

No. 35 august 2008

**Carlos Alós-Ferrer
Christoph Kuzmics**

Hidden Symmetries and Focal Points

Research Paper Series
Thurgau Institute of Economics and Department of Economics
at the University of Konstanz



**THURGAU INSTITUTE
OF ECONOMICS**
at the University of Konstanz

Hidden Symmetries and Focal Points*

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

August 2008

Abstract

This paper provides 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. We also clarify the difference between different concepts of symmetry used in the literature.

Journal of Economic Literature Classification Numbers: C72, D83.

Keywords: symmetry, focal points, Nash equilibria.

The main aim of this paper is to provide a general framework for the analysis of focal points in the tradition of Schelling (1960). Our analysis also enables us to clarify the meaning of the concept of a framed game, which is extensively used in various forms in the literature. We remark right away that the framework we develop is not limited to pure coordination games (as most papers on focal points and framed games), but encompasses arbitrary normal-form games.

Games, in this paper, are truly one-shot. 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.¹ The game, in this paper, can also come with a frame² (context, setting). The frame can be familiar to players or it can be unfamiliar as well.

We will argue that a formal definition of a focal point of a framed game has to come in two parts. First, a careful study of the symmetry structures present in the game and its frame is required. Much of this paper is devoted to the analysis of these symmetry structures. On the basis of the induced symmetry structure of a framed game, a subset of strategy profiles is identified which is consistent (in a well-specified sense) with this symmetry structure and meets certain rationality requirements. We shall call these **rational symmetric recommendations**. A multiplicity problem is unavoidable, and indeed will be exacerbated the more symmetries are broken by increasingly detailed frames. Hence, the second

*The first version of this paper was written while the first author visited the Kellogg School of Management at Northwestern University. The first author thanks these institutions for their hospitality. Financial support from the Austrian Science Fund (FWF) under Project P18141-G09 is gratefully acknowledged. We would like to thank Andreas Blume, Hans Haller, Andrew McLennan, Roger Myerson, Itai Sher, Jonathan Weinstein, and seminar participants in Northwestern University, Vancouver, and Karlsruhe for helpful comments.

[†]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

¹Suppose 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.

²The concept of frame goes back to Tversky and Kahneman (1981), but has been formalized in different ways in the game-theoretic literature.

part of the definition must appeal to a (maybe intricate) meta-norm. A focal point is then a rational symmetric recommendation which is uniquely selected by means of a such a meta-norm. We shall now explain these two parts in more detail.

Let us start with symmetries. Suppose a player approaches a game theory advisor (a consultant) asking for guidance as to how to play a particular game. Assume for concreteness that the player is dealing with a finite two-player game. The advisor will most likely start by writing down the game in one form or another, for instance using a bi-matrix form. The advisor's client could arbitrarily be called 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. Then the two players are symmetric. Similarly, the names strategies receive in the advisor's representation are arbitrary. If there are two ways of writing down the game leading to the same payoff tables although the ordering of strategies is different, then we should declare the two strategies symmetric.

It turns out that these intuitions hide a large number of subtleties. For the most part this paper is thus concerned with the question as to which strategies and which players can indeed be declared symmetric given a framed game. 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.³ 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. We will 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 will lead to different predictions than the one based on Crawford and Haller (1990). 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 allows an identification of every possible familiar frame, and, conversely, 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.

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 on 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). Any strategy profile that satisfies these three axioms shall be called a **rational symmetric recommendation**.

In some simple games, one is led to a unique rational symmetric recommendation. For general games, however, the set of such recommendations grows as the frame incorporates more and more information. Thus, as implicitly recognized by the literature, the multiplicity problem can only be solved through an appeal to some meta-norm. Meta-norms can range from fairly simple to very intricate.

Games have payoffs, either in monetary or utility terms. Suppose they are in 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 could give rise to the Pareto-optimal outcome would enable players to coor-

³Crawford and Haller (1990) use their framework to study under which conditions can players in a repeated two-player coordination game use history to coordinate in a pure equilibrium. See also Blume (2000).

dinate even if the particular game has never been encountered before. We shall call this the meta-norm of Pareto-efficiency.⁴ In most of our illustrative examples this meta-norm coincides with the meta-norm of risk-dominance and does not conflict with the meta-norm of equity.

If we say a frame is familiar we mean that players have had time to form a common meta-norm as to how compare the various labels. This meta-norm is just 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 our discussion, we will adopt the implicit assumption that there is a commonly known meta-norm in place. In fact, 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.

We will, however, not explore what these meta-norms are.⁵ Further, we will abstract from possible conflicts between alternative meta-norms. An alternative road would lead to models of incomplete information. 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. In the same spirit 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, with different types potentially being aware or unaware of some attributes strategies have.⁶ There is a commonly known distribution over these types with the caveat that if a player is unaware of a certain attribute he is also unaware that others might be aware of it. This gives rise to another incomplete information game.

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 unique label which appears twice, while two others are the only two labels that appear only once. In principle, for such cases (and after all other choice criteria, such as Pareto-optimality, fail), a meta-norm could be commonly held that tells individuals to, for instance, go for a strategy with a unique label. This is well defined, only there are two of them. So the player will have to randomly pick among these two. The meta-norm could alternatively be such that players go for the two strategies with the label that is the only one that appears twice. Again, we will not explore which meta-norm is more likely. We will simply postulate common knowledge among players of such a meta-norm. We are not aware of any study investigating salience on this level.

Finally, having formulated what we mean by a focal point we explore how focal points based on the meta-norm of Pareto-dominance fare in terms of predictive power in the relatively simple games which were explored in recent experiments by Crawford, Gneezy and Rottenstreich (2008). We find most results in that paper (with one exception) surprisingly consistent with our (normative) predictions.

I. Motivating Example

In this paper, 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. However, neither player positions nor strategy names have any intrinsic meaning. A strategy must then 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.

