructed by setting $\mathcal{I} = \mathcal{I}' \vee \mathcal{I}''$ and $\mathcal{T}' \vee \mathcal{T}'' = \{\mathcal{T}_i' \vee \mathcal{T}_i''\}_{i \in I}$. The following counterexample shows that indeed the structure of global symmetry structures is more subtle than the one of pairwise symmetry structures. The following trivial game, Game 8, already demonstrates this fact.

	L	R
U	1, 1	1, 1
D	1, 1	1, 1

Game 8.

A possible global symmetry structure is given by $\mathcal{I}' = \{\{1, 2\}\}$ (both players are symmetric), $\mathcal{T}_1' = \{\{U\}, \{D\}\}$, $\mathcal{T}_2' = \{\{L\}, \{R\}\}$, and e.g. a player identification with $\alpha_1'(\{U\}) = \alpha_2'(\{L\})$ and $\alpha_1'(\{D\}) = \alpha_2'(\{R\})$. A different global symmetry structure is given by $\mathcal{I}' = \{\{1\}, \{2\}\}$ (players are not symmetric; hence the player identification is irrelevant), $\mathcal{T}_1'' = \{\{U, D\}\}$, and $\mathcal{T}_2'' = \mathcal{T}_2' =$

$\{\{L\}, \{R\}\}$. If we construct the joins of the respective partitions, we obtain $\mathcal{I}' \vee \mathcal{I}'' = \{\{1, 2\}\}$, $\mathcal{T}_1' \vee \mathcal{T}_1'' = \{\{U, D\}\}$, and $\mathcal{T}_2' \vee \mathcal{T}_2'' = \{\{L\}, \{R\}\}$. But players 1 and 2 cannot be symmetric and have a different number of symmetry classes, hence this cannot be part of a global symmetry structure.

It is easy to see that a given pair $(\mathcal{I}, \mathcal{T})$ might admit several compatible player identifications. That is, there might be several alternative global symmetry structures $(\mathcal{I}, \mathcal{T}, \alpha^1), \ldots, (\mathcal{I}, \mathcal{T}, \alpha^K)$ sharing the same $\mathcal{I}$ and $\mathcal{T}$, but where players are declared symmetric following different mappings α^k. An example is game M_2 where, as already discussed, the players might be declared symmetric in two qualitatively different ways. This also holds (in a less relevant way) for Game 8. In such a case, the join of all the alternative global symmetry structures $(\mathcal{I}, \mathcal{T}, \alpha^1), \ldots, (\mathcal{I}, \mathcal{T}, \alpha^k)$ is a well-defined symmetry structure, due the lattice structure. However, this join, say $(\mathcal{I}^*, \mathcal{T}^*, \alpha^*)$, might in general incorporate new symmetries not captured in $(\mathcal{I}, \mathcal{T})$. The following definition captures this concept.

Definition 3.5. The **completion** of a global symmetry structure $(\mathcal{I}, \mathcal{T}, \alpha)$ is the global symmetry structure given by the join of all structures of the form $(\mathcal{I}, \mathcal{T}, \alpha')$.

For example, in the framed Game 8, the completion of the two global symmetry structures with symmetric players and the symmetry classes corresponding to the frames is the coarsest symmetry structure where all two strategies of each player are symmetric. We say that a global symmetry structure is **complete** if it is equal to its completion, and **incomplete** if not, or, equivalently, if there exist alternative player identifications. Obviously, the coarsest global symmetry structure is always complete, since its completion cannot be strictly coarser. The following result is then immediate.

Proposition 3. *If $(\mathcal{I}, \mathcal{T}, \alpha)$ is the coarsest global symmetry structure, there exists no alternative player identification α' such that $(\mathcal{I}, \mathcal{T}, \alpha')$ is a different global symmetry structure.*

This observation clarifies why the approaches of Nash [20] or Harsanyi and Selten [15] do not need a specification of player identifications. They restrict to structures considering *all* possible symmetries, and hence implicitly work with coarsest symmetry structures. For those, it is not necessary to specify player identifications.

3.3. Global symmetry structures and frames

Analogously to Subsection 2.4, there is a correspondence between global symmetry structures and frames, where not only strategies, but also players are labeled (e.g. players 1 and 5 are men, all others are women). We first extend the concept of basic frame to accommodate this possibility. It is important to focus on the interpretation of a frame as reporting on universally observable, objective characteristics. In particular, each consultant will be able to observe the labels of all players and all strategies.

Definition 3.6. A **frame** is a pair (L, L_0), where (i) L is a basic frame, i.e. $L = (L_i)_{i \in I}$ and L_0, where $L_i : S_i \to \mathcal{Z}_i$ for each $i \in I$ for some arbitrary sets of labels $\mathcal{Z}_i$, (ii) $L_0 : I \to \mathcal{Z}_0$ is a mapping assigning players to labels from some arbitrary set $\mathcal{Z}_0$, and (iii) whenever $i, j \in I$ are such that $L_0(i) = L_0(j)$, we have that $\mathcal{Z}_i = \mathcal{Z}_j$ and $|L_i^{-1}(z)| = |L_j^{-1}(z)|$ for each $z \in \mathcal{Z}_i = \mathcal{Z}_j$.

The third condition is a minimal consistency requirement. It states that two players can only be assigned the same label if their strategies are given labels from the same set in a clearly compatible way, i.e. the number of strategies labeled the same way is the same for both players. Obviously, there is no guarantee that strategies or players labeled the same way will be declared symmetric in an associated symmetry structure, because the payoff structure of the game might deliver additional information.

As before, the coarsest global symmetry structure is the strongest (coarsest) reclassification of strategies and players that a consultant can obtain from the game, based on payoffs alone, hence, is thus "frameless".

Definition 3.7. Let (L, L_0) be a frame for game Γ. The global symmetry structure induced by (L, L_0) is the coarsest symmetry structure $(\mathcal{I}, \mathcal{T}, \alpha)$ such that $\mathcal{I}$ is finer than $\{L_0^{-1}(z) \mid z \in \mathcal{Z}_0\}$, $\mathcal{T}_i(L)$ is finer than the L_i-partition of S_i for each player i, and for all symmetric players i, j and $T_i \in \mathcal{T}_i$, $T_j \in \mathcal{T}_j$ with $\alpha_i(T_i) = \alpha_j(T_j)$, it follows that for each $s_i \in T_i$ and each $s_j \in T_j$, $L_i(s_i) = L_j(s_j)$.

