Deciding Not To Decide

Ellsaesser, Florian and Fioretti, Guido

Frankfurt School of Economics and Finance, University of Bologna

10 January 2022

DECIDING NOT TO DECIDE

Florian Ellsaesser
Frankfurt School of Economics and Finance

Guido Fioretti
University of Bologna

Gail E. James
Gail James contributed a unique series of cognitive maps with her PhD thesis at University of Colorado, Boulder, 1996. We would like to have her as a co-author. If anyone knows where she is, please contact us.

Abstract

Sometimes unexpected, novel, unconceivable events enter our lives. The cause-effect mappings that usually guide our behaviour are destroyed. Surprised and shocked by possibilities that we had never imagined, we are unable to make any decision beyond mere routine. Among them there are decisions, such as making investments, that are essential for the long-term survival of businesses as well as the economy at large.
We submit that the standard machinery of utility maximization does not apply, but we propose measures inspired by scenario planning and graph analysis, pointing to solutions being explored in machine learning.

Keywords: Gail E. James, Uncertainty, Cognitive Maps, Machine Learning, Scenario Planning, Sense-Making, Bounded Rationality

JEL Classification: B49, D99, C02

We wish to thank Jochen Runde and Jean Czerlinki for helpful comments and remarks on previous versions of this manuscript.

Introduction

Sometimes, unexpected events destroy certain causal relations that used to provide a few firm signposts in spite of all uncertainty involved in managing a business. A global pandemic in the XXI century, as well the 2008 financial crisis are recent, yet by no means unique examples. The sentiment of lacking any sort of map to navigate life can have a substantial impact on forward-looking activities such as investments or strategic alliances, up to the point of questioning the very structure of one's own business model. At such times, decision-makers may simply decide not to decide. And this "wait and see" attitude does not stem from careful comparison of alternatives and probabilities of consequences, but rather from inability to carry out such a comparison. Uncertainty, in this case, concerns the set of possible events rather than the probabilities that link them to one another.

There are fundamentally two reasons why uncertainty may be different from taking expectations. The first one is relevant when probabilities cannot be reliably measured because the sample size is too small or, in the limit, no sample exists at all. From a purely conceptual point of view, this is not an issue for subjectivist interpretations of probability theory (De Finetti, 1931; Ramsey, 1931). Furthermore, by means of sub-additive probabilities it is possible to account for insufficient sampling and yet retain the whole formal machinery of probability theory (Gilboa, 1987; Schmeidler, 1989).

However, Ellsberg's paradox is there to stress that purely formal solutions do not satisfy practical purposes. Ellsberg (1961) asked his readers to consider two urns A and B. Suppose to know that A entails black and white balls in equal proportions, whereas all you know about B is that it entails black and white balls. Ellsberg remarked that, although the Principle of Sufficient Reason suggests to attach probability ½ to extract a white (or a black) ball from both urns, no-one would experience the same amount of uncertainty. The reason is that case A is equivalent to having tossed a coin infinitely many times, whereas B amounts to expressing oneself on a sample of size zero. Sub-additive probabilities come to rescue, for this extension of the theory allows probabilities not to sum up to unity if information is less than complete (Gilboa, 1987; Schmeidler, 1989). In the limit, one may assign probability zero to *both* extracting a white or a black ball from urn B, in which case taking expectations would rightly suggest decision-makers to select urn A. Nevertheless, sub-additive probabilities are not the last word on Ellsberg's paradox.

One problem with this patch to the received wisdom is that, in reality, one may be confronted with urns of type B only. If all options are such that little or no experience is available, attaching very low sub-additive probabilities to all consequences is of little help. In spite of sub-additive probabilities, we are back to the case upon which Knight (1921) had called attention. In the case he famously labelled "uncertainty" as opposed to "risk," probabilities cannot be a reliable guide simply because none is there.

We submit that, in this case, decision-makers analyse the structure of their mental representation of available possibilities, looking for areas where alternative courses of action do not fan out into widely different consequences. We shall document this practice in the context of Scenario Planning, pointing to the existence of numerical measures of the intricacy of structures that can be used to assess the strength of this sort of uncertainty.

The second reason why probabilistic expectation may be useless runs much deeper. It is due to a sort of uncertainty, eventually qualified as "Keynesian," "fundamental," "epistemic," "ontological," or "true" uncertainty[1] (Runde, 1990; Davidson, 1991; Dunn, 2001; Dequech, 2004; Lane and Maxfield, 2005), that arises when decision-makers fear that something may happen, that they are not even able to figure out. This sort of uncertainty is likely to be there if something that had not been conceived actually materialized, turning an "unknown unknown" into a "black swan" (Faulkner, Feduzi and Runde, 2017). It is worth remarking that no uncertainty is there insofar the unknown remains unknown. This sort of uncertainty rather arises at the point in time when something that had not been imagined suddenly enters the set of possibilities that a decision-maker envisages, particularly if this novel possibility is such that it upsets the network of causal relations in her mind.

Just as sub-additive probabilities can formally extend utility maximization to samples of zero size, incomplete preferences are eventually called in in order to extend the framework of utility maximization onto the problem of unknown unknowns (Eliaz and Ok, 2006; Ok, Ortoleva and Riella, 2012; Galaabaatar and Karni, 2013). Such a point of view may be questioned, for incomplete preferences in any case concern known possibilities whereas the above sort of uncertainty arises out of fear that something may happen, that one is not even able to figure out. Or it may be accepted, for this sort of uncertainty arises when a novel item enters the possibility set, that upsets existing preferences.

We do not take sides, because we deem that the fundamental problem runs much deeper. In our opinion, the flaw of all the above patches is that technical perfection does not imply usefulness. Which is the purpose of utility functions that cannot direct choices, because certain preferences are missing? With sub-additive probabilities and incomplete preferences utility maximization can be formally extended to encompass exotic types of uncertainty, but greater generality is obtained at the price of giving away any practical usage for this theory. With sub-additive probabilities all close to zero and no preference over key alternatives, expected utility maximization reduces to an empty bark.

We previously submitted that the uncertainty deriving from missing samples can be observed and measured by resorting to structural measures on possible causal linkages between alternatives and consequences. Quite similarly, we submit that also the uncertainty deriving from the surprising, destructive arrival of "unknown unknowns" can be observed and measured out of structural measure. In particular, we shall illustrate this possibility on a series of cognitive maps that have been traced before, during and after a destructive event.

Finally, we point to the possibility opened up by machine learning for carrying out such analyses automatically on huge amounts of data. It will happen in a foreseeable future, and it will open up unprecedented possibilities to observe non-probabilistic uncertainty, react to it, and elaborate solutions.

The rest of this paper is organized as follows. The ensuing section proposes to measure the level of uncertainty due to insufficient sample size out of the structure of alternative scenarios. The subsequent section illustrates a unique series of cognitive maps where an unexpected "unknown unknown" has been recorded, proposing simple structural measures. Albeit expected utility maximization is useless in these situation, human reactions are not irrational and are eventually captured by machine learning algorithms, which are the subject

[1] While Davidson (1991) contrasts "epistemic" to "ontological" uncertainty, taking the latter as a synonymous of probabilistic uncertainty, Lane and Maxfield (2005) employ the term "ontological uncertainty" in pretty much the same sense as Davidson's "epistemic" uncertainty.

of our penultimate section. Finally, conclusions attempt to assess the difficulties and prospects of the techniques illustrated hitherto.

The Complexity of Scenarios

Scenario Planning emerged among business strategists as a procedure to become aware of available options through extensive discussion and intentional search for non-obvious possibilities that may upset the received conventional wisdom (Schoemaker, 1995; Van der Heijden, 2000; Chermack, 2004; Roxburgh, 2009; Ramírez, Österman and Grönquist, 2013; Erdmann, Sichel and Yeung, 2015). The outcome of this exercise is a set of scenarios that have the purpose of preparing strategists to non-trivial future contingencies.