To fix ideas, let us consider a few variants of the following highly symmetric game, which are inspired by Schelling (1960, p.56 and p.296).

⁴Variations of this meta-norm are called the “Principle of Coordination” in Gauthier (1975), Bacharach (1991), Bacharach (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).

⁵An interesting experimental investigation into possible meta-norms for familiar labels is given by Mehta, Starmer and Sugden (1994a, 1994b), who explore the nature of salience (of labels in familiar frames) in matching games.

⁶Binnmore and Samuelson (2006) investigate an evolutionary model in which players recurrently face games, in which strategies come with two attributes and ask the question under which circumstances a meta-norm emerges that uses only one or both attributes when paying attention to both is more costly than to just one.



Figure 1. Four players around a table.

The situation is depicted in Figure 1(a). There are 4 players around a table, each indicated by a \circ . On the table there are 4 objects, each indicated by a \square . Players cannot communicate. They all look alike, as do the objects (or they differ in a nondescript way only). Each player has to choose one object. If they all choose the same they get \$12 each, otherwise they get \$0. The game is played once.

If this is the whole description of the game it is difficult to see how players could possibly manage to coordinate their actions to choose the same object except by chance. Suppose a player talks to a game theory consultant before playing. The player informs the consultant about the game. The consultant may write the game down in matrix form, identifying the four objects (strategies) as A, B, C , and D . Strategy A might, for instance, stand for the player's top-left hand object. Strategies B, C , and D are then the objects found from A in, say, a clockwise order. The payoff matrix is then given by $u(s_1, s_2, s_3, s_4) = (12, 12, 12, 12)$ whenever $s_1 = \dots = s_4$ and $(0, 0, 0, 0)$ otherwise. For a two-player version, this is given by Game 1 below.

	A	B	C	D
A	12,12	0,0	0,0	0,0
B	0,0	12,12	0,0	0,0
C	0,0	0,0	12,12	0,0
D	0,0	0,0	0,0	12,12

Game 1

	A	B	C	D
A	12,12	0,0	0,0	0,0
B	0,0	12,12	0,0	0,0
C	0,0	0,0	12,12	0,0
D	0,0	0,0	0,0	x, x

Game 2

The consultant should realize that there is a great deal of arbitrariness in her representation. First, she might as well have written the game down in a very different way, e.g. with A standing for the bottom-right strategy, and she would still obtain the exact same game form. Second, if another player had come to her she would perhaps have called him player 1 and written down the game in yet another way, still obtaining the exact same game form. Hence, the consultant should realize that if she recommends A this is completely arbitrary, as A only makes sense in her depiction of the game. Hence, the only realistic recommendation she can give is to play all 4 strategies with equal probability (which results in coordination with probability $(\frac{1}{4})^3$ only), while at the same time predicting that other players will receive the same recommendation. In doing so, the consultant is fulfilling three “axioms”. First, symmetric strategies are treated equally. Second, symmetric players are treated equally. Third, a form of rationality is respected, because ultimately she is recommending (and predicting) a Nash equilibrium.

Now suppose there is one object which pays $x > 0$ instead of \$12 if all players coordinate in choosing it. All other objects still pay \$12 in case of coordination. Lack of coordination still pays \$0. For the two-player case, this situation can be written as Game 2 above, where the consultant has chosen to identify the strategy (or object) which could provide a payoff of $\$x$ by D . This is, of course, still arbitrary. However, the strategy is now uniquely identifiable by the simple fact that it can pay out an amount that no other strategy can pay out.⁷ The other 3 strategies are still completely symmetric. A consultant can now treat strategy D differently from the rest. If we assume that a consultant, being a game theorist after all, only recommends Nash equilibria but again takes all symmetries into account, we

⁷Thus minute differences in the payoff structure induce completely different symmetry structures, potentially giving rise to very different predictions. See Section VI for a discussion of how our symmetry structures do thus quite well in predicting play in the experiments documented in Crawford et al. (2008) with subtitle: “Even minute payoff asymmetry may yield large coordination failures.”

obtain the following possible recommendations: $(0, 0, 0, 1)$ (i.e. probability 1 on D), but also $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0)$ (equal mixing over the first 3 strategies) and even another one given by $(p, p, p, 1 - 3p)$, where p is chosen such that the opponent is exactly indifferent between all strategies (i.e. $p = \frac{x}{12+3x}$). In contrast to the previous case we now have multiple possible recommendations, a situation we expect to be typical. If we insist in further pinning down play, in this case we must rely on the existence of some meta-norm, e.g. picking the Pareto-dominant recommendation if there is one. For the two-player case, this would yield the recommendation of $(0, 0, 0, 1)$ if $x > 4$ and $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0)$ if $x < 4$.

Now consider a variant of the initial situation where there is additional information represented by a frame, i.e. strategies can come with labels or objectively distinct characteristics. Suppose, for instance, that one of the four \square 's is replaced by a \blacksquare . This situation is depicted in Figure 1(b). For the two-player case, in labeled matrix form we obtain Game 3 below.

	\square A	\square B	\square C	\blacksquare D
A	12,12	0,0	0,0	0,0
B	0,0	12,12	0,0	0,0
C	0,0	0,0	12,12	0,0
D	0,0	0,0	0,0	12,12

Game 3

	\blacksquare A	\blacksquare B	\square C	\square D
A	12,12	0,0	0,0	0,0
B	0,0	12,12	0,0	0,0
C	0,0	0,0	12,12	0,0
D	0,0	0,0	0,0	12,12

Game 4

Note that it is sufficient to label only one player's strategies. This game, in a way, is very similar to the second game. Here also, strategy D is identifiable, while the other three are not. Thus we obtain the same three predictions (with $p = \frac{1}{4}$ for the third one). The meta-norm of always choosing the Pareto-dominant recommendation when there is one would allow us to predict that D , i.e. object \blacksquare , will be chosen by all players. Note that all this can be done even if we have never seen objects \square and \blacksquare in our lives before. The simple fact that one is uniquely identifiable (the odd-man out as e.g. Binmore and Samuelson (2006) call it) allows us to predict coordination on it.