Again, this global symmetry structure is always well defined by Theorem 4, with an analogous argument to the one used for frames and pairwise symmetry structures. In particular, note that if two global symmetry structures fulfill the properties stated in the definition, so does their join.

As in Subsection 2.4, the mapping which takes each frame to the global symmetry structure it induces is onto, that is, for every global symmetry structure there exists a frame which rationalizes it.

Theorem 5. *For any global symmetry structure there exists a frame such that the induced global symmetry structure is the original one.*

Proof. Fix a global symmetry structure $(\mathcal{I}, \mathcal{T}, \alpha)$, and construct the frame as follows. Let $\mathcal{Z}_0 = \mathcal{I}$. Define $L_0(i) = J$ where $J \in \mathcal{I}$ is such that $i \in J$. For each $J \in \mathcal{I}$, choose an arbitrary $j \in J$ and let $\mathcal{Z}_i = \mathcal{T}_j$ for all $i \in J$. For each $i \in J$, define $L_i(s_i) = T_j$ where $T_j \in \mathcal{T}_j$ is such that $\alpha_j(T_j) = \alpha_i(T_i)$ with $s_i \in T_i$. The L_i-partitions just reproduce $\mathcal{T}_i$, and analogously for L_0. Note that, given $z = T_j \in \mathcal{T}_j = \mathcal{Z}_j$, $L_j^{-1}(z) = T_j$ and $L_i^{-1}(z) = T_i$ with $\alpha_j(T_j) = \alpha_i(T_i)$. Since α agrees with $(\mathcal{I}, \mathcal{T})$, it follows that $|T_i| = |T_j|$, which establishes the third condition in the definition of frame. □

3.4. Equal treatment of symmetric players

Now we are finally ready to spell out the last axiom, to be added to rationality and equal treatment of symmetric strategies. From this point on, we understand the latter axiom (and the new one) to refer to global, rather than pairwise, symmetry structures.

Axiom 3. A recommendation $x \in \Theta$ satisfies the axiom of **equal treatment of symmetric players** for a global symmetry structure $(\mathcal{I}, \mathcal{T}, \alpha)$ if, whenever two players i and j are symmetric then, $x_i(T_i) = x_j(T_j)$ whenever $\alpha_i(T_i) = \alpha_j(T_j)$, $T_i \in \mathcal{T}_i$, $T_j \in \mathcal{T}_j$.

If the consultant cannot distinguish between two player roles, then he must give equivalent recommendations to those two player roles. The interpretation is similar to that behind equal treatment of symmetric strategies. If two player roles are symmetric, there is no information in the payoff matrix or the observable, exogenous labels which allows to distinguish them. Hence, when writing down a description of the game, there is an inherent randomness on the labels that a given consultant will attach to each of those possible player roles. Hence, in writing a recommendation, a consultant must necessarily treat them equally.

Definition 3.8. A recommendation $x \in \Theta$ is a **rational symmetric recommendation** with respect to global symmetry structure $(\mathcal{I}, \mathcal{T}, \alpha)$ if it satisfies the axioms of rationality, equal treatment of symmetric strategies and equal treatment of symmetric players.

Again, a recommendation satisfying all our axioms always exists.

Theorem 6. *For every finite normal form game and every global symmetry structure, there exists a rational symmetric recommendation.*

Since a Nash equilibrium fulfilling equal treatment of symmetric strategies and players for a global symmetry structure also fulfills those properties for a finer structure, it is enough to show the result for $\mathcal{T}^*$. Rational symmetric recommendations for that particular structure are identical to Nash's [20] "symmetric equilibrium points". Hence the result follows from Nash [20, Theorem 2].[12] For the sake of completeness, we include a short proof in Appendix A which builds on the proof of Theorem 3. The only (minor) added difficulty is showing that symmetric players have symmetric best responses, given a profile which respects equal treatment of symmetric strategies and players.

3.5. A definition of focal points

A minimal requirement for a strategy profile to be considered a focal point is that it is a Nash equilibrium satisfying the symmetry restrictions inherent in the possibly framed game. That is, a focal point must be a rational symmetric recommendation. In some cases this minimal requirement suffices for the identification of a focal point. If a framed game has a unique rational symmetric recommendation, then this strategy profile must be the focal point of the game. This is e.g. the case for any unframed two-player matching game M_k, where the recommendation is to randomize uniformly among all strategies. In such a case, consultants will advise in favor of this unique recommendation, and actual play will then conform to it.

Typically, however, the set of rational symmetric recommendations is not a singleton. The way these recommendations differ (which might be rather subtle) can be phrased in terms of payoff matrices and labels alone. Hence, consultants can avoid mis-coordination by appealing to a commonly known norm on how to describe and rank certain qualitative characteristics of recommendations. We call this a **meta-norm**, because it is not a norm of behavior specific to any particular game, but rather a general principle of how to choose among recommendations for all games.

We do not want to argue that such a meta-norm is always in place, but rather that it *could* be in place. If it is, actual play will conform to the rational symmetric recommendation picked by this meta-norm. A meta-norm can be defined as a mapping assigning e.g. a real number to each

[12] Note, however, that Theorem 3 does *not* follow from Nash's result. The reason is that, by Proposition 1, the coarsest global symmetry structure is *finer* than the coarsest pairwise symmetry structure, and, as Game 3 shows, it can be strictly finer.

strategy profile of the game, depending on all familiar aspects of the game. Hence, a meta-norm defines a ranking which could be interpreted as **salience**.

A simple example in which a meta-norm is needed is the two-player game of pure coordination with two pure strategies, in which coordinating on one strategy provides a higher payoff (to both players) than coordinating on the other strategy. In this case there are three rational symmetric recommendations (the two pure Nash equilibria and the unique mixed one). A simple meta-norm that has bite here is that of Pareto-dominance. If it were common knowledge that players in such cases only consider playing rational symmetric recommendations that are Pareto-undominated (among all such recommendations) then play in this game would indeed conform to this prediction.

The same meta-norm has bite in the somewhat more subtle two-player matching game M_3 with the following frame. Let the game be played twice in succession with perfect monitoring. Suppose we are now in the second stage and that both players happened to play the same strategy in the first stage. Thus, the frame in the second stage is such that two of three strategies (for both players) are labeled "not yet played" and one (the same for both players) is labeled "played". Thus, in the second stage, the so framed game has three rational symmetric recommendations: the "played" strategy for both players, randomizing uniformly among both "unplayed" strategies for both players, and randomizing uniformly among all pure strategies for both players. The meta-norm of Pareto-dominance is sufficient to identify a focal point, as the "played" equilibrium provides higher payoff than the other two. Note that the focal point of the so framed game M_3 is very different from the focal point of the unframed game M_3.

