Simple models of knowledge

preliminary version

Elias Thijsse
ITK, Tilburg University, the Netherlands

E–mail: thysse@kub.nl

February 13, 1992

Contents

1 Why go partial in knowledge representation?

The possibility of representing knowledge in finite Kripke models was studied in [Th91b]. The main reasons for the present research stem from results reported in the earlier paper. Finite representation in classical possible worlds models (i.e. Kripke structures with a bivalent truth assignment) turned out to be possible under certain conditions. For example, some piece of **S5**-knowledge α can be characterized by a 'global miniature' M (i.e. M verifies precisely all the **S5** consequences in each world) iff α is *introspective* ($\alpha \vdash_{\mathbf{S5}} K\alpha$).

However, though the existence of classical miniatures encourages further research in this direction, some drawbacks of total models point at the need of partiality in model-theoretic knowledge representation. For, though the word 'miniature' indicates a tiny thing (reflecting our initial intention), the classical miniatures are by no means small.[1] To be more specific, we will calculate the size of one type of classical **S5**–miniatures below.

Moreover, most positive results were obtained for **S5**, which is, in some sense, the simplest modal logic. For the epistemic logics **S4** (which is somewhat closer to human knowledge) and $\mathbf{S5_{(m)}}$ (which is proper for the case of many agents reasoning according to **S5**) there are no equally positive results. In particular, one proof in [Th91b] can easily be generalized to show that *incomplete* **S4**-knowledge cannot be modelled in finite classical Kripke models. A similar negative result for $\mathbf{S5_{(m)}}$ follows essentially from [FHV91]: it is shown there that non-empty finite Kripke structures can model some and only contingent common knowledge.

Finally, we notice that modal systems such as **S4** do not account for the way in which human beings deal with knowledge — real agents are not perfect reasoners, therefore they will not know everything that follows from their knowledge. These observations lead to the following central questions:

- Can we improve upon the complexity of the representation of **S5**-knowledge? *A priori*, partial models seem proper to diminish the size of the miniatures.

- Can we represent the knowledge with respect to other epistemic logics, such as **S4** and $\mathbf{S5_{(m)}}$ by means of partial models?

[1] Perhaps this shift parallels the history of the word 'minim' (the next item in Longmans' International Reader's Dictionary): formerly a very short note of music, now quite a long one. Also, there is no consensus on the etymology of the word 'miniature', but the relation to Latin 'minim' is at least one of the possible sources.

- Can we represent the type of knowledge that is closer to the way in which human beings think, i.e., logically spoken, knowledge that accords to a weaker epistemic logic?[2]

The answers to these points depend on the kind of evaluation.[3] Assuming a non-falsification perspective on valid consequence, we will show that the total miniatures are among the smaller ones. But under a verification perspective, the gain of partiality is more substantial since the miniatures will generally be much smaller and will follow the rules of somewhat weaker logics.

Complexity of classical miniatures

Total miniatures soon become very large. In fact, the smaller the relative amount of information, the larger the model will be. Some of the worst cases are those of complete ignorance with respect to a number of propositional variables. Assume, for example, that the only information is Kp and the system is totally ignorant with respect to (only) three other atoms. Then the miniature will consist of 1024 worlds divided over 255 components. In general, if there is no information about r atoms, whereas the other atoms are completely known (i.e. either Kp_i or $K\neg p_i$ for, say, $i = 1, \ldots, n - r$), the number of worlds and components of the (smallest!) miniature is superexponential in r. More precisely,

Proposition 1 (size of miniature for simple ignorance)
A classical miniature modelling complete ignorance of r propositional variables and complete knowledge of the other variables, has $\sharp C_r = 2^{2^r} - 1$ components and a total number of $\sharp W_r = 2^{2^r + r - 1}$ worlds.

Proof: First notice that the miniature is isomorphic to the model that characterizes zero information with respect to r atoms (simply drop the uniform specification for the known p_i out of the worlds). This model consists of all non-empty tight submodels[4] of the largest tight model for r atoms, which contains 2^r worlds. So

[2] cf. [FH88], [Th91a] and [Th92].

[3] [Th90b] discusses various kinds of *non-schematic* (i.e. not requiring closure under substitutions) consequence relations in partial (modal) logic. A complementary overview of various kinds of typically modal, schematic consequence relations for total models is [FHV90].

[4] See [Th91b].

there are $2^{2^r} - 1$ components. The total number $\sharp W_r$ of worlds in this model can be calculated by an easy combinatorial argument:[5]

$$\sum_{i=1}^{2^r} i \cdot \binom{2^r}{i} = \sum_{i=1}^{2^r} 2^r \cdot \binom{2^r - 1}{i - 1} = 2^r \sum_{i=0}^{2^r - 1} \binom{2^r - 1}{i} = 2^r \cdot 2^{2^r - 1} = 2^{2^r + r - 1}.$$

■

This 'worst case' analysis is even of some practical importance: a relational database can be relatively empty, that is, the number of atomic (predicate logical) formulas may be quite large, whereas the number of known facts small, and the miniature consequently gigantic (if $r \sim 100$, $\sharp W_r \sim 10^{10^{30}}$). As we will see, especially with such *simple ignorance*, partial miniatures have a dramatically better performance: given the right perspective, the model will consist of just one (1) world!

2 F–miniatures

First we will give the definition of F–miniature, 'F' for non-falsification or falsifiability. Recall from [Th90b] that $M, s \not\dashv \varphi$ means that φ is not false in s according to model M, $M \not\dashv \varphi$ that $M, s \not\dashv \varphi$ for *every* s in M and $D \not\dashv \varphi$ that $M, s \not\dashv \varphi$ for *every* model M and situation s such that $M, s \not\dashv \delta$ for *all* $\delta \in D$. So the slash in the consequence relation of relative falsifiability has a fixed meaning and does *not* indicate non-consequence.

