

Mapping Methodical Change in Safety Culture

Kaupo Viitanen^{a*}, Teemu Reiman^b, Carl Rollenhagen^c, and Nadezhda Gotcheva^d

^a VTT Technical Research Centre of Finland Ltd, Espoo, Finland

^b Lilikoi Consulting, Lohja, Finland

^c Royal Institute of Technology (KTH), Stockholm, Sweden and Vattenfall AB, Stockholm, Sweden

^d VTT Technical Research Centre of Finland Ltd, Tampere, Finland

Abstract: Evoking positive changes in safety culture or accentuating its positive characteristics has been a long-lasting challenge in safety-critical organizations. While a variety of methods exist to facilitate this process, an understanding of the main factors to be taken into consideration when leading methodical change in safety culture is often limited. In this paper the authors describe twelve principles of safety culture change which attempt to summarize the essential good practices of safety culture change. Examples of failed safety culture change initiatives are then analyzed from the perspective of these principles to provide a tentative proof-of-concept of their usefulness. We propose that these principles can be utilized by safety (culture) practitioners by serving as a reference when leading methodical change in safety culture.

Keywords: safety culture, safety culture improvement methods, organizational change

1. INTRODUCTION

Evoking positive changes in safety culture or accentuating its positive characteristics has been a long-lasting challenge in safety-critical organizations. A wide variety of models that describe the characteristics of a good safety culture have been developed (for reviews, see e.g., [1, Ch. 11.3], [2, Ch. 2]). To accompany the models, a plethora of practical tools has been proposed (e.g., [3]–[6]), which allegedly have a positive influence on safety culture. However, the understanding of how the systematic use of these tools is actually supposed to influence safety culture and safety often remains implicit. Furthermore, the tool descriptions seldom explicate the extent to which they are applicable in a given context, and if not, how they need to be tailored to better fit specific contextual requirements. It is thus still unclear what the main factors are that need to be taken into consideration when attempting methodical change in safety culture.

To address this shortcoming, the authors developed a set of principles of safety culture change. The principles crystallize the essential good practices of leading safety culture change and implementing safety culture improvement tools. They cover the following general topics: understanding the nature of sociotechnical systems and how culture change occurs in them, setting goals in complex, nested and conflicted systems, identifying and leveraging various types of interactions within the system, and acknowledging the existence of different time scales when initiating culture change. In this paper we will describe the principles of safety culture change, and elaborate the underlying assumptions and theories that serve as the basis for the principles. We will then utilize the principles as an analysis framework in the context of safety culture improvement activities to illustrate how the potential neglect of these principles can result in failed or insignificant safety culture improvement. We propose that by following these practical principles, methodical safety culture change initiatives can be implemented with a more refined understanding of the sociotechnical system, thus avoiding superficial and mechanistic applications of the concept of safety culture. We also propose that these principles can be utilized in steering safety culture change activities, and support, for instance, practitioners in choosing which safety culture change tools or practices they should (or could) use, and how. The principles can also help explain why a safety culture change effort has failed or proved ineffective, or

* Corresponding author. Email: kaupo.viitanen@vtt.fi

conversely, identify the leverage points in the system that can help drive a positive change in safety culture.

2. PRINCIPLES OF SAFETY CULTURE CHANGE

One of the main issues in improving safety culture is successfully implementing the practical methods and tools available to practitioners as a part of an overall strategy of methodical safety culture change. To provide support for this effort, the authors developed a set of principles (Table 1) which intend to summarize the essential good practices of leading safety culture change and implementing safety culture improvement tools. The principles were developed in a researcher workshop as a result of an exercise which aimed to map the elements of safety culture and the mechanisms of its change (for further details see [7]). The principles direct attention towards understanding the nature of sociotechnical systems and how culture change occurs in them, setting goals in complex, nested and conflicted systems, identifying and leveraging various types of interactions within the system, and acknowledging the existence of different time scales when initiating culture change. In this chapter we summarize the principles, which will serve as an analysis framework for the examples described in chapter 3.

