THE QUALITATIVE VERIFICATION OF QUANTITATIVE UNCERTAINTY

Simon PARSONS*, Alessandro SAFFIOTTI**

* Advanced Computation Laboratory
Imperial Cancer Research Fund, P.O. Box 123,
Lincoln's Inn Fields, London WC2A 3PX, UK

** IRIDIA, Université Libre de Bruxelles
50 av. Roosevelt, CP 194-6
B-1050 Bruxelles, Belgium

* Department of Electronic Engineering
Queen Mary and Westfield College
Mile End Road, London, E1 4NS, UK

ABSTRACT

Qualitative methods for reasoning under uncertainty may be helpful in situations where quantification of uncertainty is not appropriate. We demonstrate another use for qualitative models. The qualitative analysis of a quantitative model of uncertainty will reveal the qualitative behaviour of that model when new evidence is obtained. This qualitative behaviour can be compared with the specifications to identify those situations in which the model does not behave as expected. We report the result of two experiments performed on a small fragment of a real test-bed by using a probability model, and a model based on the Dempster-Shafer theory of evidence.

1. Introduction

Recently, there has been considerable interest in the qualitative representation of reasoning under uncertainty in networks, including qualitative probabilistic networks [2,13] as well as qualitative possibilistic and evidential networks [3,4]. Such work has been aimed at determining the impact of new evidence in situations in which full numerical results may not be obtained due to incomplete or imprecise knowledge. In this paper we suggest a different use for qualitative methods. Since the qualitative behaviour of a system may be established from quantitative knowledge, we can qualitatively analyse any quantitative model. This analysis may then be used as a simple means of verifying that a system behaves as intended by the knowledge engineer that built it. It also provides a means for guiding the correction of any faults that may be found.

The basic method of our analysis is as follows. When we find new evidence about the state of a variable we update our prior values to take account of the evidence. When using the model, we are interested in the new value obtained after updating. When verifying the behaviour of the model, however, we are interested in checking that this updating corresponds to that described by the domain expert whose knowledge is captured in the model. Since the expert's knowledge is often expressed in the form "If we observe *e* then it is more likely that *h* is the case", we may be more interested in knowing the way in which the values change than in the values themselves. Given the equations that relate two uncertainty values *val1* and *val2*, expressed in some formalism, we can establish an expression, in terms of numerical uncertainty values, for the derivative $\frac{\partial \mathrm{val1}}{\partial \mathrm{val2}}$ that relates the two quantities. This expression allows us to determine the qualitative value of the derivative, written as $\left[\frac{\partial \mathrm{val1}}{\partial \mathrm{val2}}\right]$, that is whether the derivative is positive, written as [+], negative, [–], or zero, [0]. The sign of the derivative indicates the direction of change of *val1* when *val2* increases. As $\Delta x = \Delta y \frac{\partial x}{\partial y} + \Delta z \frac{\partial x}{\partial z}$, we can get the effect of several successive pieces of evidence by combining the effect of each alone.

To validate a given model, we compute $\left[\frac{\partial \mathrm{val(h)}}{\partial \mathrm{val(e)}}\right]$ for every interesting hypothesis *h* and piece of evidence *e*, and then compare these values with the knowledge expressed by the expert. In the rest of this paper we demonstrate our qualitative analysis technique for a probabilistic and a Dempster-Shafer model on a small example extracted from a real application. A longer report [5] includes all the computations, extends the analysis to possibilistic models, and describes a debugging procedure.

2. Problem description

The problem under study is a simplified version of fault diagnosis in electricity networks. This problem was originally used by Saffiotti and Umkehrer [8] to investigate the use of different formalisms to model uncertainty. We consider the fragment of an electricity network shown in Figure 1. This fragment comprises four substations, linked by three lines L1, L2 and L3. The substation in the middle includes S1, a big conductive bar, known as a busbar, used for connect-

