CHAPTER FIVE

REFLEXIVE FRAMEWORKS

In this chapter we develop a framework, called a reflexive observer framework, in which the objects of perception of an observer O are themselves observers having the same X, Y, E, and S as has O. We display the relationship between reflexive observer frameworks and environments for extended semantics. We illustrate the definition of reflexive observer framework with several examples.

1. Mathematical notation and terminology

The examples of reflexive observer frameworks given in this chapter make use of several mathematical concepts from group theory. In this section we collect basic terminology and notation for the convenience of the reader.[1]

A ***topological group*** G is a group that is also a topological space and satisfies (i) the map $G \to G$ which sends every element to its inverse is a homeomorphism and (ii) the map $G \times G \to G$ describing the group operation is continuous. A ***measurable group*** G is a group that is also a measurable space, such that the maps in (i) and (ii) above are measurable. Every topological group is also a measurable group, with respect to the measurable structure associated to the topology (cf. **2**–1).

If $\mathcal{H}$ is an equivalence relation on a set G, then the set of all equivalence classes is called the ***quotient set*** of G by $\mathcal{H}$ and is denoted by $G/\mathcal{H}$. The map $\pi: G \to G/\mathcal{H}$ which assigns to each $g \in G$ the equivalence class to which g belongs is called the ***canonical map***. If G is a topological space, then $G/\mathcal{H}$ has a canonical topology: the ***quotient topology*** is the finest topology on $G/\mathcal{H}$ which makes the canonical map π continuous. If G is a measurable space, then $G/\mathcal{H}$ has a canonical measurable structure: the ***quotient measurable structure*** is the largest σ-algebra on $G/\mathcal{H}$ which makes π measurable.

Let $(G, \cdot)$ be a group with subgroup $(H, \cdot)$, and a an arbitrary element

[1] For more background we suggest Gilbert (1976).

of G. The set $Ha = \{ha|\ h \in H\}$ is a ***right coset*** of H in G. The set $aH = \{ah|\ h \in H\}$ is a ***left coset*** of H in G. The relation of belonging to the same left coset is an equivalence relation on G; similarly for right cosets. A subgroup $(H, \cdot)$ of a group $(G, \cdot)$ is called a ***normal subgroup*** of $(G, \cdot)$ if $g^{-1}hg \in H$ for all $g \in G$ and $h \in H$. If $(H, \cdot)$ is a normal subgroup of $(G, \cdot)$, the left cosets of H in G are the same as the right cosets of H in G. In this case the set of cosets $G/H = \{Hg|\ g \in G\}$ has a natural group structure induced by that of G, i.e., $(Hg_1) \cdot (Hg_2) = H(g_1 \cdot g_2)$.

If $(G, \cdot)$ and $(H, *)$ are two groups, the function $f\colon G \to H$ is called a ***group morphism*** or a ***group homomorphism*** if $f(a \cdot b) = f(a) * f(b)$ for all $a, b \in G$. A bijective group morphism is called a ***group isomorphism.*** If $f\colon G \to H$ is a group morphism, then the ***kernel*** of f, denoted by $\mathrm{Ker} f$, is the set of elements of G that are mapped by f to the identity of H; $\mathrm{Ker} f$ is a normal subgroup of G.

A group $(G, \cdot)$ ***acts on the left*** on the set M if (1) there is a function $\psi\colon G \times M \to M$ such that, letting $gm = \psi(g, m)$, we have $(g_1g_2)m = g_1(g_2m)$ for all $g_1, g_2 \in G, m \in M$, and (2) $\imath m = m$ if $\imath$ is the identity of G and $m \in M$. (G ***acts on the right*** if condition (1) is replaced by $\psi(g_1g_2, m) = \psi(g_2, \psi(g_1, m))$; in this case we write $\psi(g, m) = mg$. All actions here are left actions unless otherwise stated.) If G acts on M, we say that M is a G-set. If M is a topological (respectively measurable) space, the action is said to be ***continuous*** (respectively ***measurable***) if for all $g \in G$ the map $m \mapsto gm$ is a continuous (respectively measurable) map from M to M. The set of elements of G that fix $m \in M$, i.e., $\{g \in G|\ gm = m\}$, is called the ***stabilizer*** of m and is denoted Σ_m; each stabilizer is a subgroup of G. If each $m \in M$ is stabilized only by the identity $\imath$ of G, we say that G acts ***faithfully*** on M. G ***acts transitively on*** M if for every $m_1, m_2 \in M$ there exists $g \in G$ such that $gm_1 = m_2$. M is a ***principal homogeneous space*** for G if G acts both transitively and faithfully on M. The set of all images of an element $m \in M$ under the action of a group G is called the ***orbit*** of m under G, and is denoted by Gm; $Gm = \{gm|\ g \in G\}$. The orbits are the equivalence classes for an equivalence relation on M; two elements of M are in this relation precisely when they are in the same orbit. The quotient set for this relation is therefore the set of distinct orbits; it is denoted M/G.

Let G act measurably on M. A measure μ on M is called G-invariant if, for every measurable set A of M, $\mu(A) = \mu(gA)$ for any $g \in G$. If G acts on X, and $E \subset X$, then E is an ***invariant subset*** for the action if $GE = E$.

2. Definition of reflexive observer framework

We now begin to study "participator dynamics" or, more properly, "participator dynamical systems on reflexive observer frameworks." The phrase ***reflexive observer framework*** refers to a structure for the set $\mathcal{B}$ of objects of perception and for the configuration map $\Phi: \mathcal{B} \to X$ of an environment (**4**–4.4). In this chapter we introduce reflexive observer frameworks and the subclass of ***symmetric observer frameworks***; we study this subclass because it is natural and mathematically tractable. Dynamics enters the picture in the next chapter. We will find that, in the context of this dynamics, the question of true perception can be treated in a principled manner; we discuss this in chapter eight. The dynamics underlies a general-purpose theory of interaction which is nondualistic and which employs a hierarchical analytic strategy (cf. **4**–5). X_t will appear as one aspect of this dynamics.

We begin with an observer $O = (X, Y, E, S, \pi, \eta)$. We want to construct a model of an environment $(\mathcal{B}, \Phi)$ for O, as per **4**–4. The nondualism of the model results, as stated before, from the assumption that ***the objects of perception are observers.*** We take $\mathcal{B}$ to be some set of observers whose X, Y, E and S are the same as that of O. Then Φ assigns to every such observer an element of X. A reflexive observer framework furnishes the relationship between an observer $B \in \mathcal{B}$ and the element $\Phi(B)$ as follows. There is given a map Π which assigns to each $e \in E$ a map from X to Y; thus for $e \in E$ we have

$$\Pi(e): X \to Y.$$

If $\Phi(B) = e \in E$, then we require that the perspective map of the observer B is $\Pi(e)$, so that in this case $B = (X, Y, E, S, \Pi(e), \eta)$ for some η. But if $\Phi(B) = x \notin E$, we require nothing. The word "reflexive" indicates that each $e \in E$ represents both a distinguished configuration and a set of observers which perceive the distinguished configurations of E—namely the set of those observers $B \in \mathcal{B}$ whose perspective is $\Pi(e)$. This set of observers is represented by the preobserver $(X, Y, E, S, \Pi(e))$.

Once this structure is in place, we notice that the original observer O plays no role other than to specify the X, Y, E and S. And for this purpose, any of the observers in $\mathcal{B}$ serve equally well. In fact, we think of a reflexive observer framework as providing an environment simultaneously for all the observers in $\mathcal{B}$; each of these observers has the same set of objects of perception, namely $\mathcal{B}$ itself. We now present these ideas formally.

Notation 2.1. Given two measurable spaces $(X, \mathcal{X})$ and $(Y, \mathcal{Y})$, we denote the measurable maps from X to Y by $\mathtt{Hom}(X, Y)$.

Definition 2.2. Let $(X, \mathcal{X})$ and $(Y, \mathcal{Y})$ be fixed measurable spaces. Let $E \subset X$ and $S \subset Y$ be measurable subsets. A ***reflexive observer framework on X, Y, E, S*** is an injective map $\Pi\colon E \to \mathtt{Hom}(X, Y)$ such that for each $e \in E$, $\Pi(e)$ is surjective, and $\Pi(e)(E) = S$.

Terminology 2.3. We denote a reflexive observer framework by

$$(X, Y, E, S, \Pi).$$

If Π has been fixed, we write $\pi_e = \Pi(e)$, so that we can use the notation

$$(X, Y, E, S, \pi_\bullet)$$

to represent the framework in this case; the subscript "•" represents a variable on E. In this way the reflexive family is displayed as a family of preobservers parametrized by E. X is called the ***configuration space of the framework.*** Y is called the ***premise space of the framework.*** E is called the ***distinguished configurations of the framework.*** S is called the ***distinguished premises of the framework.*** Sometimes we drop the word "observer" and use the expression ***reflexive framework.***

Π will have, in general, some additional structure. If, for example, X and Y are topological spaces and all the maps π_e are continuous, then Π might be continuous for some suitable topology on the set of continuous maps from X to Y. However such restrictions do not belong in the general definition.

We give a concrete example of a reflexive framework at the end of this section (and in the next section we present classes of formal examples). We first make more explicit the connection between reflexive frameworks and environments.