More interesting perhaps is the case represented as the framed Game 4, in which two \square 's are replaced by \blacksquare 's. Assuming that \square and \blacksquare are clearly different we could still have two possibilities. Suppose first that players (or consultants), while not familiar with the game exactly, are quite familiar with these objects \square and \blacksquare and they all agree (and know that they do) that in such cases \blacksquare comes before \square (or the other way round). Thus the possible recommendations in such cases from our point of view, assuming that we are game theorists without exact knowledge of whether \blacksquare comes before \square or the other way around, but we acknowledge the fact that there might well be such a ranking, would be $(\frac{1}{2}, \frac{1}{2}, 0, 0)$, $(0, 0, \frac{1}{2}, \frac{1}{2})$, and $(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$. If however there is no such universally accepted ranking of \square and \blacksquare , i.e. the frame is unfamiliar to the players, our only possible prediction would again be $(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$ as in the very first case. This is because the two labels are undistinguishable on the basis as to how often they appear. So no-one can tell whether one should go for \square or \blacksquare , and, hence, mis-coordination is again rather more likely.

II. 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 $u : S \mapsto \mathbb{R}^I$ 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 will 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. Building on these conventions, whenever we want to single out an opponent j of i , we will also write $u_i(s) = u_i(s_i | s_j, s_{-i,-j})$, where $s_{-i,-j} \in S_{-i,-j} = \times_{k \neq i,j} S_k$, and analogously for $u(s)$. 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 $i = 1, 2$. 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 Janssen (2001, Propositions 1 and 2) and Casajus (2000, Theorem 5.6) are restricted to matching games.

In order to avoid having to write down large payoff matrices it is convenient to use a simplified representation of such games of pure coordination. A game of pure coordination with k strategies $S_1 = S_2 = \{s^1, \dots, s^k\}$ shall be denoted by $\text{Diag}(u(s^1, s^1), u(s^2, s^2), \dots, u(s^k, s^k))$ with the implicit understanding that all off-diagonal entries in the payoff-matrix are zero.

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 V for the case when labels are unfamiliar.

III. 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, which is closely related to the one used by Crawford and Haller (1990), 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. These constraints 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 IV, 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.

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

Game 5

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

Game 6

Let us call “player 1” the one choosing rows, and “player 2” the one choosing columns. Clearly, it is impossible to distinguish between both of player 1’s strategies in Game 5 and one should, hence, call the two strategies symmetric; actually, one can say that they are **duplicates**. This gives rise to the simplest and strongest notion of symmetric strategies.

DEFINITION 1: *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 $s_{-i} \in S_{-i} = \times_{j \neq i} S_j$.*

Strategies A and B of player 1 in Game 5 are clearly duplicates, but most games do not have many duplicates. In fact, when we write down the reduced normal form derived from an extensive form game we omit duplicates.

Consider now Game 6 (which is just M_2). Exactly as in Game 5, other than the arbitrary name tags “A” and “B”, and the arbitrary fact that the same tags have been 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 he has written down as “A” might have been called “B” by a consultant advising his client’s opponent.

To reflect this observation, a weaker notion of symmetry is given below. To state this definition 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 $s \in S$, it will prove convenient to introduce the notation $\rho_{-i}(s_{-i}) = (\rho_j(s_j))_{j \neq i}$.

DEFINITION 2: *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 $s_{-i} \in S_{-i}$.*

Of course duplicates are weakly symmetric with the identity functions as relabelings. The converse is not true. Strategies A and B in Game 6 are not duplicates, but they are weakly symmetric for both players (just consider the relabeling given by $\rho_{-i}(A) = B$ and $\rho_{-i}(B) = A$).

Given the simple examples above, defining symmetry of, say, strategies seems a simple matter of identifying whether two strategies obtain the same payoff vector after some relabeling of opponents’ strategies. This, however, 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 7

A relabeling of player 2’s strategies L and R in Game 7 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 following three-player game shows.

	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 8

The transposition of L and C for player 2 in Game 8 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 L and C symmetric in a stronger sense.

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 identify (tell apart) strategies L and C for player 2 (e.g. 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 R 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 3: 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

$$(1) \quad u(s_i | s_{-i}) = u(s'_i | \rho_{-i}(s_{-i}))$$

for all $s_{-i} \in S_{-i}$ (where $\rho_{-i}(s_{-i}) = (\rho_j(s_j))_{j \neq i}$). 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 for player i .

This definition corresponds to the 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 for a player symmetric, the condition $\rho_j(T_j) = T_j$ restricts us to relabelings which only exchange strategies of other players within the same symmetry class of those players (in Game 7, 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 trivially a pairwise symmetry structure, albeit not necessarily the most interesting one. 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 that game. Let $f : U^\Gamma \rightarrow \mathbb{R}^n$ be an arbitrary mapping. Then 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 . This means we can, for instance, transform the rock-scissors-paper game 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, in fact, 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 that 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 1 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 Γ has a unique coarsest pairwise symmetry structure \mathcal{T}^* .*

One remark is in order. 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 9

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.

B. 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 II.

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 4: 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 10

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 structures, 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.

C. 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. Throughout we shall fix a game Γ and a (pairwise) symmetry structure thereof. Barring additional information on the consultant's part other than the game itself, we will typically think of the symmetry structure as the coarsest one. If the game is provided with a frame, as explained in the previous subsection we are provided with a different symmetry structure. The results below hold true independently of the symmetry structure, hence for games with and without frames.

A recommendation of a consultant is simply a mixed strategy profile $x \in \Theta$.