A more subtle meta-norm is needed in the two-player matching game M_2, in which the two strategies are labeled differently, say, "played" and "not yet played". Pareto-dominance is not sufficient to identify a focal point as two of the three rational symmetric recommendations provide exactly the same payoff (which is higher than that of the third one). Thus, the fact that one strategy is labeled "played" must be used directly in the meta-norm. In fact it is assumed in Crawford and Haller [10], who study repeated pure coordination games, that in such cases players will choose the previously played strategy. They thus appeal, implicitly, to the meta-norm in which among all Pareto-undominated recommendations, those involving previously played strategies are chosen.[13]

To cover all possible cases a meta-norm would have to be rather intricate. In many cases, however, a partially specified meta-norm, as e.g. Pareto-efficiency, will suffice. There are, however, many possible meta-norms one could consider. One could, for instance, base alternative meta-norms on the concepts of risk dominance or payoff equity. Many other partially specified meta-norms are possible. The step from a normative theory to a positive, descriptive one, would require consensus as to what this meta-norm should be. Alternatively, one could develop an evolutionary argument (as in e.g. Binmore and Samuelson [4]) where play evolves over time, so that even though some games are only played once, there is sufficient time and incentive for players to "find" a commonly known meta-norm.

Most examples of meta-norms belong to one of two classes. The first corresponds to meta-norms as Pareto-efficiency or risk-dominance, which, among the rational recommendations derived from the frame, select on the basis of payoffs alone. The second kind corresponds to meta-norms which are based completely on the salience of the labels and not on the payoffs.

[13] This was pointed out by Goyal and Janssen [13].

For instance when players are asked to coordinate on either "heads" or "tails" it is well documented (e.g. in Schelling [24] and Mehta et al. [18,19]) that "heads" is generally considered more salient.

In any case, meta-norms can be very intricate, and we do not expect that a unique meta-norm has emerged which enables coordination on a single (and particular) Nash equilibrium in all games. However, if we had a commonly known meta-norm in place, for a given game it would pick up a particular Nash equilibrium, which would then be commonly and thus mutually known (see Aumann and Brandenburger [1]) among players. Our point here is that such a very intricate meta-norm covering all (framed) games could well emerge or be taught, *and parts of it are probably already in place*. Given such a commonly known meta-norm, we can define the concept of a focal point.

Definition 3.9. For a given game and frame, a rational symmetric recommendation which is uniquely selected among all rational symmetric recommendations according to some meta-norm is a **focal point** with respect to this meta-norm.

This definition has a normative character in two respects. First, we do not expect every player to be able to immediately uncover all the symmetries (and asymmetries) of a game. *If* those are uncovered, however, we believe they would immediately become a guide for actual play. Second, we do not expect a universal meta-norm to be already in place. *If* such a meta-norm is established for a family of games, however, it will act as a selection device. The result is a focal point. For simple enough games, the approach will coincide with a more descriptive one, because symmetries are easily uncovered and clear meta-norms are in place. For more general games, our approach identifies the equilibria which should become focal points as players become more sophisticated and universal conventions are developed.

If one is interested in how possibly boundedly rational people behave in framed games in labs and the real world nowadays, one would have to model both the possible bounds to rationality (as, for instance, in Blume and Gneezy [7]) and the lack of a commonly held meta-norm by introducing beliefs and types, leading to models of incomplete information. In this spirit Sugden [25] proposes a model in which each strategy receives a different label according to some probability distribution with some correlation among players due to a shared culture. Similarly, one could also model uncertainty about which meta-norm is relevant as an incomplete information game. Also Janssen [17] and Casajus [8, Section 5] investigate incomplete information games with their variable universe matching games, which are based on the variable frame theory of Bacharach [2,3].

4. Payoff transformations

In our study of symmetry concepts for games, we have implicitly adopted a cardinal approach, in the sense that strategies can be told apart by the fact that they are associated to different payoffs, and hence, one might argue, those quantities have a meaning in themselves. In other words, payoffs are understood to have a commonly accepted meaning (among players) and can be used to differentiate between strategies (as well as players). Some notions of efficiency and most notions of equity rely on a cardinal interpretation of payoffs. Thus, the possible meta-norms would change if payoffs cannot be given a cardinal interpretation. The various notions of symmetry in this paper, however, are quite robust to interpreting payoffs more flexibly as just one possible way of summarizing players' preferences over outcomes.

The first observation is that, as long as we restrict ourselves to strategy symmetry (i.e. fixing players' identities), adopting a cardinal approach is less restrictive than it might seem at first glance. Consider a game $\Gamma = [I, (S_i, u_i)_{i\in I}]$, and suppose we transform the payoffs of this game while keeping the players' preferences over randomized outcomes unchanged, i.e. using affine increasing transformations. We obtain a transformed game $\Gamma' = [I, (S_i, v_i)_{i\in I}]$, where $v_i(s) = \alpha_i + \beta_i u_i(s)$ for all $s \in S$ and $i \in I$, for some $\alpha_i, \beta_i \in \mathbb{R}$ with $\beta_i > 0$.

Since such payoff transformations only use the payoffs themselves as inputs and are injective, strategy symmetries can be neither broken nor created. Further, since such transformations also leave the best-response structure unchanged, it follows that every rational strategy-symmetric recommendation of the original game is also a rational strategy-symmetric recommendation of the transformed game and vice versa. We can collect these observations in the following result.

Proposition 4. *Let games Γ and Γ' have the same player and strategy sets, and let the payoffs of Γ' are derived from the payoffs of Γ through increasing affine transformations. Then, the sets of* (*pairwise and global*) *strategy symmetry structures of Γ and Γ' are identical, and the sets of rational strategy-symmetric recommendations* (*for a given pairwise or global strategy symmetry structure*) *of Γ and Γ' coincide.*

The above result implies that adding a requirement that consultants' recommendations should be invariant with respect to player-specific increasing affine transformations of payoffs (a la Harsanyi and Selten [15]) is redundant, i.e. does not change the set of pairwise strategy symmetry structures and their induced rational strategy-symmetric recommendations.[14]

Payoff transformations as above also keep unchanged the notion of Pareto-efficiency. Other notions of efficiency (e.g. maximizing the sum of payoffs) and most notions of equity, are only unchanged if the payoff transformations are the same for all players, which is not required in the above proposition. Thus, while the symmetry structure might be unaffected by allowing such payoff transformations, the meta-norm one might want to appeal to in order to choose among recommendations might have to be different even if we allow only for such payoff transformations.