Scenarios constitute a network of concepts linked to one another by causal relations. Indeed, the network representation of scenarios is nothing but their authors' cognitive map (Goodier et al., 2010; Amer, Jetter and Daim, 2011; Jetter and Schweinfort, 2011; Alipour et al., 2017).[2] For instance, Figure (1) illustrates a possible cognitive map for the effects of uplifting sanctions on Iran oil revenues, depending on the role played by renewable energy sources as well as shale oil (loosely inspired by Alipour et al., 2017). This map highlights that lifting sanctions to Iran may make this country increase its oil production, but that depending on other factors stagnant oil production is equally possible. Notably, the value of this exercise is in highlighting non-obvious outcomes which may not be apparent at first sight.

[2] These Authors generally adopt a probabilistic interpretation of scenarios. However, the connection they make between scenarios and cognitive maps remains equally valid if probabilities are not used.

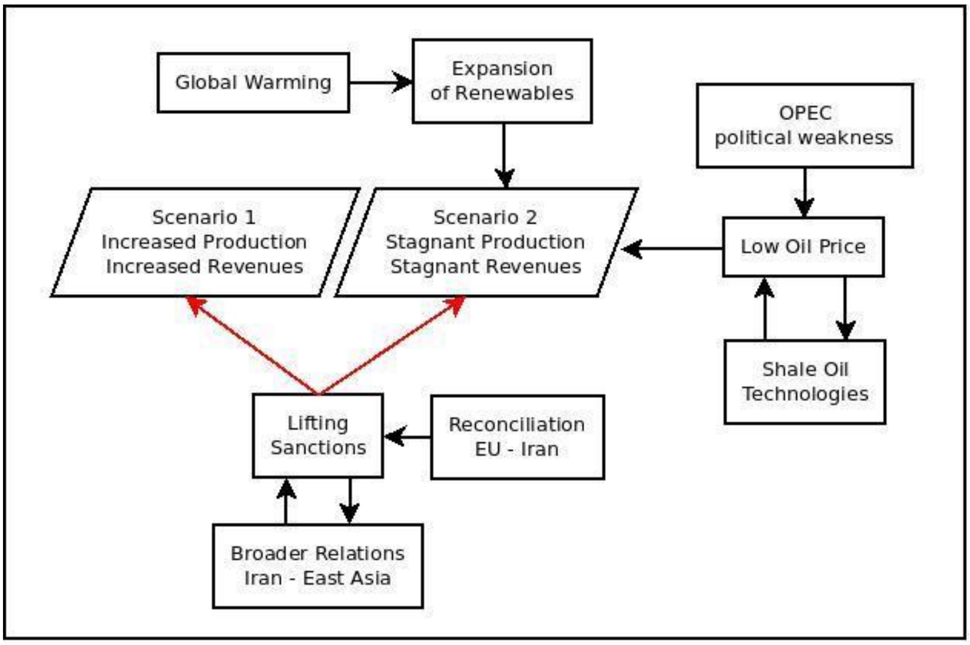

Figure 1. One and the same action - lifting sanctions to Iran - may lead to quite different outcomes depending on many other factors, such as growing availability of renewable energy sources or shale oil. One-to-many causal relations are highlighted in red. Loosely inspired by (Alipour et al., 2017).

One possible objection is that scenario planning eventually boils down to comparing expected values, and indeed many practitioners attach probabilities to the scenarios that they envisage in order to identify a most likely scenario upon which decisions should be based. However, others oppose this usage of scenario planning because scenario planning has the purpose of eliciting non-obvious possibilities whose probabilities are impossible to assess, whose structure of relations can be analysed nevertheless (Wilson, 2000; Goodwin and Wright, 2001; Wright and Goodwin, 2009; Ramirez and Selin, 2014). For instance, in the relatively trivial example of Figure (1) decision-makers may not even attempt to estimate the likelihood of the two scenarios of increasing oil production and stagnant oil production, for the value of this exercise rather resides in the very fact of realizing that non-trivial interactions are possible - in this case, swift development of renewable energy sources, or shale oil - that may reverse the outcome of lifting sanctions to Iran. In this way scenario planning is used to deal with the first sort of uncertainty that we discussed in the introduction, the one that arises when possibilities are correctly envisaged but no reliable estimate of their probabilities can be made.

When probabilities cannot be assessed, the added value of scenario planning resides in its ability to show how possibilities are linked to one another, for the very structure of causal relations provides a deeper understanding on what may happen in comparison to a mere list of possible outcomes. More specifically, uncertainty arises when a one-to-many structure appears, as it happens with the causal relations highlighted in red in Figure (1).

More in general, the structural features of cognitive maps can be captured by appropriate indicators. Henceforth, in order to keep matters tractable we shall assume that cognitive maps are made of alternatives and consequences. For instance, the cognitive map of Figure (1) would be reduced to one alternative (lifting sanctions) and two consequences (increasing or stagnant oil production, respectively). In order to stress their cognitive nature we shall adopt March and Simon's terminology of *Evoked Alternatives* (EA) and *Perceived Consequences* (PC) henceforth (March and Simon, 1958).

Figure (2) explains intuitively why formalization can add value to these considerations. On the left (a), relations between evoked alternatives and perceived consequences are one-to-one. This is the simple world where one perceives exactly which consequence follows from each of the evoked alternatives. In the middle (b), the extremely complex world where one perceives that any consequence can follow from each of the evoked alternatives. On the right (c) the somehow intermediate situation where, in spite of several one-to-many relations, one knows that certain evoked alternatives can only lead to a subset of the perceived consequences.

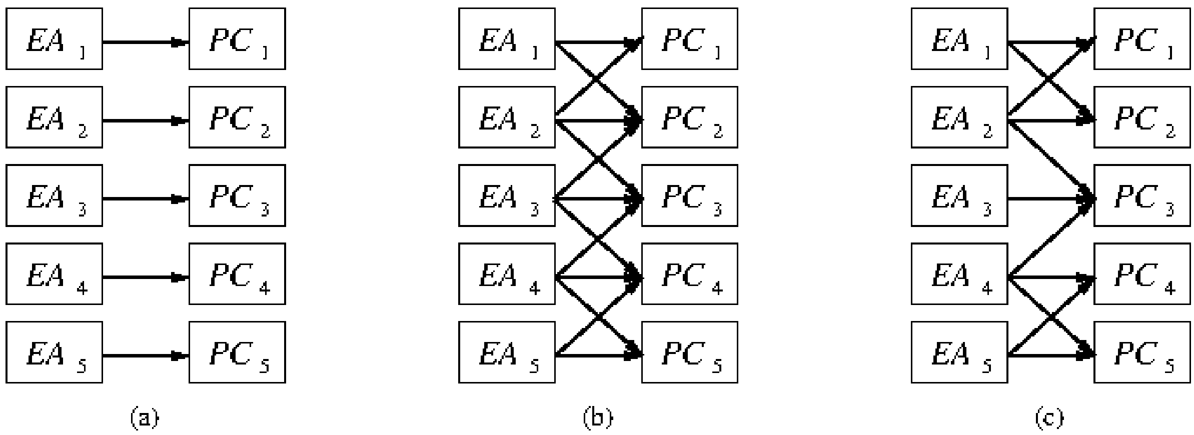

Figure 2. Three stylized cognitive maps linking evoked alternatives (EAs) to perceived consequences (PCs). In (a), a one-to-one mapping with complexity $C = 0$. In (b), the most confusing a one-to-many mapping where anything may happen. In this case, $C = 3$. In (c), an intermediate case where $C = 2$.

