The Nuclear Power Plant: Our New “Tower of Babel”?

Julie Jebeile

► To cite this version:

Julie Jebeile. The Nuclear Power Plant: Our New “Tower of Babel”?. Business Ethics and Risk Management. Ethical Economy (Studies in Economic Ethics and Philosophy), pp.129-143, 2014. hal-02105521

HAL Id: hal-02105521

<https://hal.science/hal-02105521>

Submitted on 23 Apr 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Nuclear Power Plant: Our New “Tower of Babel”?

About the Author

Julie Jebeile is a PhD student in philosophy of science at IHPST, Paris. In her work, she questions the explanatory value of scientific models that are used both as representations of natural and social systems, and as inferential tools. In particular, she examines how the intensive use of computers in modeling may impact the justification of model results and our search for understanding of the modeled phenomena. Julie also has an interest in the epistemological issues that arise with the distribution of knowledge within epistemic communities, particularly in the research and design offices of civil nuclear engineering where she worked as a neutronics engineer before turning to the philosophy of science.

Short abstract

On July 5, 2012 the Investigation Committee on the Accident at the Fukushima Nuclear Power Stations of the Tokyo Electric Power Company (TEPCO) issued a final, damning report. Its conclusions show that the human group - constituted by the employees of TEPCO and the control organism - had partial and imperfect epistemic control on the nuclear power plant and its environment. They also testify to a group inertia in decision-making and action. Could it have been otherwise? Is not a *collective of human beings*, even prepared in the best way against nuclear risk, *de facto* prone to epistemic imperfection and a kind of inertia?

In this article, I focus on the group of engineers who, in research and design offices, design nuclear power plants and model possible nuclear accidents in order to calculate the probability of their occurrence, predict their consequences, and determine the appropriate countermeasures against them. I argue that this group is prone to epistemic imperfection, even when it is highly prepared for adverse nuclear events.

1 Introduction

On July 5, 2012 the Investigation Committee on the Accident at the Fukushima Nuclear Power Stations of Tokyo Electric Power Company (TEPCO) issued its final, damning report¹. Apart from the precise facts for which TEPCO, the Nuclear and Industrial Safety Agency, and the Japanese government have respectively been reproached, one can draw from this report at least two worrying conclusions.

¹ The English version of this report is on the website <http://naiic.go.jp/en/>. An interim report was delivered on December 26, 2011. See <http://icanps.go.jp/eng/interim-report.html>.

The first is that the accident at Fukushima Daiichi on March 11, 2011, the date on which seism and tsunami raged on the nuclear site, was foreseeable and thus could have been avoided. According to the report, since 2006, both the operator and the nuclear safety agency knew the risks of a total power outage and the loss of seawater pumps in the case of a tsunami reaching the site level. However, studies and concrete measures of these risks were never undertaken. One of the main arguments used to justify this omission was that the probability of a tsunami of such size occurring was very low.

The second conclusion is just as alarming as the first one: a better response to the accident could have been given and thus the effects of the accident could have been lessened. As the report notably mentions, there was a lack of knowledge and preparation from the personnel on the site, and an absence of clarity and rapidity in the decisions made by the direction. The report also shows that the communication to the public and, in particular, to the residents close to the site (and therefore the most affected by the high levels of radioactive contamination) was vague and, in many cases, too late.

Both conclusions show that the human group - constituted by the employees of TEPCO and the control organism - had partial and imperfect epistemic control on the nuclear power plant and its environment. They also testify to a group inertia in decision-making and action. Could it have been otherwise? Is not a *collective of human beings*, even prepared in the best way against the nuclear risk, *de facto* prone to epistemic imperfection and a kind of inertia?

In this article, I focus on the group of engineers who, in research and design offices, design nuclear power plants and model possible nuclear accidents in order to calculate the probability of their occurrence, predict their consequences, and determine the appropriate countermeasures against them. I argue that this group is prone to epistemic imperfection, even when it is highly prepared for adverse nuclear events. I consider that a group is epistemically imperfect when it does not prevent individual errors from propagating and when it does not have optimal epistemic control on the nuclear machine. I show that difficulties stand in the way of its optimal epistemic control on such a machine. These difficulties might have consequences in case of exceptional issues. Under normal circumstances, one can consider that there are sufficient precautionary systems so that nuclear power plants run correctly².

