Chapter 10

The space where workers manage risk

Hazards are fluid, always in motion. Hazards are never fixed, frozen or stable. The things that hurt people are almost never the things we identify, because hazards appear and disappear as the worker does work. In reality the hazards that workers "manage-around" are constantly changing based upon the context and conditions of the work itself. Workers must adapt to the changing work conditions they face while executing work.

One of the major challenges our industry faces is the idea that risk is somehow fixed in time and space, permanent. We have designed processes that actively try to identify risk as a part of event prevention. We identify all hazards in the work planning stage and then mitigate the hazards in order to keep the worker away from the harm. Most of our tools to manage hazards assume the idea that hazards are fixed. Ask any of your workers and they will tell you the idea that once you identify the hazard you are done managing that hazard is not realistic. Hazard identification is a constant process, not a pre-job task.

Not all risk is the same just as not all work is the same for the worker; workers manage hazards as they appear in real time and as work happens. This process is fluid and almost refuses to be proceduralized or formalized. Hazards require both awareness and adaptability, and workers must tailor their work activities for these constantly moving targets – or maybe the workers are the targets for the hazards – either way the idea that work is ever "as planned" is a brief and fleeting notion to be sure.

I started to look at how workers successfully manage hazard--you know what I mean--observing successful work, work where there is no serious event. I started to notice that the workers don't really get the opportunity to manage the hazard like our planning assumes. What the workers seem to manage is the space between the actual work and the multiple hazards. This seemed interesting to me because it is counter to what we assume will happen when we plan work.

Remember how we discussed risk earlier in this chapter--risk is the degree to which the worker is faced with operational uncertainty. One way to reduce uncertainty, or increase certainty if you would prefer a more appreciative concept, is to manage the margin between the work and the hazard. The more space I have between the bad outcome and me, the less likely I am to have the bad outcome.

Workers tend to keep space between themselves and the hazards that can cause some type of harm. This "space keeping" in reality is an example of risk competency. Workers don't manage the work because successful work requires adaption, expertise, and creativity and workers don't manage the hazards because the hazards are constantly moving in and out of the work environment. That leaves the "space" between the actual work and the actual hazards as the discretionary place where workers create safety. That space could also be called the capacity to do safe work; in fact, let's call this "space capacity" to see how it fits.

What I have just discussed is a different way to think about how workers manage risk or more precisely the capacity between the work being done and the hazards to which the workers are being exposed, but there is more – because there is always more. The capacity doesn't seem to be just one type of capacity. In fact, depending on what the worker is doing, or better yet thinking while the worker is doing work, this safety capacity seems to take on different characteristics of protection.

There is a clear way workers manage this safety capacity as they prepare to do work. There is a much different way the workers manage this safety capacity as they are actually performing or executing the work. These multiple forms of creating safety capacity were interesting to observe in successful work iterations. My bet was these different safety capacities would also be interesting in understanding a failure of any significance.

The Three Types of Capacity

High reliability scholars have understood these different safety capacities between danger and safety for a long while. Reliability tends to match the capacity to operate safely and reliably to the potential disaster or failure that either can happen or has happened. Taking their idea of safety capacity types and applying their thinking not to the matching capacity to the unwanted outcome, but in fact applying their idea of matching capacity to the execution of work.

Now this is getting interesting.

There seems to be three different places where risk requires capacity for workers doing work. These three places are naturally present in the way work is performed. However, the three types of capacity are different, just as these three work stages are different. It might be easier if we introduce the three places that require capacity when doing work.

The worker must manage the capacity to do safe work at every stage of the work execution model. Workers must:

1. Plan to be safe. We call this prevention.
2. Perform or execute work safely.
3. Be able to recover if something fails. We will call this recovery.

Each of these three work capacity areas represents a specialized form of safety thinking that must be done by the worker in order for the worker to safely and productively accomplish work. Your organization's

workers manage these three capacities all the time; you manage these three capacities all the time. This is how work is safely completed. Every worker to some extent manages prevention, execution, and recovery – the question is does the worker manage these three capacities with equal importance and attention.