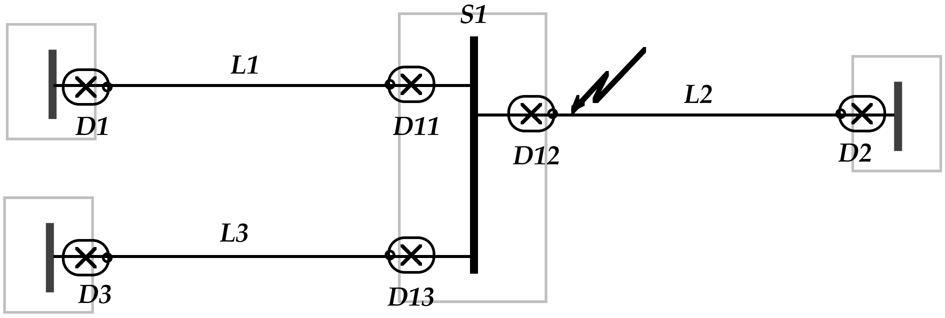

Figure 1. A fragment of an electricity distribution network.

ing more lines together. The D_is are circuit breakers, that is devices which watch the part of the network on their "hot" side, marked by a dot in the picture, for overloads, and isolate two lines if an overload is detected. If an overload occurs, a circuit breaker generates an alarm and transmits it to the control room. The alarm may be either an instantaneous alarm or a delayed alarm. Talking to domain experts revealed what the qualitative behaviour of the modelled fragment should be (no matter what formalism is used to model the uncertainty):

1) an instantaneous alarm from an outer circuit breaker should increase belief in the occurrence of a fault in the line that the breaker is on;

2) a delayed alarm from an outer breaker should increase belief in the occurrence of a fault in either the line the breaker is on, or in the busbar;

3) an alarm (of any kind) from an inner breaker should only increase our belief in the occurrence of a fault in the line the breaker is on.

The experts were also able to provide rough quantitative estimates of the uncertainty: e.g., in roughly 10% of the cases alarms are generated without faults or faults occur without alarms.

We have modelled our problem in both probability and Dempster-Shafer theory by using Shenoy and Shafer's valuation system formalism [10,11]. The tool we have used for our experiment, Pulcinella [7,8] is an implementation of valuation systems in which many uncertainty handling formalisms can be embedded. As according to the valuation system formalism, we model our problem through a set of variables, and a set of valuations linking sets of related variables. Below is a graphical representation of the model, where circles stand for variables, and rectangles for valuations.

A valuation over a set of variables expresses information about the values taken by the variables in that set, in a form that depends on the uncertainty formalism; relations among variables are expressed by valuations. Here, the *Di*'s are variables representing circuit breaker states, with possible values ok (no alarm), del (delayed alarm), and inst (instantaneous alarm); *Li*'s and *S1* represent line states, with frame {ok, fault}; and the *alarm-i*'s relate generation of alarms by breakers with states of neighbour lines. New information can be propagated through the *alarm-i*' relations to produce updated estimates of the states of the elements of the network.

In order to build the "*alarm-i*" valuations so that they behave as described above, we first split them into two groups: those referring to outer circuit breakers (*alarm-1*, *alarm-2*, *alarm-3*), and those referring to inner circuit breakers (*alarm-11*, *alarm-12*, *alarm-13*). The following tables

