

Modeling safety culture as a socially emergent complex systems phenomenon: a case study in aircraft maintenance

Mols, C; Bim, J; Passenier, D; Sharpans'kykh, OA

Publication date

2013

Document Version

Accepted author manuscript

Published in

Proceedings of the Fourth International Air Transport and Operations Symposium

Citation (APA)

Mols, C., Bim, J., Passenier, D., & Sharpans'kykh, OA. (2013). Modeling safety culture as a socially emergent complex systems phenomenon: a case study in aircraft maintenance. In s.n. (Ed.), *Proceedings of the Fourth International Air Transport and Operations Symposium* (pp. 1-15). s.n..

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Modeling Safety Culture as a Socially Emergent Phenomenon: a Case Study in Aircraft Maintenance

Abstract. Safety culture is often understood as encompassing organizational members' shared attitudes, beliefs, perceptions and values associated with safety. Safety culture theory development is fraught with inconsistencies and superficiality of measurement methods, because the dynamic and political nature of culture is often ignored. Traditionally, safety culture is analyzed by survey-based approaches. In this paper we propose a novel, systemic, interdisciplinary approach for investigating safety culture that combines multi-agent system modeling with organizational ethnography. By using this approach, mechanisms of emergence of safety culture from daily practices, operations and interactions of organizational actors can be modeled and analyzed. The approach is illustrated by a case study from the aircraft maintenance domain, based on existing ethnographic data. Using the proposed approach we were able to reproduce and explain emergent characteristic patterns of commitment to safety in the maintenance organization from this study. The model can be used for theory development and as a management tool to evaluate non-linear impacts of organizational arrangements on workers' commitment to safety.

Keywords: agent-based modeling; aircraft maintenance; organizational ethnography; power; simulation; social influence; complex socio-technical systems

1. Introduction

It is broadly recognized that safety culture plays a crucial role in shaping the safety and performance of operations in safety critical organizations (Reason, 1997) in such areas as air traffic (Mearns, Kirwan, Reader, Jackson, Kennedy and Gordon, 2013), nuclear power plant management (Lee and Harrison, 2000), and railway systems (Jeffcott, Pidgeon, Weyman and Walls, 2006). Safety culture has been defined in a variety of ways (Choudhry, Fang and Mohammed, 2007; Guldenmund, 2010). In this paper, we use the term safety culture as those aspects of organizational culture that may have an effect on safety, which is in line with Hopkins' (2006) primary focus on organizational culture and subsequent analysis of its impact on safety.

Organizational culture is a complex, context-specific phenomenon not easily harnessed in a single definition or theoretical approach (Giorgi, Lockwood and Glynn, 2015). Culture is particularly complex because its members may relate to cultural meanings in ambiguous ways (Kunda, 2009). Organizational members may engage in political ways with cultural elements such as particular basic assumptions, espoused values, technologies, and artifacts (Schein, 1990); and narratives, symbols, and rituals (Geertz, 1973). A culture is therefore not necessarily an integrated and stable whole,

but is dynamic and can be differentiated into subcultures or appear to be fragmented (Martin, 2002). This fundamental insight has not yet been thoroughly incorporated into the concept of safety culture.

Current safety cultural analytical approaches and frameworks have acknowledged limitations. They rely on linear models derived from accident research (Reason, 1990). These models, and the assumptions underlying them, are increasingly being criticized for being overly simplistic (Hollnagel, Woods and Leveson, 2006). They offer little insight in the way competing values interact in real organizations, such as the simultaneous challenges to accomplish growth and profitability as well as high safety levels (Antonsen, 2009). Current approaches to safety culture are therefore a-political, while in reality, power plays an important and complex role in the development of organizational safety cultures, as Antonsen (2009) has noted.

In this paper we present a novel, advanced research approach to investigate organizational safety culture as a complex phenomenon, emphasizing its gradual emergence in years of mostly normal operations. The approach takes a systemic view on safety modeling and analysis (Hollnagel et al., 2006) according to which safety hazards, accidents and incidents develop from complex, nonlinear interaction between diverse organizational processes. This view fits the social scientific understanding of culture described above.

The paper is organized as follows. In section 2, we describe the approach and its components of ethnographic research, multi-agent modeling, and simulation analysis. Section 3 describes the process of case selection as well as presenting details of the case study. Section 4 describes the model and how it was developed. In section 5 we present the simulation results of the model as well as virtual experiments. We conclude in section 6 with a discussion about theoretical implications, applications, limitations, and further development of the approach.

2. Approach

Modeling safety culture necessarily involves an in-depth study of a particular culture, which is usually done by ethnography. Therefore we introduce agent-based modeling and organizational ethnography, as well as how to perform agent-based modeling on safety cultural phenomena and incorporate ethnographic data.

2.1. Organizational ethnography

Because of the complexity of culture, organizational culture is typically investigated with ethnography. Ethnography involves making detailed discoveries of local informal processes and practices, understand the experiential reality of cultural members from within, and thus develop ideas about how patterns of interest emerge (Fayard & Van Maanen, 2015). Generally the method of participant observation is preferred. The researcher participates in the everyday life of the cultural group of interest to obtain an insider's understanding, while also maintaining an intellectual distance or preventing to 'go native'. In the case of auto-ethnography, a member of a culture may be trained as ethnographer and enabled to study the culture by reflecting on its common sense meanings with cultural outsiders (Doloriert and Sambrook, 2012). The data that are thus gathered are field notes of observations and conversations, transcripts of interviews and collections of meaningful documents. The analysis process is quite

particular to the research context which means there are hardly standard procedures for analysis. The results are conveyed through context-rich, or ‘thick’, qualitative descriptions (Geertz, 1973).

2.2. Agent-based modeling

The emergent nature of cultural phenomena is intuitively congruent with the multi agent paradigm. According to this paradigm, cultural patterns, seen as systemic properties of a multi-agent system, emerge and develop over time from many distributed local interactions of agents that represent organizational actors. Agent modeling can complement ethnographic theory building because it can serve as a formal test of its logic, and lead to new theoretical propositions (Harrison, Lin, Carroll and Carley, 2007).

A specification of a multi-agent system model comprises: 1) A description of agent types and structural relations between them; 2) A specification of local properties of each agent. Such properties comprise both internal (cognitive) properties and behavioral properties, i.e., temporal input-output relations of some complexity; 3) A specification of relations between agents, such as communication and power relations. 4) A specification of the environment.

This definition is in line with the work of Weiss (1999). More details on the specification of multi-agent systems are provided in Bosse, Jonker, van der Meij, Sharpanskykh and Treur (2009) and in Sharpanskykh (2008).

2.3. Agent-based modeling of safety culture

The combination of agent modeling and ethnographic data, that our approach implies, is unusual in the sense that it serves theory development. The safety cultural concept of interest, such as commitment to safety, is refined through the ethnographic study. A more complex idea about what the concept means and how it unfolds in practice is thus developed. Modeling then serves to work towards more generic statements by incorporating more social scientific theories. The aim is not to simply model an observed structure or process, but to incorporate much more detailed and conceptually refined observations that can lead to theoretical development.

The process of modeling and validation of the model through simulation requires teamwork between the researchers of the interdisciplinary team. Ethnographic data are used to build an agent model from generic socio-cultural mechanisms that aims to reproduce the patterns found in the ethnographic study. In our team, two ethnography specialists were guiding in the interpretation of the data. Modeling specialists created suitable mathematical representations, implemented and instantiated the model, and produced visuals for interpretation. The tension here is between the ethnography specialists who attempt to contextualize and the modeling specialists who attempt to see concepts and dynamics as more generic. The team must ensure that members are talking about the same concepts and make an effort in understanding some of the principles and details of the other’s discipline.

The modeling process involves jointly creating a conceptual model, after which the modeling specialists continue to formalize the model in mathematical equations. To test the model’s expressiveness and value to the current literature, we also compared it with another model of power relations (Appendix A). Our agent-based model was then implemented in Java. We did not use dedicated agent modeling tools because their

conceptual models of agents and specification languages are rather restricted for our purposes (e.g., to the BDI architecture). Furthermore, a lower level implementation runs faster, which is critical for sensitivity (Appendix B) and robustness analysis (Appendix C), when many runs with many interacting agents need to be performed.

Generating and interpreting simulations with the model serves validation as well as deepening theoretical insights and throwing up new questions. In section 5 we show how we could validate our simulation results by interpreting visualizations of the simulations and making sense of them in terms of the original research question. In addition to reproducing cultural patterns in the original ethnographic study, we performed several virtual experiments. In these experiments we explored the behavior of the model in different settings and under different conditions that seemed plausible in the given setting.

3. Case study

In this section we describe the process of case selection and present the necessary empirical details of the case itself. The case narrative is the basis for the agent model.

3.1. Case selection

Although different sequences are possible, we began by selecting an empirical research domain where a relevant contribution could be made. We chose the aircraft maintenance domain. In this domain daily practices are usually hidden from view in incident investigations, and only the emergent errors or failures are found when an incident or accident occurs. Thus, investigating aircraft maintenance safety culture has the dual benefit of understanding the emergence of safety cultural characteristics over longer periods of time, and of uncovering hidden, ‘latent failures’ (Reason, 1990) of air transport systems.

Our literature analysis furthermore suggested that the commitment to safety of maintenance technicians and maintenance teams is an essential aspect of safety culture, which presumably has a strong effect on safety of maintenance operations. Commitment to safety is a sensitive topic because organizational members are confronted with dynamic market demands and complex work situations, in which rules and procedures do not always make sense and ambiguities arise (Hale and Borys, 2013; Dekker, 2014). Under the pressure of meeting schedules and satisfying customers, a good safety culture means commitment to safety is ingrained such that it prevents pushing the balance too far towards working quickly and less safe (Edwards, Davey and Armstrong, 2013). What remains relatively under-investigated however is how context and power actually shape such commitment. Thus, the main research question of our case study is: *how does the commitment to safety of maintenance technicians emerge and develop under social and organizational influences?* Below we describe the data collection, and describe the case in terms of its internal power relations and the more complex pattern unfolding over time.

3.2. Data source

As a relevant and recent source of data for our study we used the case study on aircraft maintenance safety culture by Atak and Kingma (2010). The first author of this paper

performed auto-ethnography, allowing a unique insider's view of safety culture. We conducted additional interviews, literature and materials that helped us understand local social processes, safety cultural practices, and the context and the history of the organization in question. Furthermore, we modeled the power and influence relations between the agents using theories from social science. To gather missing data required for the model, and gain a thorough understanding of the secondary data, we conducted additional interviews with the authors Atak and Kingma, with a manager still working in the organization in question, and with domain experts from commercial aircraft maintenance organizations. In addition to the interviews, we used several reports on field studies performed in existing aircraft maintenance organizations in the context of large European projects such as HILAS (2007), ADAMS (Van Avermaete & Hakkeling-Mesland (2001), and TATEM, and smaller PhD projects (Pettersen, 2008; Ward, 2006).