AXIOM 1: *A recommendation $x \in \Theta$ is **rational** if it constitutes 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 5: *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 7 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 $x_2 = (\frac{1}{2}, \frac{1}{2})$, but the best response of player 1 against such a strategy is T. Hence there exists no Nash equilibrium of this game where M and B are treated equally.⁸

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

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 just the set of all Nash equilibria, while rational recommendations which are strategy-symmetric with respect to the coarsest structure are also strategy-symmetric with respect to any structure.

This raises an interesting point. Suppose we have a framed game and its associated symmetry structure, 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.

In particular, this delivers a reinterpretation of the approach in Crawford and Haller (1990). In that paper, 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, they could as well coordinate on the other strategy from then on, which also enables coordination. However, the fact that both are possible should in fact make it difficult, if not impossible for players to actually achieve coordination. Thus coordination requires us to appeal to a more subtle meta-norm, for instance, that already chosen strategies are more salient. We shall discuss these issues of meta-norms in more detail in section VI.

IV. Global Symmetry Structures

In this section, we tackle the two difficulties advanced in Section III 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 III. 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.

A. Preliminary Concepts and Examples

We have already discussed the rationality axiom and the axiom of equal treatment of symmetric strategies. Below we will introduce a further axiom requiring symmetric players to be treated equally in the

consultant's recommendation. In order to be able to spell out this axiom, we need 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 6: *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.*

Game 6 above (and, in general, any matching game) is strongly symmetric, thus it is clear that 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 the following (still very simple) games:

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

Game 11

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

Game 12

Game 11 (which is a version of the Battle of the Sexes) 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. The consultant can certainly distinguish among both strategies. Strategy A is the one such that, if played by both player 1 and player 2, will lead to the pure-strategy equilibrium most preferred by player 1. Likewise, strategy B is the one such that, if played by both player 1 and player 2, 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 his 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 his 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 its client to play the strategy leading to his or her 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. In this paper 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.

Game 12 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 11 are symmetric in the sense that, through reordering of the strategy sets, each player can be taken to face the same payoff table. We will 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, in fact, here only leaves 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 6 is too restrictive (hence the “strong” adjective). Under the interpretation that strategy labels have no intrinsic meaning, however, we can expand the definition of player symmetry. The new definition depends on the given (pairwise) symmetry structure \mathcal{T} of the game Γ . If we take the coarsest structure \mathcal{T}^* , we obtain a definition which relies only on the payoff structure. If we consider a particular frame L , we have to refer to the corresponding symmetry structure $\mathcal{T}(L)$. Consider the following provisional definition.

DEFINITION 7: *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

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

As we will discuss below, the tuple $(\sigma, (\tau_i)_{i \in I})$ used in the last definition is related to the concept of strategic form automorphisms used by Harsanyi and Selten (1988), which in turn are a generalization of the concept of symmetry introduced by Nash (1951). A key difference that should be observed at this point is that we require the transformations τ_i to respect the structure derived from the symmetry structure. This is crucial, because in general we might deal with alternative symmetry structures.

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 5 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 11) the unique symmetry structure is such that no strategies are considered symmetric. The two players are not symmetric according to Definition 6, but are symmetric according to Definition 7. 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 12). 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 $\tau_1(\tau_2(H_2)) = H_1$ 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.

Note that, indeed, in general the definition of symmetric players depends on the (strategy) symmetry structure of the game. Consider the following (admittedly somewhat unexciting) game.

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

Game 13

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 $\{R\}$.

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

Equal treatment of symmetric players, to be stated below, will simply require a consultant's recommendation to treat symmetric players equally. This is well-defined for the examples above, e.g. in Game 11 (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. The definition of when two players i, j are symmetric relies on a particular mapping $\tau_i : S_i \mapsto S_j$. It is easy to see that, in general, 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 the axiom of equal treatment of symmetric strategies already prescribes that all strategies are treated equally. It might be the case, however, that two players can be seen as symmetric in two qualitatively different ways, giving rise to certain subtleties. The (framed) Game 14 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 14

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”.

B. Simultaneous Determination of Symmetric Strategies and Symmetric Players

There is, however, a more serious problem with the approach followed until now. Pairwise symmetry structures, as defined above, ignore symmetries between players. Symmetric players, as just defined, extract the symmetry of players from the symmetry structure. Game 15 below shows that this approach misses some possible interactions between both concepts. 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.

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

Game 15

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

DEFINITION 8: 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 $s = (s_k | s_{-k}) \in S$,

$$(2) \quad u_k(s_k | s_{-k}) = u_{\sigma(k)}(\tau_k(s_k) | \tau_{-k}(s_{-k})).$$

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 10).

As commented above (recall Game 14), 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 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 . 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 10: 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

- for every $J \in \mathcal{I}$, $\sigma(J) = J$;
- for every $k \in I$ and $T_k \in \mathcal{T}_k$, there is a $T_{\sigma(k)} \in \mathcal{T}_{\sigma(k)}$ such that $\tau_k(T_k) = T_{\sigma(k)}$ and $\alpha(T_k) = \alpha(T_{\sigma(k)})$.

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

We are now ready to present our definition.

DEFINITION 11: 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 pair of distinct strategies $s_i, s'_i \in T_i$, there exists a symmetry (σ, τ) which agrees with $(\mathcal{I}, \mathcal{T}, \alpha)$ such that $\sigma_i(i) = i$ and $\tau_i(s_i) = s'_i$.

¹⁰Recall footnote 9.

- (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 7 because one is not allowed to swap the asymmetric strategies L and R . Careful consideration reveals that, in this framework, this is already prevented by the fact that there exists no symmetry (σ, τ) of Game 7 which swaps M and B . This is because condition (2) must apply to all strategies swapped in the symmetry, and not only (as in the original definition of symmetric strategies) 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, especially in games with duplicates. That is, in general the partition part of a global symmetry structure might have several different compatible player identifications.

Consider again the framed Game 14 (the same point can be made with M_2). We have a global symmetry structure 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. We see that 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 $\tau_1(D) = H$. However, they 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 $\tau'_1(D) = F$. 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 itself and $\tau^*(A) = C$. 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 labels of player 1 and player 2 in two different ways, which leads to the existence of 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.

