Mind Soc (2015) 14:35–55
DOI 10.1007/s11299-014-0158-4

An explanatory coherence model of decision making in ill-structured problems

M. Laura Frigotto • Alessandro Rossi

Received: 5 May 2014 / Accepted: 14 October 2014 / Published online: 19 November 2014

Abstract
Classical models of decision making deal fairly well with uncertainty, where settings are well-structured in terms of goals, alternatives, and consequences. Conversely, the typical ill-structured nature of strategy choices remains a challenge for extant models. Such cases can hardly build on the past, and their novelty makes the prediction of consequences a very difficult and poorly robust task. The weakness of the classical expected utility model in representing such problems has not been adequately solved by recent extensions. In this paper we offer an explanatory coherence model for decision making in ill-structured problems. We model alternatives as sets of concurrent causal explanations of reality that act as justifications for action. According to these premises, choice is based on an evaluation of the internal coherence and the consistency of competing explanations of the available evidence. This model is psychologically grounded on causal inference and builds on the connectionist tradition of explanatory coherence. To illustrate the model, we consider the decision of investing in a new technology and we discuss how changes in the structure of alternatives may impact on the solution. We show how the final choice depends on collecting the relevant evidence, making the suitable hypotheses, and drawing the consistent causal explanations linking the two.

Keywords: Strategic decision making; Ill-structured problems; Explanatory coherence; Constraint satisfaction; Novelty

1 Introduction

In changing and unfamiliar environments, strategy makers face unexpected competitive threats and work in business contexts that vary suddenly and often unpredictably. A rising number of decision-making problems concerns situations that are novel and difficult to interpret: strategic alternatives are blurred, causal relationships linking various events are not apparent, and the chain of consequences triggered by adopting a particular decision are largely unforeseeable. In decision theory, these situations signify ill-structured problems, (Simon 1973; Hayes and Simon 1974), i.e., problems in which alternatives, states of the world, and consequences are not given and their structure is the result of creative effort (Newell et al. 1958). In ill-structured problems, representations are continuously defined and redefined in an indefinite space (Hayes and Simon 1974; Ungson et al. 1981), and they require the decision maker to elaborate assumptions or theories that offer a structure to the problem such that it enables decision making. In approaching an ill-structured problem, decision makers may highlight different alternatives and propose different solutions, each with particular strengths and weaknesses. They address alternative points of view and create arguments justifying proposed solutions. While the decision on a well-structured problem is a right or wrong answer that derives from a given problem structure, i.e. an interpretation of the problem that is self-evident, the decision taken on an ill-structured problem requires the creation of a problem structure. We interpret this as deriving from the creation of a coherent interpretation of the problem context and we model decision as deriving from the most coherent interpretation of it.

Even if expected utility theory (EUT) is still considered the reference model in economic theorizing (Thagard and Millgram 1997), decades of behavioral studies on economic choice highlighted several limitation of the EUT in explaining decisions. EUT is not suitable for ill-structured problems, in particular, because in such cases the problem components are incomplete and indefinite. Decision models such as case-based decision theory (CBDT) have been proposed for quasi wellstructured problems, where the problem space of the problem cannot be presented as a complete set of alternatives, consequences and goals, but it can partially be built by drawing analogies with past similar problem cases. However, the extant literature has not proposed a model to address cases of true ill-structured problems i.e., novel problems where components cannot be reconstructed even by analogy. Such cases require radically different assumptions in order to be understood (Ungson et al. 1981) and they would also require radically different models.

In this paper, we propose a model of decision making for novel ill-structured problems that builds on explanation-based coherence theories (Thagard 1989; Thagard and Verbeurgt 1998). Existing coherence models of decision making have focused on the case of deliberative coherence, where a decision is viewed in terms of a complex action plan which results from the assessment of competing actions and goals according to some

M. L. Frigotto (&) _ A. Rossi
Department of Economics and Management, University of Trento, Via Inama, 5, 38122 Trento, Italy
e-mail: marialaura.frigotto@unitn.it
A. Rossi
e-mail: alessandro.rossi@unitn.it

coherence principle (Thagard and Millgram 1997; Millgram and Thagard 1996). This class of models particularly fits instances of decision making characterized by competing goals and the presence of a plurality of available actions viewed as building blocks of action plans. Conversely, we build on a similar, albeit distinct, class of coherence models (explanatory coherence) to address other instances of decision making characterized by a truly novel and ill-structured nature, such as choices concerning the investment in radically new technologies or the launch of breakthrough products. We explore the dimensions of the creative effort through which an ill-structured problem is given a structure. We represent the structure of such problem representations building on established psychological literature and the connectionist tradition on explanatory coherence. To illustrate the model, we consider a choice concerning investment in a new technology, and following the same approach adopted by Millgram and Thagard (1996), we use the model to discuss how changes in the structure of alternatives may impact on the final decision and how critical the single components are for the final decision.

The remainder of the paper is organized as follows. In the next section we frame novel problems as ill-structured problems and we discuss the limits of EUT vis-a`-vis modeling such problems. Through examples, we redraw the fit between wellstructured problems and EUT, and between routine problems and quasi illstructured problems with CBDT; we consider that despite their relevance, a model for ill-structured problem is absent. In section three we argue how our model contributes to the extant literature on decision making. In section four, we discuss the theoretical foundations of the model rooted in psychological and cognitive literature and describe the modeling of decision alternatives for ill-structured problems in novel settings. Moreover, we position our model within the constraint satisfaction tradition and provide a detailed and formal description of it. In section five we illustrate how changes in the structuring of alternatives may impact on decision making. We close with two sections, the discussion and conclusions respectively, in which we derive implications for the structuring of decision alternatives in novel problems and also propose some directions for future research.

2 Novel ill-structured problems and extant models of decision making

Expected utility theory is the unrivaled paradigm for decision making under risk and uncertainty. It considers that decision makers have a clear picture of the variables associated with the problem they face: they are fully aware of the set of alternatives from which they can choose and of the states of the world that influence the consequences of possible actions. In decision theory terms (Rubinstein 1998), considering the problem A, the decision maker can choose any alternative x belonging to A, and the chosen x* cannot be less preferred than any other x belonging to A. This implies that rational agents make inferences about the environment in which they operate, an activity which is dependent on their ability to analyze the situation as well as to easily access the alternatives, the states of the world and the consequences that are relevant for the given problem. Also, they neither invent nor discover new courses of actions; that is, the chosen x* cannot be outside the set A. In other words, they deal with well-structured problems where the solution is implied in the problem structure and found through the application of an algorithm.

However, in decision making under uncertainty, alternatives and states of the world ''are neither naturally given, nor can they be simply formulated'' and ''even a comprehensive list of all possible outcomes is not readily available or easily imagined'' (Gilboa and Schmeidler 1995: 606). CBDT has been proposed as a viable model for decisions in which decision makers do not explore the whole range of alternatives, states of the world, and outcomes. Instead, CBDT assumes that they draw analogies between past cases and the problem at hand. ''From causes that appear similar'' (Hume, quoted in Gilboa and Schmeidler 1995: 608), they expect similar effects, so they apply choices close to the ones selected in previously experienced analogous cases. In the context of strategic choice, Gavetti et al. (2005) have proposed an agent-based model of analogical reasoning, which seems to share the same theoretical foundations.