3.3. *Power relations*

The company features professional power relations and dynamics that are common in the industry. At the core of operations are teams of Aircraft Maintenance Technicians (AMTs), which can generally be subdivided into junior technicians and senior flight engineers. Junior technicians usually do most of the actual handwork repairs, and senior flight engineers usually take on more of a supervisory role such as taking care of planning and drawing up work packages. The ratio of junior to senior technicians that closely cooperate in a team may be somewhere around six to one.

Supervising the seniors is a senior Maintenance Manager (MM), who is responsible for planning in terms of total man hours and work packages for incoming aircraft. The MM thus plays a key role in regulating production pressures that AMTs experience, and to which they may sometimes mount resistance. While the incentive from MM may be to increase speed and efficiency, the Quality Assurance and/or Safety Department (SD) of a maintenance organization provides a counterbalance.

The SD is there to ensure repairs are carried out in a way that does not harm flight safety as well as occupational safety of maintenance personnel. The SD is much less visible at the work floor than the MM. The SD takes on an advisory role towards management, monitors and investigates incidents, and reports to top management. Whether the SD has the power to assure safe operations then also depends on the CEO's prioritizing, which in turn is a response to the economic and operational challenges that the company is facing at a certain point in time.

We refer to Luke's (1974) three faces of power to distinguish different power dynamics resulting from the relations described above. The first face is direct power influences between agents, such as when one agent is able to make another agent do something that it would not otherwise do. This happens for example when a MM holds punitive power over an AMT. The second face is indirect power, such as setting an agenda or refraining from making a decision. When older AMTs use their experiential authority to talk about how things used to be, they exercise this power, as it draws away attention from current demands. The third face of power can be exerted by discourse, rather than individuals. This occurs when a discourse stressing production values takes hold due to changing context, such as the development described below.

3.4. Emergent patterns

The aircraft maintenance organization developed in three distinct phases from birth to maturation in a period of 9 years, beginning in 2001 and running up to 2009. In 2001, the organization experienced a takeover. First, it had to survive the challenging first phase with many ad-hoc projects and slim margins. Of the initial work force of 35, the company still employed 20 Aircraft Maintenance Technicians (AMTs) that had worked at the company before the takeover. The increasing work pressures led to resistance amongst this old guard while the company grew to employ 60 AMTs. Second, at the start of 2004, the company contracted new customers and grew explosively to a total of 260 employees by 2008. Members of the old guard gradually left the company and resistance diminished. AMTs worked overtime and there were some close calls, highlighting safety had a lower priority than productivity at this stage. Third, in mid-2008, a new technical director led the company into more mature phase. Safety and production goals were harmonized, the AMTs could develop a sound professional culture and adhering to safety standards became normal.

These three phases correspond to three different perspectives on the organizational culture (Martin, 2002), which is seen as consisting of a safety strand and a production strand. Due to its complexity, any organizational culture can at any time be seen as integrated, differentiated or fragmented, depending on the focus of the researcher. In the case of the maintenance organization, and in agreement with Martin (2002), one perspective was found to be dominant over the others in each of the consecutive three phases (Martin, 2002, as interpreted by Atak and Kingma, 2010, p. 269):

- Survival phase: *fragmentation perspective* ‘highlights ambiguity and a lack of clarity and conflicting or changing meanings in organizations’.
- Development phase: *integration perspective* ‘refers to the shared understandings in organizations’.
- Maturation phase: *differentiation perspective* ‘focuses on the existence of sub-cultures’.

During the survival phase, the old guard resisted the pressure of Maintenance Management (MM) to increase productivity by reverting to a work-to-rule resistance strategy: extensively complying with all procedures, stalling the company’s performance. The company’s Safety Department (SD), not very active at this time, was used in this way as a shield against MM’s pressure. To new AMTs, however, ‘the way we do things around here’ was not clear at this stage. Experienced and senior old guard drew new AMTs to collaborate with their resistance, while at other times MM could force decisions. There was little common understanding about the application of procedures. Confronted with ambiguity and conflicting values, the new AMTs’ experiences reveal a fragmented culture.

In the second phase, the organizational culture integrated around production values, at the expense of the safety strand of the culture. The company grew to 260 employees, with members of the old guard gradually leaving. The new AMTs continued to use the resistance strategy of the old guard against management pressure, though to a lesser extent. They were working overtime and sometimes even double shifts. Members throughout the company adopted a flexible attitude to safety trying to satisfy rising work demand enforced by MM.

In the final maturation phase, the company culture can be seen as differentiated, with several subcultures coexisting next to one another. MM still focused on productivity yet existed in relative harmony with SD, which was in turn occupied by a new, more pro-active team that was granted more influence. AMTs gained autonomy to act on their judgment and sticking to safety procedures became easy compared to the previous phases. Production demands were resolved through planning schemes rather than by pressuring AMTs to work faster, allowing AMTs to develop a professional sub-culture with safety as a solid priority. For a more detailed description of the study we refer to Atak and Kingma (2010).

4. The Model

In this section we show how we translated the ethnographic findings to a more general conceptual model, and derive a formal model with specific parameters. We translate our conceptualization of power relations into a generic set of mechanisms, incorporate the given context of the case, and quantify from mostly qualitative data.

In accordance with the case description, we identified four types of agents: new AMT (a novice technician), old guard technician, MM and SD. The behavior of the technician agents was the main focus of modeling. MM and SD agents exerted performance and safety pressures on that behavior, which were driven by conflicting goals and different interests, values and norms. These pressures and the AMTs' response to them were formalized by performance and safety demands and efforts considered in Section 4.1. Power and influence relations between the agents are described in Section 4.2. The spread of attitudes to performance and safety in shifts of AMTs was modeled as a social contagion process, considered in Section 4.3. In Appendix A we demonstrate how the proposed model can be related to another well-known agent-based model of power based on social dependencies. This was done according to Burton's (2003) model docking approach.

4.1. Demands and efforts

To describe at a high level the execution of maintenance operations by AMTs the following variables are introduced, in line with the essential organizational goals:

- performance demand (pd) and performance effort (pe);
- safety demand (sd) and safety effort (se)

All these variables vary from 0 to 1.

The performance demand for a task is an aggregate of the task complexity, the situational complexity of the environment, in which the task is being executed, and of the time pressure. High performance demand is associated with the interval $[0.7, 1]$, average – with $[0.4, 0.7]$ and low – with $[0, 0.4]$.

Several empirical sources indicated that AMTs normally have a high workload and work pressure during the night and in the morning, and a low to average workload and work pressure during the day.

To represent the changes in the performance demand imposed on the AMTs as described in the case study, a correlated random walk is used with fixed mean values for all the phases. These values are linearly interpolated during the transition phases and their values are different for day and night. In the first phase the means are 0.5 and

0.7 for day and night respectively, 0.6 and 0.8 for the second phase, and return to 0.5 and 0.7 for the third phase (Figure 1).

The meaning of performance demand we derived from Atak and Kingma's narrative as well as additional interviews. Zero performance demand means no work needs to be done at all. At 1, AMTs must disregard all quality and safety rules and precautions and completely exhaust themselves to meet schedules. An average demand of 0.5 means that work can be accomplished at a normal rate and no heavy pressure is experienced. On the threshold of average to high, at 0.7, the AMTs begin to experience that they must pick up pace, skip steps, and work overtime in order to accomplish the work. We took skipping steps at the level of 0.7 as not necessarily harmful for safety and can be done based on expert judgment. Similarly, working overtime can be experienced as doing something out of professional zeal, rather than being exploited. This depends on the context: in the first phase, the latter meanings applied (unsafe violations, exploitation), and in the third phase, the former meanings (safe workarounds, professional zeal). Passing 0.7 towards the level of 0.8 in the second phase, production pressures will begin to affect health and safety negatively after some time. Working overtime becomes working double shifts and skipping steps in procedures to meet schedules becomes routine.

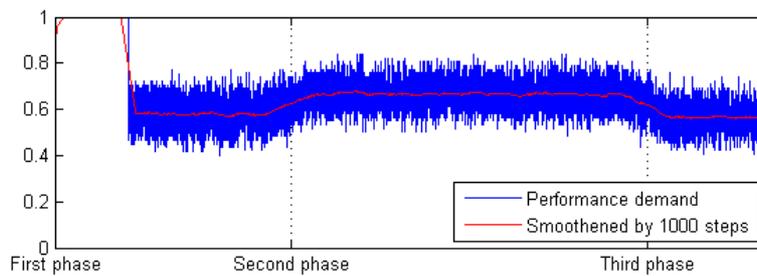


Figure 1. The dynamics of the performance demand during the three phases.

The safety demand indicates the required degree of compliance of an AMT with the safety standards of the maintenance organization and other regulatory bodies. The highest safety demand ($sd=1$) means that all the safety standards are required to be followed by the word. Multiple interpretations of standards may be possible; then the highest safety demand means that AMTs are required to choose the interpretation that by their professional judgment optimizes safety rather than efficiency or AMTs' own interests (cf. Atak and Kingma, 2010, p. 273). The minimum safety demand ($sd=0$) indicates that only a minimum set of strong safety requirements is required to be satisfied. Similarly to performance, high safety demand is associated with the interval $[0.7, 1]$, average – with $[0.4, 0.7]$ and low – with $[0, 0.4]$.

The dynamics of the safety demand imposed on the AMTs is defined in accordance with the case study description. For the first two phases a constant safety demand of 0.8, a high degree of adherence to procedures, was specified. In the last phase the safety demand lowered to 0.7 (Figure 2), reflecting the more mature organization and the autonomy over safety-related decisions granted to a developing professional culture of AMTs.

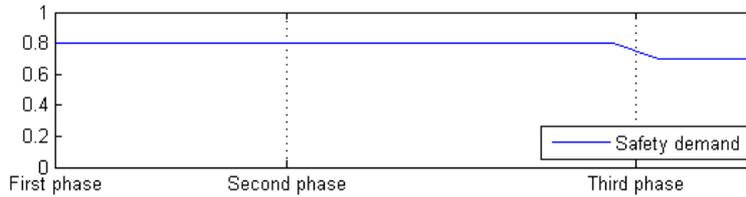


Figure 2. The dynamics of the safety demand during the three phases.