In summary, given a global symmetry structure, there might exist several player identifications compatible with it (with the obvious abuse of language here). This phenomenon is typical of games with duplicates but not restricted to it. Indeed, if we consider Game 6 (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 will indeed prove below), and thus is of no relevance for Nash (1951). However, as long as 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 . This is obviously the case if no players are symmetric at all. 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.

C. 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 10 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. In other words,

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 15 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 16

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 is given by $\tilde{\mathcal{T}}_1 = \{\{a_1, a_2, a_3\}, \{a_4, a_5, a_6\}\}$ 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 indeed 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 globally symmetric 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 16 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 \hat{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{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.

D. 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 of course 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, and then it is just a matter of showing that this structure can be translated to global symmetry structures (technically, one obtains a lattice-isomorphism). This leads to the following theorem.

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 14. 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. Due to the lattice structure, 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. However, this join, say $(\mathcal{I}^*, \mathcal{T}^*, \alpha^*)$, might incorporate new symmetries not captured in $(\mathcal{I}, \mathcal{T})$.

¹¹A similar example can be constructed with a trivial game as e.g. Game 13.

DEFINITION 12: Let $(\mathcal{I}, \mathcal{T}, \alpha)$ be a global symmetry structure. The **completion** of $(\mathcal{I}, \mathcal{T}, \alpha)$ 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 14, 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

We would like to point to an interesting connection with the work of Blume (2000), who uses a group-theoretic approach to formalize partial languages. The latter are defined as subgroups of permutations 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.

E. Global Symmetry Structures and Extended Frames

Analogously to Section III.B, we can give a correspondence between global symmetry structures and extended frames, where not only strategies, but also players are labeled (e.g. players 1,3, and 17 are men, the rest are women). Again, the coarsest global symmetry structure delivers the strongest (coarsest) reclassification of strategies and players that a consultant can obtain from the game, based on payoffs alone. Hence, this structure is “frameless”.

DEFINITION 13: 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 14: 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 $\{L_0^{-1}(z) \mid z \in \mathcal{Z}_0\}$, $T_i(L)$ 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 III.B, 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.

¹²This explains why 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.

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.

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.

F. 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 14 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 15: *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 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.

V. 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 attempt to provide a fully fledged theory of unfamiliar frames and

their induced symmetry structures, but rather restrict ourselves to a series of observations made through examples. We will exclusively use 2-player matching games, M_k , as these are the simplest games in which the points we want to make can be made. We shall assume throughout that the strategy spaces are denoted by the set of the first k letters of the roman alphabet. Since we are dealing with matching games, we shall 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 the examples in this section we will write \mathcal{T} for the partition of any one of the two players in the considered symmetry structure.

Recall the version of the motivating example given as Game 4. The game is a matching game (except payoffs are multiplied by 12), 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. For instance \blacksquare might be more salient than \square . 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 Game 4, in which players could coordinate on one label versus another and the induced symmetry structure would have to be the coarsest one. That is, the symmetry structure induced by an unfamiliar frame is quite different (always 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 (see for instance Game 3).

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 $\mathcal{T}_i^* = \{S_i\}$, $i = 1, 2$, 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 $T_i, T'_i \in \mathcal{T}_i$ such that $|T_i| = |T'_i|$ then we must have that there is a $\tilde{T}_i \in \mathcal{U}_i(\mathcal{T})$ with $T_i \cup T'_i \subset \tilde{T}_i$. 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 first Game 6, i.e. the M_2 matching game with two players and two strategies. Consider 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. This means that both 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 the property that $\mathcal{U}(\mathcal{T}) = \mathcal{T}'$.

For the M_2 matching game we thus find that whatever the frame, if it is unfamiliar, it does not refine the coarsest symmetry structure of the game. This is, however, not always the case. Consider the 2-player matching game M_3 with three strategies. Consider all possible frames (for only one player). 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 (three) remaining frames, two strategies receive the same label, while the third one receives a different label. Now this different 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 of course external to the considered symmetry structure. To make this point clear, consider the matching game with four strategies, M_4 , and suppose we are given an 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 $\mathcal{U}(\mathcal{T})$ could 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 the 2-player matching game with five strategies, M_5 . Suppose we have two frames. Frame L_1 labels strategies A and B as \odot , strategies C and D as \oslash , 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 $\mathcal{U}(\mathcal{T}(L_1)) = \mathcal{U}(\mathcal{T}(L_2))$. 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 separately.

Finally, consider the matching game M_6 , with six strategies. Consider 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 $\mathcal{T}(L_2) = \{\{A, B, F\}, \{C, D, E\}\}$ which gives rise to $\mathcal{T}(L_1) \wedge \mathcal{T}(L_2) = \{\{A, B\}, \{C\}, \{D, E\}, \{F\}\}$. 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 $\mathcal{U}(\mathcal{T}(L_1)) \wedge \mathcal{U}(\mathcal{T}(L_2)) = \mathcal{T}^*$. However, the symmetry structure induced by the two unfamiliar frames together is given by $\mathcal{U}(\mathcal{T}(L_1) \wedge \mathcal{T}(L_2)) = \{\{A, B, D, E\}, \{C\}, \{F\}\}$. We thus have that here $\mathcal{U}(\mathcal{T}(L_1) \wedge \mathcal{T}(L_2)) \prec \mathcal{U}(\mathcal{T}(L_1)) \wedge \mathcal{U}(\mathcal{T}(L_2))$. I.e. 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.

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.

VI. Focal Points

In this section we try to use the, so far normative, theory of rational symmetric recommendations, as a positive or descriptive theory of play in finite normal form games with symmetries, and compare our predictions with recent experimental work.