My work falls within an epistemology of human groups facing nuclear risk. The method I have adopted is not common and might surprise the philosopher insofar as my reflections build on work in the science of sociology. Indeed, the question of the efficiency of human organizations facing nuclear risk has been extensively discussed in the sociology of organizations. In this article, I take stock of the debate and examine the results of

² My intention is not to give the impression that nuclear power plants are always unsafe.

previous studies before presenting my own philosophical criticism. This allows me both to highlight an important scientific debate, and to situate my work within the existing scientific framework.

The political component of my work, if there is one, is to draw decisional bodies' attention to the following point: the organization of human resources, based on a division of labor and distribution of knowledge, constitutes a risk factor. This risk factor must be seriously taken into account when making decisions concerning future energy steps.

2 Normal Accidents and High Reliability Organizations

Can a human organization, implemented within a nuclear power plant, adequately face a nuclear risk? This question has been addressed by sociologists of organizations. More generally, sociologists have examined the efficiency of human organizations in high-risk enterprises, e.g. chemical plants, air traffic control stations, ships, dams, nuclear weapons, and space missions. This topic divides sociologists into two rival schools of thought. On one side are the proponents of the *Theory of Normal Accidents*, who think that nuclear power plants are hazardous enterprises which inevitably lead to accidents. According to these theorists, nuclear accidents are to be expected. They argue that better human organization would allow one to reduce the accident risk but not to remove it completely (Perrow, 1984; Sagan, 1993). On the other, the proponents of the *Theory of High Reliability Organizations* (HROs) claim, on the contrary, that, under specific conditions, enterprises, such as nuclear power plants, display organizational specificities which explain their very high level of reliability (Rochlin, Roberts and La Porte, 1987; Rochlin, 1993).

In what follows, I first present the theory of normal accidents. I then present the conditions under which an organization is said to be highly reliable by the theory of high reliability organizations. In section 3, I show that, even under these conditions, a group is prone to epistemic imperfection. In this way, I propose arguments in favor of the theory of normal accidents. Nevertheless, I consider that the theory of high reliability organizations is correct when it claims that organizational improvements must be made in order to increase security. However, I claim that the organizational improvements of the theory are not optimal and need to be revised.

2.1 Theory of Normal Accidents

Charles Perrow is one of the leading figures of the theory of normal accidents. He argues that, in high-risk technologies, “no matter how effective conventional safety devices are, there is a form of accident that is inevitable” (Perrow 1984, p. 3). Unavoidable failures, either mechanical or human, can occur anytime during the processes of design or

operation. They can also be induced by the system's external environment³. Accidents in high-risk technologies are qualified as normal not because they are frequent (fortunately, they are relatively rare) but because they are unavoidable: their occurrence is an inherent property of high-risk enterprises. In the same vein, Perrow writes: "It is normal for us to die, but we only do it once" (1984, p. 5).

According to Perrow, high-risk systems commonly present two particular properties, which are hard to get around, and which favor the occurrence of an accident. On the one hand, they display a complex interactivity which complicates the functioning of the system and makes it difficult to understand; while on the other hand, the tight coupling between their components prevents a quick recovery to a normal state of functioning in the case of an accidental situation.

A system displays a complex interactivity when its components serve several functions at the same time. The complex interactions are defined by Perrow as "those of unfamiliar sequences, or unplanned and unexpected sequences, and either not visible or not immediately comprehensible" (1983, p. 78). They are opposed to what he calls linear interactions, which can be predicted and understood and, therefore, can be controlled. Take, for example, a pressurized water reactor⁴. In such a reactor the steam generator carries out two roles at once. It is used not only to absorb heat given off by nuclear reactions in the core, but also to produce steam that is necessary for the rotation of the turbines. In the case of a steam generator breakdown, the nuclear core overheats, as the excess of heat in the primary cooling loop is not properly absorbed anymore. There is a risk of reaching the critical point of heat flux, which will then destroy the fuel cladding, and thereby create a breach of the first containment barrier, causing the release of fuel material into the reactor core. Further, if the steam generator malfunctions, the secondary system can become too cool for producing enough steam. The steam generator has a functional relation with two distinct components - the reactor core and the turbines. If it fails, its functions, namely cooling the reactor and producing steam, cannot be ensured. During the accident at Three Mile Island, a breakdown of the main feedwater pumps in the secondary loop was responsible for a failure in the steam generator. After the reactor shut down, it continually generated decay heat. As the auxiliary secondary pumps also failed, the core was not correctly cooled. Its temperature and pressure increased. In addition to that, small failures in the safety system occurred. Thus, the pressurizer relief valve of the primary loop, which opened up automatically in order to avoid a drastic increase in pressure, could not be closed in the sequel. And yet, the warning light

³ The different stages during which failures can occur are called the the DEPOSE components (*Design, Equipment, Procedures, Operators, Supplies and materials, and Environment*) by Perrow (1984, p. 8).