It is possible to define suitable measures of the *Complexity* of the graphs illustrated in Figure (2). One such measure (Casti, 1989; see also Appendix A) ascribes zero complexity to case (a), maximum complexity to case (b) and intermediate complexity to case (c). Its rationale is that case (c) exhibits some structure that moderates the uncertainty stemming from the fact that several consequences are perceived for most of the alternatives that are evoked. Structure limits the number of consequences that are perceived for at least certain alternatives, hence a lower value of complexity.

We submit that such measures of structural complexity could be used as a proxy for the uncertainty caused by missing samples. More specifically, boundedly rational agents could be assumed not to make any decision if one such measure exceeds a threshold.

When Unexpected, Novel, Destructive Possibilities Materialize

Let us now focus on the second reason why human beings may decide not to decide, namely when they become aware that some destructive possibility may materialise, which is not among those they are currently envisaging. In this case, inaction arises out of mistrust concerning the causal relations upon which they used to base their judgements.

Such sentiments make perfect sense once unexpected and destructive events have been experienced, such as major economic crises, epidemics or wars at the macro level, or technological breakthroughs at the microeconomic level. Such experiences may suggest that unthinkable novelties may appear in a world that is governed by unknown laws, if it is governed at all.

We submit that a series of cognitive maps can capture such states of mind. Cognitive maps were first introduced into social science by Robert Axelrod (1976) and later on applied to businesses by Anne Sigismund Huff and her co-workers (Sigismund Huff, 1990; Sigismund Huff, Huff and Barr, 2000; Sigismund Huff and Jenkins, 2002). In their simplest version, cognitive maps are network representations of world views whose nodes are concepts linked to one another by causal relations. Cognitive maps are generally obtained by analysing texts or recorded speeches, such as letters to shareholders or interviews.

Suppose that cognitive maps are elicited first before and then after one such disruptive event occurred, documenting how decision-makers made sense of their novel environment. We submit that the appearance of an “unknown unknown” would reflect into sudden structural changes of these cognitive maps, which could be documented by means of numerical measures.

Collecting cognitive maps entailing disruptive events is extremely difficult, for decision-makers typically hide moments of despair. No letter to shareholders, just to make an example, will ever list pending threats to conclude that management has no clue about how to proceed.

We accessed a unique series of cognitive maps capturing one such desperate magic moment that occurred to BioTech companies when they suddenly realized that they would never become Big Pharmas. A fairly long series of cognitive maps has been extracted from the written reports of a dedicated consulting company, 1986 to 1993, that includes that moment (James, 1996; See also Appendix B). Notably, the fact of analysing the reports of a consulting company allowed to capture the desperate recognition of an "unknown unknown" that single firms normally do not disclose.

In this case, the "unknown unknown" arose in 1990. Up to 1989 most biotech companies were convinced that they would grow up to become able to produce and market their own drugs. For the time being they had to stipulate strategic alliances with pharmaceutical companies, but this was seen as a temporary arrangement. By contrast, pharmaceutical companies were entering strategic alliances with biotech companies in order to acquire their technology. In 1990, biotech companies suddenly realized that many contract that they had signed entailed "poison pills" designed to squeeze their knowledge and profits (James, 1996).

Figures (3) and (4) illustrate a portion of the biotech companies' cognitive maps in 1989 and 1990, respectively (James, 1996). Note that in 1990 the block *Poison Pills* enters the map, destroying previous linkages.

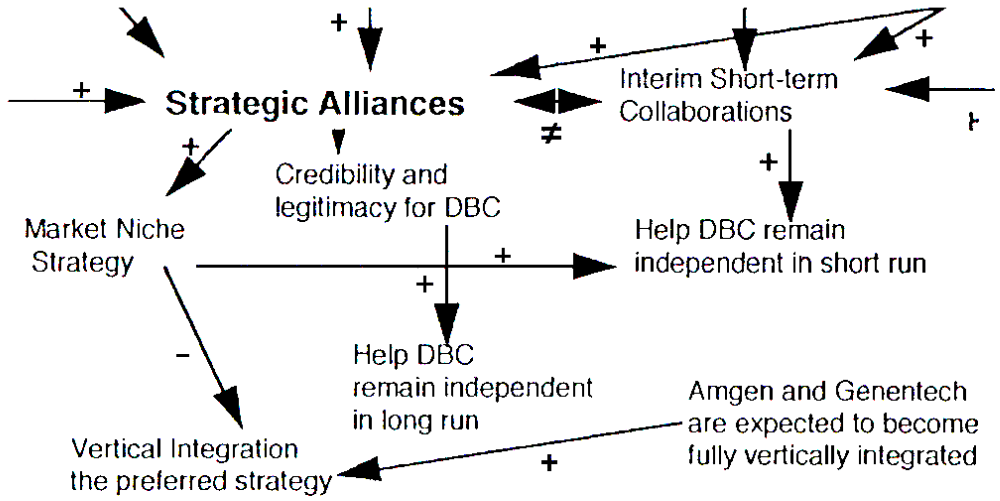

Figure 3. A portion of biotech companies' cognitive maps in 1989 (James, 1996). Incoming arrows stem from parts of the map not shown in this figure.

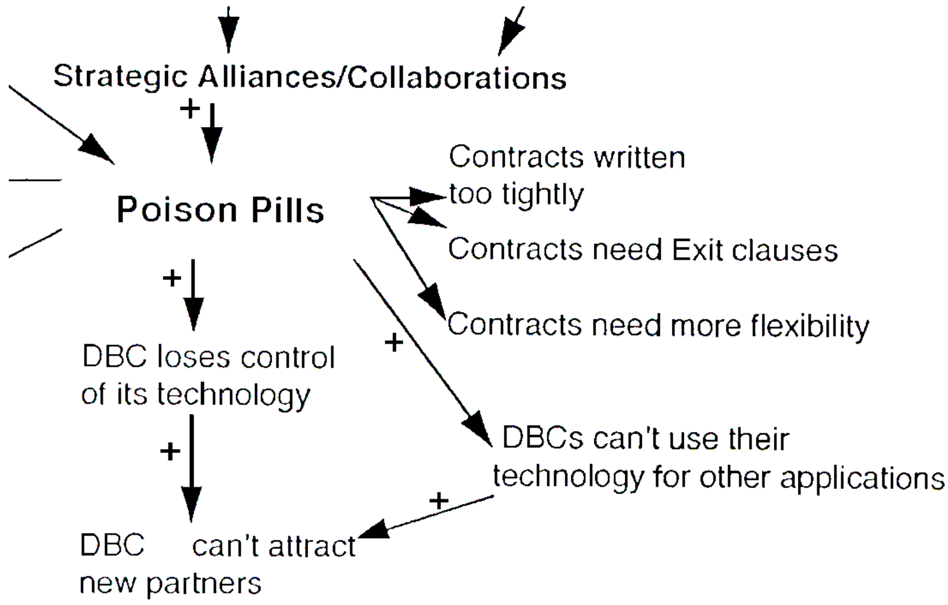

Figure 4. A portion of biotech companies' cognitive maps in 1990 (James, 1990). Incoming arrows stem from parts of the map not shown in this figure.

The rest of the story is that since 1991 biotech companies started caring about legal details, and at the same time realized the difficulties involved in drug production and distribution. Conversely, pharmaceutical companies realized that independent biotechs would guarantee a degree of exploration that in-house, hierarchically organized research could not attain. Thus, since 1991 the cognitive maps of biotech companies stabilized, providing again a reliable orientation to decision-making.