As we have seen, a reflexive framework identifies each distinguished configuration $e \in E$ with a perspective. The notation $(X, Y, E, S, \pi_\bullet)$ of 2.3 suggests another way to interpret reflexive frameworks. Suppose we begin with a set $\mathcal{B}$ of observers all of which have the same X, Y, E, S. $\mathcal{B}$ might be, for example, the set of ***all*** observers with these X, Y, E, and S (and with arbitrary perspective maps and conclusion kernels). Then if $e \in E$ is given, we can interpret the notation (X, Y, E, S, π_e) to mean the subset of $\mathcal{B}$ consisting of all those observers whose perspective map is π_e for that particular e. If we want to identify $\mathcal{B}$ explicitly in this notation we write $(\mathcal{B}; X, Y, E, S, \pi_e)$ or just $\mathcal{B}_e$. The elements, if any, of this set $\mathcal{B}_e$ are individuated only by their conclusion kernels. Let $\mathcal{B}_E$ denote the subset of $\mathcal{B}$ consisting of those observers whose perspective map is one of the π_e's, so that

$$\mathcal{B}_E = \bigcup_{e \in E} \mathcal{B}_e. \tag{2.4}$$

Then $(X, Y, E, S, \pi_\bullet)$ denotes the partition of $\mathcal{B}_E$ into the sets $\mathcal{B}_e = (X, Y, E, S, \pi_e)$ for $e \in E$. If we must make $\mathcal{B}$ explicit we write $(\mathcal{B}; X, Y, E, S, \pi_\bullet)$ or just $\mathcal{B}_\bullet$. We summarize:

2.5. Let (X, Y, E, S, Π) be a reflexive framework. An alternate notation for the framework is $(X, Y, E, S, \pi_\bullet)$. Let also be given a set of observers $\mathcal{B}$, all having the same X, Y, E, and S. Then we can interpret the notation $\mathcal{B}_\bullet = (X, Y, E, S, \pi_\bullet)$ to mean the partition of the subset $\mathcal{B}_E$ of $\mathcal{B}$ into the sets $\mathcal{B}_e = (X, Y, E, S, \pi_e)$, $e \in E$. This context fixes a meaning for the preobserver (X, Y, E, S, π_e): it is a particular set of observers—namely $\mathcal{B}_e$.

We can now state formally the basic connection between reflexive frameworks and environments.

Definition 2.6. Let $(\mathcal{B}, \Phi)$ be the environment of an extended semantics for an observer $O = (X, Y, E, S, \pi, \eta)$, where $\mathcal{B}$ is a set of observers all having the same X, Y, E, and S as O, and where Φ is the configuration map of the extended semantics,

$$\Phi \colon \mathcal{B} \to X.$$

Let $(X,\ Y,\ E,\ S,\ \pi_\bullet)$ be a reflexive framework on X, Y, E, and S. Suppose that if $\Phi(B) = e \in E$ (i.e., if $B \in \mathcal{B}_e$) then π_e is the perspective of B. We then say that *the reflexive framework supports the environment of the extended semantics.*

In this case, each observer $B \in \mathcal{B}_E$ has its perspective determined by its configuration $\Phi(B)$. 2.4 becomes

$$\mathcal{B}_E = \Phi^{-1}(E), \tag{2.7}$$

indicating that $\mathcal{B}_E$ is the set of distinguished objects of perception (cf. **4**–4.4). By contrast, there need be no relation between the perspective of an observer in $\mathcal{B} - \mathcal{B}_E$ and its configuration in $X - E$.

If a reflexive framework supports an environment $(\mathcal{B}, \Phi)$ then, with notation as above, we assume that the map

$$\Xi: \mathcal{C} \to X$$

is bijective, where $\mathcal{C}$ is the set of states of affairs for the semantics. Then the subsets $\mathcal{B}_e$ of $\mathcal{B}$ defined above play the role of the equivalence classes $\mathcal{B}_c$ of $\mathcal{B}$ associated, as in **4**–4, to the distinguished states of affairs $c \in \mathcal{C}$; in fact $\mathcal{B}_c = \mathcal{B}_e$ when $\Xi(c) = e$.

Suppose that $O = (X, Y, E, S, \pi, \eta)$ is an observer, and $(R, \mathcal{B}, \Phi, X_t)$ is an extended semantics for O, where $\mathcal{B}$ is a set of observers with the same X, Y, E, S as O. Suppose that no reflexive framework is given at the outset, but that the map Φ has the following property: for each $e \in E$, all observers in Φ^{-1} (e) have the same perspective. Then we can construct a reflexive framework which supports the environment $(\mathcal{B}, \Phi)$: we simply construct the map Π which defines the framework by letting $\Pi(e)$ be the perspective of any observer $B \in \Phi^{-1}(e)$.

Terminology 2.8. If a reflexive framework supports an environment $(\mathcal{B}, \Phi)$, we call the elements of $\mathcal{B}$ the *observers in the framework.* We view $\mathcal{B}$ as the set of objects of perception for each observer in $\mathcal{B}$, and we take the given map $\Phi: \mathcal{B} \to X$ to be the configuration map for each observer in $\mathcal{B}$. If we are given a reflexive framework $(X, Y, E, S, \pi_\bullet)$ without specifying a particular environment which the framework supports, we still use the expression "observer in the framework" to refer to any observer having the same X, Y, E, S and having one of the π_e's for its perspective. In other words, an observer in the framework is any observer which completes a preobserver (X, Y, E, S, π_e) for some $e \in E$. We will sometimes use the expression "perspective in the framework" to refer to a point of E, i.e., we abuse language by identifying e with π_e.

This terminology emphasizes that a reflexive framework represents a family of observers that observe each other. This family can be taken to be the set $\mathcal{B}$ of any environment which is supported by the framework.

We illustrate the concept of reflexive framework with an example depicted in Figure 2.9. Here the configuration space X of the framework is the plane $\mathbf{R}^2$, a portion of which is represented by the rectangle in the figure. The

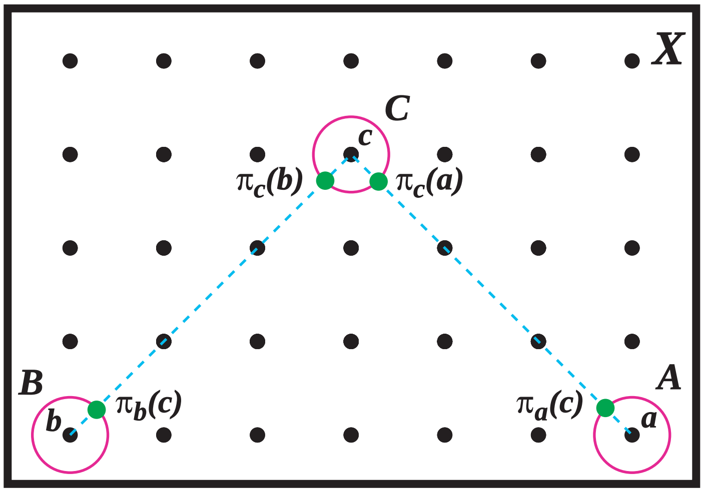

FIGURE 2.9. *A reflexive observer framework.*

distinguished configuration set is the set E of points in the plane that have integer coordinates. A few such points are represented by dots inside the rectangle. The premise space Y is the unit circle, plus one point at the center of the circle (call it s_0). We view X and Y with the measurable structures associated to their topologies (c.f. **2**–1). The distinguished premise set S is s_0 together with the set of points on the unit circle in Y which correspond to angles having rational tangents. We view Y as a measurable space where $\mathcal{Y}$ is the σ-algebra generated by the standard Borel algebra on the unit circle, along with $\{s_0\}$.

Now to describe the framework, we must assign to each $e \in E$ a measurable function $\pi_e : X \to Y$ such that $\pi_e(E) = S$. We do this as follows. To each point e of E associate a unit circle centered at e. We think of these circles as translated, but not rotated, copies of the unit circle of Y. Let $x \in X, x \neq e$. Then $\pi_e(x)$ is the point where the line determined by e and x intersects the unit circle centered at e; since this circle is a copy of Y, we view $\pi_e(x)$ as a point of Y. We define $\pi_e(e) = s_0$ in Y. If $e \neq e'$, since both have integer coordinates the line joining them has rational slope, so that the point $\pi_e(e')$ on the circle represents an angle with rational tangent. It follows that $\pi_e(E) = S$ as desired.

Figure 2.9 shows this procedure for three points of E: a, b, and c. In particular it shows (by means of green dots) $\pi_a(c)$, $\pi_b(c)$, $\pi_c(a)$, and $\pi_c(b)$ on separate copies of Y which are depicted as centered at a, b, c. Each observer in the framework has its own copy of Y, at most one point of which "lights up" at any given instant. (However, when viewing Figure 2.9 the reader should realize that the circles representing these copies are drawn *inside* X only for convenience in visualizing the maps π_a, π_b, π_c.) For example, suppose that A, B, C are three observers in the framework whose perspectives are π_a, π_b, π_c respectively. Suppose, moreover, that at a particular instant t, B and C channel with each other. Then, as shown in Figure 2.9, at that instant the point $\pi_b(c)$ lights up on B's copy of Y and $\pi_c(b)$ on C's copy. The point $\pi_c(a)$ in the figure does not light up at time t in this case. The point $\pi_a(c)$ *may* light up at time t, but if so it is not due to a channeling between A and C, for an observer interacts with at most one object of perception at any instant. It would be due to a channeling of A with some other observer whose perspective corresponds to a point of X on the same line through a as c.

We will need one further definition, which provides the syntax for the discussion of interpretation kernels in the context of reflexive frameworks.

Definition 2.10. A family of kernels $\{\eta_e\}_{e \in E}$ is called a *family of interpretation kernels* for the reflexive framework $(X, Y, E, S, \pi_\bullet)$ if, for each $e \in E$, $(X, Y, E, S, \pi_e, \eta_e)$ is an observer.

A family of interpretation kernels is a way to associate a single observer to each perspective in the framework, i.e., the observer $(X, Y, E, S, \pi_e, \eta_e)$ is associated to the perspective π_e. Equivalently, it is a way to complete each preobserver (X, Y, E, S, π_e) (for $e \in E$) to an observer.