⁴ Perrow takes instead the example of a chemical reactor in his book. But the heat process is similar to nuclear power plants, and the conclusions remain the same (1984, p. 72).

indicated the opposite to the operators. Hence, a sequence of unfortunate events and unexpected interactions led to the well-known catastrophe.

When the components of a system are tightly coupled, a small failure in one of them can quickly induce a series of problems in the other components. For example, an unforeseen increase in reactivity in the nuclear core has direct and immediate consequences on the other components with which the core is in relation. Thus recovery of the system to its normal state of functioning is more difficult. Additionally, because of the propensity of a failure to spread into the whole system, the operator cannot take the time to intervene; it is constantly in a hurry to solve the small failures. Lastly, a system that presents a tight coupling between its components requires a fine maneuverability. Thus, if the operator wants to lower the level of reactivity in the core, he can intervene, for example, by incorporating boric acid, a neutron poison, into the core. The amount of boric acid to be poured must be measured with accuracy. Because of their time-dependence and sensitivity, these systems do not allow for improvisation in the case of a problem.

Would it be sufficient to simplify the system and to avoid any tight coupling between components in order to make the system safer? In addition to the technical difficulties it would entail (see Perrow, 1984, p. 89), this solution seems to be quite counterproductive. Indeed, if we are to believe Perrow, the complex interactivity and the tight coupling of high-risk systems are the very factors that make these systems efficient. Thus, for example, the double function of the steam generator guarantees a better energy yield of the nuclear power plant. More generally, according to Perrow, “complex systems are more efficient [...] because there is less slack, less underutilized space, less tolerance of low-quality performance, and more multifunction components. [From the point of view of] design and hardware efficiency, complexity is desirable” (Perrow 1984, p. 88).

Because of their complex interactivity and tight coupling of their components, nuclear power plants are prone to accidents. Thus according to Perrow, the accident at Three Mile Island was not unexpected. He writes, concerning it, that: “[N]o matter how effective conventional safety devices are, there is a form of accident that is inevitable. This is not good news for systems that have high catastrophic potential, such as nuclear power plants, nuclear weapons systems, recombinant DNA production, or even ships carrying highly toxic or explosive cargoes. It suggests, for example, that the probability of a nuclear plant meltdown with dispersion of radioactive materials to the atmosphere is not a chance in a million a year, but more like one chance in the next decade.” (Perrow 1984, p. 4). This was a dreadful prophecy which, during last few decades, has gained credibility: at the time when Perrow wrote⁵, in 1984, only Three Mile Island (on March 28, 1979) was classed as an accident resulting in a nuclear meltdown. Since then, two

⁵ Note that Perrow has written an article in reaction to the accident in Fukushima (Perrow, 2011).

nuclear accidents of the same type have followed: Chernobyl (on April 26, 1986; six years after Three Mile Island) and Fukushima Daiichi (on March 11, 2011; twenty four years after Chernobyl).

I agree with Perrow that nuclear power plants are prone to accidents because of their complexity. Furthermore he is right when he says that they are far too complex for one to be able to anticipate, analyze and face an accident in the best way. And, as I will show in section 4, the division of labor and the distribution of knowledge, highly required by their complexity, are also sources of risk in that they hinder optimal epistemic control of the nuclear machine.

I now want to present the theory of high reliability organizations. This theory proposes four organizational improvements that have to be made in nuclear power industry. While these improvements seem appropriate in principle, as I will show in section 3, they are not efficient enough in practice.

2.2 Theory of High Reliability Organizations

The theory of high reliability organizations relies on an inference that is the inverse of the one that the normal accidents theorists make concerning the relatively low frequency of technological catastrophes. Although these catastrophes are not very frequent, for the normal accident theorists, the very fact that these accidents even arise is a sign of their normality. For the high reliability theorists, the low frequency of technological catastrophes proves that there is exceptional organization behind high-risk technologies. Roughly speaking, the disagreement between the high reliability theorists and the normal accident theorists is of the same kind as the disagreement between people who see the glass half full and those who see the glass half empty.