The limitations of EUT and CBDT surface when the present does not look like the past, thus making it difficult to imagine the constituents of decision (alternatives, states of the world, and outcomes), i.e., in the case of ill-structured problems. In such settings, the pattern of decision making cannot be adequately represented by extant theories of decision. With respect to EUT, the requirement of complete enumeration of the sets of alternatives, states of the world, and outcomes cannot be fulfilled. With respect to CBDT, strategic decisions corresponding to unique and novel cases of the past, which do not appear similar to present ones, render past experience unhelpful.

Consider the following examples. Example 1: the omelet problem (Savage 1954: 13–15). Consider making an omelet with six eggs. Five are already in a bowl. It's time to crack the sixth and pour it either into the bowl with the others or into a separate plate to examine its freshness first. This problem is typically characterized by uncertainty as it is not known whether the sixth egg is fresh or not, but it is clear that if the egg is fresh, one would prefer to pour it directly into the bowl and save on the effort to clean another dish. Conversely, if the egg is not fresh, one would prefer to pour it into the separate dish in order not to lose the other five.

Example 2: hiring a nanny (Gilboa and Schmeidler 1995: 606–607). A couple wants to hire a nanny for their child; however, they do not know how the nanny would perform once she is hired. She could be negligent, dishonest, too severe, untidy, or she might leave the job just after the child has got used to her. Uncertainty

here acts on several dimensions: there are many possible outcomes that might occur and several alternatives, and their exhaustive listing is not an easy task. In addition, if the couple was to consider all these possibilities, they would be overwhelmed by their number and complexity. To facilitate decision making in this complex situation, they might ask candidates for recommendation letters. They might consider past employers' references to assess the nanny's past competence and performance. They might then review the results in the context of their house and child and arrive at a rough idea of the behavior the candidate might exhibit, if hired.

Example 3: The launch of the personal computer in 1977[1]. In 1977 computers exist as typically huge devices that occupied a whole room. Installed only by large institutions (multinational corporations, research facilities, etc.), they are used by small groups of people for data calculations. In that year, a venture capitalist (VC) is confronted with the decision of whether and how much to invest in a new product invented by a couple of scientists with electronics skills: a new kind of computer suited for personal use. The new computer is very different from previous ones: it is smaller, cheaper (about 1,195 dollars), and includes a monitor and a keyboard. The design is attractive, and it comes equipped with artificial intelligence for playing games. However, the VC is not convinced. First, he argues that computers are complex devices and the general public is not educated enough to use them. Second, he does not see any convincing gap in the current consumer market that the product would fill. He tells the scientists: ''there is no reason for any individual to have a computer in his home'' and declines to support the project[2].

In example 1, EUT seems to offer a fair representation of decision making. Example 2 depicts a situation in which the simple consideration of the problem does not naturally suggest the elicitation of all possible outcomes or the definition of all alternatives and relevant states of the world. Moreover, the analytical construction of these sets would compound its complexity (Gilboa and Schmeidler 1995). CBDT is appropriate in this case as long as its antecedents—here synthesized in the form of recommendation letters—are available and are considered similar enough to the case at scrutiny. In example 3, similarity to past cases is so low that previous experience is negligible. On the one hand, past cases are not helpful in either suggesting viable solutions or directly supporting inference over the present situation. On the other hand, the PC evolution shows that it has been very difficult to predict the subsequent development of such a disruptive product, in that it is difficult to not only define the likelihood of such scenario but also conceive its appearance. This third example represents a strategic decision that shows the properties of ill-structured problems.

3 Positioning the model

Let's consider Gilboa and Schmeidler's (1995) refinement of Knight's distinction between risk and uncertainty, which includes a third category labeled ignorance. Risk relates to a situation in which probabilities are given; uncertainty to cases in which states are naturally defined, but the translation into probabilities is not. Ignorance refers to situations in which states are neither naturally given nor can be easily constructed (Gilboa and Schmeidler 1995: 622).

We suggest attributing such increasing difficulty in the definition of the problem elements to the novelty of the setting. The argument runs as follows: the newer the problem, the more difficult it is to identify all problem components, i.e., the more ill-structured the problem. This novelty-based scale allows us to position the extant decision-making models (Gilboa and Schmeidler 1995): EUT is associated with well-structured problems and low levels of novelty while CBDT is linked to quasi ill-structured problems and medium levels of novelty[3] EUT provides a good representation of decision making, and some generalizations of EUT (for instance, non-additive probabilities or multiple priors) have been specially developed to encompass particular aspects of uncertainty. When the problem is moderately novel but appears to be a variant of previously encountered settings, past experiences can be used to support decision making, and in these cases, CBDT seems to be a good representation of how decisions are taken.

---

[1] This example is freely built on the well-known statements made in 1977 at the Convention of the World Future Society by Kenneth Olsen, co-founder of Digital Equipment Corporation, and his later hindsight of his decisions not to invest in personal computers. Sources are: an interview with Kenneth Olsen by David Allison made on September 28–29, 1988, Division of Information Technology and Society, National Museum of American History, Smithsonian Institution, available at http://americanhistory.si.edu/collections/comphist/olsen.html#tc21, last visited February 05, 2014; Anderson (1984)

[2] About 30 years later we can clearly see that that decision was wrong. The PC has been a success and the market has grown at surprising pace. However, decision at that time was ill-structured and choice took the form of a claim and a justifying argument rather than of a calculation including estimates of risk and uncertainty.

[3] Note that CBDT is also suitable for modeling routine decisions, i.e. repeated well-structured problems whose solution has been memorized and is applied each time without processing the decision. Thus also in those cases when decision becomes rule-based, CBDT is a good model of decision.

Our model extends this sequence of models by considering ill-structured problems in which novelty is taken to the very extreme. When the decision setting is new and previously experienced cases do not offer comparable elements, decision makers need to expand their knowledge and come up with an ad hoc interpretation of the problem. In such cases, the problem needs to be constructed as no antecedent has been registered in previous experience. The model integrates CBDT in the sense that it models the construction of new cases of the repertoire. Like with EUT, our model captures the limited ability of agents to imagine the complete set of alternatives, states of the world, and outcomes and further the debate in the literature on decision making by assuming that solutions—in the form of decision settings (EUT) or stored cases (CBDT)—are not given.

4 The model

4.1 Framework

We model decision making of ill-structured problems by building on Paul Thagard's explanatory coherence theory (1989, 1992b, 2000a, b). We expand the range of numerous applications offered on scientific revolutions (Thagard 1989), scientific discoveries (Thagard 1998a), medical diagnoses (Thagard 1998b), adversarial problem solving (Thagard 1992a), juror's decisions (Thagard 2004), group decision making (Frigotto and Rossi 2012) by adapting this model to provide an explanation to novel, ill-structured decision problems. We maintain that decision makers deal with novel and ill-structured problems by producing competing causal explanations of events that act as interpretations of reality. According to such premises, decision makers select the interpretation of reality that maximizes the coherence of competing explanations of the available evidence.

This view of decision makers as solvers that build explanations of problems finds ample support in the psychological literature. Several studies in psychology (Keil 2006 for a review) indicate that explanation is a ubiquitous and fundamental element of human cognition. For example, children's tendency, even at the young age of three, to ask 'why' of every phenomenon stems from the need to give sense to the world around them in order to be able to act in it. For the same reason, lay people as well as experts use explanations when faced with new and complex situations. Through explanations, people explore different phenomena, attempting to uncover the reasons of why an effect occurred. They are willing to see the events they face as the result of causes, the identification of which is necessary to understand what is going on and how they may intervene so as to reach their goal. In other words, through explanations, people build the knowledge they do not have in other more rapid and immediate forms (e.g., stored knowledge as cases to draw analogy from, patterns or principles that activate rules or routines) and in this way provide an elementary structure to ill-structured problems that makes choice possible. Such rudimental structure is built on internal validity, i.e., explanatory coherence, while through experience it may also be tested for external validity.