When the performance effort reaches a critical value called *critical performance point* (*cpp*), it starts interacting with the safety effort. This assumption is based on several evidences provided in Atak and Kingma (2010), indicating that some safety prescriptions were not followed by AMTs because of a high time pressure (performance demand). It is reflected in the model by the assumption that the higher the performance effort of an AMT agent, the less the maximum amount of safety effort it would be able to deliver. This relation reflects the well-known dilemma between performance and safety goals in safety-critical organizations, and is also in line with observations of operation execution in many maintenance organizations. To formalize such a relation between the performance effort and the maximum amount of safety effort of an AMT (i.e., the limit on safety effort from above), the logistic function $maxsft(pe)=1-1/(e^{-w1 \cdot pe} + w2)$ with $w1=25$, $w2=20$ shown in Figure 3 was chosen. This function determines *cpp* close to 0.7, corresponding the lower bound of the high demand interval. The high steepness of the function reflects that with the increase of *pe* above *cpp*, the maximum *se* degrades rapidly, i.e., every subsequent increase of *pe* by Δpe occurs at a rapidly increasing cost for the maximum *se*. In Appendix B, sensitivity analysis results are provided for different values of $w1$ and $w2$ corresponding to different values of *cpp*.

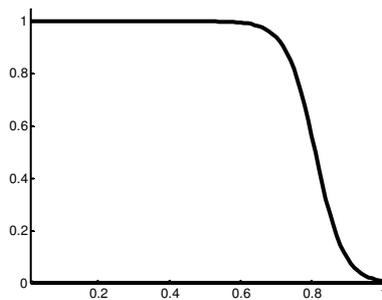


Figure 3. Change of the maximum safety effort (vertical axis) depending on the performance effort (horizontal axis).

Note that an AMT may not necessarily contribute the maximum possible amount of safety effort. The AMT's effort is determined as the result of social influences and contagion processes described in the following Sections 4.2 and 4.3.

The initial values of the efforts were drawn randomly in the beginning of the simulation from uniformly distributed ranges defined as follows: for old guard $se \in [0.7, 1]$ and $pe \in [0.4, 0.7]$; for new AMTs $pe \in [0.5, 0.8]$ and $se \in [0.5, 0.8]$.

4.2. Power and influence relations

In the model, MM is the main source of the performance demand imposed on AMTs, and SD is the main source of the safety demand. Besides the MM's and SD's influences, the AMT's commitment to performance and to safety is also shaped by influences from their peer AMTs. Such influences may be exerted explicitly by communication or may be transferred implicitly by observation.

To specify influence relations between agents, the French and Raven's power model (Raven, 1992) was used. This model introduces the following bases of power:

- *reward power*: its source is the ability to control legitimate reward and its strength increases with the magnitude of rewards;
- *coercive/punishment power*: its strength depends on the magnitude of the negative valence of the threatened punishment multiplied by the perceived probability that the punishments can be avoided by conformity;
- *legitimate power* of i over j stems from internalized values in j which dictate that i has a legitimate right to influence j and j has an obligation to accept this influence;
- *expert power* is the ability to administer information, knowledge or expertise to another agent; its strength varies with the extent of the knowledge;
- *referent power* of i over j has its basis in the identification of j with i , i.e., a feeling of oneness of j with i ;
- *informational influence* or *persuasion* is based on an information or logical argument that the influencing agent could present to the target to implement a change.

Different power bases may be correlated. For example, a source of legitimate power may be the value that one must adhere to certain types of expertise, thus enabling an agent's expert power base if the agent holds this expertise. An agent's ability to persuade may further increase his/her standing as an expert. Power bases should therefore not be thought of as independent variables, but conceptual tools to quantify the ethnographic textual descriptions and the sentiments conveyed by them. This necessarily involves some (inter-)subjectivity.

The strength of power-based influence of agent i on agent j is represented by parameter γ_{ij} with the range $[0, 1]$. For each influence relation between the agents from the case study a range of values was identified as shown in Figure 4. The actual influence values used in the simulation discussed in Section 5 were drawn from the uniform distributions defined by these ranges. In the following these influence relations are discussed more in detail.

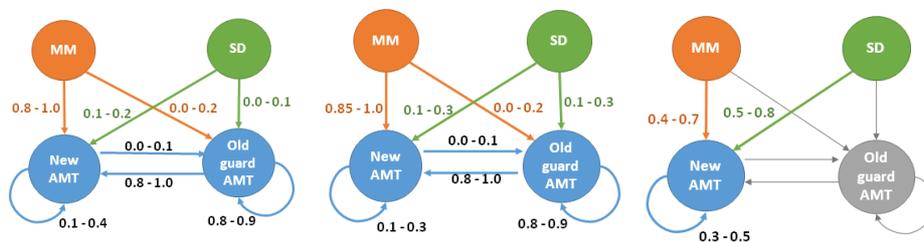


Figure 4. The influence relations between the agent types in the three phases.

According to the case description, in the first phase MM was pushing for production goals. The main power bases involved were strong coercive/punishment power, strong reward power and weak legitimate and expert power. MM had these power bases due to its formal position in the organization. For example, MM had the ability to control promotion of AMTs and to decide whether or not AMTs with a short term contract would be hired again in the future. The case study indicates that the new AMTs were more influenced by managers than the old guard (0.8-1.0 vs. 0.0-0.2) during the first phase. The old guard used their knowledge of safety regulations to resist the MM's pressure on production goals. Thus the old guard could structurally undermine both strong and weak power bases. This is reflected in high legitimate power between members of the old guard (0.8-0.9) that reduces the power of MM.

The ranges are determined by reasoning from within the specific situations that the agents can encounter. For example, we established a range of (0.8-1.0) power of MM over new AMTs for the following reasons. The influence level of 1.0 means that in some instances, new AMTs do not see how they could resist MM demands at all. New AMTs had weak legitimate power, as they are unaware of behavioral norms like how procedures should be executed, if those norms existed at all. At the same time, MM had very high coercive and reward power over newcomers. In the context of a fragmented culture, where there is no clear norm of 'how to do things', MM can play out any uncertainty about rules and always has the last word. MM may also hold some expert power towards newcomers, as the MM is himself a senior engineer and newcomers may not feel they can question this knowledge.

The lower level of 0.8 on the other hand reflects that in some situations, the legitimacy and expert power base of MM may be undermined by members of the old guard. A top-down management style may also create negative referent power. AMTs identify with one another and MM becomes the out-group, resulting in a desire *not* to behave in a manner seen characteristic of MM. That is, AMTs then resist behavior such as respecting work packages and trying to meet performance demands. This together creates a lower level that we determined at 0.8, rather than, for example, 0.7, which would be at the threshold of a medium to strong power level. The reason for this is that the power of MM over new AMTs never reaches close to average. New AMTs generally feel at the mercy of MM during the first phase as clear norms are lacking and they have little individual or shared experience to resort to.

According to the case description, the old guard acted as role models for the new AMTs, which is represented by high referent, expert and legitimate power towards new AMTs. In situations where MM is absent, the old guard could pressure new AMTs to do the complete opposite of their instructions, while the lack of norms allows them to come up arbitrarily with instructions that new AMTs should follow. This is reflected in the upper level of 1.0. The lack of an integrated culture may also undermine their own legitimacy; the old guard also does not control MMs reward and coercive power on new AMTs, which together is reflected in a lower limit of 0.8, and thus a range of (0.8-1.0). The role of SD during the first phase was only marginal. SD had very limited legitimate power to influence both new AMTs (0.1-0.2) and the old guard (0.0-0.1) by imposing safety values on them.

In the second phase the size of the old guard was decreasing, by the end of the phase all of them had left the organization. MM was still the main driver behind the production goals. The MM's influence was still based on the same power bases as in the first phase and additional referent power. In the second phase MM also included former technicians, who were perceived as role models by new AMTs. Besides the new

referent power, the legitimate power of MM had increased due to the struggle of the organization for survival. This created a context in which a discourse about meeting schedules became commonly accepted as the reality to reckon with. From the MM's position, it became easier to argue why tight plans should be met and working overtime should be considered normal. Therefore, the MM's influence strength on new AMTs had increased in comparison with the first phase as well (0.85-1.0 vs. 0.8-1.0). We take 0.85 rather than 0.9 because the effect of legitimate power, while significant, is impacted negatively by the gradual development of new AMTs' experience and knowledge of company rules and business arrangements. This allows new AMTs to develop their own informational and expert power base, and undermine coercive and legitimacy power bases. Yet because the company grows so quickly, we estimated these effects remain very small on average for the duration of the growth phase.

Even though the main focus during the second phase was on production goals, SD became more involved and present in the organization. However, their behavior was rather reactive - they responded to occurrences and supported production goals. Their legitimate power had increased a little. SD thus could thus always exert a trace of influence, translating in power level of 0.1. On occasions – such as after an incident or during a safety campaign – influence would be slightly higher, but still very weak, resulting in the range (0.1-0.3).

In the third phase a new proactive SD was formed. Management pressure on production goals was not dominant anymore and AMTs were able to work according to procedures and safety rules. The legitimate power of MM decreased and SD even had coercive and punishment power. For example, in this phase SD was able to initiate the suspension of the license of an AMT who did not work according to procedures. This new power of SD over the AMTs is reflected in the higher influence strength (0.5-0.8). In their turn, the AMTs had more freedom to work according to safety standards and had more legitimate power against management pressure. Note that there is no mechanism that establishes a new 'old guard' as time passes by and new AMTs become more experienced. This is a theoretical possibility but there were no data to substantiate this process here.

To determine combined influences of groups of agents (such as the management or old guard) based on the individual influences of the group members, the Latane's (1981) dynamic theory of social impact is used. According to this theory, the strength of influence γ_{Gi} of group G of N agents on agent i is determined by:

$$\gamma_{Gi} = N^\beta [\sum_{k=1..N} (\gamma_{ki} / I_{ki}^2) / N], \quad (1)$$

where γ_{ki} is the strength of influence of group member agent k on agent i ; it is defined for the different phases of the case study as discussed above;

$I_{ki} \in [0, 1]$ is the immediacy of agents i and k , i.e., their closeness in space and time; we assume that $I_{ki}=1$, i.e., the agents interact with each other without intermediate agents.

β is a constant used for compensation for the group size; from empirical studies $\beta = 0.3..0.5$. In our study $\beta = 0.4$.