Among the high reliability theorists, there are three groups of sociologists whose working assumptions are slightly different (see Sagan 1993 for a presentation). Despite their methodological differences, they all agree that the dangers of high-risk technologies may be avoided by human organizations if they are designed and managed in an appropriate manner. These organizations are called high reliability organizations. According to Sagan, “the common assumption of the high reliability theorists is not a naive belief in the ability of *human beings* to behave with perfect rationality; it is the much more plausible belief that *organizations*, properly designed and managed, can compensate for well-known human frailties and can therefore be significantly more rational and effective than individuals” (1993, p. 16). Furthermore, high reliability theorists think organizations can avoid the intrinsic danger of high-risk enterprises.

The high reliability theorists agree in their identification of four factors that must be improved in order to ensure the safety of dangerous technologies. At first sight, it seems that the improvement of these factors results in better security. But, as I will show, this is

not enough. And moreover, some of these factors may contribute to epistemic imperfection of the group. They are listed by Scott Sagan (1993) and I briefly take them up here before examining some of them in section 3.

First, according to the high reliability theorists, safety and reliability must be considered by the decision-making bodies, namely the political elites and the organization's leadership, as priorities of the first order. It is only in this way that all the necessary financial means will be utilized with the aim of improving the level of safety in the enterprise. Additionally, in promoting safety and reliability above all, the decision-making entities are able to send a clear message to everyone, thus creating a climate of trust within the group.

Second, redundancy, either technical or human, must absolutely be optimized. Redundancy is the multiplication of independent channels of communication, decision-making, and implementation. Thus an employee can be led to check the tasks of a colleague. Redundancy aims, in principle, to compensate for possible human errors - including mistakes made due to irrational behavior or cognitive limitations.

Third, one must reduce individual errors by establishing a policy of decentralization of decision-making authority, a culture of reliability, and constant training of the personnel. The decentralization of decision-making authority must allow the individuals that are closest to the concrete problems to quickly and appropriately respond to dangers. The culture of reliability is established with precise rules and procedures given to the personnel. This culture allows employees that are far from the decision-making centers to act in compliance with what is expected from them by the authority. Lastly, a constant training of the personnel, punctuated by on-the-job formations and simulations of emergency crisis, must break their routine, which is often favorable to slackening and negligence, and put employees in an unusual situation, thereby forcing them to remain awake and reactive in their decisional and operational powers.

Fourth, the organization must show a strong capability to learn by accelerating the adjustment of its procedures in the case of an emergency. Agents have to learn through constant process of trial and error, which allows for retention of the most efficient operations and an abandon of those that are not performing as well. This process can be performed experimentally (e.g. with crisis exercises) or by working with simulation and imagination.

In what follows, I will show that there are difficulties of an epistemic kind that arise within high reliability organizations, when these four precautions are taken. In other

words, these four precautions are not enough to prevent accidents and that is where my view differs from the theory of high reliability organizations⁶.

3 Epistemic Imperfection within High Reliability Organizations

The four precautions advocated by the high reliability theorists are conceived mainly to apply to human organizations dedicated to the operation of nuclear power plants. I assume in this work that they should also apply to organizations in charge of their design. In order to face nuclear risk it is necessary to have good organization on the site, but it is also necessary to have it in the design and research offices in order to correctly predict and prevent possible accidents. Thus, the precaution of redundancy, the decentralization of decision-making authority, and the culture of reliability should, in principle, help engineers to properly perform studies of nuclear power plants, and therefore to gain epistemic control on the nuclear machine.

However, as I would like to show, it is not enough to distribute the tasks, to increase the number of controls, to decentralize the decision-making center, or to create a culture of reliability in order to improve the epistemic control of the machine. Furthermore, as I argue, although these practices seem in principle to be genuine improvements, some of them may, in implementation, cause a loss of optimal epistemic control on the machine. I will show that, when one applies these practices, they sometimes have undesirable effects. Some of these effects are of psychological kind but have real epistemic impacts. In particular, they may prevent the group from catching individual errors. In some cases they also contribute to the isolation of individuals within their specialized work, and to a loss of know-how that the group accumulates over time. My intent is to proceed in the way in which Miranda Fricker does in her work: I propose an epistemological analysis of some elements, which have a psychological foundation, but can be given an epistemic interpretation. In addition to the psychological effects, there are economical constraints that are specific to any profit-making enterprise. These constraints sometimes have (harmful) epistemic consequences on the daily work of the engineers. However, the epistemic consequences of the psychological effects and the economical constraints are not systematic. In some cases they can be dangerous.