Definition 1 *M is an* **F–miniature** *for D iff M is finite and* $M \not\dashv \varphi \Leftrightarrow D \not\dashv \varphi$ *holds for each* φ.

Or, equivalently, M is an F–miniature iff

- M is finite,
- $M \not\dashv D$,
- $M \not\dashv \varphi \Rightarrow D \not\dashv \varphi$.

Until further notice we will concentrate on models with an equivalence accessibility relation. Moreover, we assume the models to be coherent, i.e.

[5] Alternatively, the number $\sharp W_r$ may be understood as follows: each of the 2^r state–descriptions occurs in $2^{2^r - 1}$ components (i.e. arbitrary subsets of other state–descriptions attached to it).

truth-value assignment can be defined or undefined, but not overdefined. Then the logic of the F–inferences is essentially **S5** without the (propositional and modal) *ex falso* rules (but **with** the rules of *tertium non datur*). This logic will be called **S5*** henceforth.[6]

We can give a syntactic criterion for the existence of F–miniatures. Both result and proof resemble the analogous case for classical miniatures.

Theorem 1 *D has an F–miniature iff* $D \vdash_{\mathbf{S5^*}} \bigwedge\!\!\!\bigwedge K[D]$.

Proof: (Cf. the proof of theorem 2 in [Th91b].) Notice that **S5*** induces a relation of derivational equivalence $\vdash\!\dashv$, which corresponds to F–equivalence (i.e. $\varphi \vdash\!\dashv \psi$ iff $M, s \mathrel{\equiv\!\!|} \varphi \Leftrightarrow M, s \mathrel{\equiv\!\!|} \psi$ for all M, s). With respect to this equivalence the logic is finite, as can be established by either a syntactic normal form argument or a semantic argument. Moreover, this logic also has the FMP. So the miniature is the finite disjoint union of finite counter-examples to (equivalents of) non-consequences. Using the generation lemma, it is easy to prove that this is an F–miniature for introspective D. ■

This result demonstrates that for the usual sets of data, such as epistemic formulas of the form $K\alpha$, model theoretic representation is feasible. It does not display the form of the miniature, nor how to achieve a *minimal* model. In fact, what does an F–miniature look like? In some cases an F–miniature may be a total model, as the following examples will illustrate.

Example 1 (complete information)
Let Prop $= \{p, q\}$ *and* $D = \{Kp, K\neg q\}$. *The minimal F–miniature for D is the singleton model*

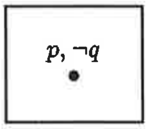

It is easily checked that this is an F–miniature for D, and it cannot be made smaller, since dropping the only point or omitting, for example, 'p' would give $M \not\mathrel{\equiv\!\!|}\ K\neg p$.

Example 2 (simple ignorance) *For the same atoms, let* $D = \{Kp\}$. *The minimal total F–miniature for D is the model:*

[6]Cf. [Th90b] and [Th92]. **S5*** is characterized by the set of rules $\mathbf{MrL^*} \cup \{\mathbf{T_r}, \mathbf{5_r}\}$, where $\mathbf{T_r}$ and $\mathbf{5_r}$ stand for $K\varphi \vdash \varphi$ and $\neg K\varphi \vdash K\neg K\varphi$ and their respective contrapositives.

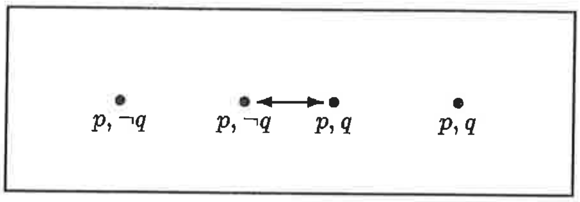

So, the idea that invoking non-falsification always produces a reduction of the model (by dropping either positive or negative information) turns out to be wrong. We can, of course, *add* partialized components, or, more precisely, despecifications modulo finite equivalence.[7] In the present case this amounts to possibly copy worlds within a component, reconnect them to the component and finally omit specifications. Partialization will lead to larger models, but, more importantly, the resulting miniatures will be equivalent.

Proposition 2 *Adding partialized components to an F–miniature results in another F–miniature (for the same information).*

Proof: given some F–miniature M for D with component N, let $N' \dot{\sqsubseteq} N$, then $M \oplus N'$ is equivalent to M, for there are finite L, L' such that $N' \equiv L' \sqsubseteq L \equiv N$, and so (1) $M \oplus N'$ will be finite; (2) suppose for some $\delta \in D$ and s in M: $N', s \Dashv \delta$, then by equivalence and persistence $N, s \Dashv \delta$, thus $M, s \Dashv \delta$, which contradicts $M \not\Dashv D$. So we obtain $N \not\Dashv D$, and therefore $M \oplus N' \not\Dashv D$; (3) if $M \oplus N' \not\Dashv \varphi$ then surely $M \not\Dashv \varphi$, and so $D \not\Dashv \varphi$. ■

The proposition licenses an optimalization procedure: an F–miniature can be minimalized by dropping components which are partializations of other components. So, a minimal F–miniature will usually consist of more or less total components.

Now for those data that can be modelled by a total miniature, we may also consider whether we can supply a truly partial miniature by *replacing* a total component by its different partializations.

Example 3 *The information $D = \{Kp, K\neg q\}$ was covered by the total model of example 1. Now consider splitting the single world into the pair constituting the model M:*

[7] So, $N_1 \dot{\sqsubseteq} N_2$ iff for some finite N_3, N_4: $N_1 \equiv N_3$, $N_2 \equiv N_4$ and $N_3 \sqsubseteq N_4$. Here $\sqsubseteq$ expresses extension of valuation for the same frame. Persistence w.r.t. valuation extension was shown in [Th90a].