4.3 Modeling social contagion

Provided a demand for performance or safety by an influencing agent, an AMT agent decides to which extent and how fast to satisfy this demand by delivering its performance effort.

The AMT i 's performance effort pe_i for the MM's performance demand pd_{MM} is determined by the following social contagion equation (Deffuant et al., 2000):

$$dpe_i = \alpha_{MM,i} \gamma_{MM,i} (pd_{MM} - pe_i) dt \quad (2)$$

Here $\gamma_{MM,i}$ is the influence of MM on i , $\alpha_{MM,i}$ is the rate of change parameter, which depends on the agent's openness to change and the expressiveness of the influencing agent/group. Since MM has a high expressiveness in all phases, $\alpha_{MM,i}$ is taken 0.8.

Similarly the AMT i 's safety effort se_i for the SD's safety demand sd_{SD} is determined by

$$dse_i = \alpha_{SD,i} \gamma_{SD,i} (sd_{SD} - se_i) dt \quad (3)$$

Here $\alpha_{SD,i}$ is taken 0.4 for the first two phases, when SD was rather passive and 0.8 for the last phase with the proactive SD.

The influence of AMT agent j on the performance and safety efforts of its peer AMT agent i is defined by the following social contagion process:

$$\begin{aligned} dse_i &= \sum_j \alpha_{ji} \gamma_{ji} (se_j - se_i) dt \\ dpe_i &= \sum_j \alpha_{ji} \gamma_{ji} (pe_j - pe_i) dt \end{aligned} \quad (4)$$

Here α_{ji} is taken 0.7 in all the phases, reflecting intensive explicit and implicit interaction between the AMTs.

If by applying the equations (2)-(4) a point (pe, se) is obtained that lies above the maximum safety effort function from Figure 2, the closest point on the function is chosen. In such a way mutual inhibition of safety and performance goals is captured.

Furthermore, two additional constraints are defined to reflect that the old guard had used safety as a shield against the productivity push of MM in the first phase and did not deliver performance effort above a certain threshold (pt):

$$\begin{aligned} pe &< pt \\ se &> \text{maxsft}(pt) pe^2 / pt^2 \end{aligned} \quad (5)$$

where pt is performance threshold setting maximal performance the old guard can reach and maxsft is the maximum safety effort function defined above (Figure 3). For our model we have chosen $pt = 0.75$ for new AMTs corresponding to a high performance demand and $pt = 0.5$ for the old guard, which reflects their opposition to a high performance demand. Essentially, the safety effort of the old guard is determined by fitting a quadratic function between $[0, 0]$ and the point on the maxsft curve for the level of the performance threshold.

The frequencies of interaction of the agents are defined as follows. Interaction between AMTs occurs within a shift every 5 minutes and between shifts during shift changes. MM influences AMTs every 10 minutes in every phase, since MM function as supervisors of AMTs and are relatively close to them. The SD's influence on AMTs occurs once a shift in the first phase, five times a shift in the second phase and 15 times

a shift in the last phase. This represents the low, indirect participation of the safety department in daily operations at the initial phases towards a more active participation in the last phase.

5. Simulation Results

In this section we discuss analysis of the simulation results. In particular we show how the model can be validated through an interpretation of the emergent patterns regarding commitment to safety. Validation is here seen as the interpretation of generated simulation data and matching the pattern to the ethnographic description. This validation criterion states that if the simulations create patterns that match the empirical case, then the model must be to some extent correct. With alternative settings, we can check if the results are intuitively sensible and construct hypotheses that can be tested with newly gathered data. We therefore performed simulations using the model from Section 4 (discussed in Section 5.1) and several variants of this model (discussed in Section 5.2).

5.1 Simulation results of the proposed model

Based on the model presented in Section 4, 100 discrete simulation runs were performed with a static time step representing one minute of time in the real world. The simulation comprises three distinct phases for which most of the simulations parameters differ in accordance with the case description. In addition, there are two transition phases between the main phases and those parameters that vary between the phases are gradually changed in transition phases. The length of all the phases separately as well as the overall length match the lengths of periods with distinct safety cultural characteristics as identified in the case study.

All agents were put in an agent pool from which they were being selected for shifts. There were three shifts with equal length during one day. The ratio between the size of a shift and number of agents was constant during the whole simulation. Only in the third phase every shift was split into two which worked separately and had no influence on each other. This simulated that some of the AMTs had worked on line maintenance and some on base maintenance. During the whole simulation the number of agents entering and leaving the company is maintained as reported in the case study.

The developed model was validated by comparing the simulation results with the patterns found in the case study (Atak and Kingma, 2010). The simulation results for one random simulation run are represented by graphs of AMTs' performance and safety efforts (Figure 5) provided the MM's performance and SD's safety demands (Figures 1 and 2). The simulation results aggregated over all 100 simulation runs are provided in Figure 5. Commitment can be interpreted as the difference of management demand and AMT performance, but not in all situations. For example, when management demand is very low, equally low AMT performance does not necessarily reflect high commitment. We chose not to include another complex parameter for commitment to safety, but rather interpret commitment straight from the performance and demand graphs.

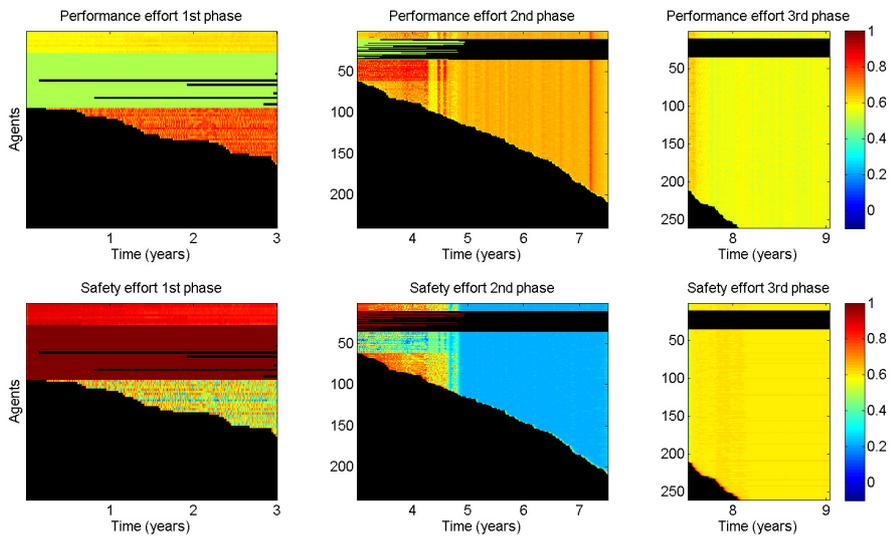


Figure 5. Visualization of performance and safety efforts of the AMTs during the three phases.

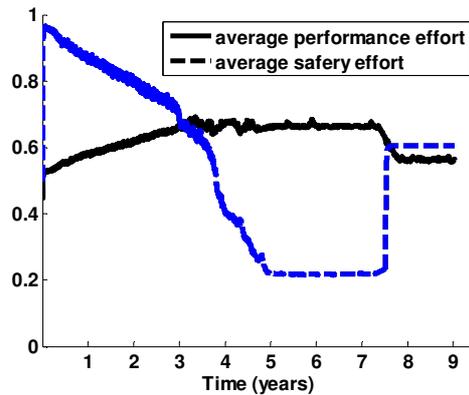


Figure 6. Performance and safety efforts in the three phases averaged over all agents in 100 simulation runs. Standard deviation at each point is less than 0.03.

The rows in the graphs in Figure 5 represent the state of a single agent's safety or performance effort. Efforts range from zero (blue) to one (red). Black corresponds to the value of -0.1, which is an empty row, i.e., the agent is not employed. In the graphs for the first phase, at time point 0, there are 35 agents present, a work force that gradually grows until the 8-year time point in the third set of graphs, where it reaches 260. Old guard AMTs occupy row 15 through 35 in Figure 5. They are gradually leaving the company starting from the first phase to roughly the first year of the second phase. The row representing their efforts abruptly turns black when an old guard AMT leaves, while empty rows below are gradually being filled with new AMTs.

In the first phase, conflicting management demands, resistance strategies and power differences between agents give rise to fluctuating, unevenly distributed safety efforts (Figure 5). This pattern corresponds to the fragmentation perspective,

highlighting ambiguity and inconsistent views amongst cultural members. Commitment to safety at this stage reflects the influence of the old guard's 'work-to-rule' resistance against management performance demands, by extensive compliance with safety rules (reflected by very high safety efforts in Figure 5). New AMTs are highly influenced by the more experienced old guard and therefore are greatly drawn in by this behavior. At the same time new AMTs experience a stronger pull from MM than the old guard does. They work harder than the old guard, compensating to some extent for the work they leave undone due to their resistance to MM (Figure 5).

As the core of this resistance—the old guard—leaves the company, it enters a second phase. Commitment in the first phase appears to have been of a superficial character. In the graph, old guard members leaving the company are represented as values that go to -1. As almost all members of the old guard have left, the safety effort graph shows how all the AMTs' safety efforts drop to dangerously low levels (Figure 5).

The transition between the first and second phase is initiated by changing power relations between MM and AMTs and increasing performance demands. The changes accompany the shifting mentalities and arrangements accommodating new, larger contracts (see section 4). Old guard AMTs continue to resist performance pressures, but new AMTs are less affected by them than by the increasing persuasiveness of MM. This reflects the observed exertion of power, where AMTs could not ignore the company's economic challenges and MM legitimized its attempts to pressure the work force by reminding technicians of this fact (Atak and Kingma, 2010, p. 271). As more and more new AMTs were recruited and the old guard AMTs' numbers decreased, the old guard AMTs was losing the struggle to resist. When only the last two members of the old guard remain in the company, safety efforts reach a turning point, drop and stabilize at a significantly lower level. Since the phasing out of the old guard is a relatively gradual process, the transition between phase one and two takes a relatively long period to complete.

The safety efforts observed in the second phase, fluctuating around 0.25, can be interpreted as AMTs regularly making shortcuts to speed up the work, also when the flight safety consequences of procedure shortcuts may not be known. Atak and Kingma (2010, p. 275) describe an event, where an incorrect aircraft part had been ordered and MM forces an AMT to install an old part, rather than wait for the correct new part. Another type of practice related to low values of safety could be acceptance of unhealthy working schedules, such as the reported overtime and double shifts that can lead to clouded judgment.