Before going further, I should draw to the attention of the reader the fact that

⁶ Normal accidents theorists also address these factors (see Sagan p. 36-43 for a summary of their arguments).

organizational rules are generally documented internally within the enterprises in charge of designing the nuclear power plants; they are not publicized. Nevertheless, in France, where nuclear power represents approximately three quarters of the electric production, these rules are partly determined by the Order of August 10, 1984 related to the quality of the design, construction and operation of basic nuclear installations. This order has been abrogated and replaced by the Order of February 7, 2012.

3.1 Unavoidable Individual Errors

According to Sagan's analysis, the high reliability theorists consider well-designed and well-managed organizations to be potentially more rational than individuals can be. Consequently, in their view, the collective is able to overcome errors potentially made by individuals. At first sight, that seems to be a fair judgment because the division of labor, the specialization of workers and the systematic check of studies seem to serve this purpose.

For example, the division of labor is obvious as a rational choice of organization. It seems preferable to subdivide the work into different tasks and to attribute each of them to the most competent person rather than to give all of them to the whole group. Further, one expects from the verification of a study, performed by a different person than its author, to increase the chance of removing residual errors of judgment or calculation. If two people, independently of each other, carry out reasoning, there ought to be a low probability that they both make precisely the same error. Further, because a "checker" should pay close attention to possible errors, he thus has a good chance of finding them.

However, in the context of civil nuclear industry, the organizational properties of the collective display retroactive effects on the individual which lead to error.

First of all, in design and research offices, the tasks assigned to engineers – in neutronics, thermohydraulics, metallurgy or mechanics – are often precise, repetitive and procedural. Even if they require a high level of qualification, their ultra-specificity can weaken the interest they have for their work, and their concentration, and thus can lead them to make mistakes.

Further, in order to detect possible individual errors, organizational rules compel that each study be performed, controlled, verified, and endorsed by four distinct agents: the author, the technical referent, the checker, and the hierarchical manager. Therefore the rules respond to the search of redundancy, decentralization of decision-making authority, and the culture of reliability. However, this produces two negative side effects: on the one hand, it minimizes the responsibility of each individual; on the other hand, it increases the trust one puts in other members of the collective.

In addition to this second effect, there is also the phenomenon of *epistemic injustice* highlighted by Fricker (1998). That is, individuals give credit to the most powerful, and therefore, not necessarily to the best. In the industrial context, it is not always easy to identify the most reliable informants. A great number of personal variables (such as the age of one's coworkers - which is not necessarily proportional to their level of competence) can exert influence. Generally, in a design and research office, the most powerful people are the hierarchical managers, or the employees whose self-confidence produces a *de facto* form of authority. As for the best, those are the employees whose competence is likely to bring reliable and efficient solutions to technical issues that are addressed to them.

The trust in other members is increased when one has no time to deliberate by oneself. There are indeed budgetary constraints that concern the authors of studies, but these also affect checkers and technical referents (hence, let us note in passing, the interest remains to grant a special budget to young recruits). Likewise, one can easily figure out that a hierarchical manager has little time to endorse a study performed by one member of her team and sign it. They are obliged to trust their quick judgment that cannot concern only the quality of the study. This trust, which is sometimes distorted, does not favor the catching errors.

Lastly, the nuclear risk is a less concrete notion for engineers in offices than for operators in nuclear power plants, for example. If the employees tend to forget that they actually work for a high-risk enterprise, they might relax their vigilance.

3.2 The scientific isolation of individuals

We have seen that the division of labor, the specialization of workers and the systematic verification of studies does not allow for the complete avoidance of individual errors. Further, these things also prevent optimal epistemic control. Optimal epistemic control requires not only reliable information about the machine and correct predictions about its possible behaviors, but also a systemic understanding of the machine, i.e. an understanding of the whole system down to the smallest but relevant details⁷.