A. A Definition of Focal Points

Our focal points are special rational symmetric recommendations resulting from the appropriate global symmetry structure induced by the framed game at hand. If there is a unique rational symmetric recommendation we have a much sharper prediction than that of players using rationalizable strategies (Bernheim (1984) and Pearce (1984)) derivable from the assumption of common knowledge of rationality. For instance, in Game 1 we obtain the prediction of $\frac{1}{4}$ probability for each of the four strategies, while the set of rationalizable strategies is the set of all mixed strategy profiles. In the more likely case of multiple rational symmetric recommendations we can not immediately expect players to be able to coordinate on one of them. Hence, it seems that we are back to predicting rationalizable strategies or some such set.

However, even if the game and its frame are completely unfamiliar to players, this does not have to be the case for every aspect of the game. A readily understood aspect of a game is that it provides payoffs, be it in monetary or utility terms. If players have experienced money (or utility) often enough before, they should be able to come up with a common salience ranking over tuples of monetary payoffs (or utility) to the various players. Individuals might well, for instance, consider more money more salient. Hence, in games with a unique Pareto-efficient rational symmetric recommendation players might well find this recommendation particularly salient. A similar argument could be made for risk dominance.¹³ Another money-based salience ranking might favor rational symmetric recommendations that are particularly equitable. In any case we expect that if many different games were played recurrently over a long time by many different people, eventually **meta-norms** will emerge telling all players what the appropriate salience ranking is in the particular game at hand, even if the game itself has never been seen before.

A **meta-norm** can be defined to be a mapping from all familiar aspects of a game into e.g. the real line, that is, defining a ranking which can be translated as **salience**. We refrain from providing a general definition formally here. Such a definition would, by necessity, be notationally involved, but in many contexts a simplified version suffices. For instance in matching games the meta-norm of choosing the Pareto-dominant rational symmetric recommendations is particularly innocuous. The meta-norm of choosing the risk-dominant rational symmetric recommendation coincides with it, while the meta-norm of choosing the most equitable one has no bite in these games. Indeed the meta-norm of Pareto-dominance has appeared a lot in the literature, especially for matching games (recall footnote 4).

For some games meta-norms will have to be a more general function into a salience ranking not only from payoff tuples, but also from other familiar aspect 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.¹⁴

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 V. 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 meta-norm that dictates them to choose, in such games, one of the two strategies which have a distinct label, i.e. they would choose $(0, 0, \frac{1}{2}, \frac{1}{2})$. Of course, an alternative meta-norm could prescribe to choose one of the two strategies which have the same label, i.e. choose $(\frac{1}{2}, \frac{1}{2}, 0, 0)$.

In any case, meta-norms can be very intricate, and we do not expect that a unique meta-norm has emerged so that players can achieve 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

¹³In an experimental study of 2×2 coordination games, Straub (1995) finds more support for risk-dominance than for Pareto-dominance. Individuals played the game recurrently for 9, 10, or 18 rounds and often eventually learned to play the risk-dominant strategy. In fact, many individuals played the risk-dominant strategy even in the first round. Since they knew that they would play the game again, albeit with different players, this still differs from a true one-shot game.

¹⁴The two experimental studies by Mehta, Starmer and Sugden (1994a, 1994b) generally explore the nature of salience (of labels in familiar frames) in matching games exclusively. Their studies are nicely complementary to this paper as they provide an investigation of what determines salience of different labels, something this paper is agnostic about.

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. Given such a commonly known meta-norm, we can define the concept of a focal point.

DEFINITION 16: *For a given game and frame, whether unfamiliar 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.*

The “symmetry” part of this definition corresponds to our three axioms, and, as our analysis shows, can be made as formally precise as one wishes. The “meta-norm” part of the definition is necessarily loose in that the details of what the meta-norm is might vary greatly for identifiably different games. In principle, one could write down a complete meta-norm such that applying it to all rational symmetric recommendations (given the coarsest global symmetry structure induced by the game’s frame) leaves us with exactly one of these rational symmetric recommendations for every possible game, which by definition is a Nash equilibrium. One could write a book detailing this 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. For a descriptive theory, these books would probably not be of much use as yet, as individuals would first have to be taught the normative theory of how one should behave or would simply need a lot of time to form such a meta-norm. In the meantime we will have to contend, for positive purposes, with partially specified meta-norms and the resulting room for, by necessity, at least partially behavioral explanations put forward in the literature.

B. Examples from the Experimental Literature

We conclude this section with a few examples from the experimental literature which demonstrate how focal points fare as a positive theory of individuals’ behavior.

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?”

Let us suppose that one can only specify amounts in multiples of one dollar. Then each player has 101 strategies. None of these strategies are symmetric in any symmetry structure of this game 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 would have to be treated equally. This, in fact, leaves only a single rational symmetric recommendation, in which each claim exactly \$50. Note that this is without any appeal to a meta-norm. There are many other Nash equilibria of the game, but in all of these both players claim different amounts (adding up to \$100). Indeed Schelling (1960) reports that players do claim mostly exactly \$50.

Another example is given by Blume and Gneezy (2000), who have subjects play this game in a very carefully designed experimental setting. We here discuss two of their four treatments. Each of two players is given a plastic disk (a physical object) with three (or nine) sectors of exactly equal shape and size in 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 player, 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 the two players in the first stage.¹⁵ Assume that players chose sectors such that there are 2

¹⁵If both players choose the same sector in stage one, the analysis is similar but easier.