In the end, we have a series of eight cognitive maps 1986 to 1993, with one stable period 1986-1989, one stable period 1991-1993, and a disrupted cognitive map in 1990. A few simple metrics can be explored in order to identify indicators that are able to single out the disrupted 1990 map from those of the stable periods that precede and follow it.

Figure (5) illustrates the number of concepts and its average (black), the number of linkages and its average (green) as well as the ratio between the number of concepts and the number of linkages (bottom line), which is remarkably constant over time. Notably, the number of linkages (green) exhibits a marked drop with respect to its average in 1990. Also the number of concepts (black) drops in 1990, though possibly less sharply.

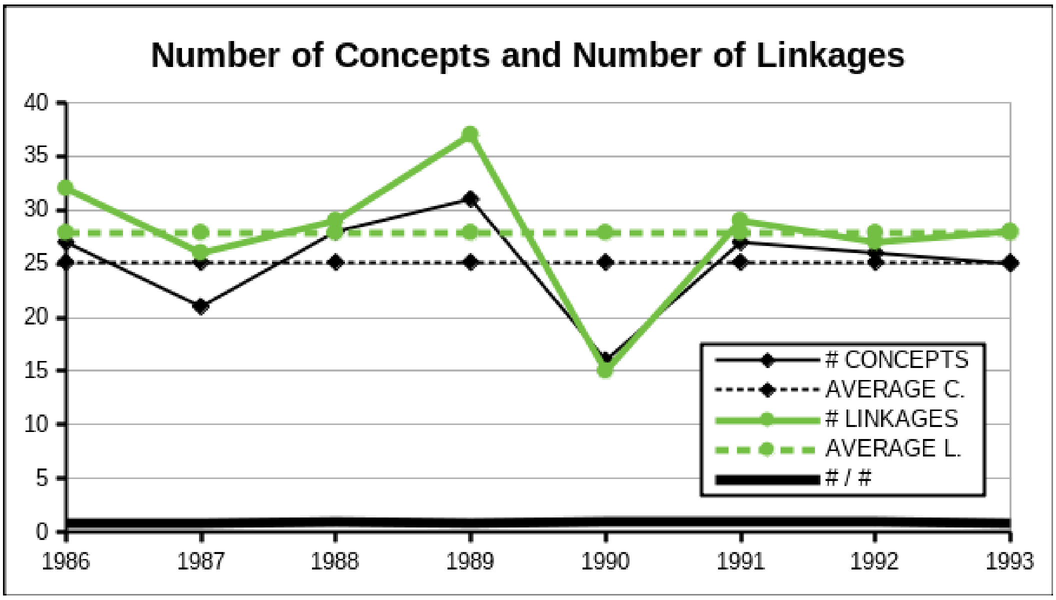

Figure 5. A few metrics of James (1996) series of cognitive maps. The number of concepts (solid black) compared to its average (dotted black), the number of linkages (solid green) compared to its average (dotted green), as well as their ratio (thick black). Notably, the ratio of two variable metrics is extremely stable, suggesting high correlation between the number of concepts and the number of linkages.

This series of cognitive maps is unique in its ability to grasp the magic moment of despair when an “unknown unknown” appears. Thus, it is impossible to make anything beyond an educated guess on possible indicators of the uncertainty that arises when decision-makers fear that possibilities might materialize, that they had not even been able to figure out. Both the variation of the number of linkages and the variation of the number of concepts appear to be good candidates, possibly with an initial preference for the first one.

Using machine learning to automatically generate cognitive maps

In the previous sections we illustrated quantitative analyses that can be carried out on scenarios and cognitive maps that are, unfortunately, seldom available. Manual reconstruction of possibilities, be it based on interviews, speeches or written text, takes time and effort. However, machine learning is opening up unprecedented possibilities to make such analyses feasible with reasonable effort.

In the last 10 years there has been a rapid development in Natural Language Processing (NLP), which uses machine learning algorithms to allow computers to process and "understand" language extracting the semantics of texts (Ruder, 2022). As a result, machine learning algorithms have surpassed human performance on a number of NLP tasks such as aspects of speech tagging or question and answering (Stanford NLP Group, 2022).

These developments built on previous advances in knowledge representation, which historically had been one of the key purposes of "good old fashioned" AI with the aim of building general or domain-specific knowledge graphs. A knowledge graph is a directly labelled graph consisting of nodes, also known as concepts, edges or linkages and sometimes, but not always, labels indicating the nature of these linkages. Henceforth we shall focus on knowledge graphs representing causal networks of concepts, hence we shall take terms "knowledge graph" and "cognitive map" as synonyms.

Knowledge graphs can be built from texts to represent knowledge in terms of relations between entities. Once knowledge has been represented as a knowledge graph, the graph can be traversed to draw relevant conclusions such as, e.g., whether *What'sApp* belongs to *Facebook* or *Joe Biden* is the *President of the United States*. Since 2019 there has been a revival of the attempt to combine statistical models that extract the meaning of texts with automatic knowledge representation in the form of graph neural networks (Zhang *et al.*, 2020).

Machine generation of cognitive maps from text processes the natural language through a pipeline composed by three main steps. The first step is pre-processing. The second step is concept recognition, which focuses on identifying the concepts which will be connected in the cognitive map. The third and last step is relational identification, where the linkages are determined.

Step 1: pre-processing.
The text is split into individual sentences and tokenized. Tokenization involves splitting up the entities of a sentence, words and punctuation into individual components. In our case, these components are embedded as vectors and punctuation is removed. Next the problem of co-reference is addressed, which is the task of finding all the expressions that refer to the same entity in the text (Clark and Manning, 2016). For instance, pronouns must be substituted by the nouns they refer to: "X is a public company. It made a loss in 2021" becomes "X is a public company. X made a loss in 2021."

Step 2: Concept recognition.
In order to recognize concepts, a machine must be trained on a meaningful network of concepts. For instance, one may train it on the network of 6.4 million concepts entailed in Wikipedia (Brank, Leban and Grobelnik, 2017). Whilst not every concept a CEO might talk about might be entailed in Wikipedia, the concepts described on Wikipedia are generally regarded as a reasonable baseline.

Step 3: Linkage extraction
Once the concepts have been identified, the linkages must be identified. The algorithm needs to identify what concept is related to another concept and how they are related. For this purpose, we take a pre-trained model. The BERT transformer is currently the model with the higher accuracy (Devlin *et al.*, 2019). Figure 6 illustrates these steps.

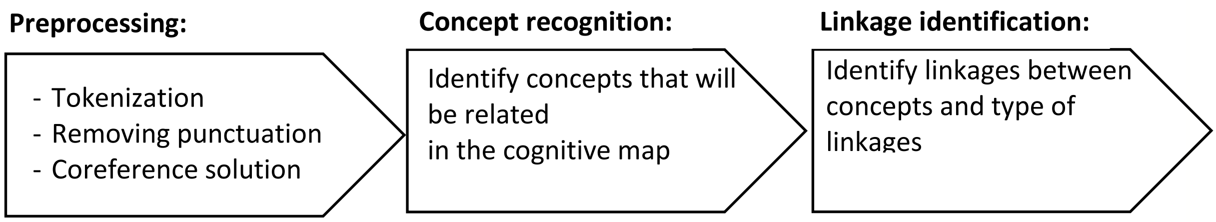

Figure 6. The pipeline for automatically generating cognitive maps from texts.

Extracting cognitive maps by means of machine learning algorithms has a potential to support new streams of research. One advantage of machine-based cognitive map generation is that concepts and links are always extracted exactly in the same way, no matter which text is submitted or which researcher is doing the job. Furthermore, a larger amount of documents can be processed because of lower costs.