Consistently low safety efforts, significantly lower than demanded by the SD (0.8), can in the second phase be interpreted as revealing low commitments to safety. The uniform distribution of low levels of safety efforts indicates that practices such as those described above are seen as normal and legitimate throughout the organization. The ethnographic data reveal that 'a new discourse on safety which stressed a flexible and practical attitude' took hold of the organization at this time and was used to justify deviations from safety rules (Atak and Kingma, 2010, p. 272). AMTs' relatively low power towards MM, and MM's increasing persuasiveness to be flexible with safety standards, gives rise to an organizational culture integrated around production values. This can be observed in the performance efforts by AMTs responding in a uniform manner to fluctuations in managers' performance demands.

In the transition towards the third phase, managerial demands and power relations are harmonized. Both safety and performance demands are lowered, and the SD gains

power relative to MM, which loses some power (see section 4). The resultant power balance shifts in favor of the SD, signifying the shift towards a serious safety policy by a new CEO, in a more comfortable, mature market position. Safety levels quickly rise to acceptable levels, reflecting the observation that the influence of the new policy was felt ‘from the very first day’ (Atak and Kingma, 2010, p. 272). Performance efforts remain at a higher level slightly longer, accounting for the work that had accumulated under the previous, higher performance demands and takes some time to complete.

The new balance of power and work pressure allows AMTs’ safety efforts to stabilize, yielding an impression of strong, internalized commitments to safety. Safety and performance pressures no longer impose conflicting demands on AMTs, evident when one compares the demand values to the maximum efforts represented in Figure 5. Safety efforts are uniformly distributed amongst the 260 AMTs at a level of approximately 0.6, meaning more than-average concern with safety where standards are normally met and exceptions coordinated. Safety efforts hardly fluctuate, even when performance efforts do (approximately between 0.5-0.55). This stabilization and uniform distribution signifies a professional subculture of AMTs where safety practices are shared and technicians are given leeway to exercise discretion over safety-critical decisions. Production pressures exerted by management still fluctuate, but hardly perturb the uniformly distributed safety efforts.

To summarize, under the social and organizational influences an emergent pattern of commitment unfolds in the three distinct phases of organizational development. Commitment is first superficial, owing to the old guard’s resistance in the context of a fragmented culture in the survival phase. The weakness of commitment shows up in the second phase, when the organizational culture integrates around production values. In the final, maturation phase, commitment takes on the stable character. Such commitment is expected of a more autonomous professional culture in a diversified cultural context where competing core goals coexist in relative harmony.

5.2 Results of ‘what-if’ simulations of model variants

To test the usability and scope of the model, we simulated three variants of the original model from Section 4. These variants are each realistic scenarios that could have an interesting impact on the patterns of commitment. The results of these simulations are considered in this section and compared to the simulation results of the original model.

Model variant 1- Weak opposition culture

In the first model variant, in contrast to the original model, old guard represents a loosely connected group, which does not oppose MM strongly. This variant is interesting because it helps illuminate the role of the old guard in the development of the culture. The following changes were made in the original model:

- influence of the old guard on themselves is [0.1, 0.4] (was [0.8, 0.9]);
- performance threshold (*pt*) of the old guard is 0.75 (was 0.5).

Using this model variant 10 simulation runs were executed. The performance and safety efforts averaged over all the agents and the simulation runs of model variant 1 were compared with the simulation outcomes of the original model (Figure 7) by using paired sample t-test with 5% significance level.

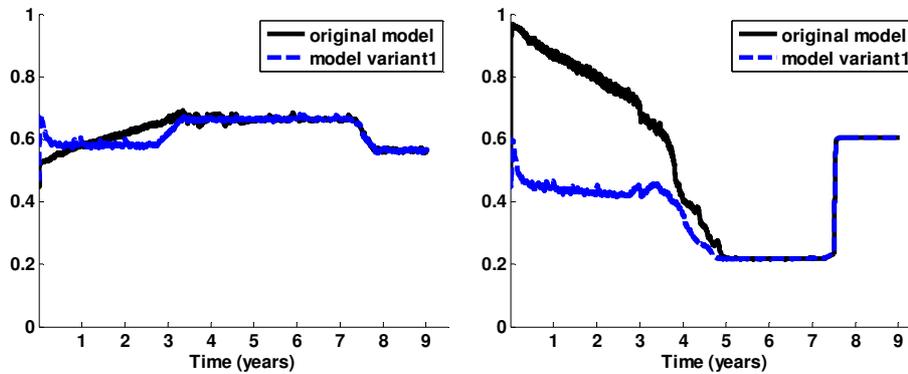


Figure 7. Performance efforts (left) and safety efforts (right) in the three phases averaged over all agents in 10 simulation runs for the original model and model variant 1. Standard deviation at each point is less than 0.05.

In phases 1 and 2 significant differences between the models were established both for the performance and safety demands. The null hypothesis of the paired t-test was supported for the phase 3.

The safety effort of the AMTs in phase 1 of the model variant 1 is significantly lower than in the original model. This can be explained by a low opposition of the old guard and a low degree of their mutual reinforcement. Furthermore, it can be observed that because of the loose relations in the shifts in model variant 1, the performance effort is not reinforced by the team members and remains relatively constant throughout each phase. In the original model the performance effort grows steadily in the first phase. In general, tight relations between agents and mutual reinforcement of their states may have positive, as well as negative amplification effects.

After the transition period the behavior of AMTs stabilizes in the second phase in both models. These results show that a resistant old guard may not only have a positive role to play in terms of safety, but also have a less detrimental effect on production than it might seem.

Model variant 2 – Conflict between safety and performance demands

In the second model variant, in contrast to the original model, SD has a high influence on the AMTs during the first and second phase. However, the performance and safety demands are both high and not harmonized, as in the third phase. This is theoretically interesting because it may help to corroborate the managerial insight that pressing on performance is not always helpful.

The following changes were made in the original model:

- the influence of SD over both new AMTs and the old guard is [0.7, 0.9]

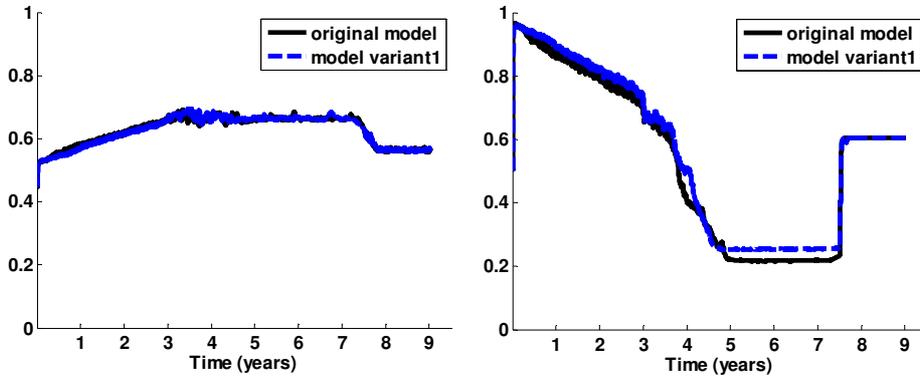


Figure 8. Performance efforts (left) and safety efforts (right) in the three phases averaged over all agents in 10 simulation runs for the original model and model variant 2. Standard deviation at each point is less than 0.05.

The paired t-test showed the models have significant differences w.r.t. the performance effort in the second phase and w.r.t. the safety effort in the first and the second phase. It means that the SD's impact on the AMTs is statistically significant, however is not very high in terms of the magnitude (Figure 8). One can conclude that enforcing high safety and performance demands at the same time without their mutual adjustment does not actually help to improve safety.

Model variant 3 – Random composition of shifts

In the third model variant a random composition of the shifts was introduced by a random permutation of the agent pool every time when all the agents were used for the shift composition. This is an interesting variation because it is a strategy used by airline companies, for example, to prevent negative group dynamics in flight crews.

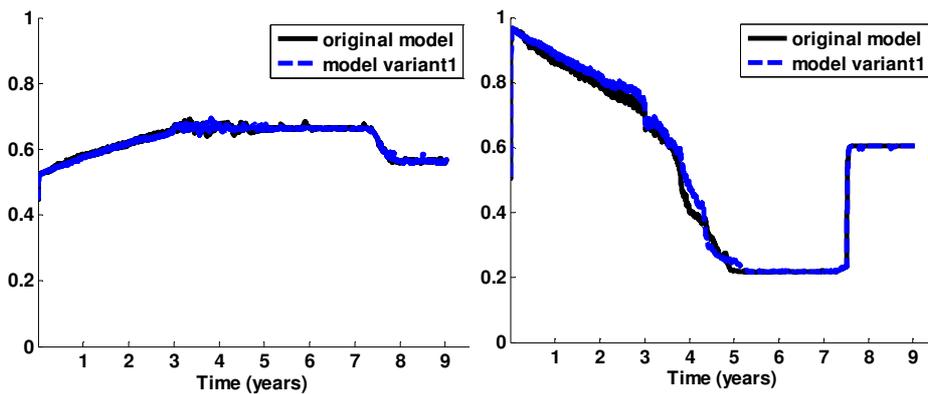


Figure 9. Performance efforts (left) and safety efforts (right) in the three phases averaged over all agents in 10 simulation runs for the original model and model variant 3. Standard deviation at each point is less than 0.05.

The paired t-test showed that there were no significant differences between the models w.r.t. the performance effort. However, there were significant differences w.r.t. the safety effort in the first and the second phases. The safety effort was on average higher in model variant 3, meaning that shift rotation might have a positive effect on safety. However the magnitude of this effect in the simulation is not very high, because the shifts did not differ much from the beginning. The more differences the shifts have, the higher will be the effect. In the future more variations between the shifts will be explored.

6. Conclusions and Discussion

This study presents an innovative approach that can be used to address current questions regarding safety culture that attempt to get beyond linear, static approaches and that take a complex, political view on culture more seriously. Our paper showcased an example from the aircraft maintenance domain addressing the research question how the commitment to safety of maintenance technicians emerges and develops under social and organizational influences. Commitment to safety is a key aspect of organizational safety culture, but current survey-based methods offer limited insight in underlying mechanisms of development of such cultural properties. We propose an interdisciplinary approach, combining organizational ethnography with formal agent modeling and simulation experiments. As a demonstration of this approach, we developed a formal model from the ethnographic case study by Atak and Kingma (2010). We used this model to simulate, bottom-up, with a generic set of mechanisms operating throughout the simulation, the observed emergent global pattern of commitment to safety of aircraft maintenance technicians, developing over a period of 9 years. Theories of social power and influence were used to derive a formal model and supplementary data were gathered through interviews and other studies on aviation maintenance operations.