Consider now global symmetry structures, which incorporate player symmetry as well. Strictly speaking, allowing for player-specific affine transformations does change global symmetry structures by changing player symmetries. Consider, for example, the following game where payoffs are in monetary terms:

	A_2	B_2
A_1	40, 3	0, 0
B_1	0, 0	30, 4

In the coarsest global symmetry structure for this game, players (and strategies) are not symmetric. Thus, all 3 Nash equilibria are feasible rational symmetric recommendations. Indeed, it is likely that A_1, A_2 would be played if one took this game to an experimental lab. However, if we allow for affine payoff transformations, the game can be transformed into the symmetric Battle of the Sexes, Game 6. In the coarsest global symmetry structure of this game, the mixed Nash equilibrium is the only feasible rational symmetric recommendation.

[14] The set of strategy symmetry structures also remains unchanged if one allows for more general payoff transformations, as long as they are injective mappings.

However, it can be argued that the player asymmetry in the game above arises because we insist on a monetary interpretation of payoffs, and the difference between 40 for player 1 and 4 for player 2 could alternatively be captured by labels in a frame. That is, the previous global symmetry structure can be recovered by labeling players in the symmetric Battle of the Sexes; it is, however, not the coarsest one anymore. If one adopts this approach, we should consider the players to be symmetric because they are symmetric in the symmetric Battle of the Sexes, which is a transformation of the game above.[15] This possibility could be easily accommodated into our framework.

Following Harsanyi and Selten [15], as done e.g. in Casajus [8], one could go further and require that two games with the same best-response structure should be played in the same way. For instance, one might require invariance of solutions to transformations which change payoffs differently depending on the strategy profile of the opponents, i.e. of the form $\tilde{u}(s_i|s_{-i}) = \beta_i(s_{-i})u(s_i|s_{-i}) + \alpha_i(s_{-i})$ with $\alpha_i(s_{-i}), \beta_i(s_{-i}) \in \mathbb{R}$. It would be possible to incorporate such a requirement into our framework. However, the resulting changes in the definition of symmetry would lead us away from our objective.

To see the effect of allowing such payoff-transformations, consider the following game.

	A	B
A	$3, 3$	$-2, 2$
B	$2, -2$	$-1, -1$

This game can be transformed, by applying best-response-preserving transformations, into the matching game M_2. Hence the modified concept would have to declare both strategies symmetric in the original coordination game. In effect, allowing for arbitrary best-response-preserving transformations for a given game creates a large number of new symmetries and will be equivalent to an application of our concept to a transformed (and, one might say, simplified) game. However, this possibility in turn affects the possible meta-norms which one might use in a definition of focal point, e.g. Pareto efficiency becomes a void concept as coordination games are turned into matching games.

The examples in this section point out that introducing more sophisticated invariance requirements is not inconsequential. However, doing so is nevertheless feasible in our theoretical structure. For example, we could use Harsanyi and Selten's [15] symmetries in place of our Definition 3.1. Fix a game Γ and consider all (payoff) transformed games Γ^v, where transformations keep the best-response structure unchanged. Call **super-symmetry structure** of Γ any global symmetry structure of any such transformed game Γ^v. This includes all global symmetry structures of Γ itself. However, for many games there may well be additional super-symmetry structures as illustrated, for instance, by the previous example. We, thus, obtain a larger lattice of super-symmetry structures, which can be analyzed analogously to our global symmetry structures.

5. Conclusion

In this paper, we take the position that the concept of focal point reflects two different considerations. The first one is **symmetry**, and boils down to the observation that strategies or players

[15] This approach was suggested by a referee.

which cannot be told apart must be treated equally. The second one is the necessity of a **meta-norm**.

The meaning of symmetry can be readily formalized. Building on concepts introduced by Nash [20], Harsanyi and Selten [15], and Crawford and Haller [10], we have shown that, given a game, there might be many alternative, internally consistent ways to describe symmetries of strategies and players. Far from being abstract objects, we show that each of this **symmetry structures** corresponds to a **frame** (or a family of equivalent frames), that is, a set of commonly observed labelings of strategies (and players) which provide additional information about their identities. The set of symmetry structures displays a very convenient mathematical structure (a lattice), and each symmetry structure can be viewed as a subgroup of a certain group of game automorphisms (symmetries in the sense of Nash [20]). The coarsest such structure corresponds to the unframed game, that is, it captures all symmetries derived from the payoff matrix of the game. The lattice structure of the set of symmetry structures provides a rich framework where the questions posed in both the theoretical and the experimental literature on focal points can be developed and, we believe, better understood.

We deal with two different concepts of symmetry. The first one, based on Crawford and Haller [10], builds on pairwise comparisons of strategies. The second, closer to Nash [20], builds on global symmetries of the game. The pairwise concept is simpler to apply and delivers most of the intuitions we want to capture. It is, however, unsatisfactory for complex games and the proper modeling of player symmetry. Indeed we show that the predictions delivered with the global version might differ from the ones arrived at with the pairwise one.

We show that, given a symmetry structure (pairwise or global), there are always possible rational symmetric recommendations, which are Nash equilibria treating symmetric strategies and symmetric players equally. As more and more information about a game is collected, the frame becomes more detailed, the information structure becomes finer, and the set of rational symmetric recommendations grows, enabling more and more equilibria. Hence, focal points cannot, in general, be defined through symmetry alone, because the attempt to provide more information (through frames, histories, etc.) will result in an enlarged set of possible predictions. As a consequence, a meta-norm (e.g. Pareto-efficiency, risk-dominance, or equity) is necessary to explain why certain outcomes might be seen as focal.

Appendix A. Proofs

A.1. Some concepts from lattice theory

We will rely on the following concepts and elementary facts from Lattice Theory. We refer the reader to Davey and Priestley [11] or Grätzer [14] for details.

A set X endowed with a partial order $\leqslant$ is a **lattice** if both the **meet** $x \wedge x' = \inf\{x, x'\}$ (i.e. the greatest lower bound) and the **join** $x \vee x' = \sup\{x, x'\}$ (i.e. the least upper bound) exist, for every $x, x' \in X$. A lattice is **complete** if both the meet $\bigwedge S = \inf S$ and the join $\bigvee S = \sup S$ exist, for every subset $S \subseteq X$. If X is finite, joins and meets of subsets can be obtained by mere iteration.

