

Chapter 12

EPISTEMIC GROUP STRUCTURE AND COLLECTIVE AGENCY

While we have looked extensively at individual agents and their interaction, a further basic feature in rational agency is the formation of collective entities: groups of agents that have information, beliefs, and preferences, and that are capable of collective action. Groups can be found in social epistemology, social choice theory, and the theory of coalitional games. Some relevant notions occurred in the preceding chapters, especially, common knowledge – but they remained a side theme. Indeed, the logical structure of collectives is quite intricate, witness the semantics of plurals and collective expressions in natural language, which is by no means a simple extension of the logic of individuals and their properties. This book develops no theory of collective agents, but this chapter collects a few themes and observations, connecting logical dynamics to new areas such as social choice.

12.1 Collective agents in static logics

Groups occur in the epistemic logic of Chapter 2 with knowledge modalities such as $C_G\varphi$ or $D_G\varphi$. But the logic had no explicit epistemic laws for natural group forming operations such as $G_1 \cup G_2$, $G_1 \cap G_2$.²⁷⁵ Actually, two logics in this book did provide group structure. One is the epistemic version *E-PDL* of propositional dynamic logic in Chapter 4, where epistemic program expressions defined complex ‘collective agents’ such as $i ; (?p ; j \cup k)^*$. Another was mentioned in Chapter 2: the combined topologies of van Benthem & Sarenac 2005. Even so, epistemic logic still needs a serious extension to collective agents: adding common or distributed knowledge is too timid. Here, group structure and information may be intertwined: for instance, membership of a group seems to imply that one knows this.²⁷⁶ Moreover, groups also have beliefs and preferences, and they engage in collective action. In all these cases, generalization may not be straightforward. Collective attitudes or actions may reduce to behaviour of individual group members, but they need not. One sees this variety with collective predicates in natural language, such as “Scientists agree that the Earth is warming”, or “The sailors quarreled”. There is no canonical dictionary semantics

²⁷⁵ The indices in the standard notation are always concrete sets of agents. Adding an explicit abstract *group algebra* to epistemic logic would be an interesting generalization.

²⁷⁶ Also, groups are often held together by *trust* between agents, a delicate epistemic feature. Cf. Holliday 2009 for a *DEL*-style analysis of the dynamics of reported beliefs, testimony and the building of trust, based on various new soft upgrades beyond those studied in Chapter 7.

for what these things mean in terms of individual predication. Finally, there is a temporal aspect. Groups may change their composition, so their structure may be in flux, with members entering and leaving – and dynamic and temporal logics come into play.

12.2 Group knowledge and communication

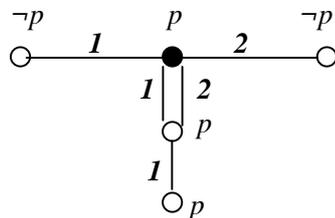
One dynamic perspective on groups is that they must *form* and stay together by being involved in shared activities. This fits well with logical dynamics, and we will pursue some illustrations in this chapter, first with knowledge, then with belief.

For a start, the recurrent static notion of common knowledge in this book is not Heavensent: it is the result of doing work, such as making public announcements. Chapter 3 raised the issue what information groups can achieve through internal communication. We discuss this a bit further here, though a general theory is still beyond our reach:

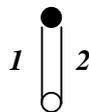
‘Tell All’: maximal communication Consider two epistemic agents in an information model \mathbf{M} , at an actual world s . They can tell each other things they know, cutting down the model. Suppose they are cooperative. What is the best correct information they can give?

Example The best agents can do by internal communication.

What is the best that can be achieved in the following model? ²⁷⁷



Geometrical intuition suggests that this must be:



This is correct. For instance, 1 might say “I don’t know if p ”, ruling out the rightmost world, and then 2 “I don’t know either”, ruling out the leftmost and bottom world. They could also say these things simultaneously and get the same effect. ■

For simplicity, in what follows, we stick to *finite models* \mathbf{M} where each epistemic relation is an equivalence relation. Clearly, any sequence of updates where agents say all they know must terminate in some submodel that can no longer be reduced. This is reached

²⁷⁷ To make things more exciting, one can mark worlds with unique proposition letters.