(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 some clockwise order. As Blume and Gneezy (2000) point out there is a unique sector with a unique label. This has to do with the fact that there is both an even (2) and an odd (5) number of sectors between the chosen sectors. Consider first the two chosen sectors. Both players only know that they have been chosen in the first stage, but they do not even know by whom. These two sectors must receive the same label "chosen in stage 1". The two sectors between them on the one side also must receive identical labels of "adjacent to a chosen sector in the even section". Then there are two sectors "adjacent to a chosen sector in the odd section" and another two with the label "one apart from a chosen sector". But this leaves a unique sector "in the middle between the chosen sectors (on the odd side)". The symmetry structure is thus $\mathcal{T}_1 = \mathcal{T}_2 = \{\{1, 4\}, \{2, 3\}, \{5, 9\}, \{6, 8\}, \{7\}\}$. If we choose recommendations according to 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 might just be 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). In another interesting experiment, however, Blume and Gneezy (2008) show that this is not necessarily due to players not recognizing that there is a unique label, but also due to players fearing that their opponent does not realize this. Thus, in such cases play will be more readily explained by 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).

The framed games experimentally investigated in Crawford et al. (2008) are similar to the ones in Blume and Gneezy (2000) but simpler. We shall look here at the second set of games in that paper, called "pie games". The discussion for their "X-Y games" would be very similar. All games in Crawford et al. (2008) are games of pure coordination with three strategies, say A, B, C, and were played only once by each subject. They all share the same frame L with labels $\blacksquare, \blacksquare, \square$, i.e. A and B always receive the same label. In the experiment what we call \blacksquare and \square here, in fact, denote the labels "colored" and "uncolored sector" of the pie, respectively. Some effort went into making sure that these are the labels. There are 8 pie games in total, which we will 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. The names of the games are taken from Crawford et al. (2008), where S stands for symmetric, A for asymmetric, M for moderate, and L for large.

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. The meta-norm of Pareto-efficiency (also that of risk-dominance or equity) yields the focal point C. Indeed 94% of subjects play C.

Games 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 11), or coordinating on C. The latter then becomes a focal point (with respect to the meta-norms of Pareto-efficiency, equity, as well as risk-dominance). Indeed in the experiment roughly 80% and 90% of players choose strategy C in games AL1 and AM1, respectively.

Game S2 is the game of pure coordination $\text{Diag}((6, 6), (6, 6), (5, 5))$. The symmetry structure is the same as for game S1, and again the meta-norm of Pareto-efficiency (also that of risk-dominance or equity) provides us with the focal point C. The experimental results for this game are perhaps the most surprising of all 8 pie-games in Crawford et al. (2008). 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 here. On the one hand they find colored slices (here label \blacksquare) more salient than uncolored ones and on the other they find Pareto-optimal rational symmetric recommendation more salient. Somehow they cancel each other out.

Game 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

■, while B is identified through its unique payoff possibility of 5. Thus, also A is identified. Symmetric rational recommendations are here thus all symmetric Nash equilibria of the game. The game has no focal points according to neither of the three meta-norms of Pareto-efficiency, equity, or risk-dominance. To us this seems a hard problem for subjects. The result is a frequency of choices given by 14%, 21%, and 64%, respectively. This is a little more systematic than complete confusion, which we would not have been surprised by here. Ex-post one could rationalize the majority of choices by assuming players use the meta-norm of Pareto-efficiency in conjunction with the meta-norm that uncolored is more salient than colored, which somewhat contradicts the ex-post explanation of results in game S2.

Game AM2¹⁶ is the game of pure coordination given by $\text{Diag}((5, 6), (6, 5), (6, 5))$. The symmetry structure here has both players identified, $\mathcal{I} = \{\{1\}, \{2\}\}$, and the strategy symmetry structure is the trivial one. This is so because strategy C is identified by its label □, and player 1 is identified through the fact that only he gets a payoff of 6 when two Cs are chosen. Then A is identified by the fact that A gives player 1 a different payoff than strategy B does. Thus, also B is identified. This game, much like game S3 has no focal points with respect to any of the three meta-norms of Pareto-efficiency, equity, or risk-dominance. Again this is a hard problem for subjects. The result in the experiment is indeed confusion (for details see Crawford et al. (2008)). Players could hold a common meta-norm here, which would allow them to coordinate on one symmetric rational recommendation, but all such meta-norms are somewhat intricate and, without having seen games like these frequently, it is not at all obvious what it should be. One could, for instance, hold the meta-norm that in such cases players should go for the strategy with an outcome (unrankable in Pareto-efficiency, risk-dominance, and equity) which is more beneficial for the player who does not so well in the uniquely labeled strategy. This would make players choose A. But many other meta-norms are possible.

VII. 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**, be it of players or strategies, 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**.

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

¹⁶The analysis and results for games AM3 and AM4 are essentially identical to this one.

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. Proofs

A. 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 simply 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.*

B. Proofs from Section III

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

LEMMA 1: *Let S_i be a finite set.*

- (a) *Let T_i and T'_i be partitions of S_i . If T'_i is coarser than T_i , then every set $T_i \in \mathcal{T}_i$ can be written as a (finite) disjoint union of the sets which form T'_i .*
- (b) *The finest partition coarser than two partitions T_i and 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 T'_i or 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 T''_i be a partition coarser than both \mathcal{T}_i and \mathcal{T}'_i . Let $T''_{i,\ell}$ 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

$$T''_{i,\ell} = \bigcup_{r=1}^{\ell} T_{i,r} = \bigcup_{r=1}^{\ell'} T'_{i,r}.$$

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

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

Let $T_i \in \mathcal{T}_i$, and let $s_i, s'_i \in T_i$. We need to show that there is a set of 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 1(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 1(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.

The fact that \mathcal{T} is the finest symmetry structure which is coarser than both \mathcal{T}' and \mathcal{T}'' follows by construction and the definition of the relation “to be coarser than”. Hence \mathcal{T} is the join of \mathcal{T}' and \mathcal{T}'' . We have thus shown that 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. The trivial symmetry structure is clearly a bottom element, and hence Fact A3 implies that symmetry structures form a lattice.

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

PROOF OF THEOREM 3: Let $\beta_i : \Theta_{-i} \rightarrow \Theta$ 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 $x_i(s_i) = x(s'_i)$ 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 (1) holds. Then, making extensive use of the $-i$ notation for product spaces,

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

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. This completes the proof.