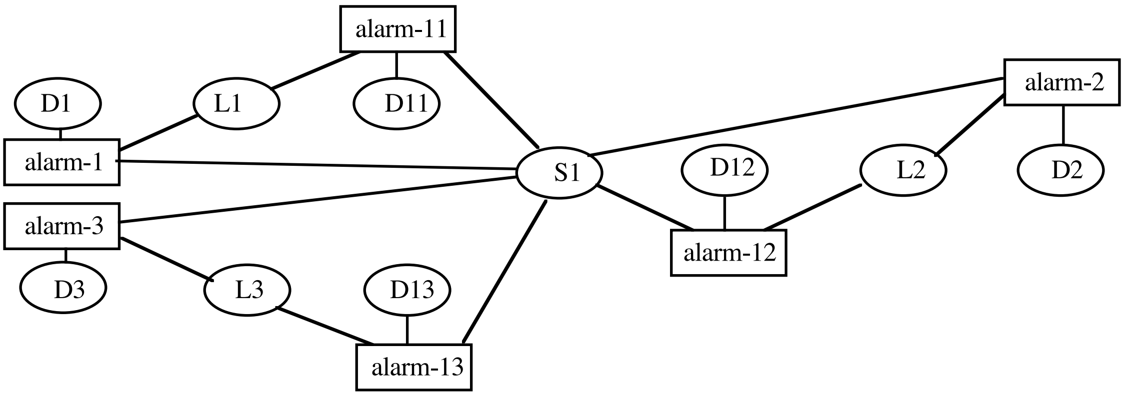

Figure 2. The valuation system for the distribution network

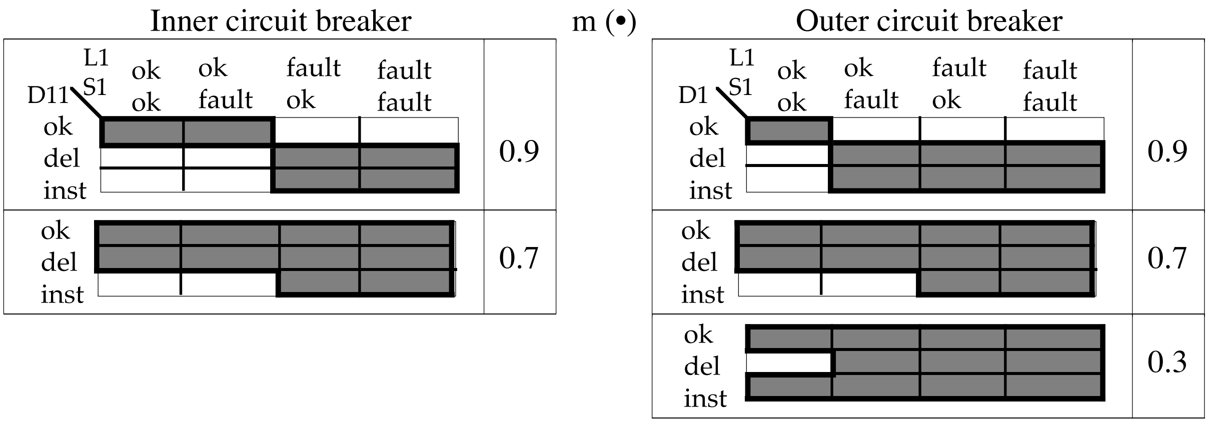

Figure 3. Mass assignments for the "alarm" relations

show the values for the two classes of valuations when probability values are used—valuations in this case are joint probability distributions.

Inner breaker			*P* (•)	*Outer breaker*		
D11				*D1*		
ok	del	inst	*L1*, *S1*	ok	del	inst
0.45	0.05	0.05	ok, ok	0.9	0.1	0.05
0.45	0.05	0.05	ok, fault	0.05	0.6	0.05
0.05	0.89	0.89	fault, ok	0.05	0.2	0.89
0.05	0.01	0.01	fault, fault	0.001	0.1	0.01

To build a Dempster-Shafer model for our problem [9,12], we represent the bits of our information by the basic mass assignments shown in Fig. 3, that are then combined by Pulcinella into one mass assignment using Dempster's rule[1].

A more detailed analysis of this example is given in [8]. In the rest of this paper we will investigate whether these models of uncertainty correctly encode the behaviour described above.

3. Qualitative analysis of the problem

It is helpful to reformulate the problem using a causal network representation. We use a network representation, similar to that of Pearl [6], where two nodes are joined by a directed arc if and only if the variable represented by the node at the end of the arc is directly dependent upon the variable represented by the node at the beginning of the arc. Thus the problem information of Section 2 may be represented by the network of Fig. 4.

Due to the rules of differential calculus, we need only consider changes in sub-networks of the following form, with y = 1, 2, 3.

[1] The combined assignments are fairly intricate, and are reported in the full paper [5]. That report also deals with the case where possibility theory is used.

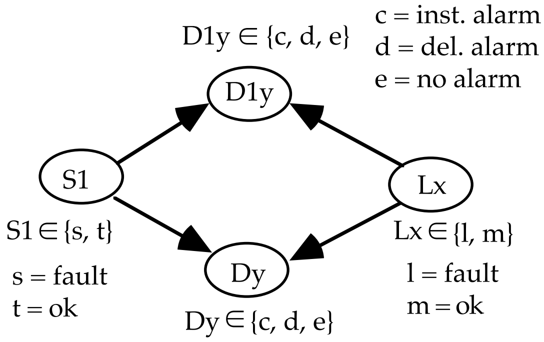

The change at S1 is the sum of all the changes due to all the Dy, and that at Lx the sum of the changes due to the two relevant Dy. In the qualitative analysis we look at changes in value of S1 and Lx given changes in value of a single Dy. If we wanted to assess the impact of several alarms we could sum the impact of the individual alarms using qualitative arithmetic [1].