However, hand-based extraction of cognitive maps is unlikely to disappear altogether. Human-based extraction allows more detailed, accurate and fine-grained analyses, whist the automatic generation allows a more consistent analysis over larger corpora. It is conceivable that, in the future, human extraction will still be used to calibrate the parameters of machine-based extraction.

Parameter identification is not trivial for machine-based extraction. Consider for instance a parameter such as entropy, which determines how strict the algorithm is in selecting concepts. If this parameter is set too loosely, concepts that are not appropriate for the text may be identified. If it is set too strict, relevant concepts may be missed. A somewhat related problem exists when cognitive maps are generated by different researchers, but human researchers adapt to the text they are analysing, whereas algorithms apply the same parameter value to texts that may be very different from one another. On the whole, machine-based extraction of cognitive maps entails more mistakes than human-based extraction, but can be applied to huge numbers of documents that are made available on the Internet or elsewhere.

In principle, we can figure out a world where businesses and Governments process texts available from sources ranging from letters to shareholders to blogs in order to extract cognitive maps that reveal perceived causal linkages and point towards future directions, for a wide variety of aims.

Consider, for instance, a likely future where machines analyse millions of documents looking for sudden breaks as the one identified by James (1996). One obvious consequence would be that certain actors – those who gear the machines – are able to detect such breaks earlier. One further possibility, already available for simpler tasks (Pathak *et al.*, 2017), is that machines suggest solutions for such breaks, such as the one the biotech industry found since 1991. It's not magic, it is simply a consequence of exploring automatically a huge number of possibilities.

Conclusions

Non-probabilistic sorts of uncertainty have been conceptualized since decades, but the prevailing attitude has often been that nothing it is impossible to formalize, measure, or make it the subject of normative prescriptions. We hope to have contributed to show that, quite to the contrary, practical steps can be undertaken in this direction.

According to our point of view, the structure of causal linkages between concepts is key to highlight and measure non-probabilistic sorts of uncertainty. By traditional means, this would be a terribly lengthy, and costly exercise to carry out. However, machine learning algorithms can usefully contribute to speed up the analysis of documents that are becoming available in increasing amounts and speed. Possibly, measurement of uncertainty out of automatized analysis of online texts will be one day just as common as the questionnaires that are distributed to entrepreneurs in order to ascertain the overall level of confidence and uncertainty.

Several weaknesses will certainly concur to limit the applicability of the approach we have outlined. Above all, the tendency of single individuals and organizations to hide the “magic moments of despair” and the consequent need to rely on assessments by third parties is likely to pose a major obstacle. However, machine learning algorithms are able to access documents of any sort over the Internet, including sources that are far from being official.

Appendix A: The Measurement of Complexity

This appendix illustrates the concept of complexity expounded in Casti (1989). Please consult the original reference for further details.

A *simplex* is the convex hull of a set of $(n+1)$ independent points in some Euclidean space of dimension n or higher. These points are its *vertices*. A 0-dimensional simplex is a point, a 1-dimensional simplex is a segment, a 2-dimensional simplex is a triangle. Henceforth, higher-dimensional simplices will not be considered.

The convex hull of any non-empty subset of the $(n+1)$ points that define a simplex is called a *face* of the simplex. In particular, 0-dimensional faces are the vertices of a simplex, 1-dimensional faces are segments that connect vertices. Two simplices are connected if they have a common face. A set of (at least) pairwise connected simplices is a *simplicial complex*.

Henceforth, a cognitive map is graphically represented as a simplicial complex. Evoked Alternatives are simplices whose vertices are the Perceived Consequences to which each Evoked Alternative is connected. Figure (A1) illustrates the simplicial complexes that correspond to cases (b) and (c) in Figure (2).

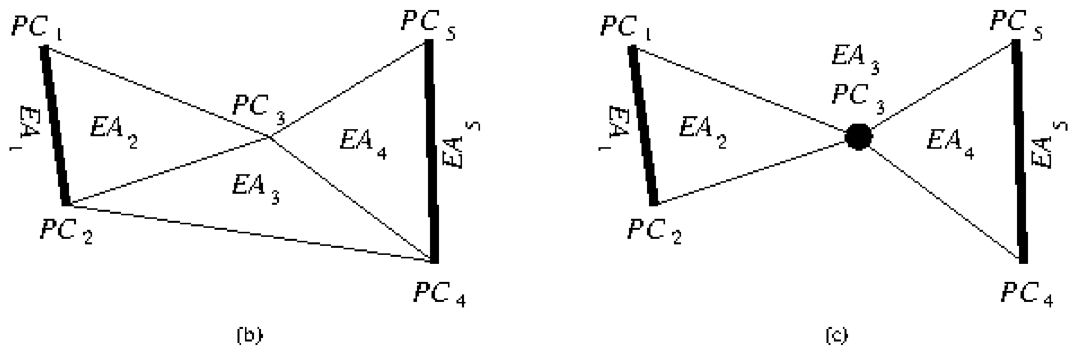

Figure A1. Left, the simplicial complex corresponding to case (b) of Figure (2). Right, the simplicial complex corresponding to case (c) of Figure (2). Segments representing 1-dimensional simplices are thicker than segments representing the faces of 2-dimensional simplices (triangles).

If the connections between categories of actions and categories of results are all one-to-one as in case (a) of Figure (2), then the simplices are isolated points so no simplicial complex exists. In this case, complexity is zero. By contrast, in case (b) and (c) the simplicial complexes illustrated in Figure (A1) yield a complexity greater than zero.

Two simplices are connected if they have at least one common vertex. Two simplices that have no common vertex may nonetheless be connected by a chain of simplices having common vertices with one another. Let us say that simplices EA_i and EA_j are $q-connected$ if there exists a chain of simplices $EA_u, EA_v, \dots EA_w$ such that $q = min\{l_{i,u}, l_{u,v}, \dots l_{w,j}\} \geq 0$, where $l_{x,y}$ is the dimension of the common face between EA_x and EA_y. In particular, two contiguous simplices are connected at level q if they have a common face of dimension q.

Let us consider the common faces between simplices and let us focus on the face of largest dimension. Let Q denote the dimension of this face. It is necessarily $Q \leq n-1$, where $Q = n-1$ means that there are at least two overlapping simplices that include all possible vertices.

Let us partition the set of simplices that compose the simplicial complex according to their connection level q. In general, for $\forall q$ there exist several classes of simplices such that the simplices belonging to a class are connected at level q. Let us introduce a *structure vector* **s** whose q-th component s_q denotes the number of disjoint classes of simplices that are connected at level q. Since $q = 0, 1, \dots Q$, vector **s** has $Q+1$ rows.

Let us define *Complexity* as:

$$C = \begin{cases} 0 & if\ all\ connections\ are\ one-to-one \\ \sum_{q=0}^{Q} \frac{q+1}{s_q} & otherwise \end{cases}$$

where the sum extends only to all terms such that $s_q \neq 0$. The complexity of two or more disconnected simplicial complexes is the sum of their complexities.

This expression takes account of two opposite effects. On the one hand, the numerator increases with the number of connections between Evoked Alternatives and Perceived Consequences. Thus, it simply measures the extent to which novel connections confuse the cognitive map. On the other hand, the denominator makes complexity decrease to the extent that cross-connections are clustered in distinct groups.

In case (b) of Figures (2) and (A1) there exists one single class of simplices connected at level $q = 0$ and one single class of simplices connected at level $q = 1$, hence $C = \frac{(0+1)}{1} + \frac{(1+1)}{1} = 3$. By contrast, in case (c) there exists one class of simplices connected at level $q = 0$ but two classes of simplices connected at level $q = 1$, hence $C = \frac{(0+1)}{1} + \frac{(1+1)}{2} = 2$.