3. Channeling on reflexive frameworks

We now make precise the term "channeling" on a reflexive framework. Recall that, in the primitive semantics, channeling denotes the presentation of an observer with a premise from an undefined probabilistic source (**4**–2.2). In the extended semantics, we speak of an object of perception "channeling" to an observer (**4**–4.4); this means that a given channeling arises from an interaction of the observer with that object of perception. Let us consider an environment

$(\mathcal{B}, \Phi)$ supported by a reflexive framework. We have a set $\mathcal{B}$ of observers—the observers in the framework—which is the set of objects of perception for each of its members. Now, according to the assumptions of extended semantics for an observer O (**4**–4.1), at each instant of time O participates in at most one channeling, implying that O interacts with but one of its objects of perception.

We make one further, and independent, assumption about channeling in this reflexive framework. This assumption supports a strategy which seeks to use direct interactions between observers as the foundation for an analysis of dynamics.

Assumption 3.1. Let $(\mathcal{B}, \Phi)$ be the environment of an extended semantics supported by a reflexive framework. With notation as above, let $A, B \in \mathcal{B}$. Suppose at time t that B channels to A, i.e., that B is the object of perception for A at that time. Then A also channels to B at time t, i.e., A is the object of perception for B. (Compare Lefebvre (1982).)

For a given instant t, let $L \subset \mathcal{B}$ denote those observers in $\mathcal{B}$ which channel at time t. For any $B \in L$, let $\tilde{\chi}(B) \in L$ be the observer with which B channels. In view of Assumption 3.1, $\tilde{\chi}: L \to L$ is a well-defined function with the property that $\tilde{\chi}^2 = \mathrm{Id}_L$ (the identity map on L); a map with this property is called an *involution* of L. (B may channel to itself: $B = \tilde{\chi}(B)$ is permissible.) We arrive at the following definition:

Definition 3.2. Let $(X, Y, E, S, \pi_\bullet)$ be a reflexive framework that supports the environment $(\mathcal{B}, \Phi)$ (**4**–4.4).

(i) A $(\mathcal{B}, \Phi)$-*channeling on the framework* is a pair $(L, \tilde{\chi})$ consisting of a non-empty subset $L \subset \mathcal{B}$ and an involution $\tilde{\chi}: L \to L$. (When there is no danger of confusion about $\mathcal{B}$ and Φ we simply say "channeling on the framework.")

(ii) Such a channeling is *elementary* if there are at most two elements in L.

With the hypotheses and notation of Definition 3.2, let $A, B \in \mathcal{B}$ and suppose that at time t, A and B channel to each other, so that $A = \tilde{\chi}(B)$. Then $\pi_{\Phi(A)}$ and $\pi_{\Phi(B)}$ denote the perspectives of A and B respectively. Thus the premise for A's perceptual inference resulting from this channeling is

$$\pi_{\Phi(A)}(\Phi(B)),$$

and similarly the premise for B's perceptual inference is $\pi_{\Phi(B)}(\Phi(A))$. More generally, we can use the $\tilde{\chi}$-notation, and summarize as follows:

3.3. With the hypotheses and notation of 3.2, let $(L, \tilde{\chi})$ be a channeling, and let $A \in L$. Then A's premise resulting from this channeling is

$$\pi_{\Phi(A)}(\Phi(\tilde{\chi}(A))).$$

Terminology 3.4. Given a channeling $(L, \tilde{\chi})$, we denote

$$D = \bigcup_{A \in L \cap \mathcal{B}_E} \{A, \tilde{\chi}(A)\}, \qquad \chi = \tilde{\chi}|_D.$$

(χ is the restriction of $\tilde{\chi}$ to D.) We call the channeling (D, χ) the ***distinguished part*** of $(L, \tilde{\chi})$.

4. Formal examples of reflexive frameworks

This section presents formal examples of reflexive frameworks. The examples do not represent a broad spectrum of types of frameworks, nor do they display an obvious relevance to everyday perception. Rather, they have been chosen to direct the exposition toward the particular subclass of *symmetric frameworks.* These we develop in the next section; they are the frameworks of primary interest in this book. In section six we develop in detail a perceptual example.

Example 4.1. Let G be a measurable group (see **5**–1), and E and H measurable subgroups of G. Denote by $Y = G/H$ the set of left H-cosets with its quotient measurable structure; H need not be normal. Let $\pi\colon G \to G/H$ be the canonical map, and let $S = \pi(E) = EH/H$. (EH denotes the set $\{eh|\, h \in H$

and $e \in E\}$, so that EH/H is the set of left cosets of H by elements of E.) Define $\Pi\colon E \to \mathtt{Hom}(G, G/H)$ as follows: for each $e \in E$, $\Pi(e)$ is the map $\pi_e\colon G \to G/H$ given by $\pi_e(g) = \pi(ge^{-1})$. We then have $\pi_e(E) = \pi(Ee^{-1}) = \pi(E) = S$ for all e (since E is a group), as required by the definition of a reflexive framework. For each $e \in E$, $\pi_e^{-1}(S) = \pi^{-1}(S) = EH$.

$$
\begin{array}{ccccc}
 & E & \subset & G & = X \\
 & \Big\downarrow{\scriptstyle \pi_e|_E} & & \Big\downarrow{\scriptstyle \pi_e} & \\
S = & HE/H & \subset & G/H & = Y
\end{array}
$$

In this particular reflexive framework the set of fibres $\{\pi_e^{-1}(y) \mid y \in Y\}$ is independent of e (as a set of subsets of X); in fact for each e the fibres are the left H-cosets in X. To change e is simply to permute the fibres.

Example 4.2. This example generalizes the previous one. Again let G be a measurable group. But now let H be an arbitrary group which acts measurably on G on the right. (Thus the elements of H correspond to bijective, bimeasurable maps from G to itself, maps which are not necessarily group homomorphisms.)

Let G/H denote the orbits in G for the action of H, and $\pi\colon G \to G/H$ the canonical map. Let E be a measurable subgroup of G, and let EH/H denote the subset of G/H consisting of those orbits which contain an element of E. For $e \in E$, define $\pi_e(g)$ to be $\pi(ge^{-1}) = H(ge^{-1})$, i.e., the H-orbit on G containing ge^{-1}. Let $X = G$, $Y = G/H$, E the given subgroup of G, $S = EH/H$, and π_e as defined above.

$$
\begin{array}{ccccc}
 & E & \subset & X & = G \\
 & \Big\downarrow{\scriptstyle \pi|_E} & & \Big\downarrow{\scriptstyle \pi} & \\
EH/H = & S & \subset & Y & = G/H
\end{array}
$$

In the case where the H in this example is a subgroup of G, acting on G by left translation, we simply recover the previous Example 4.1. However this case accounts for only a very small class of “natural” measurable actions of one group on another. In fact, the example illustrated in Figure 2.9 is of the type of 4.2, but not 4.1. In the next example we present the n-dimensional generalization of the one in Figure 2.9. In the general situation of Example 4.2 it is not true (as it was in Example 4.1) that all the maps π_e for $e \in E$ have the same set of fibres over S. This is evident from the next example.

Example 4.3. With the notation of 4.1, let $X = G = (\mathbf{R}^n, +)$ be the n-dimensional vector group; "+" denotes vector addition. Let $H = (\mathbf{R}_+, \text{multiplication})$, i.e., H is the multiplicative group of positive real numbers acting by dilation (scalar multiplication) on $\mathbf{R}^n$. Then $Y = G/H$ is the set of half-rays emanating from the origin, together with one point s_0 which is the orbit consisting of the origin itself. Since the set of half-rays is naturally identified with the $(n-1)$-dimensional unit sphere $\mathbf{S}^{n-1}$ centered at the origin in $\mathbf{R}^n$, we have $Y = G/H = \mathbf{S}^{n-1} \cup \{s_0\}$.

Now let $E = (\mathbf{Z}^n, +)$, the subgroup of points with integer coordinates, or let $E = (\mathbf{Q}^n, +)$, the subgroup of points with rational coordinates. In either case the image by π of E in Y is the same: it is the set consisting of the point s_0 together with all points on $\mathbf{S}^{n-1}$ with the property that the ratio of any pair of coordinates is rational. This set is denoted $\mathbf{S}_r^{n-1}$ in the following diagram.

$$\begin{array}{ccccc} \mathbf{Q}^n \text{ or } \mathbf{Z}^n = & E & \subset & X & = \mathbf{R}^n \\ & \Big\downarrow {\scriptstyle \pi_e|_E} & & \Big\downarrow {\scriptstyle \pi_e} & \\ \mathbf{S}_r^{n-1} \cup \{s_0\} = & S & \subset & Y & = \mathbf{S}^{n-1} \cup \{s_0\} \end{array}$$

For $e \in \mathbf{R}^n = X$ we may conceptualize π_e as follows: translate the unit sphere $\mathbf{S}^{n-1}$ (originally centered at the origin) to e. For any $v \in \mathbf{R}^n$, if $v \neq e$ take the ray from e to v, and intersect it with this translated $\mathbf{S}^{n-1}$ to obtain $\pi_e(v)$. Define $\pi_e(e) = s_0$.

Example 4.4. Here is a further generalization of Example 4.2 in which the constructions can be described without substantial change in the syntax. In its generality this example contains all the others of this section. Again, let G be a measurable group. We suppose that G has a partition in measurable subsets; denote this partition as well as the corresponding equivalence relation by $\mathcal{H}$. Let $Y = G/\mathcal{H}$ and let $\pi: G \to Y$ be the canonical map. We take for the σ-algebra $\mathcal{Y}$ the quotient measurable structure. Let J be a measurable subgroup of G with the property that $\pi(J) \subset Y$ is measurable. We set $S = \pi(J)$.