Table 1: Overview of the Safety Culture Change Principles

Description	
Principle 1	Consider the dynamics between classes of system elements
Principle 2	Select the boundaries of the system you want to change
Principle 3	Select the system elements you want to change
Principle 4	Acknowledge that safety culture is not monolithic and internally coherent, and try to benefit from this
Principle 5	Make an effort to understand what organizational members actually do to identify leverage points for safety culture change
Principle 6	Identify the assumptions embedded within safety culture change tools
Principle 7	Identify and make use of the indirect effects of safety culture change tools
Principle 8	Acknowledge that safety culture cannot be directly changed
Principle 9	Consider how power relations influence safety culture change
Principle 10	Involve the target group in safety culture change activities
Principle 11	All information acquired in safety culture activities may be useful in the future
Principle 12	Behavior or structure change may result from safety culture change activities but do not expect rapid change in values and assumptions

Sociotechnical systems consist of various classes of elements that interact with each other, thus forming the overall system. Culture is one of the phenomena that emerge in a sociotechnical system. The way in which the concept of safety culture is used can sometimes be misleading when it is viewed from the perspective of sociotechnical systems. This is due to the inconsistency and lack of an elaborated view of systems that many safety culture models suffer from: the models may refer to artifacts, values, structures, or behavior, and few (if any) to basic assumptions, but rarely do the models explicate how these classes of system elements interact and form culture. What is typically discussed under the label of safety culture does not help understand what takes place in the system and how, for instance, culture can actually be changed. Safety culture change should involve an understanding of what are the different elements of the system that are relevant to safety, how they interact with each other, and what are their temporal dynamics.

Organizational culture changes as a result of the interaction between shared values and assumptions, structures and behavior [8]. Values and assumptions define, for example, what is considered relevant, important or possible and what meanings are given to actions, decisions or perceptions [9]. In the context of safety culture, values and assumptions can, for example, relate to conceptualizations of safety, accidents, hazards, human error and safety management. Values and assumptions are often characterized as being the most inaccessible element of culture and are associated with “the core of culture”, with other elements as manifestations of it [9]. Values and assumptions can be explicit (e.g.,

organizational strategies and visions, espoused values), or implicit (e.g., unconscious taken-for-granted assumptions). Structures, on the other hand, store values and assumptions as they are reflected in the contents of administrative structures (e.g., management systems, processes, procedures rules, regulations, and other institutionalized practices) or physical structures (e.g., buildings, technologies). Behavior creates and changes structures (e.g., through managerial decisions or personnel initiatives), and values and assumptions. Behavior, like structures, also reflects values and assumptions.

Values and assumptions result from the interactions between behavior and structures, but also between other values and assumptions. If one wants to influence values and assumptions, there needs to be an overview of the elements of the system and how they interact. We believe that it is also important to acknowledge that almost any value or assumption can be safety related. Some values may be more safety-related than others, and distinguishing between them can be difficult. In general, safety culture should always be perceived in relation to other associated values and assumptions that may exist in organizations (incl. subcultures). One cannot target only one aspect of one element when leading safety culture change (e.g. pick a value or an assumption and change it), because as an adaptive entity, a sociotechnical system will reject and compensate any disturbance that it is not compatible with. Thus, other interconnected elements of the system are likely to change the system back to its original state. This often manifests itself as change resistance or unwanted side-effects. The assumption that systems consist of a diverse set of elements that define how safety culture develops forms the basis for principle 1: consider the dynamics between classes of system elements.

The boundaries of a sociotechnical system can be viewed from several perspectives, for example, vertically by separating different levels of the system, or horizontally by separating between entities at same levels. An informed selection and awareness of system boundaries is needed to understand who or what influences or is influenced by safety culture, i.e., who is the target and what is the context of safety culture change. The selection of boundaries thus affects the spatial scale (in organizational terms) of things that are taken into consideration. A narrow selection of the boundary can help focusing safety culture change activities and thus help avoid diluting them with irrelevant nuisance factors. At the same time, it can contribute to losing a big picture of how different system elements interact and create emergent patterns (cf. principle 1). Principle 2: select the boundaries of the system you want to change is based on the assumption that system boundaries define who influences and who is influenced by safety culture.

Methodical safety culture change requires making the purpose of the change clear: what system element is the target of the change: behavior, structures, or values and assumptions? It is also necessary to understand what characteristic of that element requires changing, e.g., what kind of behavior or structures is the target of the change, or which set of assumptions or values need to be influenced. Since some of the system's existing characteristics may be desirable, and some may not be, it is necessary to consider, what "safety culture improvement" specifically entails. Four types of safety culture change actions can be formulated, cf. theory of elementary actions [10]: maintaining a desirable characteristic, eliminating an undesirable characteristic, creating a desirable characteristic and preventing the emergence of an undesirable characteristic. Conceptualizing safety culture change activities from this perspective can help ensure that the organization's safety culture strategy does not only focus on one type of action without considering the other types of actions. Principle 3: select the system elements you want to change draws from the assumption that system elements have desirable and undesirable characteristics.