Appendix B: The Series of Cognitive Maps

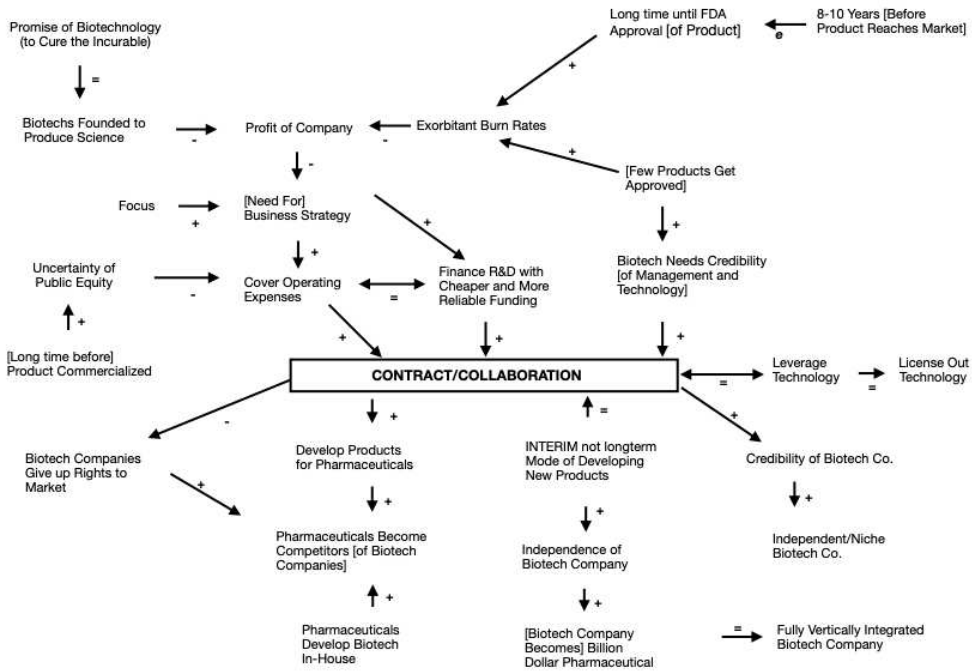

Figure B1. Biotech companies' cognitive map, 1986 (James, 1996).

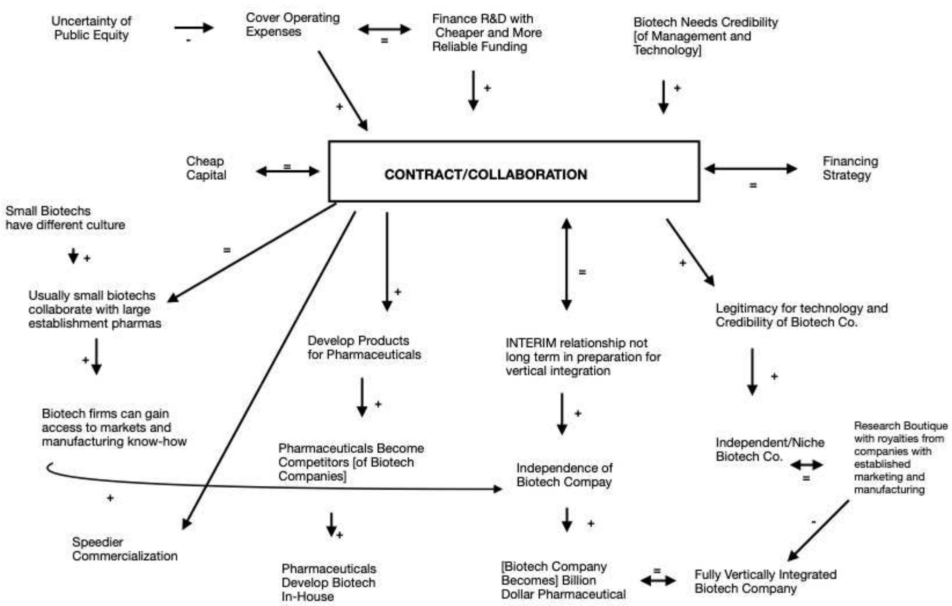

Figure B2. Biotech companies' cognitive map, 1987 (James, 1996).

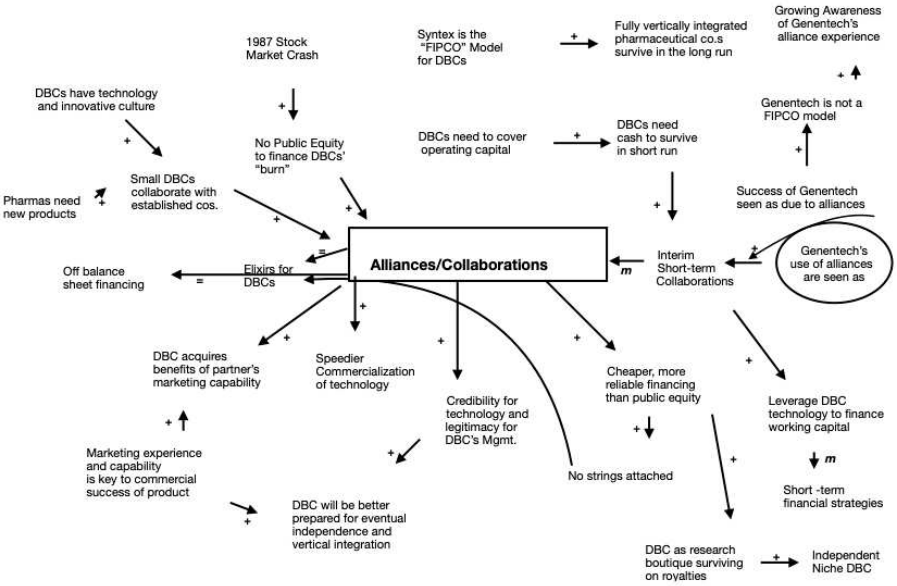

Figure B3. Biotech companies' cognitive map, 1988 (James, 1996).

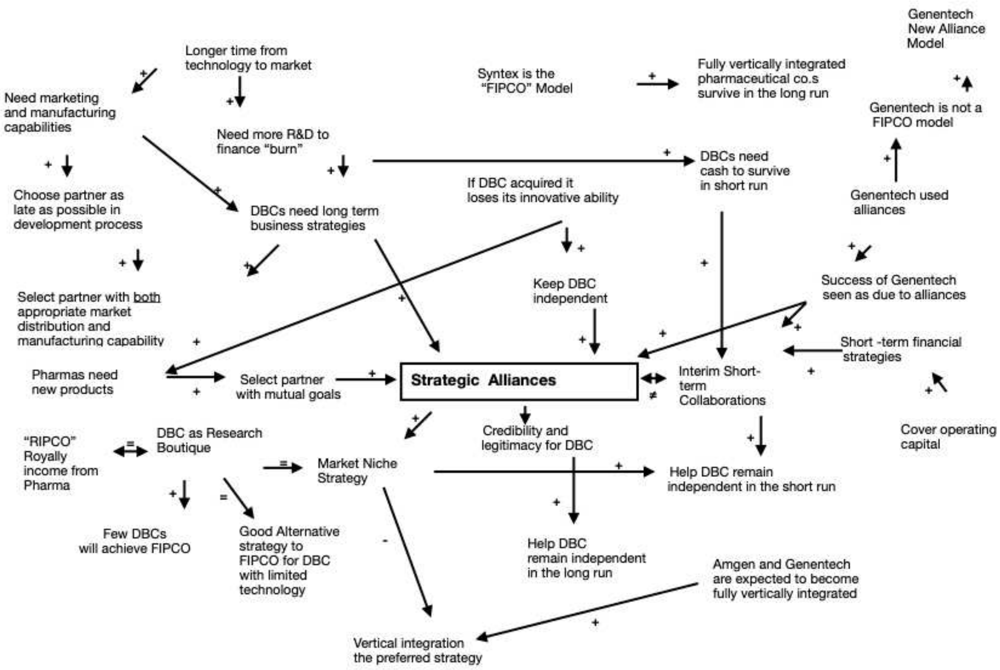

Figure B4. Biotech companies' cognitive map, 1989 (James, 1996).