As a matter of fact, in contexts of novelty, the main[4] way to build validity is internally as there is no time or opportunity to test the world and build externally valid knowledge before deciding and acting. Therefore, the criterion for choosing among various explanations is internal validity or internal coherence. Propositions must provide an internally consistent interpretation of the setting. Keil (2006), for instance, has shown that children need coherent explanations in order to make sense of situations, to be part of a world they understand and in which they can play an active role. Building on this evidence, explanatory coherence appears as an elementary pattern of reasoning employed when other more rapid and straightforward rationality shortcuts (such as the case-based type) are not available.

Thagard (1989) proposed seven principles of explanatory coherence and modelled their concurrent satisfaction as a constraint satisfaction network problem. Our model builds on such a framework and extends the scope of explanatory coherence models to novel decision making problems faced by strategy makers.

Despite the robust psychological evidence showing that internal coherence and explanations are fundamental elements of problem structuring and choice, we do not propose our model as a representation of how decisions actually takes place. Rather, we model decision in a plausible and psychologically grounded way in

---

[4] Our model refers to coherent theories of epistemic justification, usually formulated as an alternative to foundationalist views of epistemology (Gardenfors, 1990). Both coherentism and foundationalism try to answer the fundamental problem of epistemology, the regress argument that goes as follows: given some statement P, it is reasonable to ask for a justification for P, which is justified by another statement P', and so forth. Fundamentalists conclude that there is a set of statements that is self-evident and do not need further justification. All other chains of justification are based on them. For instance, Decartes' 'I think therefore I am' is the most famous example of a self-evident statement which supports the whole body of knowledge. Similarly, empiricists take observations as foundation of knowledge.

On the other hand, coherentists deny the validity of the regression argument taking the form: P'' justifies P' which justifies P etc. They consider justification a holistic process and not a chain of reasoning. P is justified because it coheres with some system of thoughts. With much less concern on the issue of truth in comparison with fundamentalists, coherentists usually identify the system with the complete set of beliefs of the individual or of a group or of a society. Within this paradigm, in our model explanations are required to be internally coherent.

Nevertheless, our model is constructed as a so-called ''discriminating coherence'' model (Thagard and Verbeurgt 1998: 16) in which elements resulting from observation or experiments are given priority. The algorithm spreads activation first to these elements that should be favored in the coherence calculation. In such way, we can have some confidence that the most coherent explanations ''also correspond to the world and are not mere mind-contrivances that are only internally coherent'' (Thagard and Verbeurgt 1998: 17).

the Carnegie tradition (March and Simon 1958; Cyert and March 1963): this model provides an account of human decision making that is "more psychologically realistic than classical decision theory" (Thagard and Millgram 1997, p. 440). It has been shown that decisions are taken for several reasons other than optimality or for none but the availability of the solution (e.g., March and Olsen 1986 garbage can model). However, several contributions suggest that decisions, especially in novel settings, are built or justified by a veil of rationality. Such rationality often takes the form of coherence packages with all the limitations that bounded rationality includes: sensemaking frames (Weick 1979, 1995), hypothetical histories or near histories (March et al. 1991), and cases (Gilboa and Schmeidler 1995). In a different setting, consultancy also relies on this thinking mechanism by promoting scenario planning even though the plethora of studies on consultancy or a research profile (Lindgren and Bandhold 2003; Ringland 2006) have not provided any robust argument on how and why this works. Our model provides theoretical support to these contributions as it formally models these logical constructions, justifies how they are constructed on the grounds of a well established debate in the psychological literature, and suggests how they may be manipulated to reduce bounded rationality limitations.

4.2 Overview

We propose a model of decision making under (extreme) novelty that is rooted in ECHO, a computational model of explanatory coherence (Thagard 1989, 1992b, 2000b; Thagard and Verbeurgt 1998), which has been used to model a large series of psychological and philosophical phenomena, ranging from individual cognitive processes to more aggregate instances. According to ECHO, coherence is modeled in terms of constraints positively linking various evidence and explanatory statements, while negative links represent inconsistency between propositions. Hence, the search for coherence in a problem representation is modeled as the attempt to reach the maximum level of satisfaction of the set of constraints.

More precisely, ECHO is an example of a constraint satisfaction model (CSM), a class of connectionist networks originally proposed by Rumelhart et al. (1986), where units (also known as nodes) represent particular traits, features, or hypotheses of some kind—in our case, elements of the problem representation—and where connections linking two units reflect the nature and strength of the underlying relationship between them. In ECHO, units involved in representing the problem are of two kinds: explanation units—or hypotheses—and observation units—or events.

Connections among units (also referred to as weights or schemata if considered as a whole) can be twofold as well: positive connections between two units (say, A and B) suggest that whenever unit A is supposed to have some state (positive activation, for instance), the same should be expected for unit B. Negative connections, on the other hand, convey competitive relationships between units, thus implying that the activation of one of the two should lead to the other unit being de-activated. In addition, the magnitude of a weight indicates the strength (or power) of the constraint linking the two units.

Such a CSM enables the study of coherence since connections can be conceived as restrictions that have to be taken into account when producing coherent structures. CSMs are characterized by their tendency to lean ("relax") toward configurations in which as many constraints as possible are satisfied. The procedure runs as follows: every unit starts with an initial activation (or input) that reflects—for event units—the fact that some evidence has been taken into consideration, or—for hypotheses units—that the agent has an initial belief in that element of the explanation. This initial activation usually does not represent an optimal state, thus, an iterative procedure updates the units' activation levels toward configurations in which as many constraints as possible are satisfied (with priority given to the constraints characterized by a larger magnitude). For a relevant class of constraint satisfaction (CS) networks (as in the case of Hopfield networks), it is possible to prove that this computational procedure (or "relaxation") converges in finite time to final activation values representing a local optimum.

4.3 Distinction from other constraint satisfaction models

Our model is different from other well-established CSMs discussed in the managerial decision-making literature (Axelrod 1984; Hutchins 1995; Marchiori and Warglien 2005). CSMs, such as Hopfield networks (1982), have been introduced to formalize the perception of visual patterns in parallel distributed processing (Rumelhart et al. 1986). On these grounds, CSMs of decision making (see for instance Hutchins 1995, Marchiori and Warglien 2005) typically conceive choice as simple dyadic categorization issues (e.g., distinguishing friends/foes in a given group of firms, assessing complementary/substitution relationships in a pool of products) occurring in repetitive, albeit noisy, settings. The decision maker is given as initial beliefs (initial units activation) a noisy observation from the field, and is endowed with a set of patterns (decision outcomes) that have been previously encountered and memorized and represent the result of a learning process. The relaxation process is aimed at representing the decision maker's ability to derive the proper memorized pattern from the noisy inputs. It mimics the mental process that makes it possible to correctly filter noise from the environmental stimuli, to identify the correct decision pattern. Even if the relation has not been made explicit in the literature, it is straightforward to see that this use of CSM is in line with the CBDT idea of having a set of cases as primitive, providing ‘‘a simple axiomatization of a decision rule that chooses a ‘best’ act based on its past performance in similar cases’’ (Gilboa and Schmeidler 1995: 605).