Traditionally, the concept of organizational culture implies unity and homogeneity of a group. Many organizational culture theories are based on this assumption, including Schein's theory [9], which views organizational culture as reflecting something that is shared. However, in reality most, if not all, organizational cultures consist of subcultures (and this is increasingly acknowledged by Schein's work, [11]), each with their unique sets of values and assumptions, sometimes conflicting and competing with each other, sometimes co-existing in harmony [12]. This means that there can be a multiplicity of (safety) cultures that are constantly created and recreated as groups of people interact with each other and their environment [13]–[17]. The existence of subcultures has implications to

methodical safety culture change: attempting to create a uniform safety culture might not be possible or even be a desirable goal. An attempt to influence safety culture from a cultural heterogeneity perspective may result in unexpected outcomes because each subculture responds to the change intervention differently. Thus, the cultural heterogeneity calls for a tailored (cf. principle 10) or at least subculture-aware approach to safety culture change.

Conflicts can also manifest within cultural entities, i.e., between or within different classes of system elements. Values and assumptions might not be in line with each other, or behavior and/or structures might not always reflect or support each other, or the values and assumptions. For example, procedures may not always match the way in which work is done in reality (structures-behavior mismatch or “Work-As-Done-Work-As-Imagined” mismatch, [18]) or the values espoused through promotion might not reflect what the values or assumptions of organizational members really are (structures-values or assumptions mismatch). The extent to which these elements agree with each other may influence the organization’s ability to manage safety, see e.g., [19]. The existence of conflicts and paradoxes in safety-critical organizations is not a negative phenomenon per se, rather, it is an inherent characteristic of a complex sociotechnical system [20]. Sometimes the heterogeneity can in fact be a necessity for safe activities: for instance, a system requires sufficient variety (e.g., in terms of interpretations) to be able to regulate the safety-critical activities or to facilitate learning [15], [21]. Principle 4: acknowledge that safety culture is not monolithic and internally coherent, and try to benefit from this is based on the assumption that systems are inherently paradoxical and conflicted.

Since culture is a group-level phenomenon, one of the essential ingredients for the formation of shared values and assumptions are the interactions between individuals. Culture forms when a group of people interacts with each other, solving externally or internally-induced problems [9]. Understanding the interaction between different organizational groups is thus essential in the study of safety culture, e.g. [19]. This suggests that the arena of interest for methodical safety culture change is where the organizational members interact: their everyday work. Principle 5: make an effort to understand what organizational members actually do to identify leverage points for safety culture change conveys the idea that being aware of the contextual problems the organizational members face in their everyday work allows for a better understanding of why certain values and assumptions are held and this can be utilized for methodical change interventions.

A wide variety of tools and methods have been developed that attempt to change safety culture. Some of the most common types of tools include safety culture trainings and seminars, influencing behavior using, for example, human performance programs, increasing the awareness of safety or safety culture through promotional materials, building employee involvement and participation, and developing organizational structures with the intent of ensuring a good safety culture [6]. The selection of documented tools and methods is quite heterogeneous and targeted at different system elements. This essentially suggests that the mechanisms of action also differ between the tools and methods. Because assumptions guide our actions, all safety culture change tools are developed and used on the basis of certain assumptions regarding how the tools function and influence the system. These assumptions can include the developers’ views regarding safety culture and its change, models of the organization or perspectives of management and leadership. Since tool development is influenced and steered by these assumptions (either consciously or unconsciously) they also get embedded in the tools. Sometimes the assumptions are explicated in tool descriptions or instructions, but often they are not. This implies a risk of a mismatch between the assumptions made during the development of the tools, and the assumptions made by the end-user regarding the tools’ intended function and effect. The mismatch can lead to the end-user utilizing the safety culture change tool in an unintended manner or otherwise render the tool ineffective. There thus needs to be a sufficient agreement between those who develop or introduce the tools and those who implement them: the end-user needs to be active in identifying the assumptions that the developer has embedded in the tools. An informed caution is necessary when adopting safety culture improvement tools, methods or approaches, even if they have been promoted as “best practice” or enforced by peer networks. Principle 6: identify the assumptions embedded within safety culture change tools reminds practitioners to take the extra step to appreciate why and how the tools were developed.

We make a distinction between direct and indirect effects of safety culture change tools. The direct effects of safety culture tools include changes in organizational structures (e.g., outcomes of tools that aim at organizational development) or behavior (e.g., outcomes of tools that aim at behavioral modification). The tools' indirect effects result from the direct effects: a change made in one part of the system may trigger adaptation in another part, sometimes leading to cascading and non-linear (much larger or much smaller than the initial, direct effect) outcomes, that sometimes can be system-wide. The indirect effects, as defined here, are not necessarily negative (as is the case with adverse side-effects), they can also be positive or neutral. In fact, the indirect effects, if well-understood, can potentially be harnessed to create effects not otherwise possible. This can be the case when attempting a change in safety-related basic assumptions (cf. principle 8). Principle 7: identify and make use of the indirect effects of safety culture change tools is based on the assumption that sociotechnical systems react and adapt to the effects of safety culture change tools.