C. 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 “ \cdot ” in G satisfying the associative property $((g_1 g_2) g_3 = g_1 (g_2 g_3))$ for all $g_1, g_2, g_3 \in G$, with an identity element ($1_G \in G$ such that $1_G g = g 1_G = g$ for all $g \in G$) and such that every element $g \in G$ has an inverse $g^{-1} \in G$ according to this operation ($g^{-1} g = g g^{-1} = 1_G$). 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 $g_1 g_2^{-1} \in H$ for every $g_1, g_2 \in H$.*

The set of groups of a subgroup have a lattice structure. The meet is particularly 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 $\ell = 1, \dots, r$.*

D. Proofs from Section IV

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 A1: *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 A4.¹⁷

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.

¹⁷If 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 14 shows.

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 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 equivalent¹⁸ 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 . The rest of the statement follows now by definition.

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

LEMMA A2: 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 A2.

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.

Define $\hat{\Theta}$ to be the subset of $\prod_{i \in I} \hat{\Theta}_i$ 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 $\alpha_i(T_i) = \alpha_j(T_j)$. It follows from definition of player symmetry that if y_i is a best response to $x \in \hat{\Theta}$, then 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, we have that $\hat{\beta}$ is convex- and upper hemicontinuous.

It remains to show that $\hat{\beta}$ is nonempty-valued. All that we need to show is that 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 (which hence use symmetric strategies in x). 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 (which is a Nash equilibrium) exists.

¹⁸To see the sufficiency, recall that the composition of any symmetries in Φ is also in Φ because Φ is a subgroup.

References

- Aumann, R. and A. Brandenburger**, “Epistemic Conditions for Nash Equilibrium,” *Econometrica*, 1995, *63*, 1161–1180.
- Bacharach, M.**, “Games with Concept-Sensitive Strategy Spaces,” 1991. Mimeo.
- , “Variable Universe Games,” in A. Kirman K. Binmore and P. Tani, eds., *Frontiers of Game Theory*, MIT Press, Cambridge, MA, 1993, p. 255276.
- Bernheim, B. D.**, “Rationalizable Strategic Behavior,” *Econometrica*, 1984, *52*, 1007–29.
- Binmore, K. and L. Samuelson**, “The Evolution of Focal Points,” *Games and Economic Behavior*, 2006, *55*, 21–42.
- Blume, A.**, “Coordination and Learning With a Partial Language,” *Journal of Economic Theory*, 2000, *95*, 1–36.
- and **U. Gneezy**, “An Experimental Investigation of Optimal Learning in Coordination Games,” *Journal of Economic Theory*, 2000, *90*, 161–172.
- and —, “Cognitive Forward Induction and Coordination Without Common Knowledge: An Experimental Study,” 2008. mimeo.
- Casajus, A.**, “Focal Points in Framed Strategic Forms,” *Games and Economic Behavior*, 2000, *32*, 263–291.
- , *Focal Points in Framed Games*, Springer, 2001.
- Crawford, V. and H. Haller**, “Learning How to Cooperate: Optimal Play in Repeated Coordination Games,” *Econometrica*, 1990, *58*, 571–596.
- Crawford, V. P., U. Gneezy, and Y. Rottenstreich**, “The Power of Focal Points is Limited: Even Minute Payoff Asymmetry May Yield Large Coordination Failures,” *American Economic Review*, 2008, p. forthcoming.
- Davey, B.A. and H.A. Priestley**, *Introduction to Lattices and Order (2nd Ed)*, Cambridge: Cambridge University Press, 2002.
- Gauthier, D.**, “Coordination,” *Dialogue*, 1975, *14*, 195–221.
- Goyal, S. and M. Janssen**, “Can We Rationally Learn to Coordinate?,” *Theory and Decision*, 1996, *40*, 29–49.
- Grätzer, G.**, *General Lattice Theory (2nd Ed)*, Basel: Birkhäuser, 2003.
- Harsanyi, J.C. and R. Selten**, *General Theory of Equilibrium Selection in Games*, Cambridge, MA: MIT Press, 1988.
- Hungerford, Thomas W.**, *Algebra*, New York: Springer Verlag, 1980.
- Janssen, M. C. W.**, “Rationalizing Focal Points,” *Theory and Decision*, 2001, *50*, 119–148.
- Mehta, J., C. Starmer, and R. Sugden**, “Focal Points in Pure Coordination Games: An Experimental Investigation,” *Theory and Decision*, 1994, *36*, 163–185.
- , —, and —, “The Nature of Salience: An Experimental Investigation of Pure Coordination Games,” *American Economic Review*, 1994, *74*, 658–673.
- Nash, J.**, “Non-Cooperative Games,” *Annals of Mathematics*, 1951, *52*, 286–295.
- Pearce, D. G.**, “Rationalizable Strategic Behavior and the Problem of Perfection,” *Econometrica*, 1984, *52*, 1029–51.
- Rose, John S.**, *A Course in Group Theory*, Cambridge, UK: Cambridge University Press (Dover Reprint 1994), 1978.
- Schelling, Thomas C.**, *The Strategy of Conflict*, Cambridge: Harvard University Press, 1960.
- Straub, P. G.**, “Risk Dominance and Coordination Failures in Static Games,” *The Quarterly Review of Economics and Finance*, 1995, *35*, 339–363.
- Sugden, R.**, “A Theory of Focal Points,” *Economic Journal*, 1995, *105*, 533–550.
- Tversky, A. and D. Kahneman**, “The Framing of Decisions and the Psychology of Choice,” *Science*, 1981, *211*, 453–458.

**THURGAU INSTITUTE
OF ECONOMICS**
at the University of Konstanz

Hauptstr. 90
CH-8280 Kreuzlingen 2

Telefon: +41 (0)71 677 05 10
Telefax: +41 (0)71 677 05 11

info@twi-kreuzlingen.ch
www.twi-kreuzlingen.ch