Fact A1. Any nonempty, finite lattice is complete.

An element x of a partially ordered set $(X, \leqslant)$ is a **top** (resp. **bottom**) or greatest (resp. smallest) element if there exists no $x' \in X$ with $x \leqslant x'$ (resp. $x' \leqslant x$) and $x \neq x'$. If a lattice is complete, the top and the bottom are given by $\sup X$ and $\inf X$, respectively.

Fact A2. Any nonempty, complete lattice has a top and a bottom.

A partially ordered set such that any two elements have a join (but not necessarily a meet) is called a **join semilattice**. In the finite case, as long as a bottom is present, existence of meets is guaranteed. An analogous result is true for the dual concept of **meet semilattice**.

Fact A3. Any finite join semilattice with a bottom (and any finite meet semilattice with a top) is actually a lattice.

A.2. Proofs from Section 2

The following lemma gives a useful characterization of the join of two partitions.[16]

Lemma A1. *Let S_i be a finite set.*

(a) *Let $\mathcal{T}_i$ and $\mathcal{T}_i'$ be partitions of S_i. If $\mathcal{T}_i'$ is coarser than $\mathcal{T}_i$, then every set $T_i \in \mathcal{T}_i$ can be written as a (finite) disjoint union of the sets which form $\mathcal{T}_i'$.*
(b) *The finest partition coarser than two partitions $\mathcal{T}_i$ and $\mathcal{T}_i'$ of S_i is given by the equivalence classes of the following relation. Two elements $s_i, s_i' \in S_i$ are related if and only if there exists a finite sequence of elements of S_i, $s_i^0 = s_i, s_i^1, s_i^2, \ldots, s_i^k = s_i'$ such that s_i^t and s_i^{t+1} are in the same symmetry class of either $\mathcal{T}_i'$ or $\mathcal{T}_i''$, for all $t = 0, \ldots, k-1$.*

Proof. Part (a) is straightforward. To see part (b), first note that the relation given in the statement is a binary equivalence relation and thus its equivalence classes define a partition, which we denote $(\mathcal{T} \vee \mathcal{T}')_i$. By construction, this partition is coarser than both $\mathcal{T}_i$ and $\mathcal{T}_i'$. It remains to show that it is the finest such partition.

Let $\mathcal{T}_i''$ be a partition coarser than both $\mathcal{T}_i$ and $\mathcal{T}_i'$. Let T_i'' be any of the sets in $\mathcal{T}_i''$. By part (a), there exist $T_{i,1}, \ldots, T_{i,\ell} \in \mathcal{T}_i$ and $T_{i,1}', \ldots, T_{i,\ell'} \in \mathcal{T}_i'$ such that

$$T_i'' = \bigcup_{r=1}^{\ell} T_{i,r} = \bigcup_{r=1}^{\ell'} T_{i,r}'.$$

A direct consequence is that no element of T_i'' can be related through the relation above to any element outside of T_i''. This proves that $(\mathcal{T} \vee \mathcal{T}')_i$ is finer than $\mathcal{T}_i''$. □

Proof of Theorem 1. Consider two symmetry structures $\mathcal{T}'$ and $\mathcal{T}''$. We will first show that the collection $\mathcal{T} = \{\mathcal{T}_i' \vee \mathcal{T}_i''\}_{i \in I}$ is also a symmetry structure.

Let $T_i \in \mathcal{T}_i$, and let $s_i, s_i' \in T_i$. We need to show that there is a set of renamings ρ_j for all $j \neq i$ such that $\rho_j(T_j) = T_j$ for all $T_j \in \mathcal{T}_j$ for all $j \neq i$ such that $u(s_i|s_{-i}) = u(s_i'|\rho_{-i}(s_{-i}))$ for all $s_{-i} \in S_{-i}$. First note that the set of renamings with $\rho_j(T_j) = T_j$ for all $T_j \in \mathcal{T}_j$ for all $j \neq i$ includes all renamings with the property $\rho_j(T_j') = T_j'$ for all $T_j' \in \mathcal{T}_j'$ for all $j \neq i$ as well as $\rho_j(T_j'') = T_j''$ for all $T_j'' \in \mathcal{T}_j''$ for all $j \neq i$. This is due to the fact that, by Lemma A1(a), each

[16] We provide this technical lemma and its proof for completeness. It is well known in abstract algebra that the set of partitions of a set, finite or not, forms a lattice, which is isomorphic to a certain lattice of permutation groups. For the original proofs see Birkhoff [5] and Ore [21]. See also Roman [22] for a more recent exposition.

$T_j \in \mathcal{T}_j$ can be written as a finite union of sets $T'_j \in \mathcal{T}'_j$ as well as of sets $T''_j \in \mathcal{T}''_j$, as $\mathcal{T}$ is coarser than both $\mathcal{T}'$ and $\mathcal{T}''$.

Suppose that there is a $T' \in \mathcal{T}'$ or a $T'' \in \mathcal{T}''$ such that s_i, s'_i are either both in T'_i or both in T''_i or both. Let us w.l.o.g. suppose s_i, s'_i are both in T'_i. By definition of symmetry structure, there is a renaming ρ_j such that $\rho_j(T'_j) = T'_j$ for all $T'_j \in \mathcal{T}'_j$ for all $j \neq i$. But by the above observation this renaming then also satisfies $\rho_j(T_j) = T_j$ for all $T_j \in \mathcal{T}_j$ for all $j \neq i$.

Now suppose that s_i and s'_i are not in the same symmetry class of either $\mathcal{T}'$ or $\mathcal{T}''$. By Lemma A1(b), there exist $s_i^0 = s_i, s_i^1, s_i^2, \dots, s_i^k = s'_i$ such that s_i^t and s_i^{t+1} are in the same symmetry class of either $\mathcal{T}'_i$ or $\mathcal{T}''_i$, for all $t = 0, \dots, k-1$. The conclusion follows from an iteration of the previous argument.