The goal we pursue with our model is different. We aim at modeling strategic decision-making processes in non-repetitive settings, and more precisely, in settings where previously accumulated experience is not readily helpful. Here, there are no cases comparable to the actual situation, and knowledge needs to be newly combined and elaborated in order to provide an interpretation of the decision setting. Interpretations are built through the definition of causal relations within observed evidence and available knowledge, expressed in the form of causal hypotheses. These relations also define the constraints that make a problem representation internally coherent. When decision makers assess two competing interpretations of a problem, they look for the interpretation that satisfies these constraints better. In this CSM, the purpose of relaxation process is not the identification of the nearest memorized pattern but that of the stable pattern characterized by the highest level of explanatory coherence.

Also, while traditional CSMs are based on networks of homogeneous units (all units are of the same type and represent traits, or features in the environment) arranged in a simple (one-layer) graph in which the relaxation rule allows us to separate units into two different groups, ECHO allows us to stylize causal relations in decision making even with the existence of different classes of units (e.g., competing sets of explanatory units and units of evidence).

4.4 Formalization

Our model is built according to the following building blocks (see Fig. 1):

(a) A decision-making setting is described as a network, in which units represent either a piece of available evidence (units of evidence) or hypotheses establishing causal explanations for one or several evidences (units of hypothesis or explanatory units). Also, special evidence units (units that are only connected with evidence units) are introduced in the model in order to clamp evidence units to positive activation values.

(b) Hypotheses are divided into two competing sets (theories A and B), formed by k and l elements that can be viewed as alternative interpretations of the world/strategic setting (e.g., the two sets can be viewed as a series of hypothetical statements supporting a go/no go decision to enter in a new market). Along with evidence units, they form a three-layer graph in which it is possible to associate a vector of the units' activations $\mathbf{s}=(s_1,\ldots,s_i,\ldots s_n)=(a_1,\ldots,a_k,b_1,\ldots,e_1,\ldots e_m)\in|\mathbf{R}^n$ (where m is the number of units of evidence and $n=k+l+m$). Let also $\mathbf{f}=(f_1,\ldots,f_i,\ldots,f_n)\in|\mathbf{R}^n$ (with $f_i=0$ *for* $i=1,\ldots,k+l$) be the activations' vector of the special evidence units.

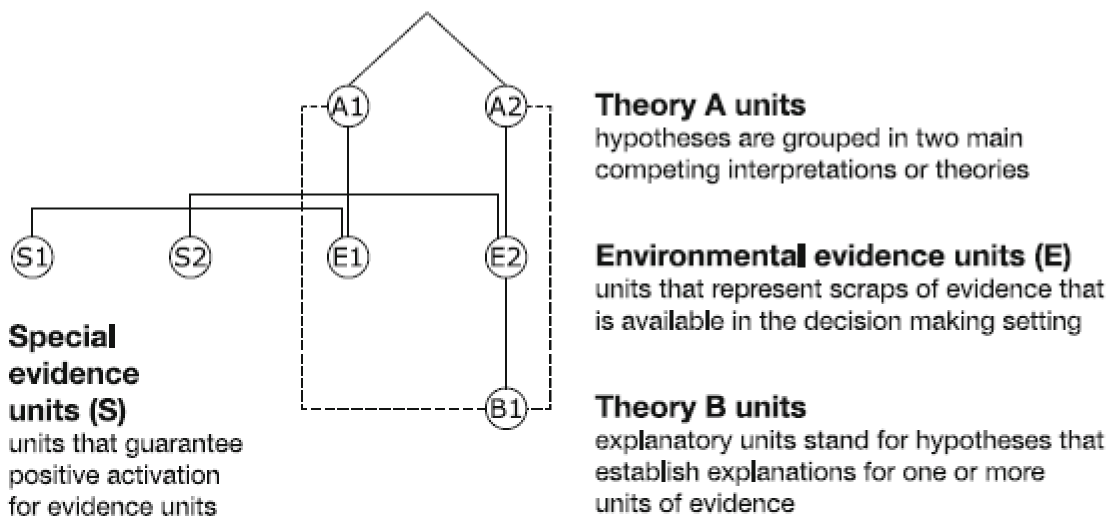

Fig. 1 A constraint satisfaction network. *Solid lines* represent positive connection between units, *dashes* represent negative (competitive) links

(c) Let W be a n x n null matrix and α the excitatory parameter. Positive connections between units of evidence and units of hypothesis represent direct causal relations (event x is causally explained by hypothesis *y*). Their intensity (or explanatory value) is coded through positive weights, so that a higher causal relationship corresponds to at a higher weight. The weights w_{ij} are updated according to the following steps: for each unit of type *e* that is causally explained by one or more explanatory units of theory C (C∈{A,B}), (1) let *i* correspond to the position of the unit *e* in **s**; (2) for each one of the *r* explanatory units of type *c* that explain s_i, let *j* correspond to the position of the unit *c* in **s** and set $w_{ij} = w_{ji} = \alpha/r$.

(d) Positive connections are also introduced between explanatory units that are causally linked to the same evidence and belong to the same theory, to introduce positive feedback between co-hypotheses. For each couple of units of type c that are co-hypotheses, (1) let *i and j* correspond to the positions of the two units in **s**; (2) set $w_{ij}=w_{ji}= \sum_{e=k+l+1}^{n} w_{ie} \cdot I$; where $I= \begin{cases} 1\ if\ wie = wje \\ 0\ otherwise \end{cases}$.

(e) We denote β as the inhibition parameter. Negative connections are introduced between explanatory units that are causally linked to the same evidence, but belong to competing theories, to account for the competitive relationship between alternative interpretations of the world. For each couple (a, b) of units that are co-hypotheses: (1) let *i* (*j*) correspond to the position of the unit of type *a (b)* in **s**; (2) set *p* as the number of evidence units that are jointly explained by s_i and s_j; (3) set q as the total number of co-hypotheses (of type a and b) that jointly explain the *p* evidence units; (4) set $w_{ij} = w_{ji} = \beta p/(q/2)$.

(f) Activation of units may take values in the (min, max) interval and are initially given weak positive activation (close to zero) in order to reflect agnostic initial beliefs on the various theories. Usually min = -1 and max = 1.

(g) Units' activation values are updated through a connectionist algorithm that is meant to increase the degree of coherence of the network in the sense that it performs a gradient-descent path toward units' activation levels that satisfy constraints more completely. At each iteration, units' activation levels are synchronously updated according to the following rule:

$$s_j(t+1)=(1-d)s_j(t)+\begin{cases} netj(max - sj(t)) & if\ netj > 0, \\ netj(sj(t) - min\ ) & otherwise \end{cases}$$

where $s_j(t)$ is the activation of unit *j* at time *t*; *d* is a decay parameter that, at each iteration, weakens the activation value of every unit and

$net_j = \sum_i WijS(t) + fi$ is the net input to unit *j*, computed as the sum of the activation of all the units weighted by the connections w_{ij} linking each of these units with unit j (including also the corresponding special evidence units f_j).