Moreover, let us assume that we have a measurable space X on which G acts measurably (on the left). Let x_0 be a distinguished point of X, and let $E = Jx_0 \subset X$. We also assume the following:

(i) G acts transitively on X.

(ii) Let $e \in E$, and $g, g' \in G$. If $ge = g'e$ then g, g' are in the same $\mathcal{H}$-class in G.

From this we now describe a reflexive framework on X, Y, E, S. Assumptions (i) and (ii) insure that we can define π_e in a manner consistent with the previous examples. In fact, for $e \in E$ and $x \in X$, let xe^{-1} denote any element $g \in G$ such that $ge = x$. Such a g exists because of (i). Then define $\pi_e(x) = \pi(xe^{-1})$. Assumption (ii) means that this definition of $\pi_e(x)$ is independent of the choice of xe^{-1}, i.e., π_e is well defined. To see that $\pi_e(E) = S$ for all e, let $e_1 \in E$, and suppose that $e = jx_0$, $e_1 = kx_0$, where $j, k \in J$. kj^{-1} is then one choice for e_1e^{-1}, so

$$\pi_e(e_1) = \pi(e_1e^{-1}) = \pi(kj^{-1}) \in \pi(J) = S.$$

Moreover it is clear that as e_1 runs over E, kj^{-1} runs over J, so that all elements of S are represented.

$$\begin{array}{cccccc} E = & Jx_0 & \subset & X & & \\ & \downarrow{\scriptstyle \pi_e|_E} & & \downarrow{\scriptstyle \pi_e} & & \\ S = & \pi(J) & \subset & G/\mathcal{H} & = Y \end{array}$$

Let Σ_e denote the ***stabilizer*** of e (i.e., the subgroup of G which leaves e fixed). In view of (i) above we may identify X with G/Σ_e; under this identification $x \in X$ corresponds to the coset $g\Sigma_e$ where $g \in G$ is any element such that $ge = x$. For g, $g' \in G$, $ge = g'e$ if and only if g and g' are in the same left coset of Σ_e. Thus (ii) above is equivalent to the assertion that every coset of Σ_e is contained in one $\mathcal{H}$-class, or equivalently that each $\mathcal{H}$-class is a union of cosets of Σ_e. We can then associate to each $e \in E$ a natural map

$$X = G/\Sigma_e \to G/\mathcal{H} = Y$$

as follows. If $x \in X$ with $x = ge$, x is identified as above with $g\Sigma_e$ in G/Σ_e which is then mapped to the element of $G/\mathcal{H}$ which represents the $\mathcal{H}$-class containing $g\Sigma_e$. But, since g here is one choice for xe^{-1}, this map from X to $G/\mathcal{H}$ is just our π_e defined above.

Example 4.4 generalizes 4.2 in two respects. First, the equivalence relation $\mathcal{H}$ on G which gives the canonical map π need not arise from the orbits of a group action. Secondly, the action of G on X need not be faithful: for example, whereas in 4.2 $X = G$, here we can have $X = G/\Sigma$ where Σ is a non-normal subgroup of G. However the action of G on X still must be transitive.

Example 4.5. We show how to get a class of generalizations of Example 4.1, where H is still a subgroup of G acting by translation, but now X is a measurable, transitive G-set for which the action is not faithful. This means that X may be identified with G/Σ, the left cosets of the stabilizer Σ of some fixed $x_0 \in X$ (as we saw in Example 4.4). We assume that the measurable structure of X is given by the quotient structure of G/Σ.

As in 4.1, let $Y = G/H$. We assume

1. $\Sigma \subset H$

and by so doing get a canonical surjective (and measurable) map $\pi\colon X \to Y$ given by $\pi(g\Sigma) = gH$.

Let J be a measurable subgroup of G and set $E = Jx_0$. We think of E as $J\Sigma/\Sigma$: the left cosets of Σ by J. Then the map π restricts to $\pi|_E\colon J\Sigma/\Sigma \to JH/H$. We set $S = \pi(E) = JH/H$.

Example 4.4 tells us that we can define the map Π of a reflexive framework if its assumption (ii) is satisfied: $g\Sigma_e \subset gH$ for all $e \in E$ and $g \in G$. If $e \in E$, then $e = jx_0$ for some $j \in J$ and $\Sigma_e = j\Sigma j^{-1}$. We therefore impose another condition.

2. For all $j \in J$, $j\Sigma j^{-1} \subset H$.

(For example, J may be contained in the normalizer of Σ, i.e., $j\Sigma j^{-1} = \Sigma$, or j may be contained in the normalizer of H.) We have

$$\begin{array}{cccccc} E = Jx_0 = & J\Sigma/\Sigma & \subset & G/\Sigma & = X \\ & \big\downarrow \pi|_E & & \big\downarrow \pi & \\ S = & JH/H & \subset & G/H & = Y \end{array}$$

The maps π_e are well-defined as follows: if $e = j\Sigma$ and $x = g\Sigma$ (i.e., $e = jx_0$, $x = gx_0$), then

$$\pi_e(x) = (gj^{-1})H.$$

If $\Sigma = \{\imath\}$ this example reduces to 4.1, and in any case it shares with 4.1 the property that all the maps π_e, $e \in E$, have the same set of fibres, namely the cosets of H (mod Σ). In order that we get a nontrivial situation, we must assume

3. $J \neq \Sigma$ (otherwise E is a singleton),
   $J \neq G$ (otherwise $E = X$ and $S = Y$),
   $J \not\subset H$ (otherwise S is a singleton).

5. Symmetric observer frameworks

All of our examples of reflexive frameworks have involved groups, although this is certainly not required by Definition 2.2. In every example in section three, X is a G-set for some group G in such a way that E is a J-set for a subgroup J of G; the actions are transitive. Moreover, in each case the maps π_e in the framework are deduced "by translation" from some fixed measurable map $\pi\colon G \to Y = G/\mathcal{H}$, $\mathcal{H}$ being an equivalence relation on G. In fact $\pi_e(x) = \pi(xe^{-1})$, where xe^{-1} denotes any element of G such that $(xe^{-1})e = x$. In other words xe^{-1} is the difference between x and e measured in terms of G. This means (using the terminology of 2.8) that the observations by an observer O with perspective e in the framework depend only on the structure of X ***relative to*** e (in the sense of the action of G.)

Consider Example 4.4. It is not misleading to think of each $e \in E$ as the center of a "frame for observation" which consists of the structure (G, Y, J, S, π) "translated" to e, where translation here refers to the action of G. This frame provides the syntax for the perceptual representations of any observer in the framework ***relative to*** O; the notion "relative" is grounded in the G-space structure of X. This is the basis for a symmetric theory of observer interaction; the symmetry in question is that of the group G. When we present the dynamics in the subsequent chapters we focus on this symmetric setting. One can certainly construct examples of reflexive observer frameworks which are not of this type and then study interaction dynamics on them in depth, but we will not do so explicitly in this book.

Definition 5.1. A ***symmetric observer framework*** is a reflexive observer framework $(X, Y, E, S, \pi_\bullet)$ for which there exists a measurable group G, a measurable subgroup $J \subset G$, and a measurable surjective map $\pi\colon G \to Y$ satisfying two requirements:

(i) G acts transitively and measurably on X, inducing a transitive action of J on E (which is automatically measurable).

(ii) For all $e \in E$ and $x \in X$, $\pi_e(x) = \pi(g)$, where g is ***any*** element in G such that $ge = x$ (i.e., $g =$ "xe^{-1}.")

The requirements on the maps $\pi_e\colon X \to Y$ in a reflexive framework, namely that π_e is surjective and $\pi_e(E) = S$, impose nontrivial conditions on the map π. However, the best way to understand the whole definition is to realize the following:

Proposition 5.2. The definition of symmetric observer framework is equivalent to the Example 4.4 of the previous section.

Proof. The fibres of the map π of 5.1 form a partition of G. The relation of joint membership in a fibre is an equivalence relation: call it $\mathcal{H}$. Then Y is identified with $G/\mathcal{H}$. Since the action of J on E is transitive, E is identified with Je_0 for any $e_0 \in E$. It remains to verify (ii) of 4.4, but this is implicit in (ii) of 5.1. ∎

Terminology 5.3. We will use the notation (X, Y, E, S, G, J, π) for a symmetric observer framework; the notation π_e will refer to the maps from X to Y defined in terms of π as in (ii) of the definition. The structure (G, Y, J, S, π) is called the ***fundamental frame*** of the framework. π is the ***fundamental map***, G and J are the ***configuration group*** and the ***distinguished subgroup*** respectively; we retain our original terminology for X, Y, E, S, namely ***configuration space, premise space, distinguished configurations, distinguished premises.*** We will frequently use the informal terminology "symmetric framework" rather than "symmetric observer framework."

In Definition 5.1 it is necessary only for group actions to exist at the level of X and E, not Y and S. Furthermore, the fundamental map $\pi: G \to Y$ need not arise in any particular group-theoretic way. Y can be $G/\mathcal{H}$ for ***any*** equivalence relation $\mathcal{H}$ on G for which the notation "$\pi(xe^{-1})$" makes sense (so that the π_e's are well-defined by (ii) of 5.1). As we have seen in 4.4, this is tantamount to saying that for all $e \in E$ and $g \in G$, $g\Sigma_e$ is contained in a single $\mathcal{H}$ equivalence class. (The equivalence relation $\mathcal{H}$ here is just the set of fibres of π.)