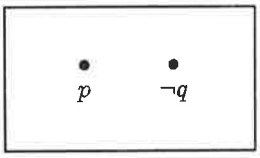

Unfortunately, M does **not** *characterize D: on the one hand $M \not\dashv K\neg p \vee Kq$, but on the other $D \not\vdash_{\mathbf{S5^*}} K\neg p \vee Kq$.*

A similar situation can be found for incomplete information which is covered by a total model:

Example 4 (simple ignorance, continued)
Reconsider $D = \{Kp\}$ (cf. example 2) A minimal F–miniature for D could be obtained from the total miniature by omitting some of the literals. Notice however that contracting the middle component to the singleton p, will not do: then $Kq \vee K\neg q$ will be non-falsified, but $Kq \vee K\neg q$ is not a consequence of Kp. Moreover, dropping both occurrences of 'p' in the central component would give the model

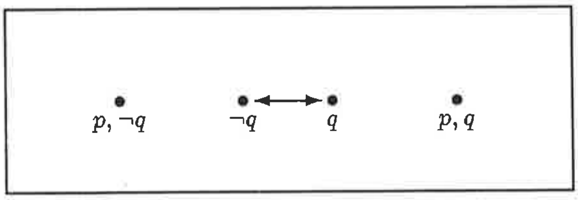

This model non-falsifies $K\neg p \vee Kq \vee K\neg q$, which is not an F–consequence of Kp. Similarly, dropping 'p' in the right-hand world of the middle component admits the non-consequence $\neg Kp \vee Kq \vee K\neg q$. In fact, as the reader may check (warning: this is tedious labour), none of the occurences of 'p', 'q' and '¬q' may be omitted without loss of characterization.

These examples suggest that minimal total F–miniatures are unique, up to isomorphism. The examples and the previous proposition may also suggest the generalization that classical (**S5**)- miniatures for some set of data D are always F–miniatures for D as well. Though tempting, the latter is not true.

Example 5 *Consider the data $D = \{Kp, K(p \rightarrow q)\}$. A total model that verifies D will also verify Kq. Consequently, essentially the only classical* **S5**-*miniature for D will be the singleton model verifying both p and q. But Kq is not an F–consequence of D, roughly because* **S5*** *does not contain* Modus Ponens. *So the singleton model is not an F–miniature for D. A small F–miniature for D is:*

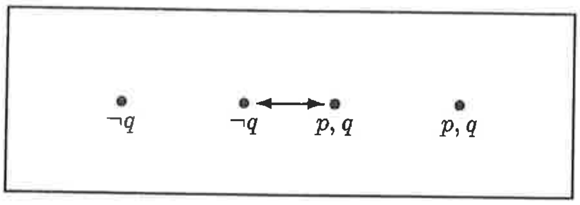

Another example of the incongruity of classical and V–miniatures may be more transparant. The point is that for inconsistent data, F–consequence and **S5**-consequence diverge widely.

Example 6 *For the data* $\{Kp, Kq, K\neg q\}$ *the minimal F–miniature is:*

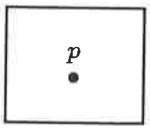

As in the previous example, there is no total F–miniature for this set of data: a total model that non–falsifies q *and* $\neg q$ *in each world should verify* q *and* $\neg q$ *in each world, thus has to be the empty model. But the empty model also non-falsifies* $K\neg p$*, which does not follow from the data in* **S5***.

¿From these comparisons between (partial and total) F–miniatures and (total) classical miniatures some generalizations are induced:

- A total F–miniature for D is also a classical miniature for D.
- If D has minimal F–miniature that is partial, it has no total F–miniature.

The first generalization is easily proved. Looking for a kind of converse notice that examples 5 and 6 show that classical miniatures may not be total F–miniatures, in fact that there is information that has a classical miniature, but no total F–miniature. So the following question emerges:

> *Is consistent information modelled by an F–miniature iff there is a classical miniature?*

¿From the left to the right this holds trivially (even without consistence): If D has an F–miniature then $D \vdash_{\mathbf{S5^*}} \bigwedge K[D]$ and so $D \vdash_{\mathbf{S5}} \bigwedge K[D]$, thus D has a classical miniature. For the other direction, notice we cannot leave out the consistency requirement; a trivial counter-example runs as follows: $p \wedge \neg p$ has the empty classical miniature, but cannot have an F–miniature, since $p \wedge \neg p \not\vdash_{\mathbf{S5^*}} K(p \wedge \neg p)$. This observation can be transformed into a genuine counter-example: Consider $\alpha = (p \wedge \neg p) \vee Kq$. α is consistent and classically equivalent to Kq, and so **S5**–introspective. However α is not **S5***–introspective, for $K\alpha$ is falsified in the left-hand world of the following model, but α is not.

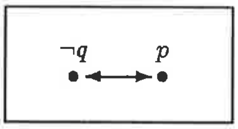

3 V–miniatures

The notion of 'V–miniature' is similar to 'F–miniature', with verification ('V') instead of non-falsification. The definitions of global verification ($M \models\!\equiv \varphi$), and relative verification (also called 'strong consequence') are obvious.[8]

Definition 2 *M is a* **V–miniature** *for D iff M is finite and* $M \models\!\equiv \varphi \Leftrightarrow D \models\!\equiv \varphi$ *holds for each* φ.

Or, equivalently, M is a V–miniature iff

- M is finite,
- $M \models\!\equiv D$,
- $M \models\!\equiv \varphi \Rightarrow D \models\!\equiv \varphi$.

As noticed in [Th90b], the set of V–consequences is usually considerably smaller than the set of normal consequences; so the V–miniature may have to represent less, which is enabled by possibility of underspecifying worlds for their propositional contents.