(h) The algorithm converges on a stable state where the units fit (in terms of coherence). This state is at a local maximum and is where the units' activation can be interpreted as the division of units of hypotheses into accepted units (with positive activation value) and rejected units (with negative activation value), eventually suggesting the existence of a "winning" theory or interpretation of the world (if all its explanatory units have positive values and all the other ones have negative values) that accounts for the available evidence better according to the explanatory coherence principle. Note that steps (*c*), (d*)* and (*e*), together constitute the matrix *W*, also known as the schemata. Note that *W* is symmetric, $w_{ii}=0$ *for i=1,…,n* and $w_{ij}=0$ for *i=1+k+l,…,n,j=1+k+l,…,n.*

It is worth mentioning that this class of constraint satisfaction network still lacks formal proof of convergence to a stationary state. Also, the relaxation rule does not guarantee maximized coherence, since the steady state might correspond only to a local maximum. However, there is a consistent body of literature that has produced results in line with empirical observations in many different domains of investigation (e.g., reconstruction of instances of legal reasoning, reproduction of many scientific debates taken from the history of science: see, for instance, Thagard 1989, 1992a, b; Nowak and Thagard 1992a, b; Eliasmith and Thagard 1997).

5 An application of the model: deciding to invest in a breakthrough technology

Deciding whether to invest in a breakthrough technology represents a typical case of decision making under (extreme) novelty to test our explanation-based model. In this category, we chose the case of the development of digital enhanced cordless telecommunications (DECT), a radically innovative technology established in the telecommunication industry[5].

In the mid-nineties, the worldwide market for mobile phone solutions was still in its infancy compared to the fixed-line phone market and was characterized by relatively higher hardware and service costs. At that time, many industry players (*in primis*, hardware producers and telecom providers) were interested in exploiting emerging market opportunities. In those days, a prospective technology in the form of digital enhanced cordless

[5] This case builds on extensive evidence collected by the authors both in terms of second-hand materials and via interviews with direct informants, experts in the telecommunication industry.

telecommunication (DECT), an European Telecommunications Standards Institute protocol for portable digital devices (Duff 2002), showed the potential for equipping home cordless phones with some basic mobile communication capabilities. Soon after the public release of DECT standard specifications, some industry operators started to ponder about investing in such technology.

The absolute novelty of the DECT system and the fact that no comparable cases were at hand to inspire a solution indicate that the decision to invest was taken in an unknown setting, characterized by information overload, no criterion for filtering good from noisy information, no powerful criterion for selection, and no previous consequences or success/failure experience.

In the first half of the nineties, portable communication devices and services were still limited to first generation (analog) mobile phones, which were, in comparison with today's handset devices, rather bulky and unreliable. Second generation (digital) mobiles, or GSM, did exist but the technology was in its early stages, with still-to-be-developed networks and relatively higher costs of hardware and services(4). The market for fixed communication, on the contrary, had already matured. There had been some interest in forms of semi-mobile communication, such as cordless fixed phones, but they showed some prominent transmission limitations even in closed environments. Given this scenario, DECT seemed to be suitable for energizing both markets. A product based on this technology could offer a more reliable and wide range of mobile connectivity to the land line market and realize the possibility of a portable phone within the range of a city. It could allow the user to make his home number portable within the metropolitan area and enable him or her to answer calls from places other than the actual home. Arguments in favor of the service launch were very powerful: the product seemed to target a promising market segment, made up of private consumers with relatively low mobility needs, interested in basic services (roaming limited to the local metropolitan area). Also, the service seemed competitively priced with respect to the existing mobile solutions and involved very low switching costs for fixed-line customers. On the other hand, arguments against the launch of the service highlighted the limitations of this technology in meeting the more sophisticated mobility needs of users, in terms of coverage and range issues and stability of the signal when travelling at high speed. On these aspects, the GSM technology appeared to be a better option provided it could be developed and was reliable and could be marketed in large volumes.

We used this case to test our model in the following way. We gathered inputs for the model by reviewing various second-hand data sources on DECT and conducting interviews with key industry informants. We distinguished between elements of evidence coming from the competitive environment or from the technological domain (in our model: statements of evidence) from arguments supporting or opposing to the investment in DECT (explanatory statements supporting, respectively, theory A—invest and theory B—do not invest). Then we screened the available documents again looking for causal explanations linking the explanatory statements with the statements of evidence. Coding was performed independently by the authors. Subsequently, every coded item was discussed and finalized as a component in the model until an agreement was reached. Figure 2 shows the components of the model. The fundamental evidence includes the following six components (Fig. 2, middle row): (E1) the market for fixed, domestic phone communication is in its maturity stage; (E2) there are unmet needs in the mobile communications market; (E3) the transmission performance of first generation mobile phones is limited; (E4) there is high penetration in the land-line market; (E5) DECT technology might enhance the home-users' portability/mobility; (E6) DECT technology offers reliable service within the metropolitan area.

The explanatory statements supporting theory A—invest in DECT—are represented by the following list (top row of Fig. 2): (A1) DECT allows greater reachability than traditional home cordless devices; (A2) switching costs from

Fig. 2 A formalization of the DECT 1 model, which support the decision to invest in the technology

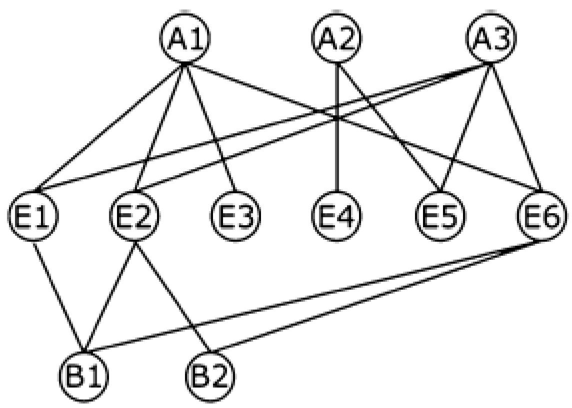

traditional landlines to DECT-based solutions are low; and (A3) DECT solutions meet the urban network communication needs. The explanatory statements supporting theory B—against investing in DECT—are the following (bottom row of Fig. 2): (B1) there will be a considerable demand for higher mobility; and (B2) mobile communication needs are not restricted to metropolitan areas and imply fast movement (e.g. users in urban areas often drive at speeds higher than the maximum speed compatible with DECT). Note that the causal explanations are drawn as links joining the explanatory statements with the causal statements, while all other links in the model (i.e., positive links between mutually supportive causal statements, negative links between competing causal statements, links between activation units and explanatory units) are not included to preserve clarity. We label this model as ''DECT 1.''

We run all the instances of the ECHO model below according to a choice of parameters that has been commonly employed in the previous literature. In particular, we set d = 0.05, α = 0.04, β = -0.06, initial activation levels for s_j=0.01 for j=1, ..., n; (fixed) activation for special evidence units f_j = 0.1 for j = 1, ..., m. The robustness of the parameters was determined by a sensitivity analysis to ensure that the qualitative results of the model do not change over a large parametric space. Taking the above-mentioned components into account, running the model resulted in positively activating the units corresponding to theory A, ''invest in DECT''. This result is consistent with the behavior of some telecommunication providers in the mid-nineties who were building DECT-based solutions, such as the Personal Handy-phone System, developed by NTT Docomo, a Japanese mobile operator and FIDO, the cordless terminal mobility network developed by Telecom Italia.