An important special case is when X is a ***principal homogeneous space*** for G. This means that G acts faithfully as well as transitively on X; in other words, all the stabilizers Σ_x are trivial. In this case, given any $x \in X$ we can identify X with a copy of G "centered at x," i.e., the element $g \in G$ is identified with $gx \in X$. Moreover when $x = e \in E$, this identification of X with G also identifies E with J. When X is a principal homogeneous G space, then for any x and e in X the element xe^{-1} is uniquely determined. In this case the condition (ii) in Definition 5.1 does not impose any requirements on π. We therefore have

5.4. Let G be a measurable group and $J \subset G$ a measurable subgroup. Let X be a principal homogeneous space for G on which the action of G is measurable. Suppose $E \subset X$ is a measurable J-invariant subset (so that the G-principal homogeneous structure of X induces a J-principal homogeneous structure for E). Let Y be a measurable space and $\pi: G \to Y$ be ***any*** measurable, surjective function; this is equivalent to saying that $Y = G/\mathcal{H}$ where $\mathcal{H}$ is an equivalence relation on G for which the equivalence classes are measurable subsets of G. Let $S = \pi(J)$. Then (X, Y, E, S, G, J, π) is a symmetric observer framework. For each $e \in E$ we define $\pi_e: X \to Y$ by $\pi_e(x) = \pi(xe^{-1})$, where xe^{-1} denotes the *unique* element $g \in G$ such that $ge = x$. We call these "principal homogeneous symmetric frameworks," or "principal frameworks" for short.

Example 5.5. This is the same as Example 4.3, where we now make explicit its structure as a principal framework: Let $G = (\mathbf{R}^n, +)$, $J = (\mathbf{Z}^n, +) \subset G$. Let $X = \mathbf{R}^n$ and $E = \mathbf{Z}^n$. We think of G acting on X by translation; it is obvious that X is principal homogenous for this action and that E is J-invariant. (In 4.3 we identified G with X at the outset, so that E is itself a subgroup of G and there is no need to introduce J. However here we are making the distinction in principle between G and X; the point is that while one can always identify a group G with a principal homogeneous G-space X, the identification is not canonical.) Let $Y = \mathbf{S}^{n-1} \cup \{s_0\}$, where $\mathbf{S}^{n-1}$ denotes the $n-1$-dimensional sphere and s_0 is a point (which we can visualize at the center of the sphere). We now define the map π. To do this, let $e_0 \in E$ denote the the origin in G. Identify the $\mathbf{S}^{n-1}$ in Y with the unit sphere centered at e_0. With this identification, for any $x \in X$, $x \neq e_0$, let $\pi(x)$ be the point of Y which is the intersection of $\mathbf{S}^{n-1}$ with the line joining e_0 and x. Let $\pi(e_0) = s_0$. It is evident that the maps π_e in 4.3 can be defined in terms of this π by the formula $\pi_e(x) = \pi(x - e)$ (where we use the additive notation "$x - e$" instead of xe^{-1}).

Finally, we elaborate the notion of a family of interpretation kernels (2.10) in the special case of symmetric frameworks.

Definition 5.6. A family $\{\eta_e\}_{e\in E}$ of interpretation kernels for the symmetric observer framework (X, Y, E, S, G, J, π) is said to be ***symmetric*** if there exists a markovian kernel $\eta\colon S \times \mathcal{J} \to [0, 1]$ such that for all $e \in E$, $s \in S$ and $\gamma \in \mathcal{E}$,

$$\eta(s, \pi^{-1}\{s\} \cap J) = 1,$$
$$\text{and} \qquad \eta_e(s, \Gamma) = \eta(s, \Gamma e^{-1}).$$

η is then called the ***fundamental kernel*** of the family η_e.

One way a family can be symmetric is as follows. Suppose we are given a symmetric observer framework (X, Y, E, S, G, J, π) and a measure ν on J. Since J acts transitively on E, given any $e \in E$ we get a surjective map $c_e\colon J \to E$ by sending $\imath$ to e. ($\imath$ denotes the identity element of J.) c_e identifies

E with the quotient space $J/\Sigma_e \cap J$, where Σ_e is the stabilizer of e in G. Let $\nu_e = (c_e)_*(\nu)$; this is the measure ν transported to E by "centering a copy of J at e."

Terminology 5.7. With the hypotheses and notation of the previous paragraph, if ν is a measure on J, the family of measures ν_e on E is called the *symmetric family of measures associated to* ν; ν is called the *fundamental measure of the family.* Concretely, if $\Gamma \in \mathcal{E}$, then $\nu_e(\Gamma) = \nu(c_e^{-1}(\Gamma)) = \nu\{j \in J |\, je \in \Gamma\}$.

Now given a probability measure ν, and its associated symmetric family $\{\nu_e\}$, we can define a family of kernels $\eta_e \colon S \times \mathcal{E} \to [0, 1]$ which are the rcpd's of the ν_e, i.e., we can let

$$\eta_e(s, \Gamma) = m_{\pi_e}^{\nu_e}(s, \Gamma)$$

(notation as in **2**–1). Another way to describe this family of kernels is as follows: Let $\eta = m_{\pi|_J}^{\nu}$, where $\pi|_J$ is the fundamental map of our symmetric framework restricted to the subgroup J of G.

$$\begin{array}{ccc} \nu & J & m_{\pi|_J}^{\nu} \colon S \times \mathcal{J} \to \mathbf{R} \\ & \big\downarrow \scriptstyle{\pi|_J} & \mathcal{J} = \text{ the Borel sets of } J. \\ & S & \end{array}$$

We have the diagram

$$\begin{array}{ccccc} \nu,\, \eta & J & \xrightarrow{c_e} & E & \nu_e,\, \eta_e \\ & \pi|_J \searrow & & \swarrow \pi_e|_E & \\ & & S & & \end{array}$$

which commutes (i.e., $\pi|_J = \pi_e|_E \circ c_e$) by definition of the π_e. From this and the fact that $\nu_e = (c_e)_*(\nu)$, it follows from the meaning of rcpd that for $\Gamma \in \mathcal{E}$ and $s \in S$, $\eta_e(s, \Gamma) = \eta(s, c_e^{-1}(\Gamma))$. Note that $c_e^{-1}(\Gamma) = \{j \in J |\, je \in \Gamma\}$; so it is consistent with our previous notation to write $c_e^{-1}(\Gamma) = \Gamma e^{-1}$.

Notation 5.8. Given a symmetric family of kernels $\{\eta_e\}$ (respectively, measures $\{\nu_e\}$), then η (respectively, ν) will always denote the fundamental kernel (respectively, measure).

Note that the measure ν never appears in the Definition 5.6; intuitively only its rcpd appears, in the form of the fundamental kernel η. Thus in order to determine ν we would need to know the measure $\pi_*\nu$ on S. The precise statement is

5.9. A symmetric family η_e of interpretation kernels, together with a measure λ on S, uniquely determine a symmetric family of measures ν_e on E (and conversely) via the relation:

$$\eta = m^{\nu}_{\pi|_J}, \qquad \lambda = (\pi|_J)_*(\nu).$$

The definition of symmetric framework expresses the role of groups in creating a theory of observer interactions which permits "relativization." The reflexive frameworks we study in this book are primarily principal frameworks. These include the framework of instantaneous rotation observers presented in the next section, the frameworks for which we develop the theory of true perception in chapter eight, those employed in the investigation of hierarchical perceptual organization in chapter nine, and the frameworks which arise in our discussion of the applications of observer theory to physics in chapter ten. However, the general theory of participator dynamics developed in chapter seven is not restricted to the principal homogeneous case.

6. Example: Instantaneous rotation

We now study one example of the visual perception of structure in three dimensions given image motion in two dimensions, namely the perception of rigid, fixed-axis motion from a premise consisting of two views of $n+1$ points. For this purpose n can be any integer ≥ 3. We think of these views as occurring in successive instants of some underlying discrete time.

Given $n+1$ points moving arbitrarily in $\mathbf{R}^3$, let $(P_0,\ P_1,\ \ldots,\ P_n)$ and $(Q_0,\ Q_1,\ \ldots,\ Q_n)$ be their positions at two successive instants of time. Let us assume that the viewer is using a moving coordinate system in which $P_0 = Q_0 = (0,0,0)$. Then this data (viz., the P_i's and Q_i's) is equivalent to the array $\mathbf{a} = (\mathbf{a}_{11},\ \mathbf{a}_{21},\ \ldots,\ \mathbf{a}_{n1};\ \mathbf{a}_{12},\ \mathbf{a}_{22},\ \ldots,\ \mathbf{a}_{n2})$ of $2n$ vectors in $\mathbf{R}^3$, where $\mathbf{a}_{i1} = P_i - P_0$, $\mathbf{a}_{i2} = Q_i - Q_0$, $i = 1,\ldots,n$. To say that $Q_0,\ Q_1,\ \ldots,\ Q_n$ are obtained from $P_0,\ P_1,\ \ldots,\ P_n$ by a rigid motion of $\mathbf{R}^3$ is equivalent to saying that $\mathbf{a}_{12}\ldots\mathbf{a}_{n2}$ are obtained from $\mathbf{a}_{11}\ldots\mathbf{a}_{n1}$ by a rotation about an axis through the origin. We call this an ***instantaneous rotation*** since two successive positions of an object in discrete time corresponds to instantaneous motion. Thus to infer an arbitrary rigid motion of $n+1$ points from two views is the same thing as inferring an instantaneous rotation of n vectors from two views.

We will define a symmetric framework $\Theta = (X, Y, E, S, G, J, \pi)$ in which the observers are instantaneous rotation observers. It turns out that in order to get the group structure here, the observer must utilize configurations which are ***pairs*** $(\mathcal{A}, \mathbf{a})$, where $\mathbf{a}$ is a $2n$-tuple of vectors in $\mathbf{R}^3$ as above, and $\mathcal{A}$ is a "reference axis":

Terminology 6.1. An ***axis*** in $\mathbf{R}^3$ is an oriented line through the origin, i.e., a line with its positive direction specified. We will denote the set of such axes by $\mathbf{A}$.