In this section we still restrict ourselves to coherent models with an equivalence accessibility relation. For relative verification the inference rules will again be somewhat weaker than the usual **S5** ones. Call this logic, which is also a variant of good-old **S5**, now with the *ex falso* rules but without *tertium non datur*: $\mathbf{S5^+}$, 'coherent verificational **S5**'.[9]

[8] See [Th90b] for details.

[9] Cf. [Th90b] and the chapter on completeness in partial modal logics in [Th92]. $\mathbf{S5^+}$ is characterized by $\mathrm{MrL}^+ \cup \{\mathbf{Tr}, \mathbf{5r}\}$.

We may repeat the question of correspondence: which syntactic or deductive qualities of D enable its verificational representation by a finite partial model? Existence of V–miniatures is warranted by the already familiar syntactic condition of (deductive) introspection:

Theorem 2 *D has a V-miniature iff* $D \vdash_{\mathbf{S5^+}} \bigwedge\!\!\!\bigwedge K[D]$.

Proof: (Cf. the proof of the previous theorem.) The condition is clearly necessary, but also sufficient. To show the latter, notice that $\mathbf{S5^+}$ is *logically finite* and has the FMP. If $\Phi = Form/\equiv$, then the disjoint union of counter-examples of non-consequences produces

$$M = \bigoplus\{N \mid N \models D, N \text{ tight and reduced } \& \ N \not\models \varphi \text{ for some } \varphi \in \Phi\}.$$

It is easily checked that M is indeed a V–miniature for D. ∎

Notice that this theorem guarantees existence of the V–miniature for introspective information, but does not produce a *minimal* V–miniature. As a matter of fact, the model produced in the proof will be usually (much) larger than the classical miniature. But in many cases smaller models can be obtained, as the examples below will demonstrate.

The gain of relative verification becomes clear in cases where only part of the propositional variables are known. Recall that for such *simple ignorance* the F–miniature amounts to a classical model of superexponential size (in the number of unknown variables).

Example 7 (simple ignorance)
Assume complete information about $p_1, \ldots, p_k$ (i.e. Kp_i or $K\neg p_i$ for each $i = 1, \ldots, k$), and complete ignorance of the rest. This set of data is modelled by the singleton miniature M: [10]

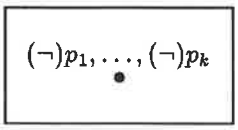

The last example was a case of 'simple ignorance': nothing was known about p_i for $i > k$, even $D \not\vdash \neg Kp_{k+1} \vee \neg K\neg p_{k+1}$. This contrasts with types of 'strong ignorance' in which, for example, there is full knowledge about $p_1, \ldots, p_k$ but $D \vdash \neg Kp_i \wedge \neg K\neg p_i$ for all $i > k$.

[10]The correctness of this and the following miniatures is proved in the appendix.

Example 8 (strong ignorance)
Let $D = \{Kp_1, \neg Kp_2, \neg K\neg p_2, \neg Kp_3, \neg K\neg p_3\}$. D is represented by the following minimal model:

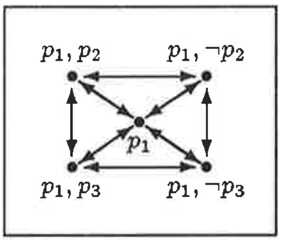

Notice on the one hand that in cases of 'strong ignorance' the information can be strengthened further: e.g. in the above example, $D \not\vdash \neg K(p_2 \vee p_3)$. Yet adding such a formula implies an increase of knowledge, so it seems fair to say that then the ignorance would have decreased.

On the other hand, in between simple and strong ignorance there are intermediate cases of semi-strong (or, partial) ignorance. Here is a paradigm.

Example 9 (semi-strong ignorance)
Assume that $D = \{Kp_1, \neg Kp_2\}$. D is minimally represented by:

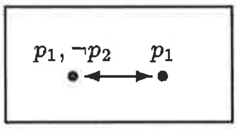

Now adding (negative) information such as $\neg Kp_i$ has a remarkable effect: it *reduces* the size of the classical miniature, but *magnifies* the size of the partial V–miniature somewhat. To wit, assume the initial information Kp_1. As we saw in the introductory section, the classical miniature for the atoms p_1, p_2, p_3 has 32 worlds (distributed over 15 components). Adding $\neg Kp_2$ leads to some reduction: for this case of semi-strong ignorance the classical miniature contains 28 worlds (12 components). The final addition of $\neg K\neg p_2$, $\neg Kp_3$ and $\neg K\neg p_3$ (strong ignorance) involves a classical miniature of 20 worlds (and 7 components). As the above examples show, the number of worlds of partial V–miniatures for these cases of simple, semi-strong and strong ignorance are 1, 2, and 4, respectively (and just one component in each case).

More generally, (semi-)strong ignorance can be captured by tight V–miniatures of polynomial, in fact even *linear* size, whereas the classical pendants need a *superexponential* amount of worlds.

Proposition 3 (size of miniatures for (semi-)strong ignorance) *The minimal V–miniature modelling (semi-)strong ignorance with respect to r atoms and complete knowledge of the others has at most $2r+1$ situations (in 1 component). A corresponding classical miniature requires at least 2^{2^r-1} worlds, divided over at least 2^{2^r-2} components (if $r>0$).*

Proof: the examples above of (semi-)strong ignorance obviate that the largest minimal V–miniature will the one for strong ignorance of r atoms, where $2r+1$ situations suffice: apart from the known literals, which occur in all situations, one containing p_i and one containing $\neg p_i$ for each unknown p_i and one containing none of these.