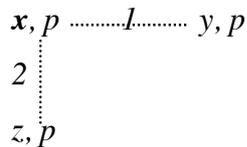
when everything each agent knows is true in every world, and hence common knowledge. It is not clear that there is a unique ‘communication core’ (henceforth, the *Core*) to which this must converge, but van Benthem 2000 proposes the set of worlds reachable from the actual world in M via each uncertainty link. These interpret distributed knowledge in the sense of Chapter 2 – so we want to know when this is reached. A related issue was raised in Chapter 11 as a property of informational protocols, namely, when it holds that

communication turns distributed knowledge into common knowledge.

When we try to make all this more precise, there are complications:

Example Problems with bisimulation.

In the following model, the communication core is just the actual world x , but all worlds satisfy the same epistemic formulas:



The reason is that there is a bisimulation contraction (cf. Chapter 2) to a one-world model. This may not look bad yet, but we will see worse scenarios below. ■

Communication does not get us to the communicative core here. An immediate response might be to reject such inflated models, working with bisimulation contractions only. This improves things to a certain extent:

Proposition On finite bisimulation-contracted models, the Core can be reached by one simultaneous announcement of what each agent knows to be true.

Proof By the contraction under bisimulation, all worlds t in the model satisfy a unique defining epistemic formula δ_t , as shown in Chapter 2. Each agent can now communicate all she knows by stating the *disjunction* $\vee \delta_t$ for all worlds t she considers indistinguishable from the actual one. This simultaneous move cuts the model down to the actual world plus all worlds reachable from it in the intersection of all \sim_i -alternatives. ■

As with the Muddy Children puzzle, things get more complicated when we let agents speak sequentially. Still, we do have one positive result:²⁷⁸

²⁷⁸ The reader can skip the following passage without loss of continuity.

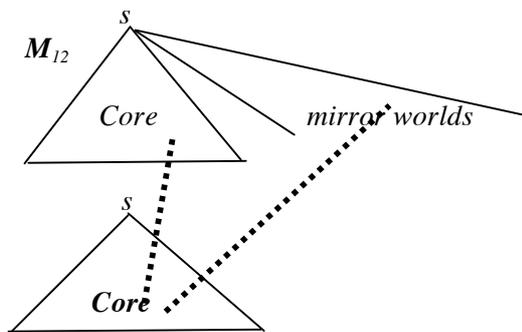
Proposition On bisimulation-contracted finite models with two agents, the Core is reached by two single-agent announcements plus one more contraction.

Proof Let agent 1 state all she knows as before. This reduces the initial model M, s to a model M_1, s with just the actual world plus its \sim_1 -successors. Now we want to let agent 2 define the set of remaining \sim_2 -successors in a similar fashion. But there is a difficulty: a new model after a *PAL*-update need not be contracted under bisimulation, as we saw in Chapter 3. And if we contract first, to make things right, we may not get to the Core, since different worlds in the core may now contract to one (after the first update, they may have come to verify the same epistemic formulas in M_1).

To get around this, let agent 2 first state the best she can. This is the set of formulas true in all her \sim_2 -successors from s (in the Core) plus all worlds having the same epistemic theory as one of these in M_1 . One formula suffices for defining such a finite set (we omit the simple argument). The result is a submodel M_{12}, s whose domain consists of the Core plus perhaps some ‘mirror worlds’ that were modally equivalent to some world in the Core in the model M_1 . Now, our claim is this:

Lemma Taking the identity on the Core and connecting each mirror world in M_{12} to all Core worlds satisfying the same epistemic formulas in M_1 is a bisimulation.