The set $\mathbf{A}$ of axes corresponds to the set of points on the unit sphere $\mathbf{S}^2$ centered at the origin: each such point determines a line through the origin, whose positive direction is taken to be the direction from the origin to the point.

The axis-body configuration $(\mathcal{A}, \mathbf{a})$ represents the motion

$$\mathbf{a}_{11} \ldots \mathbf{a}_{n1} \to \mathbf{a}_{12} \ldots \mathbf{a}_{n2}$$

"referred" to the axis $\mathcal{A}$. Since we do not detect rotational motion of points on the axis itself, we consider only those axis-body configurations $(\mathcal{A}, \mathbf{a})$ such that none of the the vectors $\mathbf{a}_{11}$, …, $\mathbf{a}_{n2}$ lie on the line through $\mathcal{A}$. By an abuse of notation, we express this assumption simply by $\mathbf{a} \notin \mathcal{A}$. Let

$$X = \{(\mathcal{A}, \mathbf{a}) | \, \mathcal{A} \in \mathbf{A}, \mathbf{a} \in (\mathbf{R}^3)^{2n}, \mathbf{a} \notin \mathcal{A}\}. \tag{6.2}$$

X is the configuration space of our framework Θ. Explicitly,

$$X = \big\{(\mathcal{A}; \mathbf{a}_{11}, \ldots, \mathbf{a}_{n1}; \mathbf{a}_{12}, \ldots, \mathbf{a}_{n2}) | \, \mathcal{A} \in \mathbf{A}, \mathbf{a}_{ij} \in \mathbf{R}^3, \mathbf{a}_{ij} \notin \mathcal{A}\big\}. \tag{6.3}$$

Let Y be the set of ordered pairs of n-tuples of vectors in $\mathbf{R}^2$, so that

$$Y = (\mathbf{R}^2)^{2n}.$$

We denote the elements of Y explicitly by

$$Y = \big\{(\mathbf{b}_{11}, \ldots, \mathbf{b}_{n1}; \mathbf{b}_{12}, \ldots, \mathbf{b}_{n2}) | \, \mathbf{b}_{ij} \in \mathbf{R}^2\big\}. \tag{6.4}$$

Let us fix a coordinate system, say (x, y, z) on $\mathbf{R}^3$. Let

$$p: X \to Y$$

be the map which forgets the axis $\mathcal{A}$, and which associates to each of the vectors $\mathbf{a}_{ij} \in \mathbf{R}^3$ its projection onto the (x, y)-plane viewed as a copy of $\mathbf{R}^2$. Thus

$$p(\mathcal{A}; \mathbf{a}_{11}, \ldots, \mathbf{a}_{n1}; \mathbf{a}_{12}, \ldots, \mathbf{a}_{n2}) = (\mathbf{b}_{11}, \ldots, \mathbf{b}_{n1}; \mathbf{b}_{12}, \ldots, \mathbf{b}_{n2}),$$

where

$$\mathbf{a}_{ij} = (x_{ij}, y_{ij}, z_{ij}), \quad \mathbf{b}_{ij} = (x_{ij}, y_{ij}). \tag{6.5}$$

We will see below how to define the fundamental map π of our framework in terms of this p.

Let E be the set of those elements of X in which the two n-tuples of vectors $(\mathbf{a}_{11}, \ldots, \mathbf{a}_{n1})$ and $(\mathbf{a}_{12}, \ldots, \mathbf{a}_{n2})$ of $\mathbf{R}^3$ are related by a rotation of $\mathbf{R}^3$ about the given axis $\mathcal{A}$. Thus

$$E = \big\{(\mathcal{A}; \mathbf{a}_{11}, \ldots, \mathbf{a}_{n1}; \mathbf{a}_{12}, \ldots, \mathbf{a}_{n2}) \in X \,|\, \sigma(\mathbf{a}_{i1}) = \mathbf{a}_{i2}; 1 \le i \le n; \text{ where } \sigma \in \mathrm{SO}(3, \mathbf{R}) \text{ is a rotation about } \mathcal{A}\big\}. \tag{6.6}$$

Remark 6.7. For $n \geq 3$, (Lebesgue) almost all n-tuples $(\mathbf{a}_{11}, \ldots, \mathbf{a}_{n1})$ of points of $\mathbf{R}^3$ do not lie in any proper linear subspace of $\mathbf{R}^3$. Since the rotation σ in 6.6 is a linear map it is therefore uniquely determined by where it sends these $\mathbf{a}_{ij}$. Moreover, the axis of the rotation σ is uniquely determined up to orientation. We conclude: For almost all points of E the σ in 6.6 is uniquely determined. Moreover, for almost all points $e = (\mathcal{A}, \mathbf{a}) \in E$ (where $\mathcal{A} \in \mathbf{A}$ and $\mathbf{a} \in (\mathbf{R}^3)^{2n}$) there is exactly one other point $e' = (\mathcal{A}', \mathbf{a}') \in E$ such that $\mathbf{a} = \mathbf{a}'$. In that case $\mathcal{A}$ and $\mathcal{A}'$ differ only in their orientation.

To recapitulate, we can think of X as the set of configurations which correspond to two successive positions of n vectors moving arbitrarily in $\mathbf{R}^3$, together with a choice of reference axis $\mathcal{A}$; "successive" refers to some particular discrete time scale. We will call such a configuration an "axis-referenced instantaneous motion of n vectors in $\mathbf{R}^3$," or just an "instantaneous motion" for short. Then E consists of those instantaneous motions which are in fact (rigid) rotations about their respective reference axes. Finally, we let

$$S = p(E) \subset Y. \tag{6.8}$$

We now have a preobserver (X, Y, E, S, p). For this preobserver, the two n-tuples of vectors in $\mathbf{R}^2$, which comprise a premise $y \in Y$, are interpreted as two successive two-dimensional projections via p of $n+1$ fixed feature points on an object moving in three dimensions. Each projection is a "view" of the object; the two n-tuples in a premise y represent the images on the observer's "retina" resulting from the two views. In other words, the interpretation of the preobserver is that the premise $y \in Y$ arises from some instantaneous motion $x \in X$ such that $p(x) = y$. (Strictly speaking there is no interpretation unless the premise y is in S. Furthermore it is observers—not preobservers—that make interpretations.)

Each point x of X includes a reference axis $\mathcal{A}$ as part of the motion it represents, even though for general points of X this motion is not a fixed-axis motion, much less a rigid fixed-axis motion about $\mathcal{A}$. Only when $p(x) = y$ is in S is it possible to infer that the instantaneous motion being viewed is a rigid rotation about its reference axis; the interpretations consistent with this inference correspond to configurations x in $p^{-1}(y) \cap E$. Is it necessary to include the reference axis as part of the configuration? Not if we simply want to describe a single instantaneous rotation observer. However it is necessary in order define the group actions of a symmetric framework. Moreover it seems clear that one's perception, when one is presented with the appropriate displays, includes a direction of rotation; this is equivalent mathematically to choosing an orientation for the axis.

For the sake of intuition, we state without proof some facts about the geometry of (X, Y, E, S, π). Details may be found in Bennett et al. (1989).

6.9. S is contained in the solution set in $Y = (\mathbf{R}^2)^{2n}$ of a family of polynomial equations (in $4n$ variables). The dimension of S is $3n+2$. Thus if μ_Y denotes Lebesgue measure on Y, $\mu_Y(S) = 0$. (This implies by **2**–3.3 that the preobserver (X, Y, E, S, p) is ideal.) The dimension of E is $3n+3$. For almost all $s \in S$, $p^{-1}(s) \cap E$ is a 1-dimensional manifold. This manifold has four connected components, corresponding to the two types of "reflections" which act on the fibre: the first is the reversal of orientation of the reference axis $\mathcal{A}$ (but leaving the points in the configuration fixed) and the second is a reflection of the entire structure about the image plane. Thus, up to choice of orientation of the reference axis and reflection in the image plane, every distinguished premise s is compatible with a one-parameter family of instantaneous-rotation interpretations. Two such interpretations for a given premise are illustrated in Figure 6.9.1. In the figure, each interpretation is represented by a system of n ellipses with the same eccentricity; the ith ellipse contains the image vectors $\mathbf{b}_{1i}$ and $\mathbf{b}_{2i}$, $i = 1, \ldots, n$. The minor axes of the ellipses in each system lie on the same line through the origin, namely the projection into the image plane of the actual axis of rotation of the corresponding interpretation. In the figure the system of ellipses for one interpretation is drawn with solid lines, and for the other interpretation with dotted lines. Note that the projected axis of rotation is the same (M) for the two interpretations. This holds true in general: For any $s \in S$, the axes of all of the distinguished configurations compatible with s project to the same line in the image plane.

Since we have defined X, Y, E and S, in order to describe the symmetric framework Θ we need to define the groups G and J and to define their actions on X and E. We also need to define the fundamental map π. For these purposes we give an alternate description of X and E in terms of which the group actions can be clearly expressed. We represent each element $x \in X$ in the form

$$x = \begin{pmatrix} \mathcal{A} & \begin{matrix} \mathbf{v}, c_{21}, \ldots, c_{n1} \\ \\ c_{12}, c_{22}, \ldots, c_{n2} \end{matrix} & \begin{matrix} h_{11}, \ldots, h_{n1} \\ \\ h_{12}, \ldots, h_{n2} \end{matrix} & \begin{matrix} l_{11}, \ldots, l_{n1} \\ \\ l_{12}, \ldots, l_{n2} \end{matrix} \end{pmatrix} \tag{6.10}$$

where we have fixed a coordinate system in $\mathbf{R}^3$, and where

$\mathcal{A}$ is an oriented line through the origin in this coordinate system, i.e., an axis;

$\mathbf{v}$ is a unit vector at the origin, perpendicular to $\mathcal{A}$;

c_{ji} are angles with $0 \le c_{ji} < 2\pi$;

h_{ji} are arbitrary real numbers; and

l_{ji} are strictly positive real numbers.