Because organizational culture constantly redevelops itself through the interaction of behaviors and structures [8], there is no “direct access” through which values and assumptions can be changed. Attempting to influence them directly is likely to be ineffective. Thus, the only way to influence values and assumptions is through changing behavior or structures. Principle 8: acknowledge that safety culture cannot be directly changed suggests that the key ingredients to methodical safety culture change are the indirect effects of behavior and structural change induced by safety culture change tools (cf. principle 7).

Being aware of power relations is crucial in leading safety culture change. However, it is often neglected in safety culture studies, see e.g., [15]. The interests of individuals in various positions of power influence the implementation and efficiency of safety culture activities. For instance, top management commitment is viewed in many safety culture models as an important prerequisite for a good safety culture, e.g., [22], [23]. One reason for this is the fact that top management is in a position of formal power in the organization. Top management commitment is a multi-faceted phenomenon that can manifest differently. A weak commitment can manifest as insufficient resources or a weak power and organizational position of safety culture experts. We believe that the key to managing organizational power relations is being aware of them and the challenges and opportunities they pose upon safety culture change activities. Principle 9: consider how power relations influence safety culture change emphasizes that methodical safety culture work cannot disregard organizational power relations as an influencing factor.

A safety culture change initiative can only be sustainable when it is compatible with the target groups. This usually requires their involvement in planning. In practice, locally tailored safety culture change can utilize organizational members that are characterized, for example, by high social connectivity (i.e., they are “hubs” who know everyone and who are known by everyone), or by having desirable safety-related competences or values. After the identification of these individuals, they can be utilized as drivers of positive influence towards safety culture by ensuring that they have sufficient opportunities to contribute to the development of a healthy safety culture. It should, however, be noted that involving the target group in safety culture activities does not mean there should be no top-down involvement. In fact, both are needed as the local group is often blind to its assumptions and may be thus hesitant to change things they consider meaningful, despite them being potentially detrimental on the system level (cf. principle 4). Principle 10: involve the target group in safety culture change activities emphasizes that to avoid using an excessive amount of resources on futile or insignificant change efforts that are either rejected or resisted by the culture, tailored and collaborative solutions are advisable.

Implementing safety culture change activities and conducting safety culture assessments usually produces a large amount of information about the organization and its culture. However, the immediate usefulness of this information can vary. For instance, some of the insights gained from safety culture-related information at a given moment of time can appear to be insignificant, irrelevant, or concern matters that cannot be changed. However, due to the dynamic and evolving nature of

sociotechnical systems, the meaning and relevance of the information can also change over time. Safety culture problems previously perceived as impossible to influence or inherent to the system can later be found to be trivial, and within the scope of common safety culture change tools. Principle 11: all information acquired in safety culture activities may be useful in the future reminds that as one carries out various safety culture activities during a longer period of time, care should be taken to avoid anchoring to initial preconceptions of what can and cannot be done to develop the system, because this may result in dismissing opportunities.

The various classes of system elements react differently to attempts to influence them. As discussed in the context of principles 1 and 8, change in values and assumptions is an indirect result of the change in behavior and structure. This suggests that the effect of intentional change in values and assumptions is likely to lag behind those in behavior and/or structures. When successfully implementing a safety culture change tool or method, the direct effect of the method should result in behavior or structure change, but most likely do not lead to significant, immediate change in values and assumptions. The change in values and assumptions is rather a result of a long-term evolution of the system. This evolution takes place naturally as the system responds to various internal and external stimuli, but can to some extent be steered by changing behaviors and structures. Principle 12: behavior or structure change may result from safety culture change activities but do not expect rapid change in values and assumptions thus suggests that methodical safety culture change requires a long-term strategy, commitment and perseverance because system elements vary in terms of temporal stability and susceptibility to change.

3. EXAMPLES OF APPLICATION

In the following subchapters we present examples of safety culture change efforts that have proven ineffective or undesirable, and briefly analyze them from the perspective of the safety culture change principles. The narratives as such are fictional but they are based on the authors' experiences from safety culture-related work in various research or consultancy projects, or as industry practitioners. The purpose of this analysis is to utilize the principles to understand and explain why a safety culture change effort had failed and how this failure could have been avoided.

