

Hidden Symmetries and Focal Points*

Carlos Alós-Ferrer[†] and Christoph Kuzmics[‡]

October 2008

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 alternative 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. The definition of a symmetry structure is given in steps of increasing complexity allowing us to precisely explain differences between existing definitions. *Journal of Economic Literature* Classification Numbers: 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 (1960). Focal points are understood as those (and all those) strategy profiles which can be viewed as the “obvious way”¹ to play a one-shot framed game, where a frame² endows strategies and players with objectively observable characteristics. 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 a given normal-form game.

*We would like to thank Andreas Blume, Vincent Crawford, Péter Esö, Hans Haller, Maarten Janssen, Andrew McLennan, Roger Myerson, Itai Sher, Jonathan Weinstein, and seminar participants in Northwestern University, Vancouver, Paris, and Karlsruhe for helpful comments. Financial support from the Austrian Science Fund (FWF) under Project P18141-G09 is gratefully acknowledged.

[†]Department of Economics, University of Konstanz, Box 150, D-78457 Konstanz (Germany). E-mail: Carlos.Alos-Ferrer@uni-konstanz.de

[‡]Kellogg School of Management, Northwestern University, Evanston IL 60208 (USA). E-mail: C-Kuzmics@kellogg.northwestern.edu

¹See e.g. Kreps (1990, pp. 405 ff.), which, among other things, motivated us to write this paper.

²The concept of frame goes back to Tversky and Kahneman (1981). It has been formalized and used in different ways in the game-theoretic literature on focal points.

The paper makes three conceptual points. First, we show that symmetry structures, which have a rich and useful mathematical structure, are in natural, essentially one-to-one, correspondence with the different possible frames of a given game. Thus studying frames is the same as studying symmetry structures.

Second, 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. We show that finer symmetry structures have larger sets of associated equilibria. As more symmetries are broken by increasingly more detailed frames, the number of possible predictions grows. Hence, if focal points are to serve as equilibrium selection devices, the second part of their definition 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**.

Third, we want to emphasize that a theory of focal points has bite beyond the case of pure coordination games (on which most of the literature has concentrated). The framework we develop encompasses arbitrary normal-form games.

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.³

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 (1960) 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. A game could also come with an unfamiliar and hence seemingly useless frame as well.

The approach in this paper is a normative one. While we believe that our analysis is helpful in predicting how players might play, we focus on characterizing solutions to a game that players “should” play. For the sake of concreteness, we imagine the following situation. A player is about to play a game which is completely new to her, and 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. For

³Suppose you and your significant other find yourselves suddenly and unexpectedly separated in a new shopping center. You do have the prior understanding that if such an event occurs you meet again at the main entrance. This shopping center, unfortunately, has two main entrances. The fact that the ensuing game of where to meet is of the same form as the game of which side on the road to drive on is of little use to both of you. The game is thus an unfamiliar one.

instance, 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 rapidly conclude that a strategy must be solely identified by its associated vector of 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. 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**.

Suppose further that all players involved in a game obtain advice from (different) game theory consultants as to how to play the game. 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. 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**.

The question of which strategies and which players can be declared symmetric for a given normal-form game is surprisingly subtle, and indeed different qualitative answers are possible. Our first answer builds on the definition of symmetric strategies in 2-player games of Crawford and Haller (1990, Appendix), extended to general n -player games.⁴ In many cases, in particular for the two-player matching games studied in (the main body of) Crawford and Haller (1990), this concept is sufficient for the analysis.

⁴Crawford and Haller (1990) 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 (2000).

We show, however, that it is in general not enough. First, strategy symmetries cannot be established independently from player symmetries. Second, Crawford and Haller's (1990) concept is based on pairwise strategy comparisons, but a global concept is needed once one moves away from pure coordination games. 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 (1951) concept of symmetries and Harsanyi and Selten's (1988) game automorphisms. The resulting concept of global symmetry structures is isomorphic 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 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.

In some simple games, the axioms sketched above lead to a unique rational symmetric recommendation. For general framed games, however, this is not the case. Thus, as implicitly recognized by the literature, the multiplicity problem can only be solved through an appeal to some meta-norm. A focal point is then a rational symmetric recommendation which is uniquely selected by means of such a meta-norm.

Meta-norms can range from fairly simple to very intricate. For a first example, one can rely on the fact that games have payoffs, in, say, monetary terms, which are familiar to all players. Within a pure coordination game (i.e. all off-diagonal payoffs are zero), familiarity with money is probably enough to ensure that players coordinate on a unique Pareto-efficient equilibrium (which is then also risk-dominant), 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.⁵

When we say a frame is familiar we mean that players hold a common meta-norm as to how to compare the various labels. This meta-norm is a ranking of relative salience (prominence, or conspicuousness) of labels, a common term in the literature on focal points. For instance, "heads" is generally considered more salient than "tails" (see e.g. Schelling (1960)). In examples we will typically assume that the meta-norm of how to evaluate labels comes lexicographically after a first appeal to a meta-norm over money, such as the Pareto-optimality criterion. Even in games with unfamiliar frames there might be aspects of the game other than payoffs that players might be familiar with. For instance, there might be a label that is only attached to a single strategy, while all other labels appear at least twice. A meta-norm could thus be commonly held, which

⁵Variations of this meta-norm are called the "Principle of Coordination" in Gauthier (1975), Bacharach (1991, 1993), Sugden (1995) and Casajus (2000), "rationality in the extended sense" in Goyal and Janssen (1996), and the "Principle of Individual Team Member Rationality" in Janssen (2001).

does not depend on the game at hand, but induces players to opt for such a uniquely labeled strategy (in cases where it is an equilibrium of the game).

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, except in a tentative exploration of recent experimental literature in Section 6. 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.

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 (2008)) and the lack of a commonly held meta-norm by introducing beliefs and types, leading to models of incomplete information. In this spirit Sugden (1995) 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 (2001) and Casajus (2000) investigate more behaviorally flavored models with their variable universe matching games, which are based on the variable frame theory of Bacharach (1991, 1993). Players play a matching game but can be of different types, which reflect their potential (un)awareness of some attributes strategies have.⁶ This gives rise to another incomplete information game.

The paper is structured as follows. Section 2 briefly introduces the notation. Sections 3 and 4 present the concepts of pairwise and global symmetry, respectively, and analyze the structure of symmetry structures, their correspondence to frames, and the existence of rational-symmetric recommendations. Section 5 provides a discussion of the idea of unfamiliar frames. Section 6 discusses our concept of focal points based on the notion of a meta-norm and tentatively explores recent experimental literature as to what this meta-norm, if one is in place, could be. Section 7 concludes. Proofs (which rely on some elementary group and lattice theory) are relegated to the Appendix.

⁶Binmore and Samuelson (2006) investigate an evolutionary model in which players recurrently face games where strategies come with two attributes. They investigate under which circumstances a meta-norm emerges that uses only one or both attributes when paying attention is costly.

2 Games and Frames

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 $u : S \rightarrow \mathbb{R}^{|I|}$ is the function whose i -th coordinate is u_i .

Following game-theoretic conventions, for all $s \in S$ we write $u_i(s) = u_i(s_i | s_{-i})$, where $s_{-i} \in S_{-i} = \times_{j \neq i} S_j$. 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 $\Theta = \times_{i \in I} \Theta_i$ the set of mixed strategy profiles. 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 $i = 1, 2$, and shall be denoted by $\text{Diag}(u(s^1, s^1), u(s^2, s^2), \dots, u(s^k, s^k))$, where k is the number of strategies. 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 (2000, Theorem 5.6) and Janssen (2001, Propositions 1 and 2) are restricted to matching games.

We now turn to frames. A **label** is any observable characteristic that can be objectively established and that consultants can attach to strategies when analyzing the game. The first examples that come to mind are neutral adjectives like “red”, “shiny”, “square”, and so on, and we will focus on such labels 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”.

Let \mathcal{Z}_i be a universal set of **labels** for each player i . A **frame** for the game Γ is a collection $L = (L_i)_{i \in I}$ where $L_i : S_i \rightarrow \mathcal{Z}_i$ for each $i \in I$. It is important to hinge on the interpretation of a frame as reporting on universally observable, objective characteristics. In particular, the consultant will be able to observe the labels $L_i(s_i)$ of all strategies of all players.

Unless otherwise stated, frames are assumed to be **familiar**. When we say a frame is familiar we mean that in addition to labels being observable and objectively distinct, players may also have a ranking of the labels in terms of their **salience**. This ranking is used by players when there is no other criterion to choose between strategies. Given that such rankings may well vary between different groups of individuals, we do not want to postulate a particular ranking but rather study the set of recommendations for

all possible such rankings. See Section 5 for the case when labels are unfamiliar.

3 Strategy Symmetry With and Without (Familiar) Frames

In this section, we present a first approximation to the idea of symmetry structures in games and their relevance for focal points. The concept we will introduce relies on two simplifications. First, we will ignore symmetry among players and concentrate on symmetry among strategies (of a given player) only. Second, we will 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 (1990).

The constraints just mentioned allow us to discuss most of the intuitions behind our approach while greatly reducing the necessary conceptual and analytical complexity. Furthermore, the resulting 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 Section 4, and building upon the intuitions developed in this section, we will discuss a **global** notion of symmetry, while simultaneously allowing for player symmetry. For some special games, such as matching games (the object of study in the main body of Crawford and Haller (1990)), our global notion of symmetry is equivalent to pairwise symmetry.

3.1 Pairwise Strategy Symmetry

In this section we first provide two simple definitions of symmetric strategies before turning to our pairwise definition of symmetry structures of games. Consider the following two trivial examples.

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

Let us call “player 1” the one choosing rows, and “player 2” the one choosing columns. It is impossible to distinguish between both of player 1’s strategies in Game 1 and one should, hence, call the two strategies symmetric, in fact, duplicates. Two strategies $s_i, s'_i \in S_i$ of player i are **duplicates** if $u(s_i|s_{-i}) = u(s'_i|s_{-i})$ for all $s_{-i} \in S_{-i}$. Most games do not have many duplicates, however.⁷

Consider now Game 2 (which is just M_2). Exactly as in Game 1, other than the arbitrary name tags “A” and “B”, and the arbitrary fact that the same tags have been

⁷In fact, in the reduced normal form derived from an extensive form game duplicates are omitted.

used for both players, there is nothing to distinguish the two strategies in this game. Likewise, there is nothing other than arbitrary names to distinguish the two player roles. When transcribing a strategic situation into a game, a consultant cannot rely on a universal convention, say, to “play A”. What she has written down as “A” might have been called “B” by a consultant advising her client’s opponent.

To reflect this observation, a weaker notion of symmetry is given below. Let a **relabeling** (or permutation) of player i ’s strategies be a bijective function $\rho_i : S_i \rightarrow S_i$. Given $\rho = (\rho_j)_{j \in I}$ and $s \in S$, denote $\rho_{-i}(s_{-i}) = (\rho_j(s_j))_{j \neq i}$.