By construction, $\mathcal{T}$ is the finest symmetry structure which is coarser than both $\mathcal{T}'$ and $\mathcal{T}''$, i.e. their join. Thus any two elements of the (finite) set of symmetry structures of a finite normal form game have a join. Thus symmetry structures form a join semilattice with a bottom (the trivial symmetry structure), and, by Fact A3, a lattice. □

Proof of Corollary 1. The (finite) set of symmetry structures is nonempty since the trivial symmetry structure exists, and forms a lattice by Theorem 1. The result follows from Facts A1 and A2. □

Proof of Theorem 3. Let $\beta_i : \Theta_{-i} \to \Theta$ denote the (mixed) best-reply correspondence of player i, and $\beta : \Theta \to \Theta$ the product correspondence given by $\beta(x) = \bigtimes_{i \in I} \beta_i(x_{-i})$. We know that the β_i, and hence β, are nonempty and convex-valued, and upper hemicontinuous. Hence, Kakutani's theorem implies existence of fixed points of β, which are (mixed) Nash equilibria of Γ. We have to show that at least one of them fulfills strategy symmetry.

Let $\mathcal{T}$ denote a symmetry structure for Γ. For each player $i \in I$, let $\widetilde{\Theta}_i$ be the set of mixed strategies x_i such that $x_i(s_i) = x_i(s'_i)$ whenever s_i, s'_i belong the same symmetry class in $\mathcal{T}_i$. Notice that $\widetilde{\Theta}_i$ is convex. Define $\widetilde{\beta}_i : \widetilde{\Theta}_{-i} \to \widetilde{\Theta}_i$ by $\widetilde{\beta}_i(x_{-i}) = \beta_i(x_{-i}) \cap \widetilde{\Theta}_i$. Thus $\widetilde{\beta}_i$ is convex-valued by definition, and upper hemicontinuous because it is the intersection of two upper hemicontinuous correspondences.

To see that it is nonempty-valued, we have to show that for any $x_{-i} \in \widetilde{\Theta}_{-i}$, there exists a best response of player i which gives the same weight to any two symmetric strategies. Fix $x_{-i} \in \widetilde{\Theta}_{-i}$. For each $j \neq i$ and each $T_j \in \mathcal{T}_j$, there exists $y(T_j) \geqslant 0$ such that $x_j(s_j) = y(T_j)$ for all $s_j \in T_j$. Let s_i, s'_i be symmetric. Then there exist renamings ρ_j of S_j (for all $j \neq i$) such that $\rho_j(T_j) = T_j$ for all $T_j \in \mathcal{T}_j$ and $u(s_i|s_{-i}) = u(s'_i|\rho_{-i}(s_{-i}))$ for all $s_{-i} \in S_{-i}$. Then, making extensive use of the $-i$ notation for product spaces,

$$\begin{aligned} u_i(s_i|x_{-i}) &= \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} x_j(s_j)\right) u_i(s_i|s_{-i}) = \sum_{T_{-i} \in \mathcal{T}_{-i}} \left(\prod_{j \neq i} y(T_j)\right) \sum_{s_{-i} \in T_{-i}} u_i(s_i|s_{-i}) \\ &= \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} x_j\big(\rho_j(s_j)\big)\right) u_i\big(s'_i \big| \rho_{-i}(s_{-i})\big) = u_i\big(s'_i \big| x_{-i}\big) \end{aligned}$$

where the third equality follows from the definition of symmetric strategies if we recall that the renamings ρ_j only permute strategies within symmetry classes T_j; thus if $x_j(s_j) = y(T_j)$, then $x_j(\rho_j(s_j)) = y(T_j)$. It follows that s_i and s'_i yield the same payoff against x_{-i}, thus either both are or neither is a best response to x_{-i}.

Consider now a pure best response to x_{-i}, s_i. Construct a mixed strategy x_i by giving identical weight to all strategies in the symmetry class of s_i. Since x_i is a convex combination of best responses, we have that $x_i \in \beta(x_{-i})$. This proves the nonemptyness of $\widetilde{\beta}_i(x_{-i})$ for all i. Hence, $\widetilde{\beta}$ satisfies nonemptyness, convex-valuedness, upper hemicontinuity, and, hence, by Kakutani's fixed point theorem there is a fixed point. Since a fixed point of $\widetilde{\beta}$ is also a fixed point of β, it is a Nash equilibrium. □

A.3. *Some concepts from group theory*

We will rely on the following concepts and elementary facts from Group Theory. We refer the reader to Rose [23] or Hungerford [16] for details.

A **group** is a nonempty set G endowed with a binary, internal operation "$\cdot$" in G satisfying the associative property ($(g_1g_2)g_3 = g_1(g_2g_3)$ for all $g_1, g_2, g_3 \in G$), with an identity element ($1_G \in G$ such that $1_G g = g1_G = g$ for all $g \in G$) and such that every element $g \in G$ has an inverse $g^{-1} \in G$ according to this operation ($g^{-1}g = gg^{-1} = 1_G$). A **subgroup** of G is a subset $H \subseteq G$ such that $1_G \in H$ and such that it is a group with the restriction of the binary operation of G to H.

Fact A4. A nonempty subset H of a group G is a subgroup if and only if $g_1g_2^{-1} \in H$ for every $g_1, g_2 \in H$.

The set of groups of a subgroup have a lattice structure. The meet is simple.

Fact A5. The intersection of two subgroups of a group is also a subgroup.

The join is more involved. The union of two subgroups is in general not a subgroup. Given a subset (not necessarily a subgroup) H of a group G, the **subgroup generated by** H, denoted $\langle H\rangle$, is defined as the smallest subgroup of G containing H. Thus the join of two subgroups H_1 and H_2 is $\langle H_1 \cup H_2\rangle$.

Fact A6. Given two subgroups H_1, H_2 of a group G, the subgroup generated by H_1 and H_2 is the set of all finite products $g_1g_2g_3\cdots g_r$ where $g_\ell \in H_1 \cup H_2$ for all $\ell = 1, \dots, r$.

A.4. *Proofs from Section 3*

We will prove Theorem 4 through a series of intermediate results. First note that the composition of symmetries (defined in the natural way) is a symmetry, and the inverse (σ^{-1}, τ^{-1}) of a symmetry (σ, τ) is also a symmetry, where $\tau^{-1} = \{\tau_i^{-1}\}_{i\in I}$. Further, the collection of identity mappings on I and S_i form a trivial symmetry. In summary, the set of all symmetries, which we will denote by $\mathrm{Sym}(\Gamma)$, forms a **group** with the operation given by symmetry composition.

Let $\mathrm{Sym}(\mathcal{I}, \mathcal{T}, \alpha)$ be the set of *all* symmetries which agree with a global symmetry structure $(\mathcal{I}, \mathcal{T}, \alpha)$. If there is a unique compatible identification, we write simply $\mathrm{Sym}(\mathcal{I}, \mathcal{T})$. Let $(\mathcal{I}^0, \mathcal{T}^0)$ denote the trivial global symmetry structure. Then $\mathrm{Sym}(\mathcal{I}^0, \mathcal{T}^0)$ is the subgroup formed by the identity symmetry only.