Here is a picture of the situation (the relation is total between the two models):



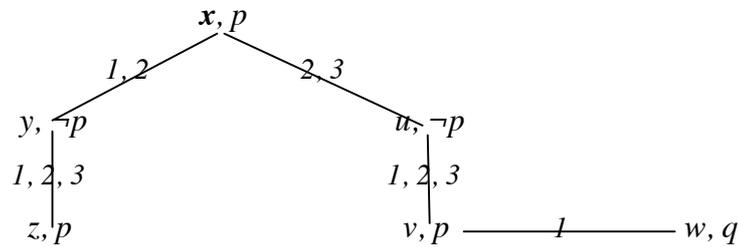
We have to check that the given map is a bisimulation from M_{12} to the free-standing model **Core** with just the core worlds. The atomic clause is clear by definition. As for the zigzag clauses, the relation \sim_1 is total in both models, and nothing has to be proved for it. Next consider \sim_2 , with a match of world x in M_{12} with a world y in **Core**. Since \sim_2 is total in **Core**, linking any given \sim_2 -successor for x in M_{12} to a suitable \sim_2 -successor of y is automatic. In the opposite direction, there are two cases. If x was in the Core, and y has an \sim_2 -successor z in **Core**, that same z also serves for x in M_{12} . But if x was a mirror world, we

argue differently. By the definition of our relation, x satisfied the same epistemic formulas in M_1 as the Core world y . Now, since $y \sim_2 z$ in the Core, M_1 , $y \models \langle 2 \rangle \varphi_z$ where φ_z defined all worlds sharing z 's epistemic theory in M_1 . But then x satisfied $\langle 2 \rangle \varphi_z$ as well, giving it a \sim_2 -successor w in M_1 satisfying φ_z . But then w was a mirror world for z in the Core, and so it stayed in M_{12} . Hence it is the \sim_2 -successor for x as needed, that got mapped to z . ■

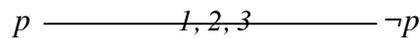
The very complexity of this argument for such a simple conclusion is ominous. With three agents, the core need not be reached at all, unless we are careful:

Example (Jelle Gerbrandy, p.c.) Different conversation, different information.

Consider the following model, with three agents, and actual world 1 :



Each world is defined by a unique epistemic formula. The Core is just $\{x\}$, and it is reached if first 1 and then 3 state all they know. But if 3 starts by announcing $\neg q$, this rules out world w that made the difference for $1, 2$ on the left and right. The new model arising then is easily seen to be bisimilar to the much simpler



and no further announcements will help. ■

Discussion: what is implicit knowledge of a group? There are subtleties here. In the final model *Core*, worlds may satisfy formulas different from their epistemic theory in the initial model: the D_G/C_G conversion that we were after applies at best to factual formulas.²⁷⁹ Also the use of bisimulation contraction works for our standard epistemic language, whereas the modality $D_G\varphi$ itself was not invariant for bisimulation (cf. Chapter 2): a case of imbalance.

But one can also argue that all this rather speaks *against distributed knowledge*, as failing

²⁷⁹ It may be distributed knowledge that no one knows where the treasure is, but inside the Core, the location may be common knowledge. Compare quantifier restriction versus relativization: $D_G\varphi$ looks only at worlds in the Core, but it evaluates the formula φ there *in the whole model*. By contrast, internal evaluation in the Core is like totally relativized statements $(\varphi)^{Core}$.

to capture the dynamic intuition of knowledge that one makes explicit by communication. Our dynamic scenarios *themselves* might be taken as yielding a more appropriate notion of *implicit knowledge* that agents can obtain by communication and contraction.

Constrained assertions The preceding scenarios were open, as long as agents said things they knew to be true. But in the puzzle of the Muddy Children of Chapters 1, 3, children could only state their knowledge or ignorance of their mud status, and in the end common knowledge resulted. This is a frequent scenario: maximal communication *with one specific assertion*. Chapters 10, 15 use this to solve games by repeated announcements of players' rationality. As we will see in Chapter 15, the limit of repeatedly announcing a formula φ (even in infinite models) is a submodel where one of two things has happened: φ is common knowledge ('self-fulfilling'), or $\neg\varphi$ is common knowledge ('self-refuting').

Beliefs and soft information Communication scenarios can play with beliefs just as well as knowledge. Even further surprising phenomena come to light then. Chapter 15 has some recent examples from Dégrémont & Roy 2009 showing how announcing differences in beliefs can switch who believes what – though in the limit, agreement will result. It also has results from Baltag & Smets 2009 on repeated announcement of soft information (cf. Chapter 7), where no worlds are eliminated, but a plausibility order gets modified.²⁸⁰ A more modest aim of communication might be creating a shared plausibility pattern.