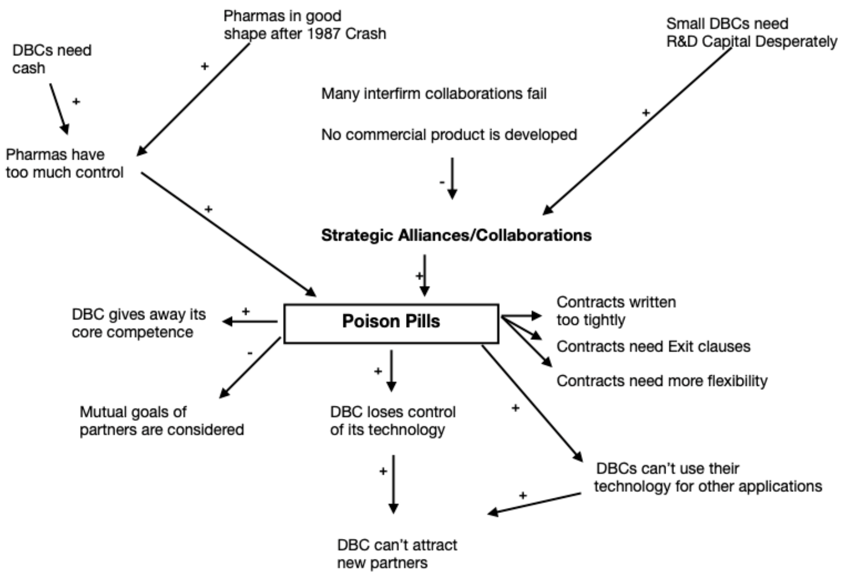

Figure B5. Biotech companies' cognitive map, 1990 (James, 1996).

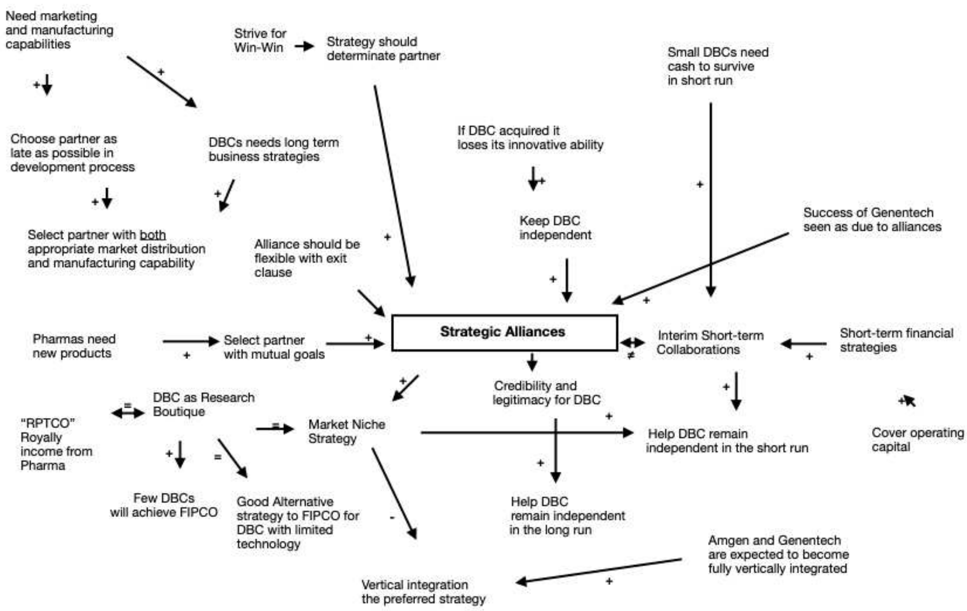

Figure B6. Biotech companies' cognitive map, 1991 (James, 1996).

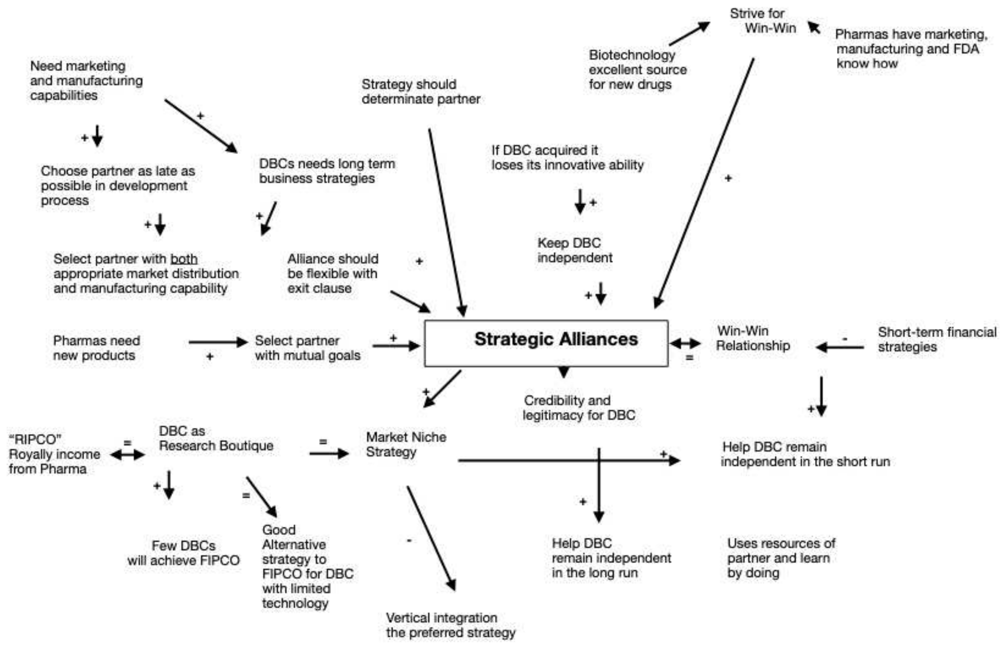

Figure B7. Biotech companies' cognitive map, 1992 (James, 1996).

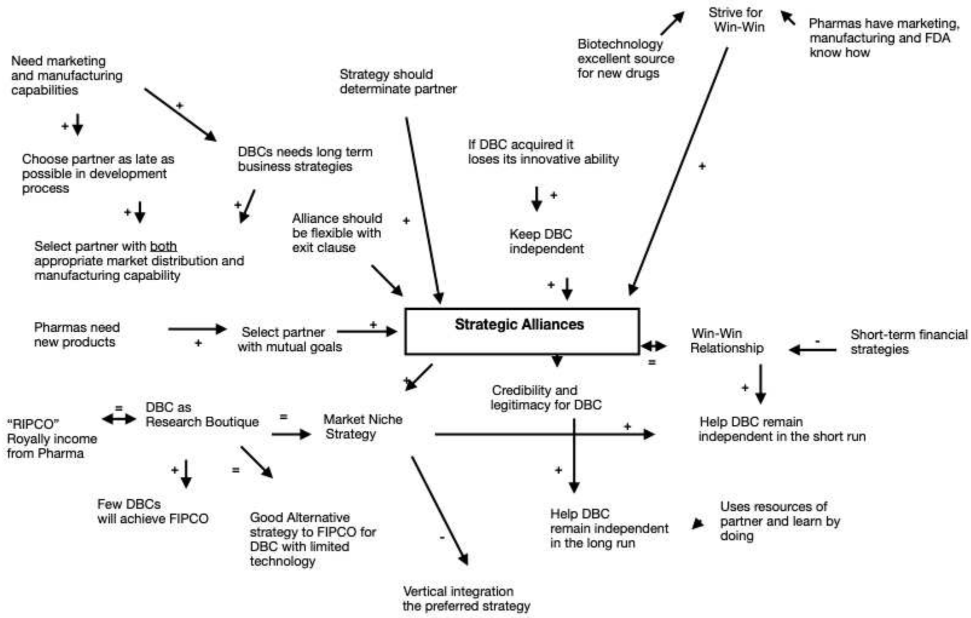

Figure B8. Biotech companies' cognitive map, 1993 (James, 1996).