6.1. Value and applicability of the approach

The merit of the approach is in offering insight in relations between social and organizational power mechanisms and emergent characteristics of safety culture, not in deriving precise numerical values. Formal modeling adds to ethnographic research that it forces researchers to make theoretical connections explicit between social relations and cultural dynamics. Ethnographic accounts will tend to remain more contextualized. The ethnographer often starts from a participatory insiders' perspective. The agent-based modeler, on the contrary, builds a model of a sociotechnical system from a complex systems perspective by identifying agents, their various local properties and interactions, taking diverse feedback mechanisms into account. By taking such a systemic approach to model development, the agent-based modeler is able to provide a feedback to the ethnographer to focus their data collection efforts, observations, and explorations on particular aspects of a sociotechnical system represented by parameters, processes, and mechanisms in the model. However, it is not necessary and even not desirable that the activities of the ethnographer are fully steered by the modeler. By keeping a more open attitude to data collection, the ethnographer may discover important aspects not considered by the modeler, which however might be important to be reflected in the model. Based on an agent-based model, diverse forms of analysis

could be performed, such as sensitivity analysis and bias and uncertainty analysis. Through such an analysis important parameters, assumptions, and uncertainties in the model can be identified, which would require refinement and further identification by ethnographic research. Thus, the systematic agent-based model development and ethnographic research should take place in interaction with each other, complementing each other in an integrated methodology.

In contrast to modeling social relations with network properties like homophily (Holzhauer, Krebs & Ernst, 2013), our model uses power influencing. This is suitable for safety culture studies because the bases of power that each agent experiences are anchored in rich qualitative context descriptions. With this approach to social relations, formalization helps to explicate contextual processes that are important in shaping a 'strong safety culture'. We thus capture a more fragile and dynamic reality than standard quantitative safety culture research methods does (Guldenmund, 2000; 2007).

The context-sensitive modeling approach may also provide managerial insights. One example is our demonstration of how lowering and harmonizing performance and safety demands can lead to a significant increase in AMTs' safety efforts, with only a small reduction of performance efforts.

6.2. Limitations

Although the tendency may exist to view multi-agent models of safety cultural phenomena as predictive, we do not claim this: the current model is explanatory. Since cultural phenomena are context-specific, the transferability of any model depends on the extent to which the dynamics it captures are recognizable elsewhere. This requires attention to any specific context of application and likely many adaptations of the model. We draw attention to transferability as a limitation, but it may also be seen as a strength. Formal modeling is a powerful tool to make assumptions explicit and theorize back and forth between what is happening in a particular setting and what we know more generally about safety culture.

That said, the current model remains relatively high level, does not include many possible feed-back interactions, and hard-codes some influences that are in reality dynamic effects. One interesting possibility is the impact of critical incidents. When commitment to safety decreases overall, safety incidents may occur and this may give voice to maintenance personnel and safety department officials to prioritize safety. The current secondary data source was not detailed enough to find empirical evidence for such feedback loops, showing where the current model could benefit from further research.

6.3. Further research

The model presented here invites improvement and extension while also inspiring further theoretical development and ethnographic research. In order to agree with the contextualizing tendency of ethnographic research, more detailed data may be gathered on the precise ways in which safety and production demands are weighed. These data can lead to more adequate models and yield interesting theoretical insights. The translation from conceptual to formal models may also be automated as described by Fuentes-Fernández, Hassan, Pavón, Galán and López-Paradez (2012).

The rewards of this approach are promising. There is a persistent—but often inappropriate—assumption in both theory and practice that safety behavior can be

regulated by mainly enforcing compliance with safety regulations (Dekker, 2014). Fundamental insights of how safety is created through culture are however still spread thin (Antonsen, 2009; Edwards, Davey and Armstrong, 2013; Woods, Dekker, Cook, Johannesen and Sarter, 2010). The current approach attempts to escape this deadlock and move towards more refined thinking about the human contribution to safety (Reason, 2008).

References

- Antonsen S (2009) *Safety culture: theory, method and improvement*. Ashgate, Farnham, UK
- Atak A, Kingma S (2010) Safety culture in an aircraft maintenance organisation: A view from the inside. *Safety Science*, 49: 268-278
- Bosse T, Jonker CM, Meij L van der, Sharpanskykh A, Treur J (2009) Specification and Verification of Dynamics in Agent Models. *International Journal of Cooperative Information Systems*, 18(1): 167-193
- Burton, RM (2003) Computational laboratories for organization science: Questions, validity and docking. *Computational & Mathematical Organization Theory*, 9(2): 91-108
- Castelfranchi, C, Micelli, M, Cesta, A (1992) Dependence relations among autonomous agents. In Werner E, Demazeau, Y (eds), *Decentralized A. I. 3*, pages 215-227. Elsevier Science Publishers B.V., Amsterdam, NL
- Choudhry RM, Fang D, Mohamed S (2007) The nature of safety culture: A survey of the state-of-the-art. *Safety Science*, 45: 993-1012
- Deci, E, Ryan, R (eds) (2002) *Handbook of self-determination research*. Rochester, NY: University of Rochester Press.
- Deffuant G, Neau D, Amblard F, Weisbuch G (2000) Mixing beliefs among interacting agents. *Adv Complex Syst*, 3: 87-98
- Dekker, SWA (2014) The bureaucratization of safety. *Safety Science*, 70: 348–357
- Doloriert, C, Sambrook, S (2012) Organisational autoethnography. *Journal of Organizational Ethnography*, 1(1): 83–95
- Edwards, JRD, Davey, J, Armstrong, K (2013) Returning to the roots of culture: A review and re-conceptualisation of safety culture. *Safety Science*, 55: 70–80
- Fayard, AL, Van Maanen, J (2015) Making culture visible: reflections on corporate ethnography. *Journal of Organizational Ethnography*, 4(1): 4–27
- Fuentes-Fernández, R, Hassan, S, Pavón, J, Galán, JM, López-Paredes, A (2012) Metamodels for role-driven agent-based modelling. *Computational and Mathematical Organization Theory*, 18(1): 91-112
- Geertz C (1973) *The Interpretation of Cultures: Selected Essays*. Basic Books, New York

- Giorgi, S, Lockwood, C, Glynn, MA (2015) The many faces of culture: Making sense of 30 years of research on culture in organization studies. *Academy of Management Annals*, 9(1): 1-54
- Guldenmund FW (2000) The nature of safety culture: a review of theory and research. *Safety Science*, 34(1): 215-257
- Guldenmund FW (2007) The use of questionnaires in safety culture research—an evaluation. *Safety Science*, 45(6): 723-743
- Guldenmund, FW (2010) (Mis)understanding Safety Culture and Its Relationship to Safety Management. *Risk Analysis*, 30(10): 1466–1480
- Hale, A, Borys, D (2013) Working to rule, or working safely? Part 1: A state of the art review. *Safety Science*, 55: 207–221
- Harrison JR, Lin Z, Carroll GR, Carley KM (2007) Simulation modeling in organizational and management research. *Academy of Management Review*, 32(4): 1229-1245
- HILAS (2007) Deliverable HILAS Del 4.16: In-field Research and Maintenance Process Mapping
- Hollnagel E, Woods DD, Leveson N (2006) Resilience engineering: Concepts and precepts. Ashgate, London
- Holzhauser, S, Krebs, F, Ernst, A (2013) Considering baseline homophily when generating spatial social networks for agent-based modelling. *Computational and Mathematical Organization Theory*, 19(2): 128-150
- Hopkins A (2006) Studying organizational cultures and their effects on safety. *Safety Science*, 44: 875-889
- Jeffcott, S, Pidgeon, N, Weyman, A, Walls, J (2006) Risk, trust, and safety culture in UK train operating companies. *Risk analysis*, 26(5): 1105-1121
- Kunda, G (2009) Engineering culture: Control and commitment in a high-tech corporation. Temple University Press
- Latane B (1981) The psychology of social impact. *American Psychologist*, 36: 343-356
- Lee, T, Harrison, K (2000) Assessing safety culture in nuclear power stations. *Safety science*, 34(1): 61-97
- Lukes, S (1974) Power: A radical view (Vol. 1). Macmillan, London
- Martin J (2002) Organizational culture: Mapping the terrain. Sage, London
- Mearns, K, Kirwan, B, Reader, TW, Jackson, J, Kennedy, R, Gordon, R (2013) Development of a methodology for understanding and enhancing safety culture in Air Traffic Management. *Safety science*, 53: 123-133
- Pettersen, KA (2008) The Social Production of Safety, Theorising the Human Role in Aircraft Line Maintenance. Dissertation, University of Illinois
- Raven, BH (1992) The Bases of Power: Origins and Recent Developments. *Journal of Social Issues*, 12(49): 227 – 251

- Reason JT (1990) *Human Error*. Cambridge University Press, Cambridge
- Reason JT (1997) *Managing the risk of organizational accidents*. Ashgate, Farnham, UK
- Reason JT (2008) *The human contribution: unsafe acts, accidents and heroic recoveries*. Ashgate, London
- Schein, EH (1990) Organizational culture. *American Psychologist*, 45(2): 109–119
- Sharpanskykh, A (2008) *On Computer-Aided Methods for Modeling and Analysis of Organizations*. Dissertation, Vrije Universiteit Amsterdam
- Van Avermaete, JA, Hakkeling-Mesland, MY (2001) Maintenance human factors from a European research perspective: results from the ADAMS project and related research initiatives. National Aerospace Laboratory (NLR)
- Ward, M (2006) *Contributions to Human Factors from three case studies in aircraft maintenance*. Dissertation, Trinity College Dublin
- Weiss, G (ed) (1999) *Multiagent systems: a modern approach to distributed artificial intelligence*. MIT press, Cambridge, MA
- Woods DD, Dekker S, Cook R, Johannesen L, Sarter N (2010) *Behind Human Error* (2nd ed). Ashgate, Farnham, UK

Appendix A. Model docking

To assess the model's expressiveness, suitability, and innovativeness, we relate it to another important modeling approach to power. Power relations between artificial agents have also been investigated in the area of multi-agent systems. Castelfranchi, Micelli and Cesta (1992) introduce different types of social dependence relations between agents, based on which social influence and power are defined. Power relations are viewed in terms of resource dependence: an agent gains power when it controls resources that another agent needs.