The large majority of the industry players were relatively more cautious about the potential of DECT technologies for mobile application. Figure 3 illustrates a decision-making process (labeled DECT 2) where such skepticisms are represented in terms of additional elements of the model (field or industry evidence, explanatory statements, causal explanations). As we can see, changes from the DECT 1 baseline model support views against the opportunity of investing in DECT. This model differs as follows: there is one more fundamental piece of evidence driving the evaluation process, namely, (E7) GSM technology is spreading very rapidly in the

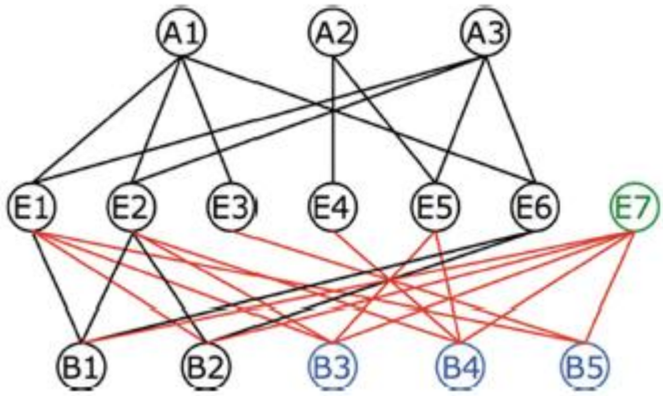

Fig. 3 A formalization of the DECT 2 model, which is against investing in the technology. In comparison with Fig. 2, the problem representation comprehends a larger number of evidence units (highlighted in *green*), explanatory elements (in *blue*), and underlying causal explanations (in *red*) (color figure online)

market. While the explanatory statements supporting the investment in DECT remain the same in the new model, there are three additional explanatory statements that support the ''do not invest'' alternative, which highlight the limitations of the DECT standard both with respect to GSM and with respect to rising sophisticated user needs. They are: (B3) each fixed line potentially hosts a plurality of users, hence concurrent use of DECT solutions might create conflict in use; (B4) intercom capabilities are unsatisfactory (handset/home communication); (B5) there are problems in perfecting the compatibility with other systems and regulation of the access to the network, which might delay the market debut. Finally, the set of causal explanations (links) is revised in accordance with the availability of additional units of evidence and explanatory units. In this case, running the model yields ''alternative B'' (do not invest) as the decision, in accordance with the widespread the industry strategy to not invest in DECT and with the operators experimenting with DECT solutions (such as the above mentioned Telecom Italia and NTT Docomo) in disinvesting early from their projects. The two instantiations of the model (DECT 1 and DECT 2) represent alternative ways to frame the investment decision that are logically constructed according to different assumptions. They result in different representations of the investment problem.

We can make twofold use of such modeling. First, as we just did, we can use it to make sense of the novelty encountered by strategy makers in the DECT challenge and is common to other breakthrough investment decisions. We can relate the different strategic reactions of the telecom players in the market to two alternative representations of the problem. Second, we can play with the problem representation and explore if strategies of incremental change in its components (e.g., adding more information or more robust explanations) result in reverting the investment decision. The result is relevant for justifying the recall of the explanation logic versus theories of lack of information or of information processing ability tout court, and for reasoning how the model can help understanding decision making.

To appreciate how the results of decision are dependent on problem structure components, we present Fig. 4, where we compare different problem representations in a cube. Starting from a baseline representation we introduce changes along three axes, representing a revision of the number and quality of, respectively, evidence units, explanation units and causal links. In particular, we examine the elements that should be included in order to change the investment decision from "yes" to "no." Thus, starting from DECT 1 problem representation (Fig. 2) that is reproduced in panel 4a), we introduce incremental changes that lead to the opposite generated by DECT 2 representation (Fig. 3), and that is positioned at the opposite corner of the cube (panel e).

Panels b, c, d of Fig. 4 represent three different scenarios integrating specific components of the DECT 2 model in the DECT1 original formulation of the problem: Fig. 4b refers to the same evidence and explanatory units of DECT 1, while including the additional causal explanations of DECT 2 (highlighted in red).

Figure 4c depicts a setting where, starting from the DECT 1 model, additional DECT 2-specific explanatory units are introduced (in blue). Finally, Fig. 4d shows the effect of integrating the additional evidence units of DECT 2 in the DECT 1 model (in green). Note that in order for such new (evidence or explanatory) units to matter in the model, the corresponding additional causal links are also introduced. Taking this into account, the 4c scenario moves towards 4g, conversely the 4d scenario moves towards 4h. Overall, these three scenarios present three strategies for partially moving from the pro-investment problem representation depicted in the DECT 1 scenario to the anti-investment representation of the DECT 2 scenario.

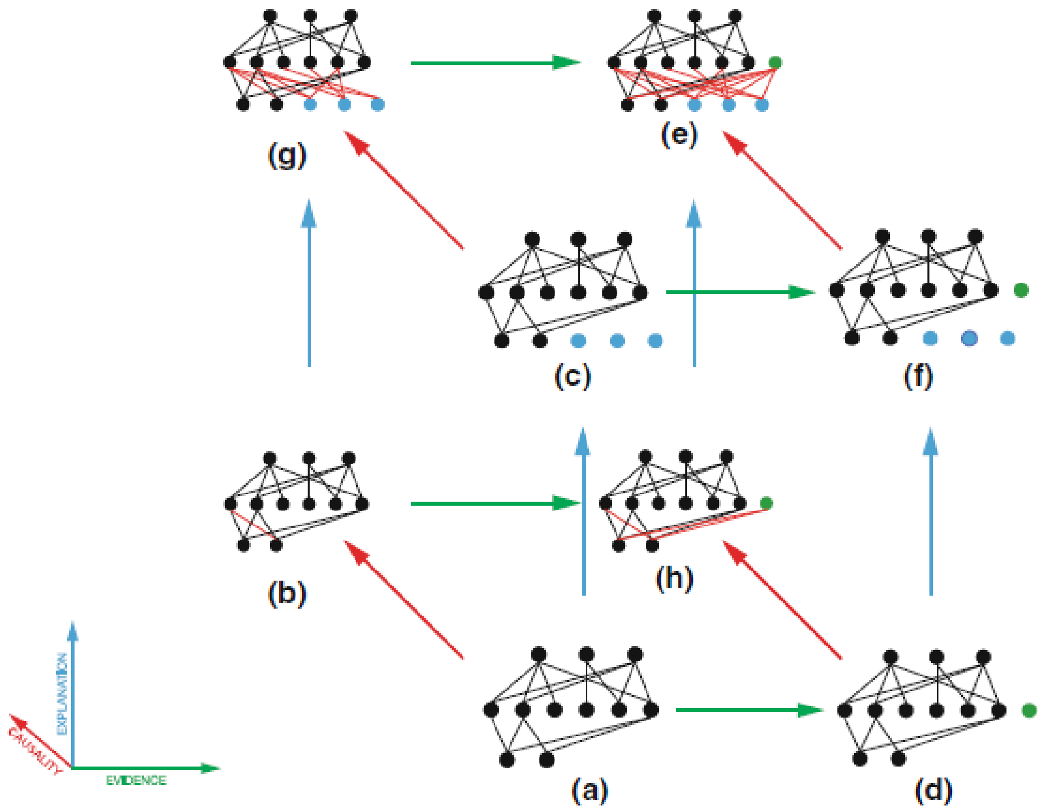

Fig. 4 The *cube* shows three different patterns for revising the problem representation. The baseline representation (DECT 1 in Fig. 1) is here depicted in (a). The revised and more complete representation (DECT 2 in Fig. 3) is depicted in (e). Different representations can be reached by revising the set of causal explanations (in *red*), the set of explanatory statements (in *blue*) and the set of evidence statements (in *green*) (color figure online)

These three strategies are: (1) look for more causal relationships between the various units, (2) look for more explanatory statements supporting a theory, and (3) look for more environmental evidence. In this respect, the DECT 2 model (Fig. 4e) can be viewed as the result of integrating these three partial strategies. These strategies represent three different patterns of action for enriching or revising the representation of the problem. In the specific case of DECT, they all prove to be ineffective, if pursued alone, in changing the idea on the investment decision because if we run the model, the units of alternative B are never activated at the steady state. These results indicate that local and incremental search strategies aimed at revising and integrating more information in the problem representation do not support a change in the decision making outcome. Conversely, a radically redefined problem representation (the one which is assumed in the DECT 2 scenario) results in a different decision outcome.