Definition 1. Two strategies $s_i, s'_i \in S_i$ of player i are **weakly symmetric** if there exist relabelings $\rho_j : S_j \rightarrow S_j$ for $j \in I, j \neq i$ such that $u(s_i | s_{-i}) = u(s'_i | \rho_{-i}(s_{-i}))$ for all $s_{-i} \in S_{-i}$.

Duplicates are weakly symmetric, with the identity functions as relabelings. The converse is not true. Strategies A and B in Game 2 are not duplicates, but they are weakly symmetric for both players (consider the relabeling $\rho_{-i}(A) = B$ and $\rho_{-i}(B) = A$).

Given the simple examples above, defining symmetry of strategies seems a simple matter of identifying whether two strategies obtain the same payoff vector after some relabeling of opponents’ strategies. This is not so. Consider the following game.

	L	R
T	6,6	7,7
M	10,10	0,0
B	0,0	10,10

Game 3

A relabeling of player 2’s strategies L and R in Game 3 renders player 1’s strategies M and B with equivalent payoff vectors. In declaring strategies M and B weakly symmetric for player 1, we are using a relabeling of strategies L and R for player 2. Those strategies, however, are not weakly symmetric for player 2, due to the differing payoffs against T. Thus the consultant can tell apart strategies L and R for player 2, and, as a consequence, he can also tell apart M and B. This example points out that the concept of strategy symmetry needs to be slightly more involved than a mere equivalence of the payoff vectors associated with a players’ strategies.

One could be tempted to call two strategies symmetric if the required permutation of strategies of the opponents only relabels weakly symmetric strategies. This is not enough, as the three-player Game 4 below shows.

The transposition of L and C for player 2 in Game 4 allows us to declare M and B weakly symmetric for player 1. Further, the transposition of W and E for player 3 allows us to declare L and C weakly symmetric for player 2. Thus, we are tempted to consider M and B symmetric in a stronger sense.

		W				E		
		L	C	R		L	C	R
T	0,0,0	1,1,1	0,0,3		T	1,1,1	0,0,0	0,0,7
M	10,10,2	0,0,0	0,0,0		M	0,0,0	10,10,2	0,0,0
B	0,0,0	10,10,2	0,0,0		B	10,10,2	0,0,0	0,0,0

Game 4

Strategies W and E are clearly identified, however, due to their payoffs against (T,R) (they are not weakly symmetric). So is strategy T for player 1. It follows that the consultant can also tell apart strategies L and C for player 2 (through their payoffs against (T,W)). But once the labels L and C acquire meaning, the consultant can also readily tell apart strategies M and B for player 1. In other words, we should not declare strategies M and B for player 1 symmetric, because to establish the symmetry we need to swap strategies L and C for player 2, and to in turn establish the symmetry of those, we need to swap strategies W and E of player 3, although those are not weakly symmetric.

It becomes apparent that one could easily be led to an infinite regress problem here, where “ k -symmetric strategies” only allow to relabel “ $(k - 1)$ -symmetric” ones. The ultimate conclusion of such an exercise is that a definition of strategy symmetry should only allow for relabelings which swap strategies of other players which are also symmetric according to the very same definition. We now introduce the concept of pairwise symmetry, which takes care of this difficulty.

Definition 2. A **pairwise 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 relabelings 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 symmetric (relative to \mathcal{T}) if they belong to the same symmetry class.

This definition is a natural generalization to n -player games of the notion of strategy symmetry introduced by Crawford and Haller (1990, Appendix) for two-player games. Note that, to declare two strategies of a player symmetric, the condition $\rho_j(T_j) = T_j$ restricts relabelings to only exchange strategies of other players within the same symmetry class of those players (in Game 3, this prevents us from exchanging the strategies of player 2, and hence from declaring strategies M and B of player 1 symmetric). The definition of symmetry is thus self-referential. Existence of a symmetry structure of any game is, however, guaranteed by the observation that the partition which consists of all singleton sets is a pairwise symmetry structure. We will refer to this as the **trivial symmetry structure**.

Before we explore the structure of symmetry structures we want to make a simple but powerful observation. Fix an n -player game $\Gamma = (I, S, u)$ and let $U^\Gamma = \{u(s) | s \in S\}$ denote the set of all payoff vectors in Γ . Let $f : U^\Gamma \rightarrow \mathbb{R}^n$ be an arbitrary mapping. Let the game $\Gamma^f = (I, S, f(u))$ be such that it shares with Γ the same player set and same strategy sets but its payoffs are transformed by f . Then, by construction, every symmetry structure of Γ is also a symmetry structure of Γ^f . Thus, for instance, the rock-scissors-paper game can be transformed into a matching game, which inherits all symmetry structures from the rock-scissors-paper game, perhaps gaining some more. If the mapping f is injective then the sets of symmetry structures in Γ^f and Γ coincide.

Given our interpretation that the game in question is such that the names of strategies at hand have no a priori meaning 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 the Appendix gives a useful characterization of the join of two partitions.

The join (least upper bound) $\mathcal{T} \vee \mathcal{T}'$ of two pairwise symmetry structures \mathcal{T} and \mathcal{T}' can be defined as the finest pairwise 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**.

Theorem 1. *For every finite game Γ the set of pairwise symmetry structures endowed with the partial order of “coarser than” forms a lattice. In particular, the join of two pairwise 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. Necessarily, this structure captures as much symmetry as actually exists in the payoff matrix of the game.

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

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.

	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 5

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” $\tilde{\mathcal{T}}$ given by $\tilde{\mathcal{T}}_1 = \mathcal{T}'_1 = \mathcal{T}''_1 = \mathcal{T}_1^* = \{\{A, B\}, \{C, D\}\}$ and $\tilde{\mathcal{T}}_2 = \{\{E\}, \{F\}, \{G\}, \{H\}\}$. 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 are duplicates. Since none of player 1’s strategies are duplicates, this is not a pairwise symmetry structure. In this example, the meet $\mathcal{T}' \wedge \mathcal{T}''$ is the trivial symmetry structure.

3.2 Symmetry Structures and Familiar Frames

The coarsest 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. In this subsection all frames are assumed familiar to all players, as addressed in Section 2.

Suppose the consultant analyzes the game in two steps. First, he extracts as much information as he can from the payoff structure alone. Thus he will arrive at the symmetry structure \mathcal{T}^* . Second, he considers the 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 frame induces a refining of \mathcal{T}^* by further partitioning the symmetry classes. Given a frame 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 counselor is in general left with some work to do to integrate the new information into a new symmetry structure.

Definition 3. Let L be a frame for game Γ . The symmetry structure induced by L , $\mathcal{T}(L)$, is the coarsest symmetry structure \mathcal{T} such that, for each player i , $\mathcal{T}_i(L)$ is finer than the L_i -partition of S_i .

Note that $\mathcal{T}(L)$ is always well defined by Theorem 1. 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), and the join of any two of its elements is also in the set. Thus 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 $\mathcal{Z}_2 = \{\blacksquare, \square\}$.

		■	■	□
		D	E	F
•	A	1,2	0,0	0,0
○	B	0,0	1,2	0,0
○	C	0,0	0,0	1,2

Game 6

In the coarsest (frame-free) symmetry structure, all strategies are symmetric, for both players. If we further 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 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.

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

Theorem 2. *For any pairwise symmetry structure there exists a 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 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 frames, then it becomes bijective. Thus we could identify frames (up to similarity) with symmetry structures, and the study of symmetry structures and the study of frames become one and the same subject.

The equivalence respects the lattice structure in the natural way. As an illustration, consider a situation where, as in Casajus (2000), Janssen (2001) or Binmore and Samuelson (2006), 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.

3.3 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 framed game (Γ, L) , i.e. for any given (pairwise) 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. From a philosophical point of view, the rationality requirement could be seen as an extension of Kant's Categorical Imperative, which, translated to the consultant's terms, reads as follows: advise your clients as if their opponents also follow your advice (or that of an equivalent consultant).

Axiom 2. A recommendation $x \in \Theta$ satisfies the axiom of **equal treatment of symmetric strategies** for 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, 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.

Definition 4. A recommendation $x \in \Theta$ is a **rational strategy-symmetric recommendation** with respect to a (pairwise) symmetry structure \mathcal{T} if it satisfies the axioms of rationality and equal treatment of symmetric strategies.

The next theorem states that a recommendation satisfying both of the axioms above always exists. We remark that, contrary to Bacharach (1993), we do not invoke the Principle of Coordination to require the players themselves to select a Nash equilibrium. Rather, the interpretation we hinge on is as follows. If consultants adhere to the axiom of equal treatment of symmetric strategies, dictated by the Principle of Insufficient Reason, then, by the next theorem, they will always find it possible to recommend a Nash equilibrium respecting this axiom.

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

The proof (see Appendix) 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 would fail if we used weak symmetry rather than symmetry. Consider Game 3 again. Strategies M and B are weakly symmetric, thus equal treatment of symmetric strategies would require that $x_1(M) = x_1(B)$. If $x_1(T) > 0$, the unique best response of player 2 is R, against which player 1 would play B, a contradiction. If $x_1(T) = 0$, it follows that $x_1 = (0, \frac{1}{2}, \frac{1}{2})$, thus player 1 must be indifferent between M and B. This implies $x_2 = (\frac{1}{2}, \frac{1}{2})$, against which player 1 strictly prefers T. Hence there exists no Nash equilibrium of this game where M and B are treated equally.⁸

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

Proposition 1. *Let \mathcal{T} and \mathcal{T}' be symmetry structures of a finite, normal-form game Γ . If \mathcal{T} is coarser than \mathcal{T}' , then any rational recommendation $x \in \Theta$ 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.

⁸The only Nash equilibria of this game are (M,L), (B,R), and $[(\frac{10}{11}, \frac{1}{11}, 0), (\frac{7}{11}, \frac{4}{11})]$.

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 (1990), 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 (1990) then rely on an additional principle, Pareto efficiency, in order to select an equilibrium. The argument above shows that such a “meta-norm” is necessary for any full definition of focal points, because refining symmetry structures results on enlarged sets of equilibria. In fact, Goyal and Janssen (1996) point out that the results of Crawford and Haller (1990) make implicit use of a second meta-norm, in addition to Pareto-efficiency. The simplest game to explain this problem is a repeated 2-player matching game where, by chance, both players choose the same strategy in the first round. Crawford and Haller (1990) argue that both players could have the repeated game strategy to stay with the chosen strategy after such an occurrence, thus achieving coordination from then on. As Goyal and Janssen (1996) argue (see also Crawford and Haller (1990, Section 5)), they could as well coordinate on the other strategy from then on. The fact that both are possible should make it difficult for players to actually achieve coordination. Thus coordination requires an appeal to a more subtle meta-norm, for instance, that already chosen strategies are more salient. We discuss these issues in more detail in Section 6.