In Burton (2003) different types of model relations are discussed, also called model docking. Such a docking may be done at the *distributional* and *relational* levels. At the distributional level the models require to produce the same (numerical) outcomes, whereas at the relational level internal components and dynamics of the models are related. In the following it is described how the models can be related at both these levels. Note that although the social dependence relations introduced in Castelfranchi, Micelli and Cesta (1992) are qualitative, it is indicated in the same paper that the specification language may be extended by incorporating quantitative degrees of dependence between agents. To enable comparison of the models at the distributional level, a three-valued scale for measuring the degree of social dependence and power of influence is introduced: 'h' corresponding to a high degree, 'm' – medium and 'l' – low.

First we demonstrate how the models can be related at the relational level.

We start by introducing acts and goals w.r.t. which social dependencies are defined. In Table A1 notations for the agents from the case study are provided which will be used for specifying social dependence relations. The goals of AMTs provided in Table A2 were identified in accordance with the self-determination theory (Deci & Ryan,

2002), which is generally accepted as a sound empirical theory of human motivation. According to this theory, humans have several types of needs, among which *competence need* is related to seeking to control the outcome and experience mastery, *relatedness* is the need to interact, be connected to, and experience caring for others, and *autonomy* is the need to be causal agent of one's own life. In addition to the goals derived from self-determination theory, g_1 through g_3 , we also listed goals g_4 through g_8 , which are derived from empirical data, i.e. the case study. Based on the case study description, in relation to the goals a set of acts is identified listed in Table A3.

x_1	new AMTs
x_2	old guard
x_3	MM
x_4	SD

Table A1. Notations for the agents from the case study

g_1	competence-related goals
g_2	relatedness-related goals
g_3	autonomy-related goals
g_4	Personal/psychological safety & security-related goals
g_5	achieve high performance
g_6	achieve high compliance to safety regulations
g_7	achieve high performance without compromising safety
g_8	avoid compromising safety

Table A2. The goals of the agents from the case study w.r.t. which dependence relations are defined

a_1	promote an AMT
a_2	hire an AMT with a short term contract
a_3	refrain from providing reprimands
a_4	ensure safety regulations are being observed by an AMT
a_5	approval of behavior of an AMT
a_6	allow for AMT professional discretion

Table A3. The acts of the agents from the case study w.r.t. which dependence relations are defined

In Table A4 social dependence relations are provided, which were identified for the first phase of the organizational development in the case study; c.f. the influence relations in Fig.4. Each social dependence relation (*S-DEP* $x y a g d$) reads as: x depends on y to degree d with respect to act a useful for realizing x 's goal g . Note that the original S-DEP relation from Castelfranchi, Micelli and Cesta (1992) was extended with the fifth argument – the degree of dependence, as was discussed previously.

To define social dependence relations, the influence relations from the model introduced in this paper (Fig.4) were taken as the basis. For each influence relation the power bases explained in Section 4.2 were related to particular goals and acts of agents. For example, the influence relation of MM on new AMT in the first phase has a strong legitimate basis to control promotions of AMTs, to re-hire AMTs with a short term contract and to provide reprimands. Both promotions and re-hiring of AMTs contribute positively to the satisfaction of their competence-related goals, i.e., they can be seen a form of professional recognition. Also refraining from providing reprimands by MM serves as a positive feedback contributing to the satisfaction of the AMTs' competence-

related goals. Furthermore, re-hiring of AMTs, as well as refraining from providing of reprimands contribute positively to the AMTs' personal safety & security-related goals. Since MM has a strong legitimate power basis on new AMTs, all these social dependence relations have a high degree (*h*) (Table A4, first column). Other social dependence relations in Tables A4 and A5 are defined in a similar manner.

<i>on MM</i>	<i>on SD</i>	<i>on old guard</i>	<i>on new AMTs</i>
(<i>S-DEP</i> $x_1 x_3 a_1 g_1 h$)	(<i>S-DEP</i> $x_1 x_4 a_4 g_1 l$) (<i>S-DEP</i> $x_1 x_4 a_4 g_4 l$)	(<i>S-DEP</i> $x_1 x_2 a_4 g_2 h$) (<i>S-DEP</i> $x_1 x_2 a_5 g_2 h$) (<i>S-DEP</i> $x_1 x_2 a_5 g_1 h$)	(<i>S-DEP</i> $x_1 x_1 a_5 g_2 l$)
(<i>S-DEP</i> $x_1 x_3 a_2 g_1 h$)			
(<i>S-DEP</i> $x_1 x_3 a_2 g_4 h$)			
(<i>S-DEP</i> $x_1 x_3 a_3 g_1 h$)			
(<i>S-DEP</i> $x_1 x_3 a_3 g_4 h$)			

Table A4. Social dependence relations of new AMTs in the first phase

<i>on MM</i>	<i>on SD</i>	<i>on old guard</i>	<i>on new AMTs</i>
(<i>S-DEP</i> $x_2 x_3 a_3 g_4 l$)	(<i>S-DEP</i> $x_2 x_4 a_4 g_4 l$)	(<i>S-DEP</i> $x_2 x_2 a_4 g_2 h$) (<i>S-DEP</i> $x_2 x_2 a_4 g_3 h$) (<i>S-DEP</i> $x_2 x_2 a_4 g_4 h$)	(<i>S-DEP</i> $x_2 x_1 a_5 g_2 l$)

Table A5. Social dependence relations of old guard in the first phase

As stated in Castelfranchi, Micelli and Cesta (1992: 10), the power of influencing is derivable from dependence relations: "If *x* is dependent on *y*'s performing a certain act in view of *p*, *y* is quite likely to have the power of influencing *x* relative to some other goal of *x*'s.". The new AMTs have a high dependence on MM and old guard (Table 4). In the case study MM uses this dependence to influence new AMTs to adopt the goal of achieving a high performance:

$$H_{i=1,2,3} (INFL-POWER \ x_3 x_1 a_i g_5 h)$$

(*INFL-POWER* $x y a g v$) relation is read as *x* has the power of influencing *y* of degree *v* if *x* can do such an act *a* that makes *y* have *g* as a goal of her own. To express power of influencing relation over multiple acts *H*-relation is used in Castelfranchi, Micelli and Cesta (1992).

In the case study old guard use the new AMTs' high dependence on them to influence the new AMTs to adopt the goal of achieving high compliance to safety regulations (i.e., opposing high performance demands):

$$H_{i=4,5} (INFL-POWER \ x_2 x_1 a_i g_6 h)$$

Since the social dependence relations of the new AMTs on SD and other new AMTs are of a low degree, they do not form a sufficient basis for power influencing.

As was discussed in Section 4.1 and illustrated in Fig.3 goals *g*₅ and *g*₆ may be in conflict with each other, especially when demands are high. Since both identified power of influencing relations have a high degree, new AMTs are pulled strongly by MM and old guard in the opposite directions. This results in a fragmented culture, which is in line with the outcomes of the model proposed in this paper.

The old guard has high mutual social dependencies, which unite this group even stronger in a self-reinforcing manner:

$$G_{i=2,3,4} (INFL-POWER \ x_2 x_2 a_4 g_i h)$$

The distinct behavior of old guard – high compliance to safety and opposition to high performance demand - can be also clearly seen in the results of the model proposed in this paper (Fig. 5).

The social dependence relations for the second phase are provided in Table A6. As management gains referent power over new AMTs in this phase, a new relation is added (*S-DEP* $x_1 x_3 a_5 g_2 h$). Furthermore, the dependence of the new AMTs on the old guard diminishes in this phase, which is reflected in the degrees of the relations in the third column. Also the mutual dependencies between the members of old guard become weaker, which is reflected in the third column of Table A7.

<i>on MM</i>	<i>on SD</i>	<i>on old guard</i>	<i>on new AMTs</i>
(<i>S-DEP</i> $x_1 x_3 a_1 g_1 h$)	(<i>S-DEP</i> $x_1 x_4 a_4 g_1 l$) (<i>S-DEP</i> $x_1 x_4 a_4 g_4 l$)	(<i>S-DEP</i> $x_1 x_2 a_4 g_2$ <i>l/m</i>)	(<i>S-DEP</i> $x_1 x_1 a_5 g_2 l$)
(<i>S-DEP</i> $x_1 x_3 a_2 g_1 h$)		(<i>S-DEP</i> $x_1 x_2 a_5 g_2$ <i>l/m</i>)	
(<i>S-DEP</i> $x_1 x_3 a_2 g_4 h$)		(<i>S-DEP</i> $x_1 x_2 a_5 g_1$ <i>l/m</i>)	
(<i>S-DEP</i> $x_1 x_3 a_3 g_1 h$)			
(<i>S-DEP</i> $x_1 x_3 a_3 g_4 h$)			
(<i>S-DEP</i> $x_1 x_3 a_5 g_2 h$)			

Table A6. Social dependence relations of new AMTs in the second phase

<i>on MM</i>	<i>on SD</i>	<i>on old guard</i>	<i>on new AMTs</i>
(<i>S-DEP</i> $x_2 x_3 a_3 g_4 l$)	(<i>S-DEP</i> $x_2 x_4 a_4 g_4 l$)	(<i>S-DEP</i> $x_2 x_2 a_4 g_2$ <i>l/m</i>)	(<i>S-DEP</i> $x_2 x_1 a_5 g_2 l$)
		(<i>S-DEP</i> $x_2 x_2 a_4 g_3$ <i>l/m</i>)	
		(<i>S-DEP</i> $x_2 x_2 a_4 g_4$ <i>l/m</i>)	

Table A7. Social dependence relations of old guard in the second phase

Thus, the following power influencing relations can be inferred in phase 2:

$$H_{i=1,3,5} (INFL-POWER \ x_3 x_1 a_i g_5 h)$$

$$(INFL-POWER \ x_2 x_3 a_4 g_6 l/m)$$

$$G_{i=2,3,4} (INFL-POWER \ x_2 x_2 a_4 g_i l/m)$$

MM clearly has the highest influencing power over new AMTs in this phase, and thus makes them to adopt goal g_5 . This is also in accordance with the results of the model proposed in this paper.

In the third phase the basis of legitimate power of MM becomes weaker and management pressure on production goals is not dominant anymore. This is reflected in the decreased degrees of dependencies in the first column of Table A8. Furthermore, SD becomes more proactive and gains a basis for coercive and punishment power. Thus, relation (*S-DEP* $x_1 x_4 a_3 g_4 m-h$) is added in the second column of Table A8. Both MM and SD gain power by allowing AMTs to make more professional judgment calls.

This contributes to satisfaction of their autonomy goals g_3 and heightens their motivation to take responsibility for the work, the demands for which are now within reasonable limits. For example, MM gives AMTs more influence in creating work packages and the SD acknowledges that not all procedure violations are unsafe, if there is sound and knowledgeable professional judgment. The strength is low-medium because this power derives only from legitimacy and persuasion power bases. Thus, the relations $(S-DEP\ x_1\ x_3\ a_6\ g_3\ l-m)$ and $(S-DEP\ x_1\ x_4\ a_6\ g_3\ l-m)$ are added.