3.1. Example 1: Construction of a Safety Culture Questionnaire

In the beginning of 1990s when safety culture was still a relatively new concept, a nuclear power company decided to construct a safety culture questionnaire and an expert was commissioned to carry out this task. Over the course of years, the questionnaire evolved as flaws were identified and subsequently corrected. Notable lessons learned during the development of this assessment tool were: a) the questionnaire was originally vague on what kind of safety it was about, which mislead the responders and increased uncertainty during analysis and interpretation of results; b) a need for at least partial contextual adaptation of the questionnaire was discovered because different departments and groups of employees deal with very different safety issues; c) too little time was invested by the power company in actually using the data acquired from the questionnaire, which resulted in its underutilization; and d) separating the analysis of safety culture questionnaires from other questionnaires (e.g., that generally cover organizational culture, etc.) was found to be artificial because many items (and concepts) correlated highly.

In this example, as the experience accumulated, a continuous improvement process is manifested as the identified flaws in the safety culture questionnaire are corrected. An issue which is apparent in this example, and often overlooked in safety-related questionnaires (and in safety communication in general), is that the concept of safety is not clearly defined. When not explicated, every group of responders may interpret it from their own perspective (e.g., some view it as referring to occupational safety, others as nuclear safety, fire safety, electrical safety or environmental safety, etc.). Different types of safety relate to different risks, hazard sources, and ways to control them, and possibly to different shared values or assumptions: for example, nuclear power plant employees who work with turbine equipment or at the office may find nuclear safety a more distant concept than occupational

safety or fire safety, which they might grasp more readily. The inherent heterogeneity of the organization and the definition of the type of safety was thus neglected in the questionnaire design, i.e., the target of the measurement was not clear (**principle 3**). Because of this, it was likely that the responders, due to their different interpretations, in fact responded a “different” questionnaire even though the questions were the same. The data can therefore be assumed to be inconsistent and drawing conclusions from it can be problematic. To correct this issue, the type of safety was subsequently explicated in questionnaire description and items.

Another way to address the issue of organizational heterogeneity discovered during the development of the safety culture questionnaire was contextual adaptation. This might involve, for instance, changing the wording of a subset of items to better meet the context of a particular target group, or having unique sets of items for each group (perhaps along with some questions that are common for all). Understanding the daily operational realities (**principle 5**) and acknowledging the influences of different groups on the questionnaire data (**principle 4**) were thus found to be crucial in a successful implementation of the questionnaire. It is also likely that the expert responsible for the construction of the questionnaire is not able to grasp the existing realities in their entirety by themselves, and involving the target groups in the design of the questionnaire is needed (**principle 10**).

The initial implementation of the questionnaire was relatively superficial in terms of its utilization: the questionnaire was expected to be useful as such without much effort put into understanding its results and providing guidance to interpret the results. This is not in line with the assumptions embedded within questionnaires as a safety culture assessment method (**principle 6**), which usually assume that an informed and possibly partially joint analysis with the target group is conducted, along with a comparison to the results of findings gained from other information sources. In connection to the latter assumption, the organization did observe a correlation between other data sources, prompting them to assume a more integrated approach to safety culture assessment instead of viewing safety culture as a phenomenon that exists in isolation from other organizational phenomena (**principle 1**). The possible root cause behind these implementation-related shortcomings was that the organizational prerequisites for implementing a safety culture questionnaire were not in place: there was not sufficient commitment (**principle 9**) and subsequently not enough resources to carry out the implementation process in such a manner that the questionnaire would be optimally useful to the organization.

3.2. Example 2: Implementation of Corrective Actions

After a nuclear power company had collected and analyzed data from a safety culture questionnaire, a couple of departments were identified where the mean scores were much lower than in the company in general. Top management required the supervisors from these departments to make quick corrective actions to remedy the situation. The high scoring departments were rewarded. Next year, the lowest scoring department scored in the highest quartile and the mean scores of the entire organization went up. The top management rewarded the entire organization.

This example illustrates a situation where a safety culture improvement activity (implementing corrective actions based on findings from assessment) resulted in changes that can potentially be detrimental to safety and at the very least rendered the insights gained from safety culture questionnaires unreliable. The main culprits in this example are the decisions made by top management, which indicate their lacking understanding of the nature of sociotechnical systems and safety culture, and the limitations of the safety culture questionnaire tool and rewarding as an intervention technique. We examine the main shortcomings in the management’s activities and relate them to the principles of methodical safety culture.