For the classical case, notice that strong ignorance now gives the *smallest* miniature. Its largest component will have 2^r worlds (state–descriptions). Each set containing more than half of these state–descriptions will necessarily satisfy both p_i and $\neg p_i$ for all unknown p_i (for if not than there can be no more than 2^{r-1} states in the set), and thus will constitute a component of the miniature. So we have that

$$\sharp C_r \geq \sum_{2^{r-1}+1}^{2^r} \binom{2^r}{i} = \frac{1}{2}\cdot 2^{2^r} - \frac{1}{2}\binom{2^r}{2^{r-1}} \geq 2^{2^r-2}$$

if $r>0$. A similar calculation for the number of worlds shows [11]

$$\sharp W_r \geq \sum_{2^{r-1}+1}^{2^r} i\cdot\binom{2^r}{i} = 2^r \sum_{2^{r-1}}^{2^r-1} \binom{2^r-1}{i} = 2^r\cdot\frac{1}{2}\cdot 2^{2^r-1} = 2^{2^r+r-2} \geq 2^{2^r-1}$$

if $r>0$. (the last estimation also holds for the borderline case $r=0$). ∎

This proposition illustrates our point that partial models are superior to classical models for at least two reasons: first, they are usually much smaller and, second, more natural since they tend to grow when information is added. Classical miniatures, on the other hand, are rather clumsy in describing knowledge. Total Kripke structures may be said to model ignorance rather than knowledge, which they are supposed to. One of the additional advantages of partial models is that, since simple ignorance does not need to be modelled, only 'relevant' propositional variables (which occur in the data) have to be taken into account. Consequently, we omit the specification of *Prop* in examples of V–miniatures.

[11] The summation may be replaced by a perhaps more insightful argument: more than half of the subsets of state–descriptions containing some particular state will contain at least half of the state–descriptions. This also yields the number $2^r\cdot\frac{1}{2}\cdot 2^{2^r-1}$.

The above examples provided V–miniatures consisting of a single component. More complex types of incomplete knowledge may require several components, however. Before discussing some more involved examples, let us find the proper analogue for proposition 2. In order to reduce its size, we need to know how V–miniatures are related to partialization. Again the relation $\dot{\sqsubseteq}$ between components of a V–miniature involves a minimalization procedure, but now in the reverse direction. So, we can now make components more complete by extending specification of worlds, after which equally specified worlds may be identified. Completization will also lead to larger models, but the resulting miniatures will be equivalent.

Proposition 4 *Adding partially completed components to a V–miniature results in another V–miniature (for the same information).*

Proof: similar to proposition 2. ■

By proposition 4 a V–miniature can be minimalized by dropping components which are essentially partial completions of other components. Therefore a minimal V–miniature will consist of 'most unspecified' components. This optimization is especially useful for more complex cases.

Example 10 (honest disjunctive knowledge)
The smallest V–miniature for $K(p \vee q)$ is:

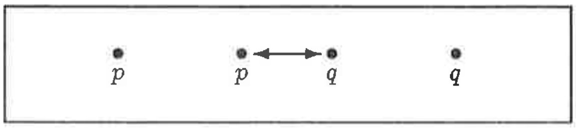

In comparison with the classical miniature[12] *the above model is still small: the classical* **S5**–*miniature is a graph consisting of 12 vertices and 6 edges, divided into 7 components.*

Although the last miniature consists of 3 components, the model as a whole represents a proper piece of information. Typically, one can succesfully declare to know *only* this or that; then the miniature has to be restricted to its central component. [HM85] discusses cases of so-called *dishonest* knowledge in which such a consistent circumscription is impossible: e.g., one cannot consistently claim to know only whether this or that. The informal explanation of Halpern & Moses is that a dishonest formula does not correspond

[12] Cf. [Th91b, example 3]

to a state of knowledge. In stead of (non–monotonically) circumscribing knowledge, we have chosen to (monotonically) describe the knowledge.[13]

Example 11 (dishonest disjunctive knowledge)
The minimal V–miniature for $Kp \vee Kq$ happens to be the model:

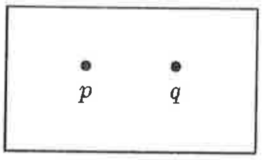

Comparison with the classical miniature[14] again points at a still considerable reduction: the classical **S5***–miniature has 7 worlds (and 2 edges, 5 components).*

All the previous examples of incomplete information revealed V–miniatures which were (much) smaller than the classical ones. This suggest the generalization that representing incomplete knowledge by V–miniatures will always be more efficient. By inspection of a stronger kind of incomplete information it shows that this generalization is not true.

Example 12 *The minimal V–miniature for $K(p \leftrightarrow q)$ is:*

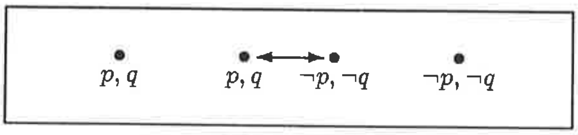

At first sight, things may even get worse in that V–miniatures may be larger than classical ones.

Example 13 *The minimal V–miniature for $\{\neg Kp, \neg K\neg p\}$ is:*

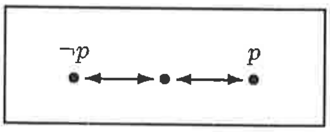

The classical miniature with respect to Prop = $\{p\}$ has one world less:

[13] See [Th91b, sections 6,7,9] for a more profound exposition and comparison of the two approaches.
[14] Cf. [Th91b, example 4]

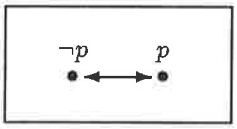

Before jumping to conclusions, note that proposition 3 indicates that larger V–miniatures are quite exceptional: for (semi-) strong ignorance the last example is the only case in which the linear growth function exceeds the superexponential one. Of course, other types of information may require larger miniatures. However, this does not give an increase of the theoretical complexity: after reduction, but even before minimalization triggered by persistence, the number of components is at most $2^{3^n} - 1$ and the total number of worlds $3^n \cdot 2^{3^n - 1}$ (cf. proposition 1). We have seen that minimalization usually cuts these numbers down considerably. Moreover, we believe that the actual absolute excess will be limited.