Optimal epistemic control cannot be met in the nuclear industry because individuals are isolated in their highly specialized work. As Perrow notes (1983, p.87), the specificity of their task, inherent to the complexity of the machine, has the effect of depriving them of a global understanding of the system. Besides, the lack of intellectual interaction between

⁷ According to a strong criterion of optimality, the group must prevent any individual errors from propagating. But this criterion is idealistic and many errors have no major impact on the accuracy of the predictions made by the group. A looser criterion can be that the group must stop any individual error that significantly impacts the predictions made by the group. But the importance of an error on the predictions can hardly be assessed *a priori*. Nevertheless we shall consider the looser criterion in this paper.

the individuals hinders the dynamics that is necessary for optimal epistemic control of the machine. The distribution of tasks reduces the possibility of interaction between individuals with different functions and specialties. Individuals generally restrict themselves to the tasks that have been assigned to them.

Furthermore, as they are subjected to budgetary constraints, the agents generally have no time to fill their gaps in the technical domains adjacent to theirs. And yet, for example, if there is a lack of interaction between the design engineers and the developers of computer programs required for the studies, an inappropriate use of these programs can be the result. Errors can occur in studies, and, because of the fluctuating confidence between workers that I have previously mentioned, they can be whisked off the control of checkers. The major problem of such a collaborative organization is that everybody thinks that if she does not know something, someone else in the firm surely does.

Let us take the example of orders made between two teams of different disciplines, two distinct epistemic communities. Engineers in thermohydraulics sometimes need neutronics inputs for their own studies. They might, for example, need to know about the nuclear reactivity in the core when doing calculations of critical heat flux. The neutronics engineer, who has to provide them with such data, often has little knowledge of thermohydraulics, and therefore has to trust his interlocutors, the thermohydraulics engineers. She will not check for her colleagues whether the question they raise is coherent or justified, or whether the input data she gets from them are plausible.

3.3 The non-conservation of accumulated know-how

Epistemic control is also in jeopardy because nothing seems to guarantee the conservation of accumulated know-how. This know-how includes not only the set of knowledge required for the achievement of the studies of accidents and design, but also the set of information about the history of the nuclear power plant in question (like, for example, the technological choices at the base of its construction). Know-how differs from mere propositional knowledge in that it includes all the “recipes” that an employee learns by doing and that are useful in performing studies, solving trivial computational issues, finding relevant information, and so on. In order to control power plants during their life-span, one often needs to come back to earlier technical studies. Thus, the organizational rules require that, for each study, a technical file that lays out the sequence of the study and the technical choices, be created and stored. In this sense they meet article 2.5.6 of the Order:

“The important activities for the protection, the technical controls, the actions of verification and assessment are the subjects of a documentation and a traceability allowing one to demonstrate *a priori* and to check *a posteriori* the respect of the defined requirements. The corresponding documents and recordings are recorded, easily accessible and readable, protected, preserved, and stored during an appropriate and justified term.”

However, these files can lack clarity, comprehensiveness, or even get lost. In these cases, one sometimes needs to directly contact the author of the file by every possible means, but it is necessary that the author still be available. Indeed, while a nuclear power plant has an average life-span of thirty years, a professional career at the same position, and in the same firm, is often much shorter. Further, people often communicate orally – both in person and by phone, and such communication is generally not stored, and thus is not systematically mentioned by authors in their study files.

This lack of continuity in the studies is aggravated by a lack of transmission: individuals often have no time to learn from each other. This is, I think, an aspect one must not underestimate. There might be a crisis in the nuclear sector at any time. Such was notably the case in the nineties in France. The nuclear industry did not, at the time, hire enough workers, and it pursued a policy of early retirement of its executives. Also, the age pyramid of the current personnel shows an ageing pattern. For several years, the nuclear industry has hired young, freshly graduated, engineers en masse, and these people are, by definition, inexperienced. The question arises - and I will keep it open: how is it possible to educate all these young arrivals when the most experienced will soon retire?

Further, there is the crucial problem of archive. Because of a comprehensible lack of physical or computer space, a great number of data and studies cannot be stored properly or are hard to access. This leads to a loss of “collective memory” in that the group loses the pieces of information it has accumulated over time. And yet, according to the article 2.1.1 of the Order, it is required that:

“The operator holds internally, in its subsidiaries, or in companies under its control [...], the technical competences to understand and appropriate in a perennial manner the foundations of these activities.”