We first analyze our problem in the probabilistic case. Every circuit breaker has three modes of operation, namely "send an instantaneous alarm", "send a delayed alarm", and "send no alarm". We can write down the probability of failure of a given fault as:

$$p(l) = \sum_{S \in \{s, t\}, D \in \{c, d, e\}} p(l, S|D)\, p(D).$$

for any Dy or D1y. This reduces to

$$p(l) = p(l|c)\, p(c) + p(l|d)\, p(d) + p(l|e)\, p(e)$$

since $p(l|c) = p(s\&l|c) + p(t\&l|c)$ as c, d and e are mutually exclusive and exhaustive. To find out how the probability of a line fault changes when we have an instantaneous alarm we write:

(1) $$\frac{\partial p(l)}{\partial p(c)} = p(l|c) + \frac{\partial p(d)}{\partial p(c)} p(l|d) + \frac{\partial p(e)}{\partial p(c)} p(l|e)$$

which, since $p(c) + p(d) + p(e) = 1$, gives

(2) $$\left[\frac{\partial p(l)}{\partial p(c)}\right] = [p(l|c) \square - \square p(l|d) \square - \square p(l|e)$$

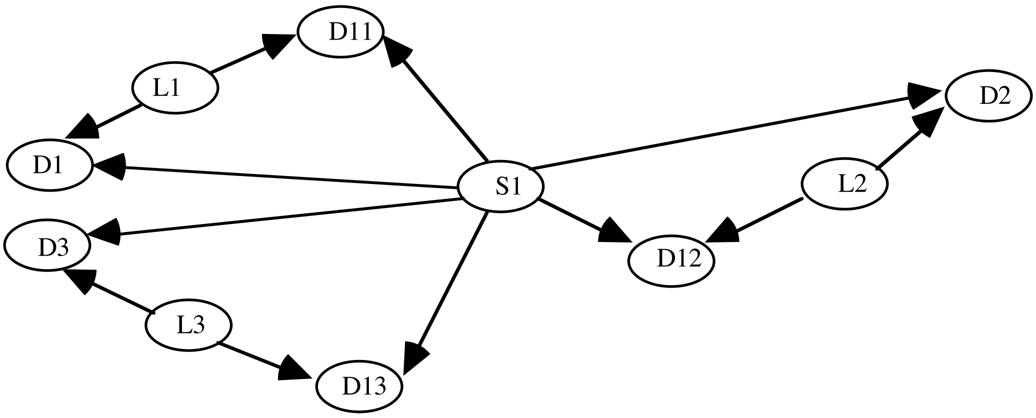

Figure 4. The causal network representation of the electricity distribution problem.

similarly for the delayed and no alarm conditions.

Probability theory strongly relates the probabilities of an event and of its complement:

(3) p(ok) + p(fault) = 1

for any line (and busbar). Hence, if $\left[\frac{\partial p(l)}{\partial p(x)}\right] = [+]$ then $\left[\frac{\partial p(m)}{\partial p(x)}\right] = [-]$ for any $x \in \{c,d,e\}$. Thus when we know how p(l) changes we can tell how p(m) changes. These results are true for both inner and outer circuit breakers, and furthermore are true for any numerical values we put into the model. Similar calculations will give us the probabilities of a busbar failure.

We now turn to considering the belief function case. Once again we start the qualitative analysis with an equation relating L to C, this time expressed in terms of belief functions:

$$bel(l) = \Sigma_{x \subseteq \{c,\, d,\, e\}} bel(l|x)\, m(x)$$

Now, recall that [9]

$$m(x) = \Sigma_{y \subseteq x} -1^{|x-y|} bel(y)$$

which tells us how m(x) varies with bel(c) so we have $\frac{\partial m(c)}{\partial bel(c)} = 1$, $\frac{\partial m(d)}{\partial bel(c)} = 0$, $\frac{\partial m(e)}{\partial bel(c)} = 0$, as well as $\frac{\partial m(c \cup d)}{\partial bel(c)} = -1$, $\frac{\partial m(c \cup e)}{\partial bel(c)} = -1$, $\frac{\partial m(d \cup e)}{\partial bel(c)} = 0$, and $\frac{\partial m(c \cup d \cup e)}{\partial bel(c)} = 1$. This in turn gives us (4):

$$\left[\frac{\partial bel(l)}{\partial bel(c)}\right] = [bel(l|c) + bel(l|c \cup d \cup e) - bel(l|c \cup d) - bel(l|c \cup e)]$$

and again we can establish similar results for the delayed and no alarm conditions. The analysis goes in a similar way for bel(m) and bel(s).