Given the extensive comparison between partial and classical miniatures in this section, we can evidently pose the same question as at the end of section 2: is existence of a V–miniature for (consistent) information equivalent to existence of a classical miniature. Actually, since $\mathbf{S5}^+$ contains the *ex falso* rules, the consistency requirement is immaterial now. Again the implication holds in one direction: if D has a V–miniature, D is $\mathbf{S5}^+$–introspective, and $\mathbf{S5}^+$ is contained in $\mathbf{S5}$, so D is $\mathbf{S5}$–introspective as well, and therefore has a classical miniature. In the other direction the implication is clearly false: $p \vee \neg p$ has a classical miniature (the same as for $\emptyset$), but no V–miniature, since $p \vee \neg p \not\vdash_{\mathbf{S5}^+} K(p \vee \neg p)$.

4 Alternative: local miniatures

A perhaps more obvious semantic approach to characterize knowledge would involve Kripke's original *local* models. The idea is to have a designated world, from which the evaluation starts, containing the facts, and the accessible worlds containing the knowledge of these facts, and the knowledge of this knowledge, etcetera.

This approach was rejected in [Th91b] because only complete information can be described in this way. More precisely, only D which were *complete* and *consistent* theories qualified. We blamed bivalence in the root world for this recalcitrant behaviour. Now, giving up bivalence, the hope of finding new possibilities for local miniatures is reviving. So, let us consider local F– and V– miniatures. To that purpose, replace M in definitions 1 and 2 by $\langle M, s \rangle$. We will discuss both kind of local miniatures separately.

Local F–miniatures

Like their global partners, the knowledge modelled by local F–miniatures does not have to be consistent. This oddity is illustrated by example 6, which also shows that if the global miniature is a singleton model, it coincides with the local miniature. Moreover, local F–miniatures do not capture honest disjunctive knowledge. For assume that, for example, $K(p \vee q)$ has a local miniature $\langle M, w \rangle$. Then $M, w \not\Rightarrow\!\!\!\!\dashv K(p \vee q)$, so $M, w \not\Rightarrow\!\!\!\!\dashv p$ or $M, w \not\Rightarrow\!\!\!\!\dashv q$, but neither p nor q are consequences of $K(p \vee q)$. The conditions for local F–miniatures appear to be very strong.

Theorem 3
D has a local F–miniature iff D is both complete and saturated.[15]

Proof: analogous to theorem 1 in [Th91b], using part of the canonical model by filtration over D + its subformulas. The suitability of saturation follows from the Henkin completeness proof in [Th90a]. ∎

Local V–miniatures

Since (global) V–miniatures are more efficient then F–miniatures, we may hope for more success in the case of local V–miniatures than we experienced for local F–miniatiures. However notice that disjunctive knowledge is still troublesome: $K(p \vee q)$ has no local V–miniature. The conditions for local V–miniatures are still very strong.

Theorem 4
D has a local V–miniature iff D is both consistent and saturated.

Proof: analogous to the previous theorem. ∎

5 Below and beyond S5: partial miniatures

Within classical logic the paradigm cases of cautious extension or variation of the **S5** system are **S4** and $\mathbf{S5_{(m)}}$. The possibilities for miniaturization of

[15] Completeness here expresses $D \vdash \varphi$ or $D \vdash \neg\varphi$, saturation $D \vdash \varphi \vee \psi \Rightarrow D \vdash \varphi$ or $D \vdash \psi$.

information turned to be extremely limited for these background logics. For the epistemic logic **S4** (advocated by Hintikka) complete information can be represented in a singleton model. In [Th91b] we conjectured that incomplete information does not have **S4**–miniatures; this claim has been verified for simple and semi-strong ignorance. The situation for $\mathbf{S5_{(m)}}$ is even worse: consistent information (whether 'complete' or not) does not have $\mathbf{S5_{(m)}}$–miniatures. Is the situation for partial logic similarly distressing? To study this in some detail, we will focus on V–miniatures in the rest of this section.

$\mathbf{S4^+}$–miniatures

Partiality slightly improves the chances for $\mathbf{S5^+}$-miniatures: both complete knowledge and simple ignorance can be modelled by singleton V–miniatures. A simple induction proof shows that $\mathbf{S4^+}$ and $\mathbf{S5^+}$ have the same inferences from this kind of information. Consequently, the miniature of example 7 still qualifies with respect to $\mathbf{S4^+}$.

For (semi–)strong information we find the opposite situation. E.g., contrasting to example 9, $\{Kp, \neg Kq\}$ has no $\mathbf{S4^+}$–miniature, for suppose M would qualify. Then $M \models \neg Kq$, so $M \models K\neg Kq$. However, $Kp, \neg Kq \not\models K\neg Kq$.

More generally, *negative* information appears to obstruct possible miniatures. More challenging are the cases of *positive* partial knowledge, such as $K(p \vee q)$. We believe that a characterizing model requires chains of unlimited length, but this has not been proven with formal rigour yet. On the other hand, the 'dishonest' formula $Kp \vee Kq$ appears to have the same miniature as in example 11.