<i>on MM</i>	<i>on SD</i>	<i>on new AMTs</i>
$(S-DEP\ x_1\ x_3\ a_1\ g_1\ m-h)$	$(S-DEP\ x_1\ x_4\ a_4\ g_1\ m-h)$	$(S-DEP\ x_1\ x_1\ a_5\ g_2\ l)$
$(S-DEP\ x_1\ x_3\ a_2\ g_1\ m-h)$	$(S-DEP\ x_1\ x_4\ a_4\ g_4\ m-h)$	
$(S-DEP\ x_1\ x_3\ a_2\ g_4\ m-h)$	$(S-DEP\ x_1\ x_4\ a_3\ g_4\ m-h)$	
$(S-DEP\ x_1\ x_3\ a_3\ g_1\ m)$	$(S-DEP\ x_1\ x_4\ a_6\ g_3\ l-m)$	
$(S-DEP\ x_1\ x_3\ a_3\ g_4\ m)$		
$(S-DEP\ x_1\ x_3\ a_6\ g_3\ l-m)$		

Table A8. Social dependence relations of new AMTs in the third phase

Based on the social dependence relations from Table A8 the following power of influencing relations can be inferred:

$$H_{i=1,3,6} (INFL-POWER\ x_3\ x_1\ a_i\ g_7\ m-h)$$

$$H_{i=3,4,6} (INFL-POWER\ x_4\ x_1\ a_i\ g_8\ m-h)$$

New g_7 and g_8 goals reflect harmonization of the performance and safety demands in the organization in the third phase. Note that both power of influencing relations have $m-h$ degree, which gives AMTs some freedom to perform operations in their own way.

The compared models can thus be formally related to each other, providing a test of our model's expressiveness and appropriateness for the topic of study. In the other model, social dependence relations are defined at a more detailed level than the power relations in our model. Such a level of detail is not deemed to be necessary for the application considered in this paper, given the nature of the secondary data that we referred to. Furthermore, social dependence relations are formalized in the other model using a qualitative predicate logic-based language, which is more suited to specify agent states and is less applicable for describing dynamics of processes. In contrast, our model is quantitative and continuous; it takes a dynamic view on the evolution of power relations in an organization. To address the current limitations of safety cultural analyses and theorizing, our approach appears more appropriate.

Appendix B. Sensitivity analysis for the function $maxsft$

In this section, the sensitivity of the patterns of commitment produced by the model to the parameters of the function $maxsft(pe)=1-1/(e^{-w1 \cdot pe + w2})$ are discussed.

For $w1=15$ and $w2=10$ and cpp close to 0.5 the function has the shape:

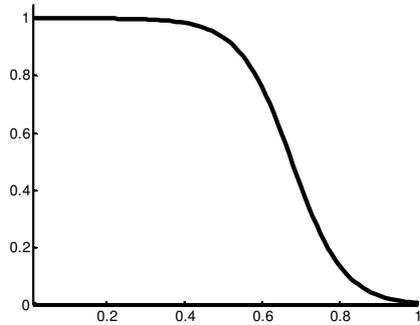


Figure B1. Change of the maximum safety effort (vertical axis) depending on the performance effort (horizontal axis) with parameters $w1=15$ and $w2=10$.

The corresponding patterns of performance and safety efforts of the AMTs for the three phases of organizational development are provided below. The agents that left the organization are indicated by dark blue.

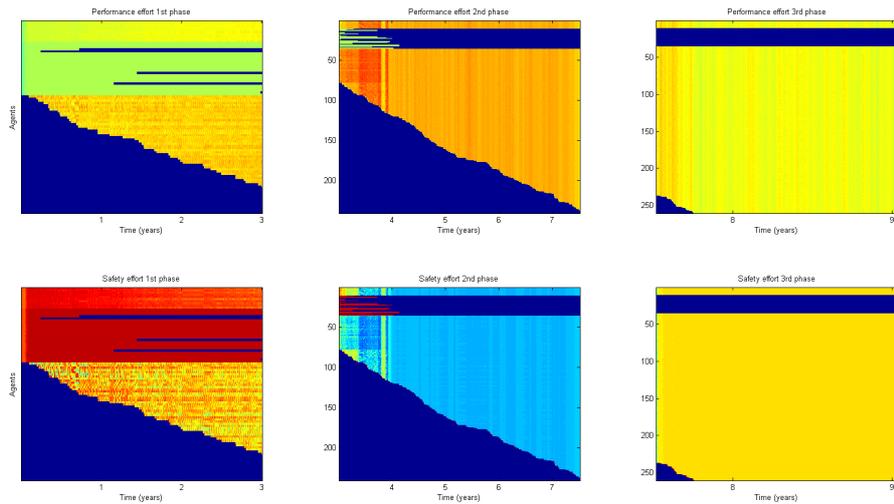


Figure B2. Visualization of performance and safety efforts of the AMTs during the three phases with parameters $w1=15$ and $w2=10$.

The essential validation findings discussed in Section 5 are also reflected in Figure B2. In the first phase commitment to safety is superficial, when safety efforts are the result of resistance of the old guard to high performance demands. Notice the difference between the high, almost uniform old guard AMTs' safety efforts (hues of red) and the new AMTs' safety efforts fluctuating under the competing influences of the management and the old guard. In the second phase, when the source of resistance is removed, safety efforts plummet across the entire workforce. In the third case, commitment becomes a stable, professional trait and is hardly impacted by the day-to-day pressures.

For $w1=35$ and $w2=30$ and c_{pp} close to 0.8 the function has the shape:

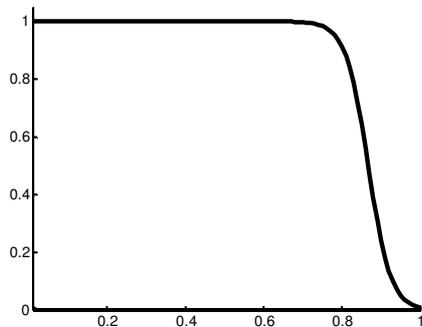


Figure B3. Change of the maximum safety effort (vertical axis) depending on the performance effort (horizontal axis) with parameters $w1=35$ and $w2=30$.

The corresponding patterns of performance and safety efforts of the AMTs for the three phases of organizational development are provided below.

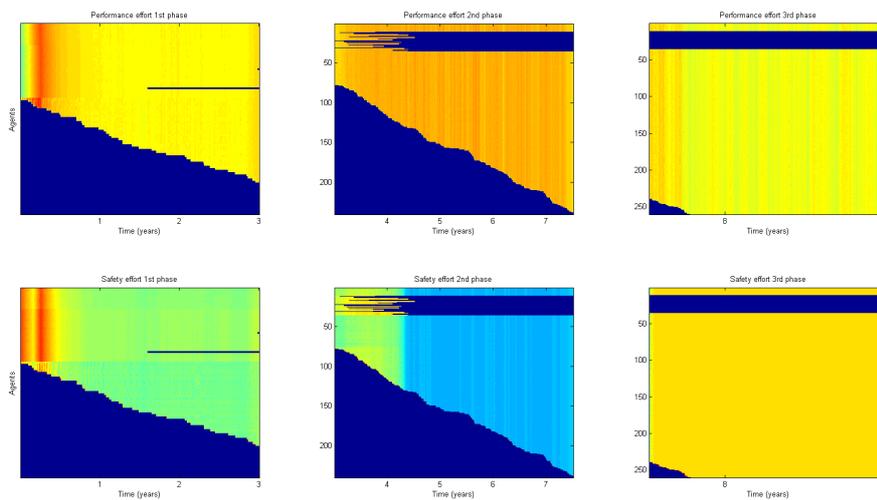


Figure B4. Visualization of performance and safety efforts of the AMTs during the three phases with parameters $w1=35$ and $w2=30$.

Also in this case the essential validation findings discussed in Section 5 can be observed in the emergent dynamics (Figure B4).

Appendix C. Model robustness

The robustness of the model was evaluated by running 200 Monte Carlo simulations, in which all essential parameters of the model were drawn from the uniformly distributed intervals, representing possible variations of agent attitudes in the maintenance organization being modeled. These intervals are defined as follows:

- the initial safety efforts of the new AMTs: [0.5, 0.8] (medium-high values);
- the initial safety efforts of the old guard: [0.7, 1] (high values);
- the initial performance efforts of the new AMTs [0.5, 0.8] (medium-high values);

- the initial performance efforts of the old guard [0.4, 0.7] (medium values);
- the standard deviation of parameter α is 0.1 in all phases of the social contagion model described in Section 4.3;
- the intervals of the influence relations between the agents were specified as in Figure 4.

The patterns of commitment to safety of the AMTs identified in Section 5 can be seen in every Monte Carlo simulation run (Figure C1). The standard deviation is the highest in the transition period from phase 1 to phase 2, which involves many organizational changes: orientation towards production goals, increase of the MM's influence on the AMTs and marginalization of the old guard. Thus, in general the patterns produced by the model are robust w.r.t. different parameter settings.

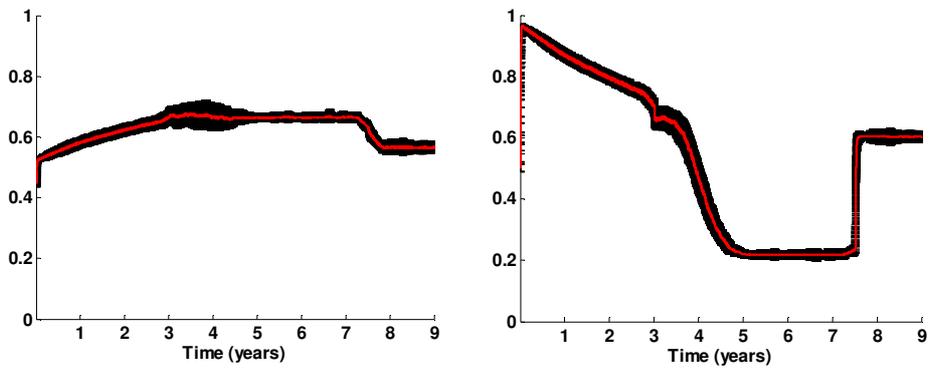


Figure C1. Mean (in red) and standard deviation (in black) of performance effort (left) and safety effort (right) calculated over all agents and 200 simulation runs.