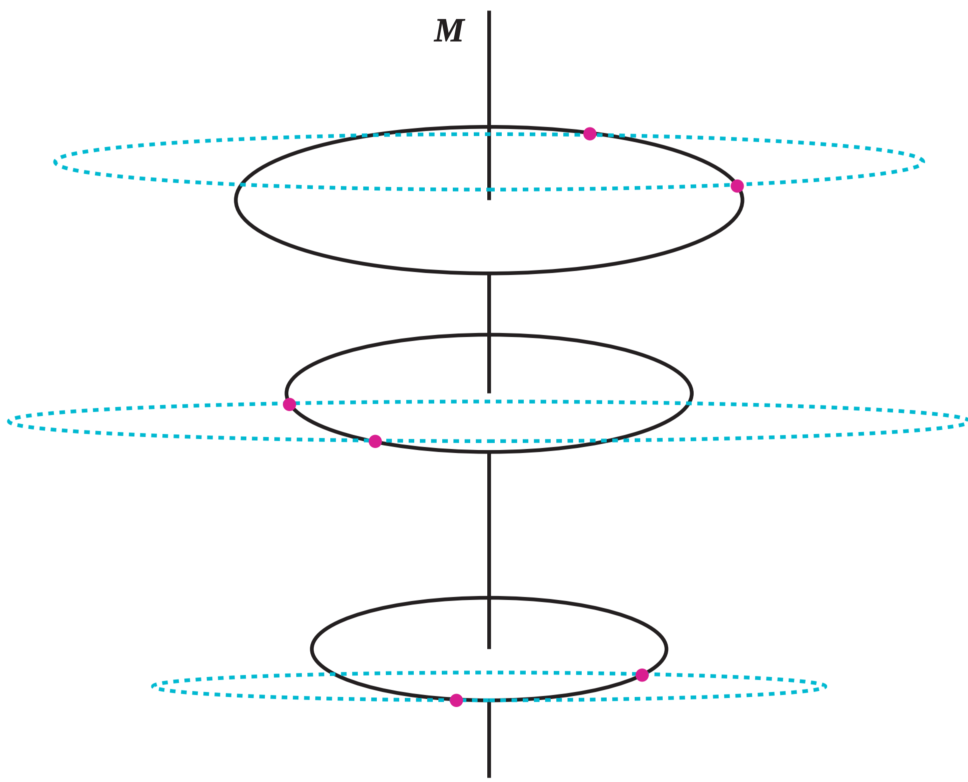

FIGURE 6.9.1. *Two rigid interpretations from the one-parameter family.*

6.10 is essentially a "cylindrical coordinate" representation of the original configuration

$$x = (\mathcal{A}; \mathbf{a}_{11}, \ldots, \mathbf{a}_{n1}; \mathbf{a}_{12}, \ldots, \mathbf{a}_{n2})$$

as follows. Let P_{ji} denote the point at the tip of vector $\mathbf{a}_{ji}$. Given $\mathcal{A}$, imagine the point P_{ji} as being connected to $\mathcal{A}$ by a vector $\mathbf{r}_{ji}$ which is perpendicular to $\mathcal{A}$ (see Figure 6.11). l_{ji} is the length of $\mathbf{r}_{ji}$; h_{ji} is the coordinate on $\mathcal{A}$ at the point where $\mathcal{A}$ meets $\mathbf{r}_{ji}$; and $\mathbf{v}$ is the direction of the projection of $\mathbf{r}_{11}$ in the plane L through the origin and perpendicular to $\mathcal{A}$. Then $c_{ji}((i, j) \neq (1, 1))$ describes the angular displacement of $\mathbf{r}_{ji}$ relative to $\mathbf{r}_{11}$; it is the counterclockwise angle between $\mathbf{v}$ and the projection of P_{ji} in L. Here the notion of "counterclockwise" is determined using the right-hand rule by the orientation of $\mathcal{A}$.

By the requirement that $\mathbf{a} \notin \mathcal{A}$ in 6.2, the vectors $\mathbf{r}_{ji}$ are all non-zero, so that the unit vector $\mathbf{v}$ and the angles c_{ji} are well-defined.

We can think of (c_{ji}, h_{ji}, l_{ji}) as cylindrical coordinates of P_{ji} with respect to the axis $\mathcal{A}$ and the vector $\mathbf{v}$.

The representation of X given in (6.10) and illustrated in Figure 6.11 shows that X is a "good" configuration space in the sense that it is coordinatized by a set of geometric descriptors, namely c_{ji}, h_{ji}, l_{ji}, which are directly adapted to the perception of the geometry of arrays of points relative to a fixed axis. In particular, the instantaneous rotations may be described within X in a very natural way as the solution set of equations which are *linear* in these coordinates.

Proposition 6.12. E is the subset of X consisting of those elements x whose representation in the form (6.10) has the following properties:

(i) $c_{12} = c_{22} - c_{21} = \ldots = c_{n2} - c_{n1}$.
(ii) $h_{j1} = h_{j2}$ for each $j = 1, \ldots, n$.
(iii) $l_{j1} = l_{j2}$ for each $j = 1, \ldots, n$.

Proof. Let e denote an element of X for which these conditions are satisfied. Let us denote by θ the common value of $c_{22} - c_{21}, \ldots, c_{n2} - c_{n1}$. For each $j = 1, \ldots, n$ denote by h_j and l_j the common values of $h_{j1} = h_{j2}$ and $l_{j1} = l_{j2}$. If also we drop the superscripts on the $c_{21}, \ldots, c_{n1}$ then we can write e in the form

$$e = (\mathcal{A}, \mathbf{v}, c_2, \ldots, c_n, \theta, h_1, \ldots, h_n, l_1, \ldots, l_n). \tag{6.13}$$

As before, let P_{j1} be the point of $\mathbf{R}^3$ whose cylindrical coordinates relative to $\mathcal{A}$ are (c_j, h_j, l_j), where the angle c_1 is measured with respect to $\mathbf{v}$ (so that $c_1 = 0$). Let σ denote the rotation about the axis $\mathcal{A}$ through the angle θ. Then

$$e = (\mathcal{A}, P_{11}, \ldots, P_{n1}; P_{12}, \ldots, P_{n2}) \in E. \quad \blacksquare \tag{6.14}$$