Lemma A2. *Given a global symmetry structure* $(\mathcal{I}, \mathcal{T}, \alpha)$*, the set* $\mathrm{Sym}(\mathcal{I}, \mathcal{T}, \alpha)$ *is a subgroup of* $\mathrm{Sym}(\Gamma)$.

Proof. It is enough to observe that the composition of two symmetries agreeing with $(\mathcal{I}, \mathcal{T}, \alpha)$ also agree with $(\mathcal{I}, \mathcal{T}, \alpha)$, and the inverse of a symmetry agreeing with $(\mathcal{I}, \mathcal{T}, \alpha)$ also agrees with $(\mathcal{I}, \mathcal{T}, \alpha)$. The proof follows then from Fact A4.[17] □

The following property follows directly from the definition of $\mathrm{Sym}(\cdot)$.

Lemma A3. *Let $(\mathcal{I}', \mathcal{T}', \alpha')$ and $(\mathcal{I}'', \mathcal{T}'', \alpha'')$ be two global symmetry structures of a game Γ. Then $(\mathcal{I}', \mathcal{T}', \alpha')$ is coarser than $(\mathcal{I}'', \mathcal{T}'', \alpha'')$ if and only if* $\mathrm{Sym}(\mathcal{I}'', \mathcal{T}'', \alpha'')$ *is a subgroup of* $\mathrm{Sym}(\mathcal{I}', \mathcal{T}', \alpha')$.

Let Φ be an arbitrary subgroup of $\mathrm{Sym}(\Gamma)$. Let $\mathcal{I}(\Phi)$ be the partition of I given by the binary equivalence relation where two players i and j are related if and only if there exists $(\sigma, \tau) \in \Phi$ such that $\sigma(i) = j$. For each player i, Let $\mathcal{T}_i$ be the partition of S_i given by the binary equivalence relation where two strategies s_i, s_i' are related if and only if there exists $(\sigma, \tau) \in \Phi$ such that $\sigma(i) = i$ and $\tau_i(s_i) = s_i'$. That these relations are indeed binary equivalence relations follows from the fact that Φ is a subgroup.

Proposition A1. *For each subgroup Φ of* $\mathrm{Sym}(\Gamma)$*, there exists a unique player identification $\alpha(\Phi)$ such that the collection $(\mathcal{I}(\Phi), \mathcal{T}(\Phi), \alpha(\Phi))$ with $\mathcal{T}(\Phi) = \{\mathcal{T}_i(\Phi)\}_{i \in I}$ is a global symmetry structure. Further, for each global symmetry structure $(\mathcal{I}, \mathcal{T}, \alpha)$, if $\Phi = \mathrm{Sym}(\mathcal{I}, \mathcal{T}, \alpha)$ then $(\mathcal{I}(\Phi), \mathcal{T}(\Phi), \alpha(\Phi)) = (\mathcal{I}, \mathcal{T}, \alpha)$.*

Proof. All we have to show is that there exists a suitable player identification. This is equivalent[18] to the statement that, for any two symmetries (σ', τ'), $(\sigma'', \tau'') \in \Phi$, whenever two players i, j are such that $\sigma'(i) = j$ and $\sigma''(i) = j$, then for each $T_i \in \mathcal{T}_i$, $\tau_i'(T_i) = \tau_i''(T_i)$. Suppose not, i.e. $\tau_i'(T_i) = T_j' \neq T_j'' = \tau_i''(T_i)$. Since Φ is a subgroup, one has $(\sigma', \tau') \circ (\sigma'', \tau'')^{-1} \in \Phi$. But this symmetry maps player j to itself and symmetry class T_j'' to T_j', a contradiction with the definition of $\mathcal{T}_j$. □

In this sense, a global symmetry structure can be identified with its associated group of symmetries.[19]

Let Φ be an arbitrary subgroup of $\mathrm{Sym}(\Gamma)$. It is immediate to see that Φ is a subgroup of $\mathrm{Sym}(\mathcal{I}(\Phi), \mathcal{T}(\Phi), \alpha(\Phi))$, where again $\leqslant$ denotes the relation "to be a subgroup of". However, equality needs not hold in general.[20]

---

[17] If we had considered the set of all symmetries agreeing with a global symmetry structure but ignored player identifications (i.e. dropped the requirement $\alpha(T_k) = \alpha(\tau_k(T_k))$), this result would not be true.

[18] The sufficiency follows because the composition of symmetries in the subgroup Φ is also in Φ.

[19] There is an interesting connection with Blume [6], who uses a group-theoretic approach to formalize partial languages, which are defined as subgroups of permutations of abstract labels, which could e.g. correspond to the labels used in a frame. Blume [6] concentrates on the issue of how partial languages facilitate coordination for repeated matching games (following Crawford and Haller [10]), and how fast learning occurs.

[20] An example can be constructed as follows. Take the game M_3 with strategies A, B, C for both players. The *alternating group* is generated by the permutation $A \to B \to C$. Using permutations in this group only one can construct a subgroup of $\mathrm{Sym}(\Gamma)$ which generates the coarsest global symmetry structure, i.e. all strategies are declared symmetric. However, this is a strict subgroup of $\mathrm{Sym}(\Gamma)$.

We are now ready to prove Theorem 4.

Proof of Theorem 4. The set of subgroups of a group is a lattice due to Facts A5 and A6. By Fact A3, it is enough to prove that any two global symmetry structures have a meet. Let $(\mathcal{I}', \mathcal{T}', \alpha')$ and $(\mathcal{I}'', \mathcal{T}'', \alpha'')$ be two global symmetry structures of a game Γ. Let $\Phi' = \mathrm{Sym}(\mathcal{I}', \mathcal{T}', \alpha')$ and $\Phi'' = \mathrm{Sym}(\mathcal{I}'', \mathcal{T}'', \alpha'')$, and consider the subgroup $\Phi' \cap \Phi''$. Construct the global symmetry structure $(\mathcal{I}^*, \mathcal{T}^*, \alpha^*)$ from $\Phi' \cap \Phi''$ as in Proposition A1. It follows from construction that this structure is finer than both $(\mathcal{I}', \mathcal{T}', \alpha')$ and $(\mathcal{I}'', \mathcal{T}'', \alpha'')$.