4 Global Symmetry Structures

In this section, we tackle the two difficulties advanced in Section 3 and provide a more comprehensive, but also more involved concept of symmetry, called **global symmetry**, which is not based on pairwise comparisons of strategies and at the same time provides a natural formalization of the idea of symmetric players. We will rely on the concept of **symmetry** (or automorphism) of a game, introduced by Nash (1951) and later generalized by Harsanyi and Selten (1988).

Again, we will consider all possible symmetry structures based on this concept in order to provide a link with the concept of frame. For the particular case of the coarsest possible symmetry structure, the concept we will deal with boils down to the concept of symmetric strategies implicit in Nash (1951) (see also Harsanyi and Selten (1988) and Casajus (2000, 2001)). As we will see, this concept immediately implies strategy symmetry under Crawford and Haller’s (1990) definition and the corresponding generalization (**pairwise symmetry**) given in Section 3. To the best of our knowledge, it is

still an open question whether the converse is true. As further motivation, we will settle this question below, providing a counterexample which shows that (even in the absence of player symmetries) pairwise symmetry does not imply global symmetry in general.

4.1 Preliminary Concepts and Examples

Below we introduce a further axiom requiring symmetric players to be treated equally in the consultant's recommendation. This requires us to provide a definition of symmetric players, and, as a consequence, symmetric games (games where all players are symmetric). Although a formal definition is rarely explicitly given in the literature, the standard (textbook) concept of a symmetric game can be captured by the following definition.

Definition 5. The game $\Gamma = [I, (S_i, u_i)_{i \in I}]$ is **strongly symmetric** if all players have the same strategy set, $S_i = S_j = S^*$ for all $i, j \in I$, and there exists a function $u^* : S^* \times (S^*)^{|I|-1} \rightarrow \mathbb{R}$ such that $u_i(s_i | s_{-i}) = u^*(s_i | s_{-i}) = u^*(s_i | s'_{-i})$ for all players $i \in I$, for all $s_i \in S^*$, and all $s_{-i}, s'_{-i} \in (S^*)^{|I|-1}$ which differ only in a permutation of the strategies among players.

Matching games as Game 2 are strongly symmetric, thus the consultant should treat both players symmetrically. For general games, however, establishing symmetry among players might be less straightforward than one might think at first glance.

Consider Game 7 below, which is a version of the Battle of the Sexes. It is not strongly symmetric, and indeed it will usually not be considered symmetric at first glance. It is certainly true that no strategies of the same player are symmetric. Strategy A is the one such that, if played by both players, will lead to the pure-strategy equilibrium most preferred by player 1. Likewise, strategy B is the one such that, if played by both players, will lead to the pure-strategy equilibrium most preferred by player 2. But, who is player 1, and who is player 2? If the consultant relies on a prescription of the form “aim for your most preferred equilibrium”, and her client's opponent's consultant does the same, the game will result in mis-coordination, hardly a desirable outcome. The mixed-strategy equilibrium is asymmetric, given by $x_1 = (\frac{3}{7}, \frac{4}{7})$ and $x_2 = (\frac{4}{7}, \frac{3}{7})$. Since the consultant cannot reliably ascribe her clients one of the names “player 1” or “player 2”, it might appear that, under the assumption that the clients' opponents will receive similar recommendations, all equilibria are unattainable. This is not the case, for the consultant might recommend her client to play the strategy leading to his most preferred equilibrium with probability 3/7, and the other one with the remaining probability. If the client's opponent receives exactly the same recommendation, play will lead to the mixed-strategy Nash equilibrium. We shall thus require a recommendation to satisfy the axiom of **equal treatment of symmetric players**, which will build upon an appropriate definition of (player) symmetry.

	A_2	B_2		H_2	T_2
A_1	4,3	0,0	H_1	1,-1	-1,1
B_1	0,0	3,4	T_1	-1,1	1,-1

Game 7
Game 8

Game 8 poses a harder problem. This game (which is, of course, Matching Pennies) is not strongly symmetric either. Indeed, while player 1 wants to coordinate choices, player 2 wants to uncoordinate them. But “coordinate” is a term which depends on an arbitrary labeling of strategies, and is hence meaningless. Suppose we swap player roles as column and row players, and reorder the strategies of player 1. Player 2 would consider himself a row player in the following game

	T_1	H_1
H_2	1,-1	-1,1
T_2	-1,1	1,-1

which is again Matching Pennies, with player 2 now wanting to “coordinate” on the diagonal. That is, both Matching Pennies as well as Game 7 are symmetric in the sense that, through reordering of the strategy sets, each player faces the same payoff table. We thus require that a recommendation for Matching Pennies must be the same for both players (after relabeling one player’s strategies) as well as the same for both strategies. I.e. the recommendation has to satisfy both axioms of equal treatment of symmetric players as well as of equal treatment of symmetric strategies. This leaves only a single feasible recommendation $(\frac{1}{2}, \frac{1}{2})$, which is a Nash equilibrium. Theorem 6 below shows that there always exists a Nash equilibrium satisfying both equal-treatment axioms.

The last two examples show that Definition 5 is too restrictive. Consider the following (provisional) definition, which depends on the given (pairwise) symmetry structure.

Definition 6. Let \mathcal{T} be a pairwise symmetry structure of Γ . Two players $i, j \in I$ are (pairwise) **symmetric** relative to \mathcal{T} if there exists a permutation of the players’ names $\sigma : I \rightarrow I$ with $\sigma(i) = j$ and there exists bijections $\tau_k : S_i \rightarrow S_{\sigma(k)}$ for each $k \in I$ such that for every $T_k \in \mathcal{T}_k$ there is a $T_{\sigma(k)} \in \mathcal{T}_{\sigma(k)}$ such that $\tau_k(T_k) = T_{\sigma(k)}$, and, for all $k \in I$ and all $s = (s_k | s_{-k}) \in S$,⁹ $u_k(s_k | s_{-k}) = u_{\sigma(k)}(\tau_i(s_k) | \tau_{-k}(s_{-k}))$.

Say that a game is symmetric if all players are symmetric according to the definition above. Coming back to our examples it is easy to see that the two players in Game 1 are not symmetric, as one player always obtains a payoff of 0, while the other does obtain a payoff of 1 in some cases. In the Battle of the Sexes (Game 7) the unique symmetry structure is such that no strategies are considered symmetric. The two players are

⁹We avoid cumbersome notation by using the convention $\tau_{-k}(s_{-k}) \in S_{-k}$. In vector notation, this involves the appropriate permutation of coordinates, i.e. $\tau_{-k}(s_{-k}) = (\tau_{\sigma^{-1}(k)}(s_{\sigma^{-1}(k)}))_{k \neq \sigma(i)}$.

not symmetric according to Definition 5, but are symmetric according to Definition 6. The bijections that allow us to reach this conclusion are given by $\tau_1(A_1) = B_2$ and $\tau_1(B_1) = A_2$, and $\tau_2 = \tau_1^{-1}$. We then have $u_1(A_1|A_2) = u_2(\tau_1(A_1)|\tau_2(A_2))$, and analogous equalities for all other payoffs.

Consider the Matching Pennies game (Game 8). The coarsest symmetry structure is such that $\mathcal{T}_1 = \{\{H_1, T_1\}\}$ and $\mathcal{T}_2 = \{\{H_2, T_2\}\}$. I.e. both strategies are symmetric for both players. If that is the case, however, then both players are symmetric. The bijections $\tau_1 : S_1 \rightarrow S_2$ given by $\tau_1(H_1) = T_2$ and $\tau_1(T_1) = H_2$ and $\tau_2 : S_2 \rightarrow S_1$ given by $\tau_2(H_2) = H_1$ and $\tau_2(T_2) = T_1$ satisfy $u_1(s_1|s_2) = u_2(\tau_1(s_1)|\tau_2(s_2))$ for all $(s_1, s_2) \in S$. Notice that $\tau_1(\tau_2(H_2)) = T_2$ and, hence, there is an implicit relabeling of strategies for player 2 within symmetry classes. Thus, given this symmetry structure the two players are indeed symmetric. The only other symmetry structure is such that no strategies are considered symmetric. According to the definition, players would not be symmetric relative to this symmetry structure.

Indeed, in general the definition of symmetric players depends on the (strategy) symmetry structure of the game. Consider the following trivial game.

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

Game 9

The two players are symmetric in the coarsest symmetry structure. If we take the symmetry structure \mathcal{T} with $\mathcal{T}_1 = \{\{U, D\}\}$ and $\mathcal{T}_2 = \{\{L\}, \{R\}\}$, however we can not find a bijection $\tau : S_1 \rightarrow S_2$ such that the symmetry classes survive the mapping. I.e. $\tau(\{U, D\}) = \{L, R\}$, which is neither $\{L\}$ nor $\{R\}$.

Equal treatment of symmetric players, to be stated below, simply requires a consultant's recommendation to treat symmetric players equally. This is well-defined for the examples above, e.g. in Game 7 (Battle of the Sexes) we are led to the unique, apparently asymmetric Nash equilibrium in mixed strategies $x_1 = (\frac{3}{7}, \frac{4}{7})$ and $x_2 = (\frac{4}{7}, \frac{3}{7})$.

There is, however, a problem with this definition. Declaring two players i, j symmetric relies on a particular mapping $\tau_i : S_i \mapsto S_j$. In general, however, alternative mappings could be used. As long as the image of a given symmetry class of player i is always the same symmetry class under all alternative mappings τ_i , this is unproblematic, for equal treatment of symmetric strategies already prescribes that all strategies are treated equally. Certain subtleties arise if two players can be seen as symmetric in two qualitatively different ways. The (framed) Game 10 below illustrates this point. The given frame induces the symmetry structure $\mathcal{T}_1 = \{\{A, B\}, \{C, D\}\}$, $\mathcal{T}_2 = \{\{E, F\}, \{G, H\}\}$. By equal treatment of symmetric strategies, recommendations must be of the form

$x_1 = (p, p, \frac{1}{2} - p, \frac{1}{2} - p)$, $x_2 = (q, q, \frac{1}{2} - q, \frac{1}{2} - q)$. There is a priori no further requirement, e.g. p might adopt any value in $[0, 1/2]$. This allows for instance for the Nash equilibrium where only A, B and E, F are played.

		■	■	□	□
		E	F	G	H
•	A	1,1	0,0	1,1	0,0
•	B	0,0	1,1	0,0	1,1
○	C	1,1	0,0	1,1	0,0
○	D	0,0	1,1	0,0	1,1

Game 10

Players are symmetric with σ being the transposition of their names and $\tau_1 = \tau_2^{-1}$ given by $\tau_1(A) = E, \tau_1(B) = F, \tau_1(C) = G$, and $\tau_1(D) = H$. A requirement of equal treatment of symmetric players then leads to the conclusion that, in a recommendation, $p = q$ must hold, which eliminates e.g. the equilibrium where player 1 randomizes between strategies A and B and player 2 randomizes between G and H .