Dynamics of communication Groups form and persist through actions of communication, and hence group knowledge and belief are fundamentally linked to this dynamics. This view uncovers a wide range of complex phenomena, of which we have only seen a few.²⁸¹ Our tentative tour suggests many new logical problems, but right now, we move on.

12.3 Belief and preference merge for groups

Our topic so far was modification of information and plausibility for single agents in an interactive process of communication. But a more radical step can be made. Consider the formation of *collective beliefs* or preferences, say, when a group forms in an encounter. Technically, group merge requires integration of separate epistemic-doxastic models that may contradict each other or involve different languages. Belief merge and social choice are two particular headings where this issue arises. Here we just present one way to go.

²⁸⁰ But there are again ugly surprises, such as *cycles* when updating plausibility orders.

²⁸¹ Apt, Witzel & Zvesper 2009 study group communication with explicit channel structure.

Andréka, Ryan & Schobbens 2002 (*ARS*) propose a model that fits well with our logics. For a start, they note that the usual accounts of creating collective relations from individual ones work with input that is too poor, viz. just a set of relations. In general, we need richer input: a graph of dominance order in the group. For this to work well, as in Chapters 7, 9, relations are reflexive transitive pre-orders, not necessarily connected:

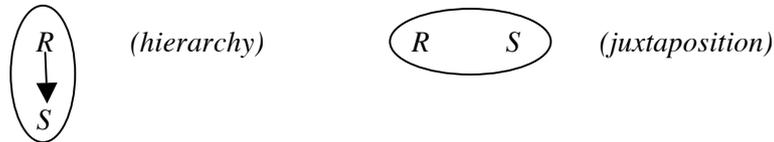
Definition Prioritized relation merge.

Given an ordered *priority graph* $\mathbf{G} = (G, <)$ of indices for individual relations that may have multiple occurrences in the graph, the *merged group priority relation* is:

$$x \leq_{\mathbf{G}} y \text{ iff for all indices } i \in \mathbf{G}, \text{ either } x \leq_i y, \text{ or there is some } j > i \text{ in } \mathbf{G} \text{ with } x <_j y \quad \blacksquare^{282}$$

Example Merging via simple graphs

Consider the following two simple graph pictures with two relations each:



Putting the relation R above S , the merged group priority orders objects lexicographically: $x \leq y$ iff either $x R y \wedge x S y$ or there is a difference in the S relation and $x R^+ y$ with R^+ the strict version of R . Putting R alongside S leads to the intersection $x R y \wedge x S y$. ■

ARS merge already occurred in Chapter 9. There, perhaps confusingly, it was used to find a preference relation for a single agent that goes by an ordered family of criteria. The latter were propositions P , and so relations in the graph were of the form

$$x \leq_P y \text{ iff } (Px \rightarrow Py).$$

Girard 2008, Liu 2008 show how this subsumes belief merge, priority-based preference, ceteris paribus logic, and agendas in inquiry (cf. Chapter 6).²⁸³

As for dynamics, there are natural operations that change and combine priority graphs:

sequential composition $G_1 ; G_2$ (putting G_1 on top of G_2 , retaining the same order inside) and *parallel composition* $G_1 \parallel G_2$ (disjoint union of graphs).

²⁸² Thus, either x comes below y , or if not, y ‘compensates’ for this by doing better on some comparison relation in the set with a higher priority in the graph.

²⁸³ Andréka, Ryan & Schobbens prove that priority graphs are universal as a preference aggregation procedure, and give a complete graph algebra. Girard 2008 has an alternative modal analysis.

As we just saw, $G_1 // G_2$ defines intersection of the separate relations for G_1 , G_2 , and $G_1 ; G_2$ defines lexicographic order (cf. the radical upgrade of Chapter 7). *ARS* is a step toward an abstract logic of group agency, and we will see one particular use in the next section.