To see that $(\mathcal{I}^*, \mathcal{T}^*, \alpha^*)$ is the meet of the two original structures, let $(\mathcal{I}, \mathcal{T}, \alpha)$ be finer than $(\mathcal{I}', \mathcal{T}', \alpha')$ and $(\mathcal{I}'', \mathcal{T}'', \alpha'')$. Let $\Phi = \mathrm{Sym}(\mathcal{I}, \mathcal{T}, \alpha)$. It follows that Φ is a subgroup of $\Phi' \cap \Phi''$, which in turn is a subgroup of $\Phi^* = \mathrm{Sym}(\mathcal{I}^*, \mathcal{T}^*, \alpha^*)$. By Proposition A1, $(\mathcal{I}, \mathcal{T}, \alpha)$ is finer than $(\mathcal{I}^*, \mathcal{T}^*, \alpha^*)$. □

Although we skip the details here, Fact A6 provides us with an explicitly computable construction of the join of two global symmetry structures $(\mathcal{I}', \mathcal{T}', \alpha')$ and $(\mathcal{I}'', \mathcal{T}'', \alpha'')$. In the join, two players are declared symmetric if and only if $\sigma(i) = j$ in a symmetry (σ, τ) which can be written as the product of symmetries which agree with either $(\mathcal{I}', \mathcal{T}', \alpha')$ and $(\mathcal{I}'', \mathcal{T}'', \alpha'')$. The construction for strategies is analogous.

We now turn to our second existence result. The proof is only sketched and given for completeness, because the result also follows from Nash [20, Theorem 2].

Proof of Theorem 6. Let $(\mathcal{I}, \mathcal{T}, \alpha)$ denote the symmetry structure. Define $\widetilde{\beta}_i$ as in the proof of Theorem 3, with the obvious change that symmetry classes belong to a global symmetry structure and not a pairwise one. Hence it is a convex-valued, upper-hemicontinuous correspondence.

Let $\hat{\Theta}$ be the subset of $\prod_{i \in I} \widetilde{\Theta}_i$ such that, whenever two players i, j are symmetric, $x_i(s_i) = x_j(s_j)$ for any $s_i \in T_i \in \mathcal{T}_i$, $s_j \in T_j \in \mathcal{T}_j$ with $\alpha_i(T_i) = \alpha_j(T_j)$. Let y_i be a best response to $x \in \hat{\Theta}$. By definition of player symmetry, defining y_j as just specified for any player j which is symmetric with i yields a best response for player j. Thus we can define $\hat{\beta}(y) = \widetilde{\beta}(y) \cap \hat{\Theta}$ and, since $\hat{\Theta}$ is convex, $\hat{\beta}$ is convex- and upper hemicontinuous.

It remains to show that $\hat{\beta}$ is nonempty-valued, that is, given $x \in \hat{\Theta}$, every player i has a best response which gives the same weight to any two symmetric strategies. This follows as in the proof of Theorem 3, the only change being that the symmetry linking two symmetric strategies might call for a permutation among symmetric players. As argued above, symmetric players always have symmetric best responses against profiles in $\hat{\Theta}$, and we conclude that $\hat{\beta}$ is nonempty-valued. Hence a fixed point exists. □

References

[1] R. Aumann, A. Brandenburger, Epistemic conditions for Nash equilibrium, Econometrica 63 (1995) 1161–1180.

[2] M. Bacharach, Games with concept-sensitive strategy spaces, Mimeo, 1991.

[3] M. Bacharach, Variable universe games, in: K. Binmore, A. Kirman, P. Tani (Eds.), Frontiers of Game Theory, MIT Press, Cambridge, MA, 1993, pp. 255–276.

[4] K. Binmore, L. Samuelson, The evolution of focal points, Games Econ. Behav. 55 (2006) 21–42.

[5] G. Birkhoff, On the structure of abstract algebras, Proc. Camb. Philos. Soc. 31 (1935) 433–454.

[6] A. Blume, Coordination and learning with a partial language, J. Econ. Theory 95 (2000) 1–36.

[7] A. Blume, U. Gneezy, Cognitive forward induction and coordination without common knowledge: An experimental study, Games Econ. Behav. 68 (2010) 488–511.

[8] A. Casajus, Focal points in framed strategic forms, Games Econ. Behav. 32 (2000) 263–291.

[9] A. Casajus, Focal Points in Framed Games, Springer Verlag, 2001.
[10] V. Crawford, H. Haller, Learning how to cooperate: Optimal play in repeated coordination games, Econometrica 58 (1990) 571–596.
[11] B. Davey, H. Priestley, Introduction to Lattices and Order, 2nd ed., Cambridge University Press, Cambridge, 2002.
[12] D. Gauthier, Coordination, Dialogue 14 (1975) 195–221.
[13] S. Goyal, M. Janssen, Can we rationally learn to coordinate?, Theory Dec. 40 (1996) 29–49.
[14] G. Grätzer, General Lattice Theory, 2nd ed., Birkhäuser, Basel, 2003.
[15] J. Harsanyi, R. Selten, General Theory of Equilibrium Selection in Games, MIT Press, Cambridge, MA, 1988.
[16] T.W. Hungerford, Algebra, Springer Verlag, New York, 1980.
[17] M.C.W. Janssen, Rationalizing focal points, Theory Dec. 50 (2001) 119–148.
[18] J. Mehta, C. Starmer, R. Sugden, Focal points in pure coordination games: An experimental investigation, Theory Dec. 36 (1994) 163–185.
[19] J. Mehta, C. Starmer, R. Sugden, The nature of salience: An experimental investigation of pure coordination games, Amer. Econ. Rev. 74 (1994) 658–673.
[20] J. Nash, Non-cooperative games, Ann. Math. 52 (1951) 286–295.
[21] O. Ore, Theory of equivalence relations, Duke Math. J. 9 (1942) 573–627.
[22] S. Roman, Lattices and Ordered Sets, Springer Verlag, Berlin–Heidelberg–Vienna, 2008.
[23] J.S. Rose, A Course in Group Theory, Cambridge University Press, Cambridge, UK, 1978, Dover reprint 1994.
[24] T.C. Schelling, The Strategy of Conflict, Harvard University Press, Cambridge, 1960.
[25] R. Sugden, A theory of focal points, Econ. J. 105 (1995) 533–550.
[26] A. Tversky, D. Kahneman, The framing of decisions and the psychology of choice, Science 211 (1981) 453–458.