Our specification of τ has been, however, rather arbitrary. In this game, there is a qualitatively different alternative. Players are also symmetric according to the bijections $\tau'_1 = \tau_2'^{-1}$ given by $\tau'_1(A) = G, \tau'_1(B) = H, \tau'_1(C) = E$, and $\tau'_1(D) = F$. Using this transformation, equal treatment of symmetric players leads to the constraint $p = \frac{1}{2} - q$. If both possibilities are taken into account, we are left with only one candidate profile where $p = q = \frac{1}{4}$. This illustrates, however, that in some cases the concept of player symmetry might allow additional strategy symmetries, not contemplated by the original symmetry structure, to “sneak in through the back door”.

4.2 Simultaneous Determination of Symmetric Strategies and Players

There is, however, a more serious problem with the approach followed until now. Pair-wise symmetry structures ignore symmetries between players. Symmetric players, as just defined, extract the symmetry of players from the symmetry structure. Game 11 below shows that this approach misses some possible interactions between both concepts.

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

Game 11

In this game, both strategies are symmetric for players 1 and 2. This in turn would allow us to declare both players symmetric, provided the two strategies of player 3 are

also declared symmetric. But, for player 3, the two strategies would not be (pairwise) symmetric according to our first definition and the approach in Crawford and Haller (1990). This is because declaring them symmetric requires swapping players 1 and 2. Since these two players are symmetric, however, there should be no objection to this.

We conclude that the concepts of symmetric strategies and symmetric players must be determined **simultaneously**. The definitions below will accomplish this objective.

Definition 7. A **symmetry** of a normal form game Γ is a tuple (σ, τ) where $\sigma : I \rightarrow I$ is a permutation of the players' names and $\tau = (\tau_i)_{i \in I}$, where, for each $k \in I$, $\tau_k : S_k \rightarrow S_{\sigma(k)}$ is a bijection, fulfilling that,¹⁰ for all $k \in I$ and all $s = (s_k | s_{-k}) \in S$,

$$(1) \quad u_k(s_k | s_{-k}) = u_{\sigma(k)}(\tau_k(s_k) | \tau_{-k}(s_{-k})).$$

The concept of symmetry of a game was introduced by Nash (1951). Harsanyi and Selten (1988) reformulated it along the lines above, and further generalized it by allowing for positive affine transformations of the payoffs. Since Harsanyi and Selten (1988) were only concerned with the best response structure of the game, this made sense for their framework. Such transformations, however, are not appropriate for our objective, because one would then declare players symmetric who can be easily told apart on the basis of payoffs alone (consider, for example, the unframed Game 6).

As commented above (recall Game 10), one needs to specify which players are symmetric and also **how** they are symmetric. The following definition will be essential to accomplish this objective.

Definition 8. 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 next definition summarizes how a symmetry should agree with a candidate symmetry structure and an identification thereof.

Definition 9. 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 $\alpha(T_k) = \alpha(T_{\sigma(k)})$.

¹⁰Recall footnote 9.

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 10. 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.

Recall that strategies M, B were not declared symmetric in Game 3 because one is not allowed to swap the asymmetric strategies L and R. Under Definition 10, this is already prevented by the fact that there exists no symmetry (σ, τ) which swaps M and B . This is because condition (1) must apply to all strategies swapped in the symmetry, and not only (as in Definition 2) to the two strategies one wishes to declare symmetric.

Let us briefly comment on the necessity of player identifications. We say that the identification α is **compatible** with $(\mathcal{I}, \mathcal{T})$ if $(\mathcal{I}, \mathcal{T}, \alpha)$ is a global symmetry structure. Whenever two players are symmetric, there is a natural correspondence between (the symmetry classes of) their strategies. However, this correspondence might not be unique. That is, in general the partition part of a global symmetry structure might have several different compatible player identifications.

Consider again the framed Game 10. A global symmetry structure is given by $\mathcal{I} = \{\{1, 2\}\}$ and $\mathcal{T}_1 = \{\{A, B\}, \{C, D\}\}$, $\mathcal{T}_2 = \{\{E, F\}, \{G, H\}\}$. Ignore for a moment the player identification requirements. Players are symmetric with $\sigma(1) = 2$ and $\tau_1 = \tau_2^{-1}$ given by $\tau_1(A) = E, \tau_1(B) = F, \tau_1(C) = G$, and $\tau_1(D) = H$. They are also symmetric with bijections $\tau'_1 = \tau_2'^{-1}$ given by $\tau'_1(A) = G, \tau'_1(B) = H, \tau'_1(C) = E$, and $\tau'_1(D) = F$. That is, we have two different symmetries (σ, τ) and (σ, τ') which agree with the global symmetry structure. The composition of these two symmetries is another symmetry $(1_I, \tau^*)$ where each player is mapped to himself and $\tau_1^*(A) = C$. However, in the symmetry structure A and C are **not** symmetric, because they have different labels in the frame. The problem is that we can identify the player labels in two different

ways, which leads to two alternative player identifications. Intuitively, in the definition above, one cannot allow for all symmetries which agree with either identification, because then the composition of those symmetries (required by the expected transitivity of the symmetry relation) would produce symmetry classes larger than specified.

The existence of multiple player identifications is typical of games with duplicates but not restricted to it. Indeed, if we consider Game 2 (M_2) with a frame where each strategy of each player has a different label, exactly the same problem appears. Player 1 and 2 can be declared symmetric by either mapping A to A and B to B or by mapping A to B and B to A . The composition of these two symmetries, however, does not agree with the specified symmetry structure, because it would force us to declare both strategies of each player symmetric.

This problem is moot for the coarsest global symmetry structure, assuming one exists (as we indeed prove below), and thus is of no relevance for Nash (1951). However, if we are interested in frames, it cannot simply be assumed away, and thus the player identification is a necessary part of the definition. Still, in many cases there exists a natural, unique player identification of a global symmetry structure, up to the specification of the sets Ω_i , as in the case where no players are symmetric. Hence, in examples where the player identification is unique, we will simply omit any reference to it.

Finally, the observation about payoff transformations that we made for pairwise symmetry structures can be repeated here. Transformations of a game's payoff vectors preserve all global symmetry structures, perhaps adding some, provided they map all permutations of a given payoff vector to the appropriate permutations of the mapped payoff vector. I.e. if we map payoff vector $(-1, 0, 1)$ to $(3, 2, 5)$ we have to map $(0, 1, -1)$ to $(2, 5, 3)$ to preserve player symmetry. This is always the case if we take a single mapping transforming all individual payoffs, rather than one transforming payoff profiles. Again if this mapping is injective the sets of all symmetry structures of the transformed game and the original game coincide.

4.3 Pairwise Symmetry Does Not Imply Global Symmetry

Suppose we force players to be considered asymmetric, i.e. we consider global symmetry structures $(\mathcal{I}^0, \mathcal{T})$ where $\mathcal{I}^0 = \{\{i\} \mid i \in I\}$. Then, for any symmetry (σ, τ) which agrees with $(\mathcal{I}, \mathcal{T})$, σ becomes the identity and all mappings τ_i become relabelings of the corresponding S_i . The identification of players becomes irrelevant and can be taken to be $\alpha^0(T_i) = T_i$ for each $i \in I$ and $T_i \in \mathcal{T}_i$. Say that such a symmetry structure is **without player symmetry**. Definition 9 then reduces to the condition $\tau_i(T_i) = T_i$ for every $T_i \in \mathcal{T}_i$, thus implying pairwise symmetry of all symmetric strategies.

Proposition 2. *If $(\mathcal{I}^0, \mathcal{T}, \alpha^0)$ is a global symmetry structure without player symmetry, then \mathcal{T} is a pairwise symmetry structure.*

Game 11 above shows that there are global symmetry structures $(\mathcal{I}, \mathcal{T}, \alpha)$ such that \mathcal{T} , taken by itself, is not a pairwise symmetry structure. A deeper question is posed by the converse of the result, that is, whether every pairwise symmetry structure is a global symmetry structure without player symmetry. 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 (1990)). Can they always be declared globally symmetric (i.e. symmetric in the sense implicit in Nash (1951))? This question was already posed (as an open question) by e.g. Casajus (2000, p.20). The following game shows that the answer is negative.

	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 12

Consider pairwise symmetry. The differences in payoffs (each coordination game block of three strategies for each player has distinct payoffs compared with the other three blocks) ensures that we can safely work with two blocks of three strategies per player. The coarsest pairwise symmetry structure has $\tilde{\mathcal{T}}_1 = \{\{a_1, a_2, a_3\}, \{a_4, a_5, a_6\}\}$ and $\tilde{\mathcal{T}}_2 = \{\{b_1, b_2, b_3\}, \{b_4, b_5, b_6\}\}$. 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 $b_1 \rightarrow b_2$ and $b_4 \rightarrow b_5$. 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, $\tilde{\mathcal{T}}$ is a pairwise symmetry structure.

Even ignoring player symmetry, this partition is **not** a global 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 $b_1 \rightarrow b_2$ and $b_4 \rightarrow b_5$. But permuting $b_1 \rightarrow b_2$ implies that the strategies of player 1 must be permuted in such a way that $a_6 \rightarrow a_5$. On the other hand, permuting $b_4 \rightarrow b_5$ implies that the strategies of player 1 must be permuted in such a way that $a_4 \rightarrow a_5$, 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 $\hat{\mathcal{T}}_1 = \{\{a_1, a_3\}, \{a_2\}, \{a_4, a_6\}, \{a_5\}\}$ and $\hat{\mathcal{T}}_2 = \{\{b_1, b_3\}, \{b_2\}, \{b_4, b_6\}, \{b_5\}\}$. 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 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.

Even though settling the question of the (non-)equivalence between pairwise symmetry concepts as in Crawford and Haller (1990) and global ones as in Nash (1951) or Harsanyi and Selten (1988) is interesting in itself, this is more than a technical point. Game 12 also illustrates that the recommendations might differ qualitatively under both approaches. Suppose that we provisionally adopt the convention that the focal points of an unframed game are the Pareto-efficient equilibria which fulfill equal treatment of symmetric strategies with respect to the coarsest symmetry structure. Adopting the pairwise approach, the coarsest symmetry structure is given by $\tilde{\mathcal{T}}$ 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, \dots, 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 prediction resulting in an expected payoff of 1. On the other hand, the global approach delivers the coarsest symmetry structure $\hat{\mathcal{T}}$, which enables coordination in the equilibrium (a_5, b_2) , with a payoff of 4. This is then the unique prediction, which is qualitatively different from the one arrived at under the pairwise approach.