12.4 Belief change as social choice

Groups are composed of agents, but agent themselves may also be groups of entities when we analyze their structure in detail. To show how this can be illuminating, we return to the belief revision rules of Chapter 7, and ask for a more principled analysis than what was offered there. More concretely, we will analyze belief revision as a process of group merge for ‘signals’. We will show how the basic rule of Priority Update is social choice between relations in an initial model M and an event model E , where the relation in $M \times E$ results from either treating the two as equally important, or taking them in a hierarchy.

Abstract setting: ordering pair objects given component orders Two pre-orders (A, R) and (B, S) are given, with possibly different domains A, B : for instance, think of a doxastic model M and an event model E with their separate domains and plausibility orders. Now we seek to order the product $A \times B$ by a relation $O(R, S)$ over pairs (a, b) .²⁸⁴

The main analogy The Priority Update Rule took the event model E to rank above the doxastic model M in terms of authority, defining the following order in $M \times E$:

$$(s, e) \leq (t, f) \text{ iff } (s \leq t \wedge e I f) \vee e < f. \quad ^{285}$$

With pre-orders, we state intuitions for four cases $x < y$, $y < x$, $x \sim y$ (indifferent), and $x \# y$ (incomparable). For vividness, we mark these cases graphically as \rightarrow , \leftarrow , \sim , and $\#$.

Intuitive conditions on plausibility update What sort of process are we trying to capture? I first choose a very restrictive set to zoom in exclusively on Priority Update. Later on I relax this, to get greater variety in update rules. The first condition says that the choice should not depend on individual features of objects, only their ordering pattern:

Consider any two permutations of A and B . Thinking of A, B as disjoint sets, without loss of generality, we can see this as one permutation π . We require the following behaviour:

²⁸⁴ In general, we only need to order a *subset* of this full product, as with *DEL* in Chapter 4.

²⁸⁵ Here $e I f$ stands for *indifference*: $e \leq f \wedge f \leq e$. By a simple computation on the rule, we get the version for strict equivalence $(s, e) < (t, f)$ iff $(s < t \wedge e \leq f) \vee e < f$.

Condition (a): Permutation invariance

$$O(\pi[R], \pi[S]) = \pi [O(R, S)].$$

This standard condition imposes a strong uniformity on possible formats of definition (cf. the accounts of logicity in van Benthem 2002B). Here is one more constraint:

Condition (b) Locality

$$O(R, S) ((a, b), (a', b')) \text{ iff } O(R|\{a, a'\}, S|\{b, b'\}) ((a, b), (a', b')).$$

Thus, we only order using the objects occurring in a pair, a form of context-independence akin to Independence from Irrelevant Alternatives in social choice.²⁸⁶

Table format Together, Permutation Invariance and Locality force any operation O to be definable by its behaviour in the following 4x4-Table:

$S \text{ on } b, b'$	\rightarrow	\leftarrow	\sim	$\#$
$R \text{ on } a, a'$	\rightarrow	-	-	-
	\leftarrow	-	-	-
	\sim	-	-	-
	$\#$	-	-	-

Here entries stand for the 4 isomorphism classes on two objects: all that matters given the invariance condition. Under certain conditions, some parts of this Table are even forced. We will fill in the same four types of entry in the Table, subject to further conditions.²⁸⁷

Choice conditions on the aggregation procedure Now we state some conditions on how the component relations are going to be used in the final result. Even though we will only be using them for a choice with two actors, they make sense more generally. The names have been chosen for vividness, but nothing is claimed for them in naturalistic terms:

²⁸⁶ Locality holds for the *radical update* $\uparrow A$ of Chapter 7, that can be modeled by Priority Update using a two-point event model with an A -signal more plausible than a $\neg A$ -signal. But Locality fails for *conservative update* $\uparrow A$ where we place only *the best A-worlds* on top in the new ordering. Checking if worlds are *maximal in A* requires running through other worlds.

²⁸⁷ *Caveat*. Strictly speaking, one might just want to put YES/NO in the slots marking whether the relation \leq holds between the pairs. In using the four types, strictly speaking, one should check that all intuitions to be stated hold for Priority Product Update as defined above.

Condition (c) Abstentions

If a subgroup votes indifferent (\sim), then the others determine the outcome.

Condition (d) Closed agenda

The social outcome always occurs among the opinions of the voters.