$\mathbf{S5_{(m)}}^+$–miniatures

Here the situation is worse. One of the points is that for a multi-agent logic there can be no complete information *without* contingent common knowledge, which can not be expressed in the simple modal language. And without an operator for common knowledge, we cannot obtain miniatures: the models will always be too strong. In fact the proof of the nonexistence of $\mathbf{S5_{(m)}}$–miniatures for consistent data can simply be transposed to the realm of partial semantics. On the other hand, simply adding an operator C for common knowledge may not solve all problems. Though it is clear that C-introspection is a necessary condition for the existence of $\mathbf{S5_{(m)}}^+$–miniatures, it is not obvious that the condition is sufficient.

A Correctness of several V–miniatures

correctness of example 7 (simple ignorance)

Proof: Clearly $M \models D$, and assume that $M \models \varphi$. Let$N, s \models D$ for some model M containing situation s, then for the generated submodel N_s also $N_s \models D$. Thus $M \dot{\sqsubseteq} N_s$, for M can be obtained from N_s by omitting all p_i specifications for $i > k$ and then identify equally specified situations. Therefore $N_s \models \varphi$, so $N, s \models \varphi$, and consequently $D \models \varphi$. ∎

correctness of example 8 (strong ignorance)

Proof: Correctness and minimality of this miniature can be shown by constructing the semi-lattice of verifying tight models, partially ordered by $\dot{\sqsubseteq}$, which has the displayed miniature as its bottom element. This, however, is a laborious exercise. An easy argument may do just as well: Notice that every tight model verifying D will consist of situations which at least contain p_1, whereas some situations should contain p_2, $\neg p_2$, p_3, and $\neg p_3$. Now the given model can be strengthened to any verifying model, and thus minimally characterizes D (cf. proposition 4 for formal justification). ∎

correctness of example 9 (semi-strong ignorance)

Proof: Here a construction of the semi-lattice of verifying components is feasible. In fact there are two ways to generate this structure. One may inspect the 8 minimal models which have p_1 in each world and possibly also p_2 or $\neg p_2$, and check whether $\neg K p_2$ holds in them. Then the obtained models are ordered for $\dot{\sqsubseteq}$, and the displayed miniature then appears to be the bottom of 3 element semi-lattice. But one may also start with the *classical* miniature for D: its components will verify D in the partial sense too.

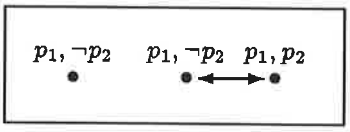

Then by weakening (i.e. partializing and possibly duplifying and identifying situations) the minimal partial miniature is obtained in the end. The latter method often turns out to be more efficient. ∎

correctness of example 10 (honest disjunctive knowledge)

Proof: it suffices to draw the graph of all (reduced) tight V–models for $K(p \vee q)$, ordered with respect to $\dot{\sqsubseteq}$ (indicated by lines), where the uphill components are more specific and thus may be omitted.

References

[FH88] R. Fagin & J. Halpern, 'Belief, awareness and limited reasoning', *Artificial Intelligence 34*, pp. 39–76, 1988

[FHV90] R. Fagin, J. Halpern & M. Vardi, 'What is an inference rule?', *Proceedings Jerusalem Conference on Information Technology* pp. 391–401, 1990. Expanded version in IBM Research Report RJ 7865, 1990

[FHV91] R. Fagin, J. Halpern & M. Vardi, 'A Model-Theoretic Analysis of Knowledge', *Journal of the* ACM 38:2, pp. 382–428, 1991

[HM85] J. Halpern & Y. Moses, 'Towards a theory of knowledge and ignorance', in Kr. Apt (ed.), *Logics and Models of Concurrent Systems*, Springer, Berlin, 1985

[HV91] J. Halpern & M. Vardi, 'Model checking vs. theorem proving: a manifesto', in J. Allen (e.a. eds.) *Principles of Knowledge Representation and Reasoning: Proceedings of the 2^{nd} International Conference*, 1991

[La87] G. Lakemeyer, 'Tractable meta-reasoning in propositional logics of belief', *Proceedings of* IJCAI87 *Milan*, Morgan Kaufmann, Los Altos CA, pp. 402–408, 1987

[Th90a] E. Thijsse, *Partial Logic and Modal Logic: a Systematic Survey*, ITK Research Report 11, Tilburg University, 1990

[Th90b] E. Thijsse, 'Partial propositional and modal logic: the overall theory', M. Stokhof & L. Torenvliet (eds.) *Proceedings of the 7^{th} Amsterdam Colloquium*, part 2, pp. 555–579, ITLI, Amsterdam, 1990

[Th91a] Thijsse, E., 'On total awareness logics (with special attention to monotonicity constraints and flexibility)', in: M. de Rijke (ed.) *Colloquium on Modal Logic*, Dutch TLI Network, Amsterdam, 1991, pp. 253 — 278. to appear with Kluwer Academic Publishers in the series *Studies in Logic, Language and Information*

[Th91b] E. Thijsse, 'Finite Kripke Models for Knowledge Representation', in: R. van de Riet & R. Meersman (eds.), *Proceedings of the 1991 Workshop on Linguistic Instruments in Knowledge Engineering (*LIKE91*)*, North-Holland, Amsterdam, to appear

[Th92] E. Thijsse, *Partial logic and knowledge representation*, dissertation, University of Tilburg, 1992 (forthcoming)