4.4 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 \mathcal{I}^0 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 the Appendix). Essentially, the set of all symmetries of a game, denoted $\text{Sym}(\Gamma)$, forms a group with the composition of symmetries defined in the natural way. It can be shown that each global symmetry structure corresponds to exactly one subgroup of this group. The group of symmetries associated to the coarsest global symmetry structure is the grand group $\text{Sym}(\Gamma)$. Global symmetry structures which coincide except for the player

identification correspond to different subgroups. The set of subgroups of a group has a lattice structure. The theorem below then follows from showing that this structure can be translated to global symmetry structures (there is a lattice-isomorphism).

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

Contrary to Theorem 1, it is not true that the 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.

Reconsider Game 10. A possible global symmetry structure is given by $\mathcal{I}' = \{\{1, 2\}\}$ (both players are symmetric), $\mathcal{T}'_1 = \{\{A, B\}, \{C, D\}\}$, $\mathcal{T}'_2 = \{\{E, F\}, \{G, H\}\}$, and e.g. a player identification with $\alpha'_1(\{A, B\}) = \alpha'_2(\{E, F\})$ and $\alpha'_1(\{C, D\}) = \alpha'_2(\{G, H\})$. A different global symmetry structure is given by $\mathcal{I}'' = \{\{1\}, \{2\}\}$ (players are not symmetric; hence the player identification is irrelevant), $\mathcal{T}''_1 = \{\{A, B, C, D\}\}$, and $\mathcal{T}''_2 = \mathcal{T}'_2 = \{\{E, F\}, \{G, H\}\}$. If we construct the joins of the respective partitions, we obtain $\mathcal{I}' \vee \mathcal{I}'' = \{\{1, 2\}\}$, $\mathcal{T}'_1 \vee \mathcal{T}''_1 = \{\{A, B, C, D\}\}$, and $\mathcal{T}'_2 \vee \mathcal{T}''_2 = \{\{E, F\}, \{G, H\}\}$. But player 1 and 2 cannot be symmetric and have a different number of symmetry classes, hence this structure cannot be part of a global symmetry structure.¹¹

Suppose a pair $(\mathcal{I}, \mathcal{T})$ admits several compatible player identifications. The join of all the resulting global symmetry structures $(\mathcal{I}, \mathcal{T}, \alpha^1), \dots, (\mathcal{I}, \mathcal{T}, \alpha^k)$ is a well-defined symmetry structure, due the lattice structure. However, this join, say $(\mathcal{I}^*, \mathcal{T}^*, \alpha^*)$, might incorporate new symmetries not captured in $(\mathcal{I}, \mathcal{T})$.

Definition 11. 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 10, 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 four 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.¹²

There is an interesting connection with Blume (2000), who uses a group-theoretic approach to formalize partial languages, which are defined as subgroups of permutations

¹¹A similar example can be constructed with a trivial game as e.g. Game 9.

¹²Thus the problems leading to the necessity of player identifications are not encountered if, as e.g. Nash (1951) or Harsanyi and Selten (1988), one restricts to structures considering all possible symmetries.

of abstract labels, which could e.g. correspond to the labels used in a frame. Blume (2000) concentrates on the issue of how partial languages facilitate coordination for repeated matching games (following Crawford and Haller (1990)), and how fast learning occurs. For matching games, one could consider the case where each strategy receives a different label, but the same label for all players. For this particular case, the subgroup of a particular partial language corresponds to one of our global symmetry structures.

4.5 Global Symmetry Structures and Extended Frames

Analogously to Section 3.3.2, there is a correspondence between global symmetry structures and extended frames, where not only strategies, but also players are labeled (e.g. players 1 and 5 are men, the rest are women). 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 12. An **extended frame** is a pair (L, L_0) where (i) L is a frame, i.e. $L = (L_i)_{i \in I}$ with $L_i : S_i \rightarrow \mathcal{Z}_i$ for some arbitrary sets \mathcal{Z}_i , (ii) $L_0 : I \rightarrow \mathcal{Z}_0$ 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 $z \in \mathcal{Z}_i = \mathcal{Z}_j$.

Extended frames provide not only a labeling of strategies, but also a labeling of players. 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.

Definition 13. Let (L, L_0) be an extended 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. As in Section 3.3.2, the mapping which takes each extended frame to the global symmetry structure it induces is onto, that is, for every global symmetry structure there exists an extended frame which rationalizes it.

Theorem 5. *For any global symmetry structure there exists an extended 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 extended frame as follows. Let $\mathcal{Z}_0 = \mathcal{T}_0$. 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 extended frame. \square

Again this result allows us to assign a (global) symmetry structure to each (extended) frame. Through the identification of extended frames which lead to the same symmetry structure, we can interpret this theorem as saying that global symmetry structures are actually the same objects as extended frames.

4.6 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, in the absence of exogenously given labels for players and strategies, the consultant cannot distinguish between two player roles, then he must give equivalent recommendations to those two player roles. The idea behind this requirement could be (in admittedly a somewhat strained way) likened to Rawls' Veil of Ignorance. It essentially requires that, in making a recommendation, the consultant treats it as a recommendation for every possible player role his or her client might end up in.

For Game 10 above, as already commented, the frame is compatible with declaring players symmetric using two different player identifications. That would lead to different requirements which of course depend on the player identification. If one wants to insist on all possible symmetries being incorporated, one needs to move "up" to the completion of the considered structure. Then one is left with the equilibria of the coarser structure where symmetry classes are merged together.

Definition 14. 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 exist.

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 recommendation for that structure are identical to Nash’s (1951) “symmetric equilibrium points”. Hence the result follows from Nash (1951, Theorem 2). For the sake of completeness, we include a short proof in the Appendix which builds on the proof of Theorem 3. The only (minor) added difficulty is to note that, given a profile which respects equal treatment of symmetric strategies and players, then symmetric players have symmetric best responses.

5 Unfamiliar Frames

In this section we provide a discussion of the symmetry structures which arise when a game is accompanied by an unfamiliar frame. We do not provide a fully fledged theory, but restrict ourselves to a series of observations, using the simplest examples, 2-player matching games, M_k , in which they can be made. For any game M_k let strategies be denoted by the first k letters of the Roman alphabet. Since we are dealing with matching games, we also assume that both players are declared symmetric throughout. Thus we can consider frames which label strategies of one player only (the other player automatically receives the same labels). Abusing notation, in this section we write \mathcal{T} for the partition of any one of the two players in the considered symmetry structure.

Consider game M_4 , with two strategies labeled \blacksquare and two labeled \square . If individuals recurrently play games with such labels, a convention could develop which translates to a ranking of salience of these two labels. In this case we would call the frame familiar and the game plus frame falls into the domain of the previous sections. However, it is conceivable that these labels have never before been encountered in any meaningful sense, i.e. no such salience-ranking has developed. We thus have an unfamiliar frame. In this case there is no way, in this game, in which players could coordinate on one label versus another and the induced symmetry structure has to be the coarsest one. That is, the symmetry structure induced by an unfamiliar frame is quite different (always weakly coarser) than the symmetry structure induced by the same, but familiar, frame.

In general, although players are all unfamiliar with the given frame of a given game, the frame might yet allow players to refine the symmetry structure of the game. Consider again game M_4 but now with three strategies labeled \blacksquare and one labeled \square . Then the symmetry structure induced by this frame is the same whether it is familiar or not. The two players are symmetric and the three strategies with same label \blacksquare are symmetric, while the remaining strategy with label \square is identified.

For a given matching game $\Gamma = M_k$, for some k , let $\mathcal{U} = \mathcal{U}^\Gamma$ denote the operator that maps a symmetry structure of that game into another symmetry structure of that game, with the interpretation that the former is the symmetry structure induced by some familiar frame and the latter is the symmetry structure induced by the same but now unfamiliar frame. Recall that for familiar frames there is essentially a one-to-one mapping between frames and symmetry structure. Although for arbitrary games frames involving different numbers of attributes might give rise to the same symmetry representation, for matching games this cannot happen and we can safely define this operator on the set of symmetry structures. The coarsest symmetry structure for 2-player matching games is given by $\mathcal{T}^* = \{S\}$, where $S = S_1 = S_2$, i.e. each player has a single symmetry class including the whole strategy space. Then for a given symmetry structure \mathcal{T} , $\mathcal{U}(\mathcal{T})$ is simply given by two conditions. First, if $T, T' \in \mathcal{T}$ such that $|T| = |T'|$ then there is a $\tilde{T} \in \mathcal{U}(\mathcal{T})$ with $T \cup T' \subset \tilde{T}$. Second $\mathcal{U}(\mathcal{T})$ must be the finest symmetry structure such that the first condition holds. Such a symmetry structure always exists. In other words, $\mathcal{U}(\mathcal{T})$ is constructed by pooling all attributes together which appear the same number of times as strategy labels. The interpretation is that, in the absence of a convention, such attributes cannot be told apart.

Consider Game 2, i.e. game M_2 and all possible frames, L , for this game. There are only two possibilities: either both strategies receive the same label or they receive different labels. If both receive the same label clearly there is no way of differentiating the two. This is true regardless whether the frame is familiar or unfamiliar. Hence, for these frames L we have $\mathcal{T}(L) = \mathcal{T}^*$ as well as $\mathcal{U}(\mathcal{T}(L)) = \mathcal{T}^*$. In the second case, in which the two strategies receive different labels, if the frame is familiar we can distinguish the two strategies and, hence, obtain $\mathcal{T}(L) = \{\{A\}, \{B\}\}$. If the frame is unfamiliar, however, labels itself have no intrinsic meaning. Thus the two strategies, although distinct, are distinct in a non-useful way. Hence, the induced symmetry structure of this frame is again the coarsest one, i.e. we have $\mathcal{U}(\mathcal{T}(L)) = \mathcal{T}^*$. This example demonstrates that not every symmetry structure can be induced by an unfamiliar frame. I.e. there are symmetry structures \mathcal{T}' such that there is no symmetry structure \mathcal{T} with $\mathcal{U}(\mathcal{T}) = \mathcal{T}'$.

For game M_2 we thus find that whatever the frame, if it is unfamiliar, it does not refine the coarsest symmetry structure. This is not always the case. Consider game M_3 and all possible frames. For all unfamiliar frames L in which either every strategy receives the same label or every strategy receives a different label, we again obtain the coarsest symmetry structure $\mathcal{U}(\mathcal{T}(L)) = \mathcal{T}^*$. In all remaining frames, two strategies receive the same label, while the third receives a different label. This label is identifiable through the fact that it is the “odd-man out” (as e.g. Binmore and Samuelson (2006) call it), i.e. through the fact that it is the only label which is attached to only one strategy. Hence, even if this frame is unfamiliar, it induces a non-coarsest symmetry structure,

given by $\mathcal{U}(\mathcal{T}(L)) = \mathcal{T}(L) = \mathcal{T}^x = \{\{x\}, \{x\}^c\}$, where x is the strategy with the unique label. Thus, even unfamiliar frames can refine the coarsest symmetry structure.