This implies *Unanimity*: “if all members of a group agree, then take their shared outcome”, but it is much stronger. Finally, consider agents who care, and are not indifferent about outcomes. An ‘over-rule’ is a case where one opinion wins over the other.

Condition (e) Overruling

If an agent’s opinion ever overrules that of another, then her opinion *always* does.

This goes against the spirit of democracy and letting everyone win once in a while, but we should not hide that this is what the Priority Rule does with its bias toward the last event.

Our main result now captures Priority Update, though with a twist. We derive that input order must be hierarchical. But we do not force the authority to be the second argument – say, the event model E .²⁸⁸ Thus our result speaks of “a”, not “the”, Priority Update:

Theorem A preference aggregation function is a Priority Update iff it satisfies
Permutation Invariance, Locality, Abstentions, Closed Agenda, and Overruling.

Proof First, Priority Update satisfies all stated conditions. Here one needs to check that the original formulation boils down to the case format in our Table. For instance, if event arguments are incomparable, this will block any comparison between the pairs, whence the last column. Also, if $e < f$, and $s < t$, it is easy to check that then $(s, e) < (t, f)$. Etcetera.

Conversely, we analyze possible Table entries subject to our conditions. Here the diagonal is clear by Unanimity, and the row and column for the indifference case by Abstentions:

S on b, b'	\rightarrow	\leftarrow	\sim	$\#$
R on a, a'	\rightarrow	\rightarrow	1	\rightarrow
	\leftarrow	3	\leftarrow	\leftarrow
	\sim	\rightarrow	\leftarrow	\sim
	$\#$	5	6	$\#$

²⁸⁸ The other option of giving priority to the first argument (the initial model M), is a conservative anti-Jeffreyan variant (cf. Chapter 8) where little learning takes place.

This leaves six slots to be filled. But there are really only three choices, by simple symmetry considerations. E.g., an entry for $\rightarrow \leftarrow$ automatically induces one for $\leftarrow \rightarrow$.

Now consider slot I . By Closed Agenda, this must be either \rightarrow or \leftarrow . Without loss of generality, consider the latter: S overrules R . Using Overruling to fill the other cases with S 's opinion, and applying Permutation Invariance, our Table is this:

S on b, b'	\rightarrow	\leftarrow	\sim	$\#$
R on a, a'	\rightarrow	\rightarrow	\leftarrow	\rightarrow
	\leftarrow	\rightarrow	\leftarrow	\leftarrow
	\sim	\rightarrow	\leftarrow	\sim
	$\#$	\rightarrow	\leftarrow	$\#$

It is easy to see that this final diagram is precisely that for Priority Update in its original sense. The other possible case would give preference to the ordering on M .²⁸⁹ ■

Weaker conditions: additional update rules Now we relax conditions to allow democratic variants where arguments count equally – in a special case, a flat epistemic product update of M and E where $(s, e) \leq (t, f)$ iff $s \leq t$ and $e \leq f$.²⁹⁰ Now Closed Agenda fails: with this rule, the above clash case $\rightarrow \leftarrow$ ends up in $\#$. Instead, we state two new principles:

Condition (f) Unanimity

If voters all agree, then their vote is the social outcome.

Condition (g) Alignment

If anyone changes their vote to get closer to the group outcome,
the group outcome does not change.

Theorem A preference merge function satisfies Permutation Invariance, Locality, Abstentions, Overruling, Unanimity, and Alignment iff it is either (a) a priority update, or (b) flat product update.²⁹¹

²⁸⁹ Both are instances of the basic ‘But’ operator of ARS, i.e., sequential graph composition.

²⁹⁰ Intersection of relations was ARS-style parallel composition: ‘And’, instead of ‘But’.

²⁹¹ There are analogies with May’s Theorem on majority voting: cf. Goodin & List 2006.