4. Validating the models

Having analysed the way in which qualitative uncertainty values are propagated through the network structure of our test case, we can use the numerical information in Section 2 to examine the qualitative behaviour of the models we have proposed for our test-bed. Our predictions will be compared with the actual behaviour of the models, as determined by running Pulcinella, in order to check the qualitative analysis.

The information that a particular alarm has arrived from a breaker is typically introduced in the model by increasing the value for the associated state, at the expense of the values of the alternative states: for example, a report of an instantaneous alarm is encoded by forcing $\Delta val(c) = [+]$, $\Delta val(d) = [-]$, and $\Delta val(e) = [-]$. These changes are related to the change in the uncertainty value of a particular fault hypothesis, say a line fault, by:

$$(9)\quad \Delta val(l) = \left(\left[\frac{\partial val(l)}{\partial val(c)}\right] \otimes \Delta val(l)\right) \oplus \left(\left[\frac{\partial val(l)}{\partial val(d)}\right] \otimes \Delta val(d)\right) \oplus \left(\left[\frac{\partial val(l)}{\partial val(e)}\right] \otimes \Delta val(e)\right)$$

where ⊗ and ⊕ denote qualitative multiplication and addition, respectively [1,5]. From the quantitative knowledge of Section 2 we can establish the qualitative derivatives using the results of Section 3. These can then be used along with (9) to establish the qualitative behaviour of the various models.

In the case of probability theory, we have, for the outer circuit breakers:

p(l&s|c) = 0.01	p(l&t|c) = 0.89
p(l&s|d) = 0.1	p(l&t|d) = 0.2
p(l&s|e) = 0.001	p(l&t|e) = 0.05

so that:

p(l|c) = 0.90, p(l|d) = 0.30, p(l|e) = 0.051.

Using these values in (2), we find the qualitative values of the derivatives that link the probability of a line fault to that of an alarm for the outer breakers:

$$\left[\frac{\partial p(l)}{\partial p(c)}\right] = [+], \quad \left[\frac{\partial p(l)}{\partial p(d)}\right] = [-], \quad \left[\frac{\partial p(l)}{\partial p(e)}\right] = [-].$$

Using these values in (5), we can predict how the probability of a line fault, p(l), changes qualitatively after different alarm reports:

Report	none	inst	delayed	no alarm
[Δp(l)]	[0]	[+]	[?]	[?]

where the [?] value indicates that it is not possible, based on purely qualitative information, to predict the direction of change. However, we can refine this prediction for a particular model by exploiting some order of magnitude information. For instance, we know that

$$\left[\frac{\partial p(l)}{\partial p(e)}\right] > \left[\frac{\partial p(l)}{\partial p(c)}\right] \approx \left[\frac{\partial p(l)}{\partial p(d)}\right]$$

and, from the probabilities in our model, we have $|\Delta p(e)| \approx |\Delta p(d)| >> |\Delta p(c)|$ when a delayed alarm is reported. Thus, when evaluating (5) to establish the change in the probability of a line fault given a delayed alarm, the third term dominates and [Δp(l)] = [+]. Similar reasoning tells us that [Δp(l)] = [–] for no alarm. We can obtain the same results for the inner breakers.

Thus for both inner and outer breakers, the probability of the line being faulty increases with both instantaneous and delayed alarms, and decreases when we known that there is no alarm, and the model responds to the specifications in Section 2. These predictions are borne out in practice when we evaluate our model in Pulcinella. The following tables show the values of p(l) after a given report has been received by an outer (a) or an inner (b) circuit breaker.