The concept of familiarity is external to the considered symmetry structure. Consider game M_4 , with the unfamiliar frame which labels A,B,C, and D as $\blacksquare, \blacksquare, \square$, and \triangle , respectively. This induces the symmetry structure $\mathcal{T} = \{\{A, B\}, \{C\}, \{D\}\}$. In the absence of a convention on the meaning of \square and \triangle , there is no way to distinguish C and D, and the induced symmetry structure is $\mathcal{U}(\mathcal{T}) = \{\{A, B\}, \{C, D\}\}$. One could note that the latter symmetry structure corresponds to a hypothetical unfamiliar frame where C and D receive the same label, and further “familiarize” it to obtain $\mathcal{U}(\mathcal{U}(\mathcal{T})) = \mathcal{T}^*$. This, however, is conceptually wrong. In $\mathcal{U}(\mathcal{T})$, the label implicitly attached to C and D has acquired meaning. Those are “the strategies labeled by symbols which occur only once”, while A and B are “the strategies labeled by the symbol which occurs twice”. Hence $\mathcal{U}(\mathcal{T})$ should be taken to be a familiar frame.

Recall that combining two familiar frames can be done by finding the meet of the two induced symmetry structures. For unfamiliar frames this is not the case. Consider game M_5 and two frames. Frame L_1 labels strategies A and B as \odot , strategies C and D as \ominus , and strategy E as \ominus , which induces $\mathcal{T}(L_1) = \{\{A, B\}, \{C, D\}, \{E\}\}$ and $\mathcal{U}(\mathcal{T}(L_1)) = \{\{A, B, C, D\}, \{E\}\}$. Frame L_2 labels A and C as \clubsuit , B and D as \diamond , and E as \spadesuit , inducing $\mathcal{T}(L_2) = \{\{A, C\}, \{B, D\}, \{E\}\}$ and $\mathcal{U}(\mathcal{T}(L_2)) = \mathcal{U}(\mathcal{T}(L_1))$. The combination of the two frames is yet another frame, in which every single strategy has a different label. Hence, its induced symmetry structure $\mathcal{T}(L_1) \wedge \mathcal{T}(L_2)$ is the trivial symmetry structure, while its counterpart is $\mathcal{U}(\mathcal{T}(L_1) \wedge \mathcal{T}(L_2)) = \mathcal{T}^*$, the coarsest symmetry structure. This demonstrates that the symmetry structure induced by a combination of two unfamiliar frames can be coarser than the meet (and even the join) of the two symmetry structures induced by the two frames separately. I.e. we here have $\mathcal{U}(\mathcal{T}(L_1) \wedge \mathcal{T}(L_2)) \succ \mathcal{U}(\mathcal{T}(L_1)) \wedge \mathcal{U}(\mathcal{T}(L_2))$.

Note furthermore that if we have another frame L_3 which is identical to frame L_1 then obviously $\mathcal{U}(\mathcal{T}(L_3)) = \mathcal{U}(\mathcal{T}(L_1)) = \mathcal{U}(\mathcal{T}(L_2)) = \{\{A, B, C, D\}, \{E\}\}$ and, hence, also $\mathcal{U}(\mathcal{T}(L_3)) \wedge \mathcal{U}(\mathcal{T}(L_1)) = \mathcal{U}(\mathcal{T}(L_2)) \wedge \mathcal{U}(\mathcal{T}(L_1))$ and yet $\mathcal{U}(\mathcal{T}(L_1) \wedge \mathcal{T}(L_2)) \neq \mathcal{U}(\mathcal{T}(L_1) \wedge \mathcal{T}(L_3))$. This demonstrates that the symmetry structure induced by a combination of two unfamiliar frames can not be written as an operation on the two symmetry structures induced by the two unfamiliar frames.

Finally, consider game M_6 and the following two frames. Frame L_1 attaches labels $\bullet, \bullet, \bullet, \circ, \circ, \circ$ to the 6 strategies and frame L_2 attaches labels $\square, \square, \blacksquare, \blacksquare, \blacksquare, \square$. We thus have $\mathcal{T}(L_1) = \{\{A, B, C\}, \{D, E, F\}\}$ and $\mathcal{T}(L_2) = \{\{A, B, F\}, \{C, D, E\}\}$ with $\mathcal{T}(L_1) \wedge \mathcal{T}(L_2) = \{\{A, B\}, \{C\}, \{D, E\}, \{F\}\}$. However, the symmetry structures induced by the same unfamiliar frames are given by $\mathcal{U}(\mathcal{T}(L_1)) = \mathcal{U}(\mathcal{T}(L_2)) = \mathcal{T}^*$ and, hence, also $\mathcal{U}(\mathcal{T}(L_1) \wedge \mathcal{T}(L_2)) = \mathcal{T}^*$. However, the symmetry structure induced by the two

unfamiliar frames together is given by $\mathcal{U}(\mathcal{T}(L_1) \wedge \mathcal{T}(L_2)) = \{\{A, B, D, E\}, \{C\}, \{F\}\}$. Thus, the symmetry structure induced by the combination of two unfamiliar frames can be finer than the meet of the induced symmetry structures of the individual unfamiliar frames. I.e. $\mathcal{U}(\mathcal{T}(L_1) \wedge \mathcal{T}(L_2)) \prec \mathcal{U}(\mathcal{T}(L_1)) \wedge \mathcal{U}(\mathcal{T}(L_2))$.

To summarize, a frame, if it is unfamiliar, induces a weakly coarser symmetry structure than the same frame does when it is familiar. The symmetry structure induced by the combination of two unfamiliar frames is not straightforward and can be either finer or coarser than the symmetry structure obtained from the meet of the two symmetry structures induced by the two unfamiliar frames separately.

6 Focal Points

6.1 A Definition of Focal Points

Suppose that, given a game and the global symmetry structure associated to an accompanying frame, there exists a unique rational symmetric recommendation. 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, it will be an easy task for consultants to advise in favor of this unique recommendation, and actual play will then conform to it. Thus one obtains a much sharper prediction than the one derived from e.g. rationalizability (Bernheim (1984) and Pearce (1984)).

Typically, the set of rational symmetric recommendations is not a singleton. By construction, however, each rational symmetric recommendation must be unique in one way or another (else they would involve playing different but symmetric strategies, a contradiction). Furthermore, again by construction, 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-coordinating on which recommendation to choose 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 from all familiar aspects of a game into e.g. the real line, that is, defining a ranking which can be interpreted as **salience**. Of course, this mapping needs to be quite detailed to cover all cases. In many cases, however, a partially specified meta-norm will suffice.

What are familiar aspects of a game? Consider a situation in which both the game and its frame are completely unfamiliar to players. Yet, there are familiar aspects of

this game. A readily understood aspect of a game is that it provides payoffs, be it in monetary or utility terms. As all games have payoffs, this is something consultants can use to coordinate their choice of a recommendation. A partially specified meta-norm, for instance, could thus say: “always choose the Pareto-undominated recommendation if there is a unique one”. Obviously this does not cover all the cases, but it is simple enough and will still provide a unique recommendation in many contexts. Indeed one could imagine that all consultants go through the same school(ing) of thought, say, being indoctrinated with the principle that one should recommend Pareto-undominated recommendations first if there is a unique one. Thus, if all players follow the advice of consultants, then indeed play will conform to the meta-norm.¹³

There are, however, many possible meta-norms one could write down,¹⁴ teach, and follow. One could, for instance, recommend the risk-undominated recommendation if there is a unique one. Another meta-norm could choose the most equitable payoff-tuple if there is a unique one. Many other partially specified meta-norms are possible. For the theory in this paper to “come true”, i.e. to transform from a normative theory to a positive, descriptive one, one would need, first, consensus among all game theorists as to what this meta-norm should be, and second, the teaching of the meta-norm to all potential players. Alternatively, one could develop an evolutionary argument (as in e.g. Binmore and Samuelson (2006)) 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 (if indeed such a process converges).

For some games meta-norms may (have to) be a more general function into a salience ranking not only from payoff tuples, but also from other familiar aspects of the game. Especially in games with familiar frames there must be a meta-norm ranking labels in terms of their salience. For instance when players are asked to coordinate on either “heads” or “tails” it is well documented (e.g. in Schelling (1960) and Mehta, Starmer, and Sugden (1994a, 1994b)) that “heads” is generally considered more salient.

A more subtle meta-norm is necessary in unfamiliar games with unfamiliar frames. Consider again game M_4 with the unfamiliar frame L with labels $\blacksquare, \blacksquare, \square, \triangle$, which we discussed in Section 5. The induced symmetry structure $\mathcal{U}(\mathcal{T}(L))$ is given by $\{\{A, B\}, \{C, D\}\}$ for both players. Both rational symmetric recommendations $(\frac{1}{2}, \frac{1}{2}, 0, 0)$ or $(0, 0, \frac{1}{2}, \frac{1}{2})$ are payoff equivalent. Hence, one might think that individuals might be confused and choose $(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$. While this is a quite plausible outcome in an experimental setting, our argument is that individuals could in principle hold a common

¹³Being well-indoctrinated economists, we of course believe that Game Theory should be taught in primary school. In that case, with teachers as consultants, the meta-norm is learned by all players at school, who, knowing that it is in place, never have an incentive not to follow it.

¹⁴One could indeed write a book detailing a meta-norm. In fact, Harsanyi and Selten (1988) is an attempt to provide exactly such a book. There are, however, other possible books waiting to be written.

meta-norm that dictates them to choose, in such games, one of the two strategies with a distinct label, leading to $(0, 0, \frac{1}{2}, \frac{1}{2})$, or a meta-norm prescribing to choose one of the two strategies which have the same label, leading to $(\frac{1}{2}, \frac{1}{2}, 0, 0)$.

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 (1995)) among players. Our point here is that such a very intricate meta-norm covering all games (without frames or with familiar or even unfamiliar frames) 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 15. For a given game and frame, whether familiar or not, 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.