References

Alipour, M., Hafezi, R., Amer, M. and A.N. Akhavan, 2017, "A new hybrid fuzzy cognitive map-based scenario planning approach for Iran's oil production pathways in the post-sanction period". Energy, 135, 851-864.

Amer, M., Jetter, A. and T. Daim, 2011, "Development of fuzzy cognitive map (FCM)-based scenarios for wind energy". International Journal for Energy Sector Management, 5, 564-584.

Axelrod, R., 1976, Structure of Decision. Princeton: Princeton University Press.

Brank, J., Leban, G. and Grobelnik, M., 2017, Annotating Documents with Relevant Wikipedia Concepts. *Proceedings of the Slovenian Conference on Data Mining and Data Warehouses* (SiKDD 2017), Ljubljana, Slovenia, 9 October 2017.

Chermack, T.J., 2004, "Improving decision-making with scenario planning". Futures, 295-309.

Clark, K.C, and Manning, C.D., 2016, Deep Reinforcement Learning for Mention-Ranking Coreference Models. Working Paper arXiv:1609.08667.

Davidson, P., 1991, "Is probability theory relevant for uncertainty? A post-keynesian perspective". Journal of Economic Perspectives, 5, 129-143.

De Finetti, B., 1931, "Sul significato soggettivo della probabilità". Fundamenta Mathematicae, 17: 298-329.

Dequech, D., 2001, "Uncertainty: Individuals, institutions and technology". Cambridge Journal of Economics, 28, 365-378.

Devlin, J., Chang, M.W., Lee, K. and Toutanova, K., 2019, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. Working Paper arXiv:1810.04805.

Dunn, S.P., 2001, "Bounded rationality is not fundamental uncertainty: A post-keynesian perspective". Journal of Post Keynesian Economics, 23, 567-587.

Eliaz, K. and E.A. Ok, 2006, "Indifference or indecisiveness? Choice-theoretic foundations of incomplete preferences". Games and Economic Behavior, 56, 61-86.

Ellsberg, D., 1961, "Risk, ambiguity and the Savage axioms". The Quarterly Journal of Economics, 75: 643-669.

Erdmann, D., Sichel, B. and L. Yeung, 2015, "Overcoming obstacles to effective scenario planning". Mc Kinsey & Company, Strategy & Corporate Finance Insights, Jun.

Faulkner, P., Feduzi, A. and J. Runde, 2017, "Unknowns, black swans, and the risk / uncertainty distinction". Cambridge Journal of Economics, 41, 1279-1302.

Galaabaatar, T. and E. Karni, 2013, "Subjective expected utility with incomplete preferences". Econometrica, 81, 255-284.

Gilboa, I., 1987, "Expected utility with purely subjective non-additive probabilities". Journal of Mathematical Economics, 16: 65-88.

Goodier, C., Austin, S., Soetanto, R. and A. Dainty, 2010, "Causal mapping and scenario building with multiple organizations". Futures, 42, 219-229.

Goodwin, P. and G. Wright, 2001, "Enhancing strategy evaluation in scenario planning: A role for decision analysis". Journal of Management Studies, 38, 1-16.

James, G.E., 1996, Strategic Alliances: Elixirs, Poison Pills or Virtual Integration? A Longitudinal Exploration of Industry-Level Learning in Biotech. University of Colorado at Boulder, Graduate School of Business, PhD Thesis.

Jetter, A. and W. Schweinfort, 2011, "Building scenarios with fuzzy cognitive maps: An exploratory study of solar energy". Futures, 43, 52-66.

Knight, F.H., 1921, Risk, Uncertainty and Profit. Boston: Houghton Mifflin Company.

Lane, D.A. and R.R. Maxfield, 2005, "Ontological uncertainty and innovation". Journal of Evolutionary Economics, 15, 3-50.

March, J.G. and H.A. Simon, 1958, Organizations. New York: John Wiley & Sons.

Ok, E.A., Ortoleva, P. and G. Riella, 2012, "Incomplete preferences under uncertainty: Indecisiveness in beliefs versus tastes". Econometrica, 80, 1791-1808.

Pathak, D., Agrawal, P., Efros, A.A. and Darrell, T., 2017, Curiosity-Driven Exploration by Self-Supervised Prediction. Proceedings of the 34th International Conference on Machine Learning. Sydney, Australia.

Ramírez, R., Österman, R., and D. Grönquist, 2013, "Scenarios and early warnings as dynamic capabilities to frame managerial attention". Technological Forecasting & Social Change, 80, 825-838.

Ramírez, R. and C. Selin, 2014, "Plausibility and probability in scenario planning". Foresight, 16, 54-74.

Ramsey, F.P., 1931, "Truth and probability". In Braithwaite, R.D. (Ed.), The foundations of mathematics and other logical essays. London: Routledge & Kegan Paul, pp. 156-198.

Roxburgh, C., 2009, "The use and abuse of scenarios". Mc Kinsey & Company, Strategy & Corporate Finance Insights, Nov.

Ruder, S., 2022, "NLP-progress". Retrievable at http://nlpprogress.com. Accessed on Jan 8th, 2022.

Runde, J., 1990, "Keynesian Uncertainty and the Weight of Arguments". Economics and Philosophy, 6, 275-292.

Schmeidler, D., 1989, “Subjective probability and expected utility without additivity”. Econometrica, 57: 571-587.

Schoemaker, P.J.H., 1995, “Scenario Planning: A tool for strategic thinking”. Sloan Management Review, 36: 25-40.

Sigismind Huff, A., 1990, Mapping Strategic Thought. Chichester: John Wiley & Sons.
Sigismund Huff, A., Huff, J.O. and P.S. Barr, 2000, When Firms Change Direction. Oxford: Oxford University Press.

Sigismund Huff, A. and Mark Jenkins, 2002, Mapping Strategic Knowledge. London: Sage Publications.

Stanford NLP Group, 2022, “SQuAD 2.0: The Stanford Question Answering Dataset”. Retrievable at https://rajpurkar.github.io/SQuAD-explorer. Accessed on Jan 8th, 2022.

Van der Heijden, K., 2000, “Scenarios and forecasting: Two perspectives”. Technological Forecsatng & Social Change, 65: 31-36.

Wilson, I., 2000, “From scenario thinking to strategic action”. Technological Forecasting and Social Change, 65, 23-29.

Wright, G. and P. Goodwin, 2009, “Decision making and planning under low levels of predictability: Enhancing the scenario method”. International Journal of Forecasting, 25, 813-825.

Zhang, Z., Zhuang, F., Zhu, H., Shi, Z., Xiong, H., and He, Q., 2020, “Relational Graph Neural Network with Hierarchical Attention for Knowledge Graph Completion”. Proceedings of the AAAI Conference on Artificial Intelligence, 34, 9612-9619. Retrievable at https://doi.org/10.1609/aaai.v34i05.6508.