Report	none	inst	delayed	no alarm
(a)	0.0065	0.5	0.18	0.0064
(b)	0.0065	0.67	0.5	0.0063

A similar verification can be performed for the probability of busbar faults given alarms from the outer breakers, which shows that the model behaves according to the specifications in Section 2. The inner breakers, however, constitute something of a surprise. We obtain:

Report	none	inst	delayed	no alarm
[Δp(s)]	[0]	[–]	[–]	[+]

which means that if we have a report of any kind of alarm in the inner breakers then the probability of a busbar fault decreases, while knowing for sure that there is no alarm means that the probability of failure increases. This behaviour, which is confirmed by the data below, computed by Pulcinella, is rather odd, and marks a departure of our model from the specifications.

Report	none	inst	delayed	no alarm
p(s)	0.000175	0.00008	0.00008	0.000176

We now turn to considering the belief function model. In order to validate this model, we first need to extract the conditional beliefs from the joint mass assignments given in Section 2. To compute the conditional of, say, L1 and S1 given D1 = del, i.e. *bel*(L1&S1|D1=del), we consider all the second rows in the joint distributions (those corresponding to D1 = del), and combine them using Dempster's rule. The full set of conditional assignments over L1 and S1 is shown in Fig. 5. From this we can establish, for instance, that for the inner circuit breakers bel(m|e) = 0.9, bel(l|d) = 0.9, and bel(l|c) = 0.97, while for the outer breaker, bel(m∩t|e) = 0.9, bel(l∪s|d) = 0.93, bel(l∪s|c) = 0.27, bel(l|c) = 0.7 and all other conditional beliefs are zero. From bel(m∩t|e) = 0.9 we know that bel(m|e) ≥ 0.9 and bel(t|e) ≥ 0.9. From (6), (7) and (8) we learn that for the inner circuit breakers we have:

$$\left[\frac{\partial bel(l)}{\partial bel(c)}\right] = [+], \quad \left[\frac{\partial bel(l)}{\partial bel(d)}\right] = [+], \quad \left[\frac{\partial bel(l)}{\partial bel(e)}\right] = [0]$$

while for the outer circuit breakers

$$\left[\frac{\partial bel(l)}{\partial bel(c)}\right] = [+], \quad \left[\frac{\partial bel(l)}{\partial bel(d)}\right] = [0] \quad \left[\frac{\partial bel(l)}{\partial bel(e)}\right] = [0]$$

showing that our model behaves as it should. This prediction is once again verified by running

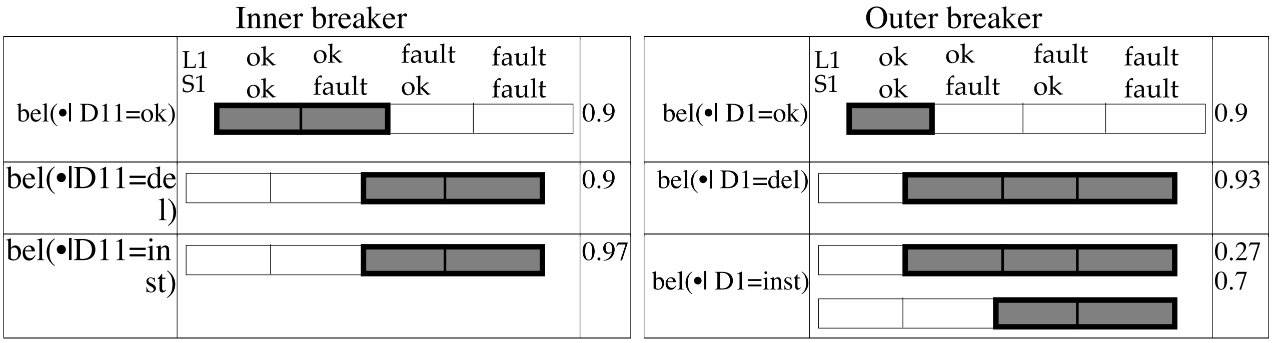

Figure 5. Conditional belief funtions corresponding to the masses in Fig. 3

Pulcinella on some sample data ((a) = outer breakers, (b) = inner).

Report	none	inst	delayed	no alarm
(a)	0	0.98	0.0	0.0
(b)	0.0	0.97	0.9	0.0

Similar results may be established for the other cases.