In summary, the work of high reliability theorists shows that it is conceivable, and even highly desirable, to do improvements of an organizational kind so that better control of the group can be possible. However, three difficulties of an epistemic type remain. These prevent the collective from reaching optimal epistemic control of the nuclear machine. These three difficulties are: i) unavoidable individual error, ii) the scientific isolation of individuals and iii) the non-conservation of accumulated know-how over time. I will now show that the division of labor and the distribution of knowledge are at the origin of these epistemic difficulties, and therefore are responsible for the epistemic imperfection of the group. And yet the theory of high reliability organizations does not seem to recognize this.

4 The Tower of Babel

A first sight, the division of labor and the distribution of knowledge are compatible with the search for optimal epistemic control. The division of labor is absolutely indispensable in the operation of complex technologies, which requires various competences and a high level of specialization from the employees. The division of labor allows the group to save time in the areas of both education and work. It takes much longer for a single individual to develop multiple skills than for several individuals to specialize in only one of these competences. Further, a single person would take much more time to perform thirty intellectual tasks than thirty persons dedicated to only one operation.

However, the division of labor and the distribution of knowledge are at the root of the bursting of knowledge and systemic understanding of the machine within the group. Thus it is worth questioning whether the nuclear power industry is, to some extent, a real Tower of Babel. In the group, each member has a specific education and particular scientific competencies. Each member possesses specific pieces of knowledge about a part of the machine, but not about the whole, and thus cannot possess a full understanding of the plant. The operators can claim a certain understanding about how things work since they effectively manipulate the machine, they know the functional roles of its parts and the effects of their operations on them (see Dretske 1994). But this understanding, gained by the manipulation of the machine, does not really help to answer what-if-things-had-been-different questions. Besides, neither the design engineer (and which one? the engineer in thermohydraulics? in neutronics?), nor the constructor, possesses full understanding of the machine, so neither is able to anticipate all possible failures or accidents. Each member holds only partial knowledge of the nuclear power plant. The one and only cognitive unit, which could claim a full understanding, is the group itself. The group must therefore be considered as a distributed and extended cognition. It is a distributed cognition in that its memory and its knowledge are split into its different members. It is an extended cognition in that a part of its memory and knowledge is not possessed by human beings but rather placed on material support, such as paper or computers. However, the pieces of knowledge of the group do not exceed the sum of knowledge of its members. We cannot even claim that the group understands how the machine works in its entirety more than only one of its members. How, then, could it foresee any possible failure or accident?

In order for this to be possible the group must be more rational than its members. But, notably, because of the mutual trust between members of the same team, it seems to be difficult to eliminate individual errors. Furthermore, its parts must be perennial, so that the “collective” memory is conserved. In other words, knowledge must be preserved and be at its disposal at any time. And yet, nothing seems to guarantee the conservation of the accumulated know-how: traceability can be difficult (due to the quality of technical files,

and problems of storage); there is a difference in scales between the length of careers and the span-life of power plants, which also contributes to the problem; and lastly, because the parts of organizations are mortal agents who get sick, go on vacation, are fired, quit, retire or die, this too contributes to the problem of continuity. Further, if the members of the group do not interact with each other, overlap of competence and knowledge is not possible. Each member becomes indispensable for the survival of the whole and such a situation is not desirable.

Additionally, in order to reduce nuclear risk, the group must not only predict any failure, but also it must offer a specific forecast in case of a problem, in a sufficient time. It has to speak as one in order to indicate the actions to follow in case of emergency. And yet this is possible only if there is an orchestra conductor who gives instructions and decisions to the whole group. This conductor cannot be an individual. It can only be the group itself. But the group often has difficulty in expressing its own decisions, as I would like to show.

If each member has his own *opinion*, there will be a divergence in the group: in one word, “chaos”. And the nuclear power plant would be a genuine Tower of Babel. The decisions must therefore appear as unified results of the work of the whole group. In order for the group to divulge them in concert, at least one of its members, not necessarily the same one every time, must be able to explicitly express them so that its decisions can be publicly known, thus allowing the group to act. Thus, the constraint of a “final single voice” is absolutely mandatory for the good management of the group. It states that an individual, who is at the end of a long chain of informants (engineers, builders, designers, experts, etc.), must be able to express the group’s decisions. This last ‘link’ of the informational chain must know enough information to formulate them. She must at least possess the minimum amount of information required for expressing it.