When the top management issued the order to quickly carry out corrective actions in low-scoring departments and rewarded high-scoring departments, the departments reacted to this decision by producing a bias in the next year’s scores, which suggests that rather than actually changing the underlying causes of the low scores, the top management’s actions concealed them further as the next year’s scores improved all-round. Effectively this means that the potential organizational flaws or

cultural issues that contributed to the low scores continued to exist, and the monitoring tool that was able to identify them was disabled. This suggests that the indirect effects of a safety culture change tool (i.e., rewarding) were misunderstood (**principle 7**). This example illustrates that the top management's selection of the target of their rewarding intervention was not informed (**principle 3**). It was directed at questionnaire scores, and consequently the scores changed. However, if they actually wanted to improve, for example, safety-related behavior or structures, or values and assumptions, not enough thought was put on understanding the root causes behind the low scores and selecting the correct target of the intervention (and perhaps a different intervention method) to address the causes. By focusing only on rewarding by questionnaire scores, the top management ended up creating and reinforcing unwanted behavior (i.e., questionnaires response bias).

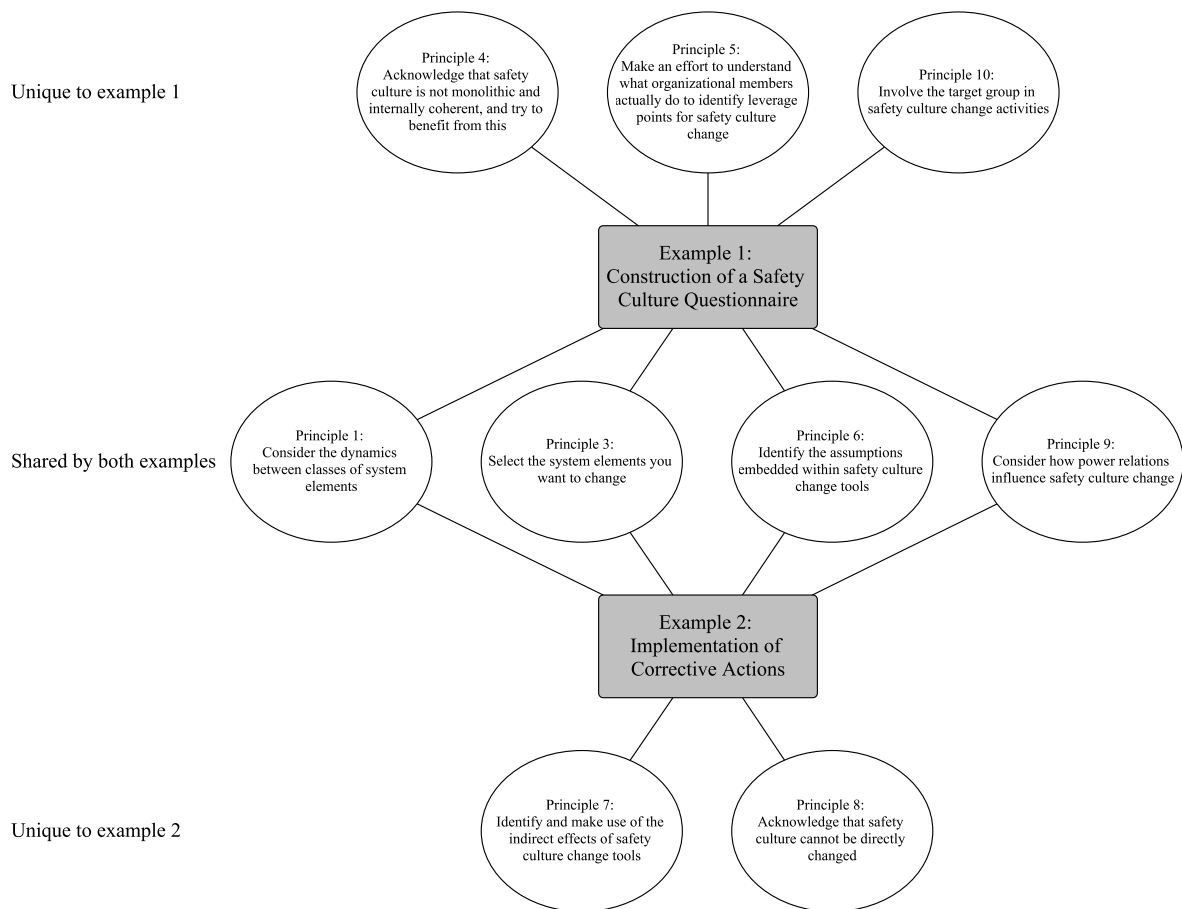
Furthermore, the top management failed to appreciate the relation and dynamics between different classes of system elements (**principle 1**) when they made the assumption that a low questionnaire scores reflect undesirable characteristics of safety culture as such (**principle 8**), and that a simple corrective action would fix this. Making such an assumption also suggests that the top management was not properly aware of the assumptions embedded in safety culture tools (**principle 6**). The questionnaire validity of safety culture questionnaires is based on the assumption that they will be used as a part of an assessment toolkit to triangulate the state of safety culture together with other, complementary data collection methods. Rewarding, on the other hand, requires a careful selection of rewarding criteria and their indirect effects.