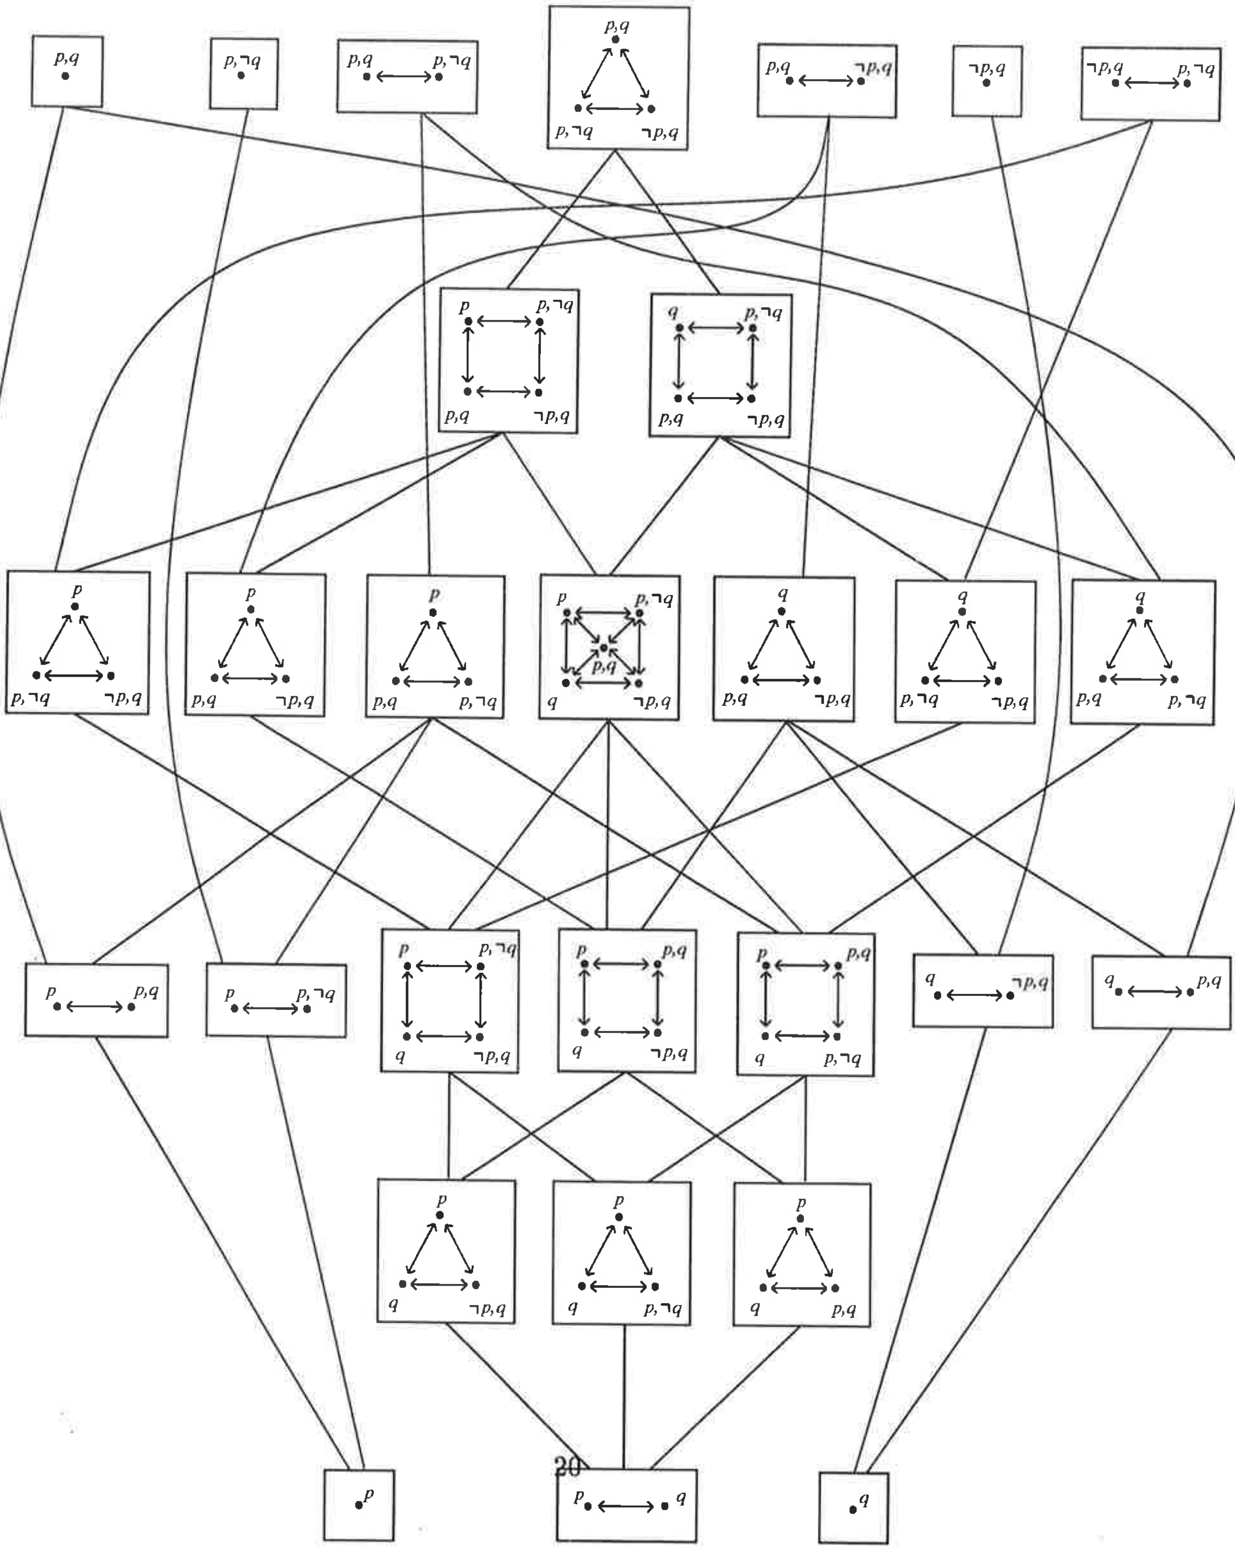

20