However, the constraint of a “final single voice” cannot be met in practice. The inherent complexity of nuclear power plants requires from employees a high level of specialization. Thus, the employees understand only a part of the big machine through their own perspective, i.e. their specialty. Further, it is difficult to say if an approximate knowledge of the machine is enough to anticipate its possible failures. Most people roughly know how a nuclear power plant works. But, in most cases, they do not have knowledge of all of the details which are required to predict failures. These details are precisely hard to identify by a social organization. An example of missing details, in the case of the Fukushima Daiichi plant, is the height of sea walls. If there had been optimal epistemic control of the plan, the group of engineers in charge of its design would have thought that the plant, built on a seafont, could be exposed to a risk of seism or tsunami. The group would then have modified the height of the sea walls and thereby protected the diesel generators from a possible drowning.

5 Conclusion

Even the production of a simple bicycle requires diversity in specialized knowledge. As highlighted by André Gorz (1989, p. 55), in the manufacturing of a bicycle, each worker possesses a small fraction of the required knowledge, but no one is in the position to appropriate the entire production process. The same can be claimed for the much more complex design and construction processes of a nuclear power plant. However, a bicycle does not involve the same environmental risks as a nuclear power plant, and thus it is worth questioning whether the division of labor and the distribution of knowledge are compatible with conditions of safety in the civil nuclear industry.

It is possible, and highly desirable, to make improvements of an organizational kind within the nuclear power industry so that the group can gain better control of, and can act more efficiently against, nuclear risk. Thus, the four precautions recommended by the theory of high reliability organizations should be taken seriously even if they do not completely prevent certain epistemic difficulties. However, one must take into account these difficulties and find solutions for preventing them. Thus, for example, human resources policies could be improved by increasing the value of technical work and experience in order to encourage engineers to dedicate themselves more fully to their work, by setting up work sessions with the explicit aim of promoting knowledge transmission between elders and youth, and by reducing the workload of engineers.

However, the division of work and the distribution of knowledge still hinder the search for an entire epistemic control on the machine. Indeed, nobody has a full and deep understanding of how nuclear power plants work: not the design engineers in R&D, nor the operator of a plant. Moreover, the American or French expert, invited to make a diagnosis in Fukushima, does not possess full knowledge about the plant, since she is a specialist in a specific area. If nobody really knows how a plant works, this is a problem. When designing and maintaining a machine require too many specialized skills, the machine is similar in all respects to a Tower of Babel.

Acknowledgements

I am very grateful to Anouk Barberousse, Jacques Dubucs, Ashley Graham Kennedy and Charles Perrow for their helpful comments on the paper. Any remaining shortcomings in the paper are of course mine. I am also in debt to Camille Lancelevée who introduced me to the sociology of organizations.

6 References

- Dretske, F. 1994. "If You Can't Make One, You Don't Know How It Works." In *Midwest Studies in Philosophy* Vol. 19, edited by Peter French, Theodore Uehling and Howard Wettstein: 468-482.
- Fricker, M. 1998. "Rational Authority and Social Power: Towards a Truly Social Epistemology," *Proceedings of the Aristotelian Society*, 19 (2): 159–177.
- Gorz, A. 1989. *Critique of Economic Reason*. Ed. Verso. First published as *Métamorphoses du travail: Quête du sens*, Galilée, 1988
- Perrow, C. 1984. *Normal Accidents: Living With High Risk Technologies*. (Revised edition, 1999). Princeton, NJ: Princeton University Press.
- Perrow, C. 2011. Fukushima and the inevitability of accidents. *Bulletin of the Atomic Scientists* 67: 44
- Rochlin, G. La Porte, T. and Roberts K. 1987 "The Self-Designing High-Reliability Organization: Aircraft Carrier Flight Operations at Sea" *Naval War College Review* 84-85.
- Rochlin, G. 1993. "Defining High-Reliability Organizations in Practice: A Definitional Prolegomenon." In K. H. Roberts, ed., *New Challenges to Understanding Organizations* (New York: Macmillan), 11-32.
- Sagan, S. 1993. *The Limits of Safety: Organizations, Accidents, and Nuclear Weapons*. Princeton, Princeton University Press.