Proof The crucial step is now that, without Closed Agenda, Slot 6 in our diagram

S on b, b'	\rightarrow	\leftarrow	\sim	$\#$
R on a, a'	\rightarrow	\rightarrow	6	\rightarrow
	\leftarrow	3	\leftarrow	\leftarrow
	\sim	\rightarrow	\leftarrow	\sim
	$\#$	5	6	$\#$

may also have entries \sim or $\#$. But Alignment rules out \sim . If S changes its vote to \sim , the outcome should still be \sim , but it is \rightarrow . So, the entry must be $\#$. But then, using Alignment once more for both voters (plus Permutation Invariance), all remaining slots are $\#$:

S on b, b'	\rightarrow	\leftarrow	\sim	$\#$
R on a, a'	\rightarrow	\rightarrow	$\#$	\rightarrow
	\leftarrow	$\#$	\leftarrow	\leftarrow
	\sim	\rightarrow	\leftarrow	\sim
	$\#$	$\#$	$\#$	$\#$

This is clearly the table for the flat update. ■

Variations Our results are just a start. For instance, dropping Overruling allows for mixtures of influence for M and E .²⁹² Also, other themes from social choice theory make sense. For instance, what corresponds to commonly made restrictions on the individual preference profiles fed into the rule?²⁹³ But perhaps the main benefit is the view itself. I am intrigued by the idea that ‘I’ am ‘we’: the social aggregate of all signals in my life.

12.5 Further directions: dynamics of deliberation

This chapter consisted of some observations showing how logical dynamics interfaces with groups as entities in their own right. Many themes need to be elaborated, but also, the list

²⁹² More might also be said about relative power of update rules in creating new relational patterns. Compare Priority Update versus Flat Product Update. Which rule is more general if we allow re-encoding of the relational arguments that provide their inputs?

²⁹³ My answer: assumptions on the *continuity* of the information streams we encounter in the world. It has often been observed that we only learn well if the universe is kind enough to us.

we considered is far from complete.²⁹⁴ ²⁹⁵ Here is one further direction to be explored which I find particularly appealing from the general viewpoint of logical dynamics:

Dynamic epistemic logic fits well with social choice theory, and it can provide two things: *informational structure*, and more finely-grained *procedure*. For the first, think of Arrow's Theorem, and the horror of a dictator whose opinions are the social outcome. But even if it is common knowledge 'de dicto' that there is such a dictator, this does no harm if there is no person whom we *know* 'de re' to be the dictator. Not even the dictator herself may know. To see the real issues of democracy, we need social choice plus epistemic logic.

As for the second aspect, social choice rules seem far removed from the practice of rational communication and debate, and intuitive notions of fairness having to do with these. One would want to study in detail how groups arrive at choices by *deliberating*, and ways in which agents then experience preference changes extending the concerns of Chapter 9.²⁹⁶ Perceived fairness resides as much in the process as in a final act of voting. But this is not to say that our dynamics has all the answers. To the contrary, dynamic epistemic logic should start looking at discussion and debate, since it is there that information update, belief revision, and preference change among agents occur at their most vivid.²⁹⁷

12.6 Literature

There is a huge literature on group behaviour in many fields: cf. Anand, Pattanaik & Puppe, eds., 2009 on social choice and economics, or Sugden 2003 in social epistemology.

²⁹⁴ A deeper study might profit from the *linguistic semantics* of individual and collective predicates. See the chapters on Temporality and Plurals in van Benthem & ter Meulen, eds., 1997.

²⁹⁵ Cf. also Dégrémont 2010, Kurzen 2010 on *computational complexity* for group activities.

²⁹⁶ This includes two processes: adjustment of individual preferences through social encounters, and joining in the formation of new groups with preferences of their own.

²⁹⁷ Our dynamic logics are still far from dealing with subtle procedural phenomena in deliberation, such as well-established rules for speaking, voting, or the dynamics of 'points of order'. Here it may be time to join forces with *Argumentation Theory* and other dialogical traditions. Logicians like Barth, Krabbe, or Gabbay have been doing this for a long time. Van Benthem 2010A is a new attempt, tying logical dynamics to the issues of procedure raised in Toulmin 1958.

Andréka, Ryan & Schobbens 2002 is our key source for belief and preference merge. The communication dynamics in this chapter comes mainly from van Benthem 2000, 2006C. Van Benthem 2009F is the first *DEL*-style treatment of belief revision as social choice.