6 Discussion

The paper proposes a model of decision making for novel ill-structured problems, which is founded on the explanationist tradition of causal reasoning. By using the case of investment in DECT technologies for testing our model, we identified in DECT 1 the logical components that supported the ''invest'' scenario, and in DECT 2, those elements which made the decision-maker opt against investing. We have played with the differences between these two representations and discussed the impact of the introduction of different components supplementing the problem representation, i.e., the availability of different information and different inferential capabilities on explanatory statements and causal explanations.

Looking back at the rapid advancements in the mobile telecommunication industry between the late nineties and the new millennium, one could easily argue that the time to market and the mobile capabilities of the phone were the two most critical elements that did not favor the adoption of DECT. At the same time, it would not be prudent to consider operators such as Telecom Italia or Docomo unwise because they decided to invest in such technologies. Rather than supporting the view that all the elements were there ''waiting to be picked up'' (Seidl 2004: 156), the DECT 1 and DECT 2 models show that investment decisions are taken according to different problem representations that are constructed and not simply given. We argue that revising an investment decision cannot be interpreted only in terms of inertia, or sunk costs logic, but rather it needs to be considered as a process of revision that needs time in order for the decision-maker to reconsider his or her view of the world and translate new intuitions and insights into components of the decision-making model.

Our model offers important implications. It may support problem solvers by serving as a tool for analyzing interpretations or scenarios and for assessing their validity. In addition, it can also be used to consider scenarios that may otherwise remain unexplored, though changes in model components, i.e., causalities and events. Further, our model provides decision makers with a tool to reason the structure of the future hypotheses and identify their fragility if one event does not take place. It not only shows how this may impact the overall internal consistency of the construct but also allows for the possibility of expanding information or the sources of evidence and causalities. In scenario planning practices, this may be particularly useful for evaluating competing scenarios built by managers. In particular, our model suggests that playing with the components of the problem structure may be helpful in improving decision making in novel settings. On the one hand, it allows decision makers to assess the robustness of an interpretation and its sensitivity to peculiar hypothesized events. On the other hand, it facilitates access to new and distant problem representations that could have hardly been reached through local search.

Despite its benefits as a decision support tool, the model we present in this paper has some limitations. First, it yields decisions that are neither guaranteed to succeed nor externally valid. Conversely it supports better decisions to the extent that it promotes an extended search in the problem representation domain. Second, the model seems difficult to apply to real ill-structured and novel problems such as investment decisions in breakthrough technologies. However, several studies have positively tested the power of the model to support decision in several contexts (Schank and Ranney 1991, 1992, 1993; Ranney and Schank 1995, 1998; Read and Marcus-Newhall 1993; Spellman et al. 1993; Pennington and Hastie 1986; Wang et al. 2006). Moreover, explanatory coherence theory used in the model is psychologically recognized (Keil 2006) as a prominent basis for structuring a problem, and problem structure is necessary to build alternatives and chose among them. Our model builds on such robustness and extends the scope of explanatory coherence models to novel decisions faced by strategy makers.

To conclude, it seems worthwhile to reflect on the possible extensions of this model. Overall, the model seems to be good for representing individual decision making in novel settings or decision making within an organization considered as a single agent. It can also be extended in order to represent group decision making. The debate on heterogeneity in teams is still controversial. Various contributions show that groups of diverse agents reach better performances than the set of the best-performing individuals (Hong and Page 2001, 2004) as they access different problem representations and explore the decision settings more broadly. On the other hand, other perspectives on diversity (we refer to the similarity-attraction paradigm—see Williams and O'Reilly 1998, or Mannix and Neale 2005, for a review) suggest that when diversity is taken too far, agents face obstacles in sharing thoughts and in negotiating meanings which lack shared background knowledge. Our model can interestingly contribute to this discussion as it allows to control for two key variables heterogeneity and communication intensity at the same time, and it can explore the relation between the two. In another paper(6), we run this extended model and results show that increases in communication intensity are effective only if heterogeneity is not too high, whereas for very heterogeneous agents, communication is confounding and they are no longer able to come up with a coherent interpretation of the problem. In this case, agents perform better alone.

Our model paves the way for new interdisciplinary directions in future research for understanding decision making. In bounded rationality, initial problem representations define an anchor to successive thinking; our model suggests ways to lessen such limits. Studies on decision making and human cognition should help clarify, as a first point, why some problem representations are chosen as primitives, i.e., why they are more attractive than others so that strategy makers naturally see them, and secondly, how attention may be diverted from them and which strategies could help in expanding the set of constructed settings and consequently, increase the probability of success. Kahneman's lecture (2002) identifies that these issues are at the edge of behavioral decision-making studies. Cognitive psychology suggests that people create their problem representations, i.e., build contexts for the interpretation of problems, by eliciting relevant information and accessing knowledge structures, not only by rapid, automatic processes—what is normally addressed as intuition—but also via more conscious symbolic manipulation. However, these phenomena have only partially been defined in terms of dynamics, processes, and shortcuts. In particular, the way in which structures of knowledge (both intuition and more elaborated reasoning) are recalled and a distinction between the strategies for deciding and the context of decision have still to be clarified.

7 Conclusions

In this paper we provided a coherence-based model for novel decisions encountered by strategy makers, which is rooted in a well-established debate in the psychological literature. We showed how the problem structure may be manipulated in order to account for different evidence or explanatory statements and causal explanations. Our model reflects the idea that a problem representation is built in the form of a set of explanations supporting some evidence according to a rule of explanatory coherence. We used the model to illustrate the decision to invest in DECT solutions, which highlights the role of the various components of problem representation. This supports our theoretical construct that representations play a crucial role in highly novel contexts where choice is a byproduct of the problem representation (Simon 1991). In this respect, our model can be seen as a tool that can assist decision making by structuring the exploration of various problem representations.

Acknowledgments

We gratefully acknowledge the comments received on earlier versions of this paper from Nicolao Bonini, Guido Fioretti, Mimmo Iannelli, Alessandro Narduzzo, Paul Thagard, Massimo Warglien, Enrico Zaninotto, and the participants of seminars held at the Faculty of Science and at the Faculty of Economics at the University of Trento. We also would wish to thank the editorial board of the journal for the engaging conversation they stimulated through the selection of reviewers. In reconstructing the innovation case study, we benefited from the expertise and advice of Sandro Dionisi, Laura Riganti, Roberto Parodi e Leopoldo Tranquilli. The usual disclaimer applies. One of the authors was supported by the Professor Claudio Dematte` Memorial Grant during this research.