Finally, the top management did not consider what effect their actions can have on questionnaire results in subsequent measurement times: because the management is an actor in high formal power, its actions and decisions will have a strong influence on the behavior of others (**principle 9**). In this example, it can be argued that during the second measurement time, the questionnaire in fact measured the management's persuasive power rather than what it had originally been intended to measure.

4. DISCUSSION

Closer examination of the examples cases demonstrated how the insights embedded in the safety culture change principles were not taken into consideration, which consequently contributed to the adverse outcome of the activity in question. We use a conceptual network to illustrate the connections between the examples and the principles that emerged in the analysis (Figure 1).

Figure 1: The Relation between Principles and Examples



The themes unique to the first example related to system heterogeneity. During the construction of the questionnaire it was not acknowledged how diverse the organization is, which led to ambiguous results. Considering the principles 4, 5 and 10 would probably have helped avoid the situation by directing attention towards organizational complexity and helping find ways to deal with it.

The themes unique to the second example related to the complex and indirect relations between safety culture change tools and culture. It was misunderstood what the questionnaire says about the organization's culture and how misplaced rewarding would influence employees. Taking insights from principles 7 and 8 might have prompted the top management to think more carefully the consequences of their actions.

Several themes are shared by both cases. First, in both cases the interactions between classes of system elements that contribute to the formation or development of culture were neglected. In first case, this manifested through the (initial) assumption that safety culture is separate from other cultural phenomena, and in the second case it was assumed that examining explicit manifestations of culture (as covered by questionnaires) are sufficient for understanding the complex and adaptive nature of organizational culture. Consulting principle 1 to understand the mechanisms through which culture creates and recreates itself and to identify the networks of behaviors, structures, and values and assumptions – or at least being aware of their existence – might have avoided the misuse of the safety culture tools.

Secondly, in both cases an uninformed selection of the target of a tool was done. In the first case, it the target of the measurement (type of safety) was not clearly defined which led to ambiguous findings. In the second case, questionnaire scores were selected to be the criteria for rewarding, which resulted in the employees optimizing their future responses according to this. Principle 3 directs attention to

explicating what exactly one wants to do to the system, and more closely explicating what are the desirable or undesirable characteristics of system elements. For instance – referring to the second case – getting low questionnaire score is not necessarily a negative issue as such, and should not be treated as one.

Another topic in common between the two cases was that the safety culture tools themselves were not fully understood. In the first example, questionnaire was initially expected to be useful without much effort put into reflecting its implications and comparing it with the results gained from other data collection tools. This indicates a lack of knowing the limitations of the method. In the second example, neither the questionnaire method, nor the influences of rewarding were understood or taken into consideration. Principle 6 reminds practitioners to always be aware of how the tools and methods they use are intended to be used and what are the prerequisites of their successful use. In the example cases, following principle 6 might have prompted a closer reflection of the tools and methods and consequently how they should be used properly.

Finally, both cases displayed issues with power relations. In the first case, top management was not committed enough to allocate a sufficient amount of resources for the questionnaire deployment, which hindered its utilization, and in the second case the top management did not realize the influence of their actions over the personnel. Principle 9 suggests that when leading safety culture change, power relations should always be taken into consideration. In these cases, this might have helped ensure, for example, that top management is involved and understands its role.

This retrospective analysis is not intended as a conclusive evidence of the validity of the principles. Empirical studies with longitudinal data collection are advisable for that purpose. In the context of this paper, we view the analysis an illustration and preliminary proof-of-concept of how the principles can be utilized in helping create an informed strategy for safety culture change. They can serve as a checklist for practitioners and remind about the good practices of leading safety culture change.

5. CONCLUSION

This summarizes good practices for practitioners working in high-risk industries who aim to develop safety culture in their organizations. We presented descriptions of and rationale behind twelve principles of safety culture change. These principles draw from theories of safety culture, organizational culture, sociotechnical systems, management and leadership, and safety management. We believe that they represent the key issues that are relevant for practitioners to acknowledge when they plan to positively influence safety culture in their organizations. The principles can be used to steer safety culture activities, to decide which safety culture tools or methods should be used and for what purpose, to identify leverage points where safety culture can be affected, or to retrospectively understand why a safety culture improvement effort has resulted in adverse, insignificant or otherwise unexpected outcomes. The principles can thus serve as a reference for improving safety culture in a methodical and informed manner. As a proof-of-concept of the potential relevance of the principles, we utilized them to analyze two examples of unsuccessful safety culture improvement efforts to identify the factors that contributed to their failures. Despite the potential limitations of the analysis design, overall it showed that the principles can be useful for examining safety culture improvement activities and pin-pointing success factors.