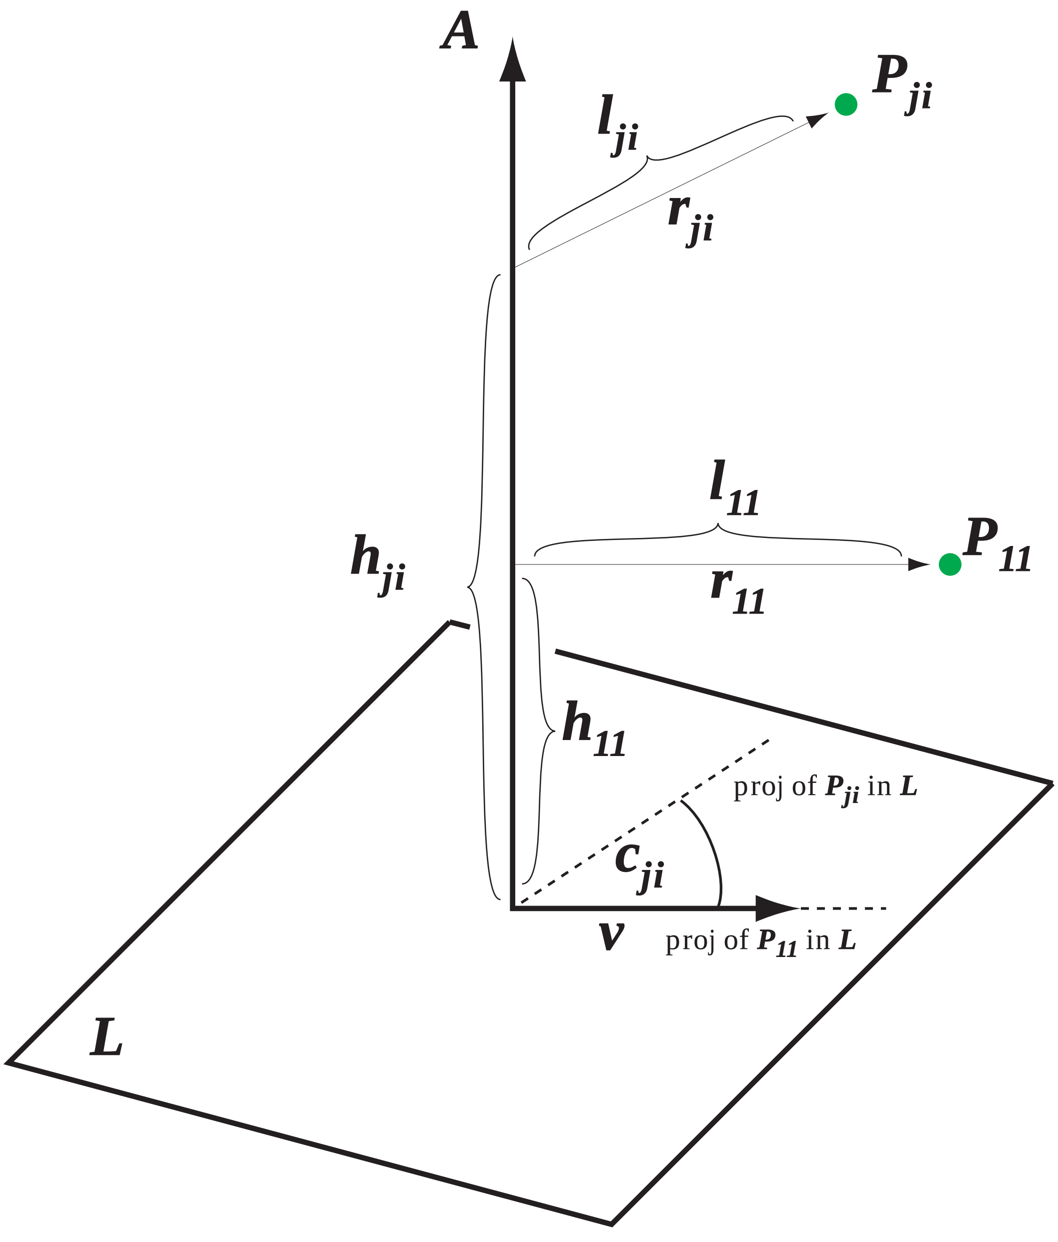

FIGURE 6.11. *Cylindrical coordinate representation of a configuration* $(\mathcal{A}, \mathbf{a}_{11}, \ldots, \mathbf{a}_{n1}; \mathbf{a}_{12}, \ldots, \mathbf{a}_{n2})$. P_{ji} *is the tip of the vector* $\mathbf{a}_{ji}$.

We now introduce the groups G and J.

$$\begin{aligned} G &= \mathrm{SO}(3,\mathbf{R}) \times (\mathbf{S}^1)^{n-1} \times (\mathbf{S}^1)^n \times \mathbf{R}^n \times \mathbf{R}^n \times (\mathbf{R}^*)^n \times (\mathbf{R}^*)^n. \\ J &= \mathrm{SO}(3,\mathbf{R}) \times (\mathbf{S}^1)^{n-1} \times \mathbf{S}^1 \times \mathbf{R}^n \times (\mathbf{R}^*)^n. \end{aligned} \tag{6.15}$$

$\mathbf{S}^1$ is the circle group, namely the additive group $\mathbf{R}/2\pi\mathbf{Z}$, and $\mathbf{R}$ and $\mathbf{R}^*$ are the additive and multiplicative real number groups respectively. Let us denote elements g of G in the form

$$g = \begin{pmatrix} \beta & & \gamma_{21}, \ldots, \gamma_{n1} & \zeta_{11}, \ldots, \zeta_{n1}n & \lambda_{11}, \ldots, \lambda_{n1} \\ & \gamma_{12} & \gamma_{22}, \ldots, \gamma_{n2} & \zeta_{12}, \ldots, \zeta_{n2} & \lambda_{12}, \ldots, \lambda_{n2} \end{pmatrix} \tag{6.16}$$

with $\beta \in \mathrm{SO}(3,\mathbf{R})$, the γ's in $\mathbf{S}^1$, the ζ's in $\mathbf{R}$, and the λ's in $\mathbf{R}^*$. We will write elements $\jmath$ of J in the form

$$\jmath = (\beta, \gamma_2, \ldots, \gamma_n, \delta, \zeta_1, \ldots, \zeta_n, \lambda_1, \ldots, \lambda_n) \tag{6.17}$$

We view J as a subgroup of G by identifying $\jmath$ in (6.17) with the element g of G given by

$$g = \begin{pmatrix} \beta & & \gamma_2, \ldots, \gamma_n & \zeta_1, \ldots, \zeta_n & \lambda_1, \ldots, \lambda_n \\ & \delta & \gamma_2 + \delta, \ldots, \gamma_n + \delta & \zeta_1, \ldots, \zeta_n & \lambda_1, \ldots, \lambda_n \end{pmatrix}.$$

We now describe the action of G on X. Let $x \in X$ be as in (6.10) and $g \in G$ as in (6.16). Then

$$gx = \left(\begin{matrix} \beta\mathcal{A} & & \beta\mathbf{v} & (c_{21} + \gamma_{21}), \ldots, (c_{n1} + \gamma_{n1}) \\ & & c_{12} + \gamma_{12} & (c_{22} + \gamma_{22}), \ldots, (c_{n2} + \gamma_{n2}) \end{matrix} \right.$$

$$\left. \begin{matrix} (h_{11} + \zeta_{11}), \ldots, (h_{n1} + \zeta_{n1}) & \lambda_{11} l_{11}, \ldots, \lambda_{n1} l_{n1} \\ (h_{12} + \zeta_{12}), \ldots, (h_{n2} + \zeta_{n2}) & \lambda_{12} l_{12}, \ldots, \lambda_{n2} l_{n2} \end{matrix} \right). \tag{6.18}$$

Here $\beta\mathcal{A}$ and $\beta\mathbf{v}$ denote the axis and vector in $\mathbf{R}^3$ which are the images of $\mathcal{A}$ and $\mathbf{v}$ under the rotation β. This induces an action of J on E which may then be described as follows: If e is as in (6.13) and $\jmath$ is as in (6.17), then

$$\begin{aligned} \jmath e = (\beta\mathcal{A}, \beta\mathbf{v}, &(c_2+\gamma_2), \ldots, (c_n+\gamma_n), \\ &\theta+\delta, (h_1+\zeta_1), \ldots, (h_n+\zeta_n), \lambda_1 l_1, \ldots, \lambda_n l_n). \end{aligned} \tag{6.19}$$

We can now see that, given any pair $(x, \bar{x})$ of elements of X, there is a unique $g \in G$ such that $\bar{x} = gx$. For suppose x is as in 6.10 and $\bar{x}$ has components $\bar{\mathcal{A}}$, $\bar{\mathbf{v}}$, $\bar{c}_{21}$ etc. Given any two pairs $(\mathcal{A}, \mathbf{v})$ and $(\bar{\mathcal{A}}, \bar{\mathbf{v}})$, each consisting of an oriented axis and a unit vector orthogonal to it, there is a unique β such that $(\bar{\mathcal{A}}, \bar{\mathbf{v}}) = (\beta(\mathcal{A}), \beta(\mathbf{v}))$. This gives us the coordinate β of the required $g \in G$. Referring to Figure 6.11 and Equation 6.18, it is clear that the remaining coordinates are fixed by the requirements that

$$\begin{aligned} \gamma_{ji} &= \bar{c}_{ji} - c_{ji} \pmod{2\pi}, \\ \zeta_{ji} &= \bar{h}_{ji} - h_{ji}, \\ \lambda_{ji} &= \bar{l}_{ji}/l_{ji}. \end{aligned}$$

Therefore, X is a principal homogeneous space for G, and E is a principal homogeneous space for J. (It follows from this that the dimension of E is the same as the dimension of J which is $3+3n$, as is easily seen from 6.15.)

To complete the description of our symmetric framework $\Theta = (X, Y, E, S, G, J, \pi)$ we need to define the fundamental map $\pi: G \to Y$; the definition uses the map $p: X \to Y$ of (6.5). Actually, there is no single canonical choice of π here, but rather a canonical set of π's, and the relations between them can be stated precisely.

For each $x_0 \in X$, we have a bijective map $f_{x_0}: G \to X$ defined by $f_{x_0}(g) = g(x_0)$. That is, we identify X with G by displaying it as the orbit, under the action of G, of a distinguished element $x_0 \in X$. Then, for $x_0 \in E$, we let

$$\pi^{x_0} = p \circ f_{x_0}: G \to Y. \tag{6.20}$$

Thus,

$$G \xrightarrow{\pi^{x_0}} Y$$

is the composition

$$G \xrightarrow{f_{x_0}} X \xrightarrow{p} Y.$$

Since $x_0 \in E$, then $f_{x_0}(J) = Jx_0 = E$ so that $\pi^{x_0}(J) = p(E) = S$ by definition of S (6.8). We summarize:

6.21. With notation as above, for each choice of $x_0 \in E$,

$$\Theta = (X, Y, E, S, G, J, \pi^{x_0})$$

is a symmetric observer framework with fundamental map π^{x_0}.

We then get a reflexive framework

$$\{(X, Y, E, S, \pi_e^{x_0}) | \, e \in E\}.$$

Recall that the perspective maps $\pi_e^{x_0}: X \to Y$ are defined from the fundamental map π^{x_0} by the formula

$$\pi_e^{x_0}(x) = \pi^{x_0}(xe^{-1})$$

where xe^{-1} denotes the unique element of G sending e to x. Suppose, for example, that $x = ke$, i.e., $xe^{-1} = k$. Then, by definition of π^{x_0} this formula may be written

$$\pi_e^{x_0}(ke) = p(kx_0). \tag{6.22}$$

This may be interpreted as saying that the use of π^{x_0} as the fundamental map of the framework means that each observer in the framework "thinks of itself" as having configuration x_0 and perspective p. To understand this, view group elements in G as indicating "displacement." Then 6.22 says the following: the premise acquired by an observer with configuration e (in the framework whose fundamental map is π^{x_0}) when interacting with an observer displaced from it by k, is the same as the premise acquired by an observer with fixed perspective $p: X \to Y$, when interacting with an observer displaced from it by k.

Finally, we note that a straightforward calculation shows that the dependency of this structure on the noncanonical choice of x_0 can then be stated as follows:

6.23. For e, x_0, $x_0' \in E$

$$\pi_e^{x_0} = \pi_{(x_0' x_0^{-1})e}^{x_0'}.$$

In chapter nine we use this framework to give a participator-dynamical interpretation of "incremental rigidity schemes" for the human visual perception of rigid motion.