5. Conclusions

We have shown that the qualitative analysis may be used to validate the quantitative model since the qualitative predictions may be compared against the opinion of the domain expert to determine whether the model has captured the expert's knowledge. In our example, this validation exposed an anomaly in the probabilistic model: the corresponding valuation system, with values as ascertained by the knowledge engineer, does not behave quite as might be expected from the description of its intended behaviour that is supplied in Section 2. It is important to notice that in every case the qualitative predictions were borne out in practice by running our models on Pulcinella, indicating that the qualitative analysis of the uncertainty handling formalisms is accurate.

The fact that the predictions are observed in practice does not mean that the qualitative analysis is redundant: quite the opposite. The qualitative equations focus on one aspect of the model — the way some values are influenced by some observations, while abstracting away from the numerical configurations of input-output values. This allows us to spot those points where the behaviour of the model does not meet the specifications without having to go through an empirical sequence of numerical tests. After the qualitative analysis has spotted an unexpected behavior, numerical tests can be carried out to assess the performance of the system at these points, and the quantitative impact of the discrepancy can be evaluated. In our case, Pulcinella has shown small quirks in the probabilistic model in the direction predicted by the qualitative analysis. It is the task of the model designer to decide whether these quirks are important enough to require a correction to the model. If so, the qualitative analysis can guide us, by indicating which are the crucial derivatives for the undesired behaviour, and which are the quantitative values they depend upon. Moreover, the qualitative analysis can tell us which are the values we can freely change without affecting other "healthy" behaviour. In the full report [5], we describe a systematic technique for debugging the quantitative values based on the results of the qualitative analysis.

The qualitative analysis is a way to analyse the behaviour of a system at a high level of abstraction. As such, it relies on weak information and produces results that, although correct, may at times be too weak to be useful. We have shown in the above how we can enrich a purely qualitative analysis by introducing some informal order of magnitude arguments.

Acknowledgements. Work by the first author was partially supported by Esprit Basic Research Action 3085 DRUMS, and a grant from British Telecommunications. The second author was partially supported by the Action de Recherches Concertées BELON founded by the Communauté Francaise de Belgique. Comments from Moisés Goldszmidt, Paul Krause, Inaki Laresgoiti, Philippe Smets, Elisabeth Umkehrer, Hong Xu and an anonymous reviewer have been very helpful.

REFERENCES

[1] Bobrow, D. G., (1984) *Qualitative Reasoning about Physical Systems*, North-Holland, Amsterdam.

[2] Druzdzel, M. and Henrion, M. (1993) Belief propagation in qualitative probabilistic networks, in *Qualitative reasoning and decision technologies* (N. Piera Carreté and M. G. Singh (eds.) CIMNE, Barcelona.

[3] Parsons, S. (1993) Qualitative methods for reasoning under uncertainty, PhD Thesis, Department of Electronic Engineering, Queen Mary and Westfield College.

[4] Parsons, S. and Mamdani E. H. (1993) On reasoning in networks with qualitative uncertainty, *Proc of the 9th Conf. on Uncertainty in AI*, Washington D.C.

[5] Parsons, S. and Saffiotti, A. (1993) The qualitative verification and debugging of qualitative uncertainty: a case study. Technical Report TR/IRIDIA/93–1, Université Libre de Bruxelles, Brussels, Belgium.

[6] Pearl, J. (1988) *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufman, San Mateo Cal.

[7] Saffiotti, A., and Umkehrer, E. (1991) Pulcinella User's Manual, Technical Report TR/IRIDIA/91–5, Université Libre de Bruxelles, Brussels, Belgium.

[8] Saffiotti, A. and Umkehrer, E. (1991) Pulcinella: a tool for propagating uncertainty by local computation. *Proc. of the 6th Conf. on Uncertainty in AI*, Los Angeles, Cal.

[9] Shafer, G. (1976) *A mathematical theory of evidence*, Princeton University Press, Princeton, NJ.

[10] Shenoy, P. P. (1991) Valuation Based Systems, Univ. of Kansas, School of Business, Working Paper 226.

[11] Shenoy, P., P. and Shafer, G. (1990) Axioms for probability and belief function propagation, in: *Uncertainty in AI 4* (R. D. Shachter, T. S. Levitt, L. N. Kanal, and J. F. Lemmer, eds.) Elsevier, North-Holland.

[12] Smets, P and Kennes, R. (1994) The transferable belief model, *Artificial Intalligence*, forthcoming.

[13] Wellman, M. P. (1990) Fundamental concepts of qualitative probabilistic networks, *Artificial Int.*, 44.