References

Anderson JJ (1984) Dave tells Ahl the history of Creative Computing. Creat Comput 10(11):66

Axelrod R (1984) The evolution of cooperation. Basic Books, New York

Cyert RM, March JG (1963) A behavioral theory of the firm. Prentice-Hall, Englewood Cliffs, NJ

Duff I (2002) Profile of Telecom Italia S.p.A. (Telecom Italia). Institute for Japanese–European Technology Studies, University of Edinburgh. http://www.telecomvisions.com/articles/pdf/telecom_italia_profile.pdf. Accessed 15 Jan 2012

Eliasmith C, Thagard P (1997) Waves, particles, and explanatory coherence. Br J Philos Sci 48:1–19

Frigotto ML, Rossi A (2012) Diversity and communication in teams: improving problem-solving or creating confusion? Group Decis Negot 21(6):791–820

Gardenfors P (1990) The dynamics of belief systems: foundations vs. coherence theories. Revue Internationale de Philosophie 44:24–46

Gavetti G, Levinthal DA, Rivkin JW (2005) Strategy making in novel and complex worlds: the power of analogy. Strateg Manag J 26:691–712

Gilboa I, Schmeidler D (1995) Case-based decision theory. Q J Econ 110:605–639

Hayes JR, Simon HA (1974) Understanding written problem instructions. In: Gregg LW (ed) Knowledge and cognition. Potomac, Lawrence Eribaum Associates

Hong L, Page SE (2001) Problem solving by heterogeneous agents. J Econ Theory 97:123–163

Hong L, Page SE (2004) Groups of diverse problem solvers can outperform groups of high-ability problem solvers. PNAS 101:16385–16389

Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. Proc Natl Acad Sci USA 79:2554–2558

Hutchins E (1995) Cognition in the wild. MIT Press, Cambridge Kahneman D (2002) Maps of bounded rationality: a perspective on intuitive judgment and choice. Nobel Prize Lect, December 8, also in Am Econ Rev (2003) 93:1449–1475

Keil FC (2006) Explanation and understanding. Annu Rev Psychol 57:227–254

Lindgren M, Bandhold H (2003) Scenario planning: the link between future and strategy. Palgrave Macmillan, New York

Mannix EA, Neale MA (2005) What differences make a difference? Psychol Public Interes 6(2):31–55

March JG, Olsen J (1986) Garbage can models of decision making in organizations. In: March J, Weissinger-Baylon R (eds) Ambiguity and command. Organizational perspectives on military decision making. Pitman Publishing, Marshfield

March JG, Simon HA (1958) Organizations. Wiley, New York

March JG, Sproull LS, Tamuz M (1991) Learning from samples of one or fewer. Organ Sci 2(1):1–13

Marchiori D, Warglien M (2005) Constructing shared interpretations in a team of intelligent agents: the effects of communication intensity and structure. In: Terano T, Kita H, Kaneda T, Arai K, Deguchi H (eds) Agent-based simulation: from modeling methodologies to real-

world applications. Postproceedings of the third international workshop on agent-based approaches in economic and social complex systems 2004. Springer, Berlin, pp 58–71
Millgram E, Thagard P (1996) Deliberative coherence. Synthese 108:63–88
Newell A, Shaw JC, Simon HA (1958) The process of creative thinking. The RAND Corporation, Santa Monica
Nowak G, Thagard P (1992a) Copernicus, Ptolemy, and explanatory coherence. In: Giere R (ed) Cognitive models of science, vol 15. University of Minnesota Press, Minneapolis, pp 274–309
Nowak G, Thagard P (1992b) Newton, Descartes, and explanatory coherence. In: Duschl RH, Hamilton RJ (eds) Philosophy of science, cognitive psychology and educational theory and practice. SUNY, Albany, pp 69–115
Pennington N, Hastie R (1986) Reasoning in explanation-based decision making. Cognition 49:123–163
Ranney M, Schank P (1995) Protocol modeling, bifurcation/bootstrapping, and convince me: computerbased methods for studying beliefs and their revision. Behav Res Methods Instrum Comput 27:239–243
Ranney M, Schank P (1998) Toward an integration of the social and the scientific: observing, modeling, and promoting the explanatory coherence of reasoning. In: Read S, Miller L (eds) Connectionist models of social reasoning and social behaviour. Erlbaum, Mahwah, pp 245–274
Read SJ, Marcus-Newhall A (1993) Explanatory coherence in social explanations: a parallel distributed processing account. J Pers Soc Psychol 65:429–447
Ringland G (2006) Scenario planning: managing for the future, 2nd edn. Wiley, Chichester
Rubinstein A (1998) Modeling bounded rationality. MIT Press, Cambridge
Rumelhart DE, Smolensky P, McClelland JL, Hinton GE (1986) Schemata and sequential thought processes in PDP models. In: Rumelhart DE, McClelland JL, the PDP Research Group (eds) Parallel distributed processing: explorations in the microstructure of cognition, vol 2. MIT Press, Cambridge, pp 45–76
Savage LJ (1954) The foundations of statistics. Wiley, New York
Schank P, Ranney M (1991) The psychological fidelity of ECHO: modeling an experimental study of explanatory coherence. In: Proceedings of the 13th annual conference of the cognitive science society. Erlbaum, Hillsdale, pp 892–897
Schank P, Ranney M (1992) Assessing explanatory coherence: a new method for integrating verbal data with models of on-line belief revision. In: Proceedings of the 14th annual conference of the cognitive science society. Erlbaum, Hillsdale, pp 599–604
Schank P, Ranney M (1993) Can reasoning be taught? Educator 7(1):16–21
Seidl D (2004) The concept of weak signals revisited: a re-description from a constructivist perspective. In: Tsoukas H, Shepherd J (eds) Managing the future: foresight in the knowledge economy. Blackwell, Oxford, pp 151–168
Simon HA (1973) The structure of ill-structured problems. Artif Intell 4:181–201
Simon HA (1991) Bounded rationality and organizational learning. Organ Sci 2:125–134
Spellman BA, Ullman JB, Holyoak KJ (1993) A coherence model of cognitive consistency: dynamics of attitude change during the Persian Gulf War. J Soc Issues 49(4):147–165
Thagard P (1989) Explanatory coherence. Behav Brain Sci 12:435–467
Thagard P (1992a) Adversarial problem solving: modeling an opponent using explanatory coherence. Cogn Sci 16:123–149
Thagard P (1992b) Conceptual revolutions. Princeton University Press, Princeton
Thagard P (1998a) Ulcers and bacteria I: discovery and acceptance. Stud Hist Philos Sci Part C: Stud Hist Philos Biol Biomed Sci 29:107–136
Thagard P (1998b) Explaining disease: correlations, causes, and mechanisms. Minds Mach 8:61–78
Thagard P (2000a) Probabilistic networks and explanatory coherence. Cogn Sci Q 1:91–114
Thagard P (2000b) Coherence in thought and action. MIT Press, Cambridge
Thagard P (2004) Causal inference in legal decision making: explanatory coherence vs. Bayesian networks. Appl Artif Intell 18:231–249
Thagard P, Millgram E (1997) Inference to the best plan: a coherence theory of decision. In: Ram A, Leake DB (eds) Goal-driven learning. MIT Press, Cambridge, pp 439–454
Thagard P, Verbeurgt K (1998) Coherence as constraint satisfaction. Cogn Sci 22(1):1–24
Ungson GR, Braunstein DN, Hall PH (1981) Managerial information processing: a research review. Adm Sci Q 26(1):116–134
Wang H, Johnson TR, Zhang J (2006) The order effect in human abductive reasoning: an empirical and computational study. J Exp Theor Artif Intell 18(2):215–247
Weick KE (1979) The social psychology of organizing. Addison-Wesley, Reading
Weick KE (1995) Sensemaking in organzations. Sage, Thousand Oaks
Williams KY, O'Reilly CA (1998) Demography and diversity in organizations: a review of 40 years of research. Res Organ Behav 20:77–140