6.2 Meta-Norms and the Experimental Literature

The objective of this paper has been to develop a normative framework to identify all possible focal points of a given framed game, and to demonstrate the need of a meta-norm selecting among the set of rational symmetric recommendations. Except for arguing that such a meta-norm could exist (be written down, be taught, be learned), we remain agnostic as to what it might actually be. One possible avenue of research calls for an empirical/experimental assessment of whether, and if so what kind of, a (probably partial) meta-norm is in place today. Indeed there are quite a few experimental studies of this nature, beginning with Schelling (1960). In this section we selectively explore this literature and whether it allows us to make some inference about the meta-norms, if any, used nowadays. We do this in increasing order of the level of intricacy the meta-norm needs to exhibit in order to choose a unique recommendation. The first example, due to Schelling (1960), is one, in which, although there are many Nash equilibria, there is only one rational recommendation, and, hence, in fact, no meta-norm is required. In the second example, taken from Blume and Gneezy (2000), the games are framed, but simple meta-norms over payoff-tuples suffice to give unique recommendations.¹⁵ Finally, the framed games studied by Crawford, Gneezy, and Rottenstreich (2008) allow best to

¹⁵The experimental studies by Mehta, Starmer, and Sugden (1994a, 1994b) require a meta-norm based both on payoff-tuples as well as labels in, however, a simple and non-conflicting way. These papers explore the nature of salience (of labels in familiar frames) in matching games. It is implicitly (and probably correctly) assumed that individuals choose according to the meta-norm of Pareto-dominance (or simply maximize the probability of coordination). However, strategies are labeled in many different ways and thus a meta-norm over labels is needed.

disentangle meta-norms that only depend on payoff-tuples from meta-norms that only depend on labels (and any combination thereof).

First, consider Game 4 in Schelling (1960, p. 61): “You and your partner (rival) are to be given \$100 if you can agree on how to divide it without communicating. Each of you is to write the amount of his claim on a sheet of paper; and if the two claims add to no more than \$100, each gets exactly what he claimed. If the two claims exceed \$100, neither of you gets anything. How much do you claim?”

Suppose that one can only specify amounts in multiples of one dollar. None of the 101 strategies of the game are symmetric in any symmetry structure as they all give rise to different possible payoffs. However, the two players are clearly symmetric. For each strategy of player 1 there is an equivalent strategy of player 2, that is, the identification α can be readily specified. Hence, in a rational symmetric recommendation both players have to be treated equally. This leaves only a single rational symmetric recommendation, in which each player claims exactly \$50. This is without any appeal to a meta-norm. There are many other Nash equilibria, in which players claim different amounts (adding up to \$100). Indeed Schelling (1960) reports that players do claim mostly exactly \$50.

We now discuss two of the four treatments in the (very carefully designed) experimental setting of Blume and Gneezy (2000). Each of two players is given a physical plastic disk with three (or nine) sectors of exactly equal shape and size on each side. Nothing distinguishes these sectors or the two sides of the disk. The game has two stages. In stage one each player is shown one side of the disk, asked to select one slice, and told that there will be a stage two. Players do not know in which order they are asked. Between the two stages chosen sectors are marked (indistinguishably) on both sides. In the second stage the disk is again shown to both players, who are now asked to choose a sector again and told that they will be paid \$10 if they both choose the same sector and nothing if they choose different sectors.

Consider the game with 9 sectors and the (most likely) case in which different sectors have been chosen by both players in the first stage. Suppose that players chose sectors such that there are 2 (and 5) sectors between them. Let us arbitrarily label strategies in the following way. Call one of the two players player 1 and let the two players' strategies be 1 for the sector he marked and 2, 3, up to 9, in clockwise order. Blume and Gneezy (2000) explain how sectors can be labeled in an objective way, a discussion which we can readily translate in terms of symmetry structures. First, the two chosen sectors must receive the same label “chosen in stage 1”, because players do not know who chose them. The two sectors between them also must receive identical labels of “adjacent to a chosen sector in the even (small) section”. Then there are two sectors “adjacent to a chosen sector in the odd (big) section” and two others with the label “one apart from a chosen sector”. This leaves a unique sector “in the middle between

the chosen sectors (on the odd section)”. The symmetry structure is thus $\mathcal{T}_1 = \mathcal{T}_2 = \{\{1, 4\}, \{2, 3\}, \{5, 9\}, \{6, 8\}, \{7\}\}$. Under the meta-norm of Pareto-dominance we must pick strategy 7 (the only uniquely labeled sector). In the experiment Blume and Gneezy (2000) find that players often play according to this theory in the three-sector game, but fail to do so in the nine-sector game. This is probably not so much due to the lack of a meta-norm, but rather due to the fact that the game is not simple and recognizing that the sector in the middle of the odd side is “uniquely unique” is not so easy (we remind the reader that, in contrast, our approach is normative). Thus, even if a commonly known meta-norm might be in place (such as Pareto-dominance), players might not follow it due to some boundedness in their rationality, which leads them to misperceive the symmetry structure. Explaining the outcome of this game in the lab, while maintaining the meta-norm of Pareto-dominance, thus requires a behavioral model.¹⁶

Perhaps the most interesting study from the point of view of this paper is that of Crawford, Gneezy, and Rottenstreich (2008), hereafter CGR. The framed games studied there allow a variety of possible meta-norms, some over payoff-tuples only, some over labels, some a combination of the two. In our terminology, their main finding is that the prevailing meta-norms (when they seem to be common knowledge) are mostly in terms of payoff-tuples. Labels seem to come in almost only lexicographically. The framed games investigated in CGR are similar to the ones in Blume and Gneezy (2000) but simpler, in terms of figuring out the symmetry structure, yet more varied in terms of payoffs. The names of the games are taken from CGR, where the first digit stands for (S)ymmetric or (A)symmetric, the second for (S)light, (M)oderate, or (L)arge payoff differences, and the third for (L)abeled or (U)nabeled.

The first set of games in CGR, called “X-Y games”, also offer a nice illustration of our framework. The games ASU and AMU are unframed battle of the sexes games with player symmetry but no strategy symmetry and the mixed-strategy equilibrium as the unique rational recommendation: play the strategy that gives you a higher payoff in case of coordination with the appropriate probability ($\frac{50}{101}$ in ASU and $\frac{5}{11}$ in AMU). The findings are that in both games roughly 38% of the players use the higher-payoff strategy, which, with a standard error of about 10% (for each game separately), is not significantly different from the recommendation. The other four games are framed, with one strategy labeled X and the other Y. X (Y) always provides a payoff of 5 to player 1 (2) if used by both players, while the four games vary in terms of what player 1 (2) obtains when both play Y (X). These payoffs are 5 in SL, 5.1 in ASL, 6 in AML, and 10

¹⁶Blume and Gneezy (2008) show that the outcome of play in this game is not necessarily only due to players not recognizing that there is a unique label, but also to players fearing that their opponent does not realize this. This emphasizes the need for the meta-norm in place to be common knowledge. Thus, actual play in such games might be more readily explained by, for instance, the Variable Universe Matching Game (VUMG) models of Janssen (2001) or Casajus (2000) or the model provided by Blume and Gneezy (2008), all of which are based on the Variable Frame Theory of Bacharach (1993).

in ALL. Given the, assumed familiar, labeling X,Y, strategies are not symmetric and, except for game SL, neither are the two players. Thus going from SL to the other three games breaks the player symmetry. In SL 76% of players chose X, while the frequency of choosing X and Y varies greatly across player roles in the other three games. Thus the meta-norm of using label X seems confined to game SL. The problem might be as follows. In a game as ASL, AML, or ALL, where the two players obtain different payoffs when playing X, a salience ranking favoring X over Y might induce a salience ranking over players, ranking the player with highest payoffs first. However, the ranking might also rank this player last. It seems that players are not sure about which meta-norm is in place and thus are confused as to which one to follow.

We look now at the second set of games in CGR, called “pie games”. Those are games of pure coordination with three strategies, say A,B,C, and were, as were the “X-Y” games, played only once by each subject. They all share the same frame L with labels ■, ■, □, i.e. A and B always receive the same label. In the experiments these labels were implemented with colored and uncolored sectors of the pie. There are 8 pie games in total, which we will briefly review here. We do not know whether subjects were familiar with this frame, i.e. held a salience ranking over color versus no color, but symmetry structures here are always the same for the familiar and the unfamiliar case. In all games below, when the two meta-norms of Pareto-efficiency and risk-dominance yield the same focal points, while the meta-norm of equity is silent, we refer to them simply as focal points without further specification.

Game S1 is the 2-player matching game with 3 strategies, M_3 , with payoffs multiplied by 5. The associated symmetry structure is given by $\mathcal{T}_1 = \mathcal{T}_2 = \{\{A, B\}, \{C\}\}$, $\mathcal{I} = \{\{1, 2\}\}$. I.e. we have player symmetry and strategies A and B are symmetric for both players. Thus we obtain the focal point C, which is indeed played by 94% of the subjects.

AL1 and AM1 correspond to the pure coordination games $\text{Diag}((5, 10), (10, 5), (5, 5))$ and $\text{Diag}((5, 6), (6, 5), (5, 5))$. Players are symmetric, while the strategy symmetry structure is the trivial one, i.e. every strategy is identified. Due to player symmetry, the rational symmetric recommendations involve mixing among A and B (as in Game 7), or coordinating on C. The latter then becomes a focal point, and indeed in the experiment roughly 80% and 90% of players chose it in AL1 and AM1, respectively.

S2 is the game of pure coordination $\text{Diag}((6, 6), (6, 6), (5, 5))$. The symmetry structure is the same as for S1, and again we obtain the focal point C. The experimental results for this game are perhaps the most surprising of all 8 pie-games in CGR. Here only 43% of subjects play C (this is in fact roughly consistent with the completely mixed rational symmetric recommendation). Ex-post one could perhaps explain this puzzle by assuming players have two conflicting meta-norms in mind. On the one hand they find colored slices more salient than uncolored ones and on the other they find Pareto-optimal

rational symmetric recommendations more salient. Somehow they cancel each other out. It thus seems plausible (this is indeed the avenue CGR then pursue) that players do not hold a commonly known meta-norm for this game, perhaps then trying to guess what norms are more likely and then following some form of level- k reasoning.

S3 is the game of pure coordination $\text{Diag}((6, 6), (5, 5), (6, 6))$. Players are symmetric, but the strategy symmetry structure is the trivial one. This is so because strategy C is identified through its label \square , while B is identified through its unique payoff possibility of 5. Thus, also A is identified. All Nash equilibria are rational symmetric recommendations, and the game has no focal points according to any of the three meta-norms of Pareto-efficiency, equity, or risk-dominance. Given that, one might expect players to choose the strategy with a unique label, C, which somewhat contradicts the ex-post explanation of results in game S2. In any case players do not quite do this. The frequency of choices was 14%, 21%, and 64%, respectively. This is just a little more systematic than complete confusion. Thus, there seems to be no commonly known meta-norm in place that would help players to coordinate play to achieve equilibrium in this game.

Game AM2¹⁷ is one of pure coordination given by $\text{Diag}((5, 6), (6, 5), (6, 5))$. It is easy to see that the symmetry structure has both players identified, $\mathcal{I} = \{\{1\}, \{2\}\}$, and the strategy symmetry structure is the trivial one. Like game S3, this game has no focal points with respect to any of the three meta-norms mentioned above. It does have a focal point (C) induced by the meta-norm of choosing the unique color. Yet, the result in the experiment is confusion (see CGR for details). Again there does not seem to be a commonly known meta-norm in place.