Acknowledgements

The work presented in this paper is based on a research project funded by Nordic Nuclear Safety Research NKS and the Finnish Research Programme on Nuclear Power Plant Safety SAFIR2018.

References

- [1] A. I. Glendon, S. Clarke, and E. F. McKenna, *Human safety and risk management*, 2nd edition. Boca Raton, FL: CRC/Taylor & Francis, 2006.
- [2] M. S. Wright, P. Brabazon, A. Tipping, and M. Talwalkar, "Development of a Business Excellence Model of Safety Culture," Health and Safety Executive, 1999.
- [3] IAEA, "Examples of safety culture practices," International Atomic Energy Agency, Vienna, 1997.
- [4] INPO, "Human Performance Tools for Managers and Supervisors. General Practices for Organizing, Planning, Monitoring, and Feedback That Promote Excellence in Human Performance," Institute of Nuclear Power Operations, 07-006, 2007.
- [5] INPO, "Benchmarking – Nuclear Safety Culture Practices," Institute of Nuclear Power Operations, 12-006, 2012.
- [6] K. Viitanen, N. Gotcheva, and C. Rollenhagen, "Safety Culture Assurance and Improvement Methods in Complex Projects – Intermediate Report from the NKS-R SC_AIM," NKS, NKS-381, 2017.
- [7] K. Viitanen, N. Gotcheva, C. Rollenhagen, and T. Reiman, "Safety Culture Assurance and Improvement Methods in Complex Projects – Final Report from the NKS-R SC_AIM," NKS, NKS-405, 2018.
- [8] T. Reiman and C. Rollenhagen, "Safety culture," in *Handbook of Safety Principles*, N. Moller, S. O. Hansson, J.-E. Holmberg, and C. Rollenhagen, Eds. Hoboken: John Wiley & Sons, 2018.
- [9] E. H. Schein, *Organizational Culture and Leadership*. San Francisco, CA: John Wiley & Sons, 2010.
- [10] G. H. von Wright, *An essay in deontic logic and the general theory of action: With a bibliography of deontic and imperative logic*. North-Holland Publishing Company, 1968.
- [11] E. H. Schein and P. Schein, *Organizational Culture and Leadership*, 5 edition. Wiley, 2016.
- [12] J. Martin, *Cultures in organizations: three perspectives*. New York: Oxford University Press, 1992.
- [13] C. Rollenhagen, J. Westerlund, and K. Näswall, "Professional subcultures in nuclear power plants," *Saf. Sci.*, vol. 59, pp. 78–85, Nov. 2013.
- [14] J. R. D. Edwards, J. Davey, and K. Armstrong, "Returning to the roots of culture: A review and re-conceptualisation of safety culture," *Saf. Sci.*, vol. 55, pp. 70–80, Jun. 2013.
- [15] S. Antonsen, "Safety culture and the issue of power," *Saf. Sci.*, vol. 47, no. 2, pp. 183–191, Feb. 2009.
- [16] N. Pidgeon, "Safety culture: Key theoretical issues," *Work Stress*, vol. 12, no. 3, pp. 202–216, Jul. 1998.
- [17] A. Richter and C. Koch, "Integration, differentiation and ambiguity in safety cultures," *Saf. Sci.*, vol. 42, no. 8, pp. 703–722, Oct. 2004.
- [18] E. Hollnagel, "Why is Work-as-Imagined Different from Work-as-Done?," in *Resilient Health Care, Volume 2 The Resilience of Everyday Clinical Work*, R. L. Wears, E. Hollnagel, and J. Braithwaite, Eds. Farnham, Surrey, UK: Ashgate Publishing, Ltd., 2015, pp. 249–264.
- [19] S. Antonsen, *Safety culture: theory, method and improvement*. Farnham, England: Ashgate, 2009.
- [20] T. Reiman, C. Rollenhagen, E. Pietikäinen, and J. Heikkilä, "Principles of adaptive management in complex safety-critical organizations," *Saf. Sci.*, vol. 71, Part B, pp. 80–92, Jan. 2015.
- [21] K. E. Weick, "Organizational Culture as a Source of High Reliability," *Calif. Manage. Rev.*, vol. 29, no. 2, pp. 112–127, Jan. 1987.
- [22] IAEA, *Leadership and Management for Safety. No. GSR Part 2*. Vienna: International Atomic Energy Agency, 2016.
- [23] INPO, "Traits of a healthy nuclear safety culture," Institute of Nuclear Power Operations, 12-012, 2012.