In sum, the (selectively chosen) evidence above seems to indicate that, nowadays, simple meta-norms such as choosing Pareto-undominated recommendations are (almost) commonly known and are effective when they provide a unique recommendation. If such simple meta-norms based on payoff-tuples alone do not provide unique recommendations, and only then, (somewhat) commonly know meta-norms over label-salience are sometimes held. However, in many cases, when simple meta-norms such as Pareto-dominance do not provide unique predictions, players do not hold a commonly known meta-norm. Still, we only have a very incomplete picture of which, if any, meta-norms are in place.

7 Conclusion

In this paper, we take the position that the concept of focal point, the “natural way to play a game”, reflects two different considerations. The first one is **symmetry**, and boils down to the observation that strategies or players which cannot be told apart, must be treated equally. The second one is the necessity of a **meta-norm**.

¹⁷The analysis and results for games AM3 and AM4 are essentially identical to this one.

The meaning of symmetry can be readily formalized. Building on concepts introduced by Nash (1951), Harsanyi and Selten (1988), and Crawford and Haller (1990), 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 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 (1951)). The coarsest such structure corresponds to the unframed game, that is, it captures all symmetries which can be 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 (1990), builds on pairwise comparisons of strategies. The second, closer to Nash (1951), 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 can not, 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.

A Appendix: 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 (2002) or Grätzer (2003) for details.

A set X endowed with a partial order \leq 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, \leq) is a **top** (resp. **bottom**) or greatest (resp. smallest) element if there exists no $x' \in X$ with $x \leq x'$ (resp. $x' \leq x$) 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.

Fact A3. Any finite join semilattice with a bottom is actually a lattice.

A.2 Proofs from Section 3

The following simple Lemma gives a useful characterization of the join of two partitions.

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 , $s_i^0 = s_i, s_i^1, s_i^2, \dots, s_i^k = s'_i$ 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 $t = 0, \dots, k-1$.

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 $T_{i,1}, \dots, T_{i,\ell} \in \mathcal{T}_i$ and $T'_{i,1}, \dots, T'_{i,\ell'} \in \mathcal{T}'_i$ 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 T_i'' . \square

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 relabelings 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 relabelings with $\rho_j(T_j) = T_j$ for all $T_j \in \mathcal{T}_j$ for all $j \neq i$ includes all relabelings 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 $T_j \in \mathcal{T}_j$ 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 relabeling 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 relabeling 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. \square

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. \square

PROOF OF THEOREM 3: Let $\beta_i : \Theta_{-i} \rightarrow \Theta$ denote the (mixed) best-reply correspondence of player i , and $\beta : \Theta \rightarrow \Theta$ the product correspondence given by $\beta(x) = \times_{i \in I} \beta_i(x_{-i})$. We know that the β_i , and hence β , are non-empty 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 $\tilde{\Theta}_i$ be the set of mixed strategies x_i such that $x_i(s_i) = x_i(s'_i)$ whenever s_i, s'_i belong to the same symmetry class in \mathcal{T}_i . Notice that $\tilde{\Theta}_i$ is convex. Define $\tilde{\beta}_i : \tilde{\Theta}_{-i} \rightarrow \tilde{\Theta}_i$ by $\tilde{\beta}_i(x_{-i}) =$

$\beta_i(x_{-i}) \cap \tilde{\Theta}_i$. Thus $\tilde{\beta}_i$ 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 $x_{-i} \in \tilde{\Theta}_{-i}$, there exists a best response of player i which gives the same weight to any two symmetric strategies. Fix $x_{-i} \in \tilde{\Theta}_{-i}$. For each $j \neq i$ and each $T_j \in \mathcal{T}_j$, there exists $y(T_j) \geq 0$ 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 relabelings 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,

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

where the third equality follows from the definition of symmetric strategies if we recall that the relabelings 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} . If neither is a best response, then in all best responses x_i the strategies s_i and s'_i have identical weight zero. If both are best responses then so is any convex combination between them. In particular then there is also $x_i \in \beta(x_{-i})$ with $x_{i,s_i} = x_{i,s'_i}$. This proves the nonemptiness of $\tilde{\beta}_i(x_{-i})$ for all i . Hence, $\tilde{\beta}$ satisfies non-emptiness, convex-valuedness, upper hemicontinuity, and, hence, by Kakutani's fixed point theorem there is a fixed point. Since a fixed point of $\tilde{\beta}$ is also a fixed point of β , it is a Nash equilibrium. \square

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 (1978) or Hungerford (1980) for details.

A **group** is a nonempty set G endowed with a binary, internal operation “.” 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 $1_Gg = g1_G = g$ 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 $g_1 g_2 g_3 \cdots g_r$ where $g_\ell \in H_1 \cup H_2$ for all $\ell = 1, \dots, r$.

A.4 Proofs from Section 4

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 $\text{Sym}(\Gamma)$, forms a **group** with the operation given by symmetry composition.

Let $\text{Sym}(\mathcal{I}, \mathcal{T}, \alpha)$ 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 $\text{Sym}(\mathcal{I}, \mathcal{T})$. If $(\mathcal{I}^0, \mathcal{T}^0)$ is the trivial global symmetry structure, then $\text{Sym}(\mathcal{I}^0, \mathcal{T}^0)$ is the subgroup formed by the identity symmetry only.

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

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 A4.¹⁸ \square

Let Φ be an arbitrary subgroup of $\text{Sym}(\Gamma)$. 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 $\text{Sym}(\Gamma)$, 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

¹⁸If 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, as Game 10 shows.

a global symmetry structure. Further, for each global symmetry structure $(\mathcal{I}, \mathcal{T}, \alpha)$, if $\Phi = \text{Sym}(\mathcal{I}, \mathcal{T}, \alpha)$ 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 equivalent¹⁹ to the statement that, for any two symmetries $(\sigma', \tau'), (\sigma'', \tau'') \in \Phi$, 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 . \square

Hence, the set of global symmetry structures of a game and the set of subgroups of symmetries are lattice-isomorphic. In this sense, a global symmetry structure can be identified with its associated group of symmetries. This observation yields the following.

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 $\text{Sym}(\mathcal{I}'', \mathcal{T}'', \alpha'')$ is a subgroup of $\text{Sym}(\mathcal{I}', \mathcal{T}', \alpha')$.

This allows us to prove Theorem 4.

PROOF OF THEOREM 4: The set of subgroups of a group is a lattice due to Facts A5 and A6. The conclusion follows by Proposition A1 and Lemma A3. \square

It is worth remarking that 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 (1951, Theorem 2).

PROOF OF THEOREM 6: Let $(\mathcal{I}, \mathcal{T}, \alpha)$ denote the symmetry structure. Define $\tilde{\beta}_i$ 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 $\prod_{i \in I} \tilde{\Theta}_i$ 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 $\hat{\beta}(y) = \tilde{\beta}(y) \cap \hat{\Theta}$ and, since $\hat{\Theta}$ is convex, $\hat{\beta}$ is convex- and upper hemicontinuous.

¹⁹The sufficiency follows because the composition of symmetries in the subgroup Φ is also in Φ .

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. \square

References

- AUMANN, R., AND A. BRANDENBURGER (1995): “Epistemic Conditions for Nash Equilibrium,” *Econometrica*, 63, 1161–1180.
- BACHARACH, M. (1991): “Games with Concept-Sensitive Strategy Spaces,” Mimeo.
- (1993): “Variable Universe Games,” in *Frontiers of Game Theory*, ed. by A. K. K. Binmore, and P. Tani, p. 255–276. MIT Press, Cambridge, MA.
- BERNHEIM, B. D. (1984): “Rationalizable Strategic Behavior,” *Econometrica*, 52, 1007–29.
- BINMORE, K., AND L. SAMUELSON (2006): “The Evolution of Focal Points,” *Games and Economic Behavior*, 55, 21–42.
- BLUME, A. (2000): “Coordination and Learning With a Partial Language,” *Journal of Economic Theory*, 95, 1–36.
- BLUME, A., AND U. GNEEZY (2000): “An Experimental Investigation of Optimal Learning in Coordination Games,” *Journal of Economic Theory*, 90, 161–172.
- (2008): “Cognitive Forward Induction and Coordination Without Common Knowledge: An Experimental Study,” mimeo.
- CASAJUS, A. (2000): “Focal Points in Framed Strategic Forms,” *Games and Economic Behavior*, 32, 263–291.
- (2001): *Focal Points in Framed Games*. Springer.
- CRAWFORD, V., AND H. HALLER (1990): “Learning How to Cooperate: Optimal Play in Repeated Coordination Games,” *Econometrica*, 58, 571–596.
- CRAWFORD, V. P., U. GNEEZY, AND Y. ROTTENSTREICH (2008): “The Power of Focal Points is Limited: Even Minute Payoff Asymmetry May Yield Large Coordination Failures,” *American Economic Review*, p. forthcoming.

- DAVEY, B., AND H. PRIESTLEY (2002): *Introduction to Lattices and Order (2nd Ed)*. Cambridge University Press, Cambridge.
- GAUTHIER, D. (1975): “Coordination,” *Dialogue*, 14, 195–221.
- GOYAL, S., AND M. JANSSEN (1996): “Can We Rationally Learn to Coordinate?,” *Theory and Decision*, 40, 29–49.
- GRÄTZER, G. (2003): *General Lattice Theory (2nd Ed)*. Birkhäuser, Basel.
- HARSANYI, J., AND R. SELTEN (1988): *General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, MA.
- HUNGERFORD, T. W. (1980): *Algebra*. Springer Verlag, New York.
- JANSSEN, M. C. W. (2001): “Rationalizing Focal Points,” *Theory and Decision*, 50, 119–148.
- KREPS, D. M. (1990): *A Course in Microeconomic Theory*. Princeton University Press, Princeton.
- MEHTA, J., C. STARMER, AND R. SUGDEN (1994a): “Focal Points in Pure Coordination Games: An Experimental Investigation,” *Theory and Decision*, 36, 163–185.
- (1994b): “The Nature of Salience: An Experimental Investigation of Pure Coordination Games,” *American Economic Review*, 74, 658–673.
- NASH, J. (1951): “Non-Cooperative Games,” *Annals of Mathematics*, 52, 286–295.
- PEARCE, D. G. (1984): “Rationalizable Strategic Behavior and the Problem of Perfection,” *Econometrica*, 52, 1029–51.
- ROSE, J. S. (1978): *A Course in Group Theory*. Cambridge University Press (Dover Reprint 1994), Cambridge, UK.
- SCHELLING, T. C. (1960): *The Strategy of Conflict*. Harvard University Press, Cambridge.
- SUGDEN, R. (1995): “A Theory of Focal Points,” *Economic Journal*, 105, 533–550.
- TVERSKY, A., AND D. KAHNEMAN (1981): “The Framing of Decisions and the Psychology of Choice,” *Science*, 211, 453–458.