The three forms of risk capacity are all different in importance and complexities, yet each of the three capacities must be present when performing high-consequence work. The three capacities are not equal, at least not in the mind of the person doing the work. In fact depending on what the worker is doing or thinking, the worker may be spending time thinking about execution only. At other times the worker may be completely involved in prevention and planning the work. For high-risk work with uncertain outcomes, it is a good bet the worker is spending time thinking of escape routes and protections needed if the work fails.

Yet all three of these types of risk capacity are equally important. Spending time on one type of capacity and not spending time on the other types of capacity is like giving away two-thirds of your protection. You made not need all three types of capacity, but if you do need this capacity you want the capacity to be there waiting for you. I know this might sound complicated or unrealistic, but I would invite you to really think about how you perform safe work, I will absolutely guarantee that you will see these three types of protective capacity present is stable and safe work.

It is like when you invest in retirement fund. You don't put all your money in risky stocks. Part of your money is in secure and stable money-market funds. Some of your money is in cash. You spread the capacity (your money) not the risk – the risk is there whether your 17 dollars is invested or not.

This observation that we don't manage risk, we manage the capacity we expose to risk (or really hazard) really changed the way I started looking at fatalities and serious events. Perhaps what we have is not bad workers making bad decisions, nor do we have bad work (the hazards

again) that is trying to kill our workers. Perhaps what we have here is an absence of capacity at one of the three safety capacities.

If we have all our capacity on prevention, we would have little capacity for work execution and even smaller capacity for recovery. I would guess the transverse of this would be true. If we put all our eggs in the recovery basket, we would have little prevention capacity in order to safely do our work.

It strikes me that you could represent this idea using a picture of a scale – you know, like the scales of justice. A scale would show the need for a balanced approach to managing capacity at the organizational level with one huge exception – a scale normally is used to show a balance between two values and we have three capacities. That has led me to represent this idea in my organization as a bow tie model. You can show a balance in the sides of the bows and the middle seems perfect to represent the capacity to execute work.

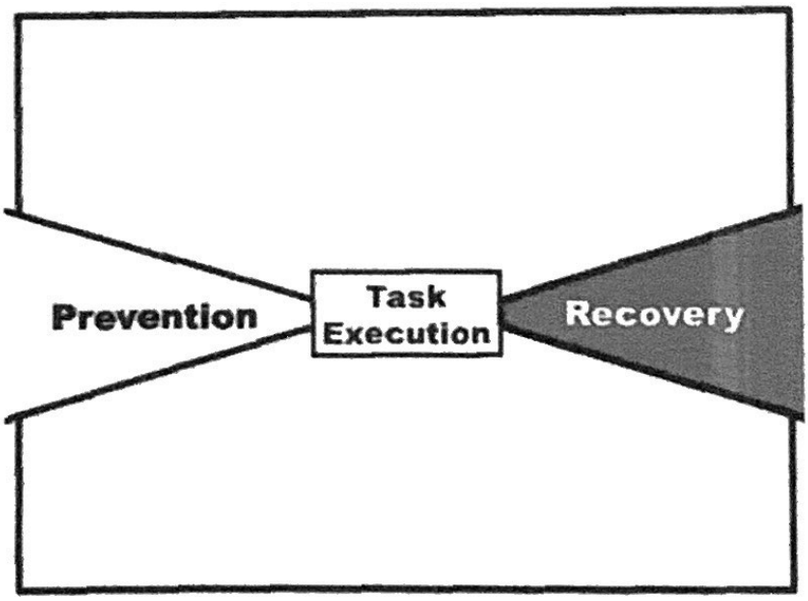

However, our discussion thus far has worked hard to make the claim we give greater preference to one of these three types of safety capacity over the other two. In fact, I will stake this books reputation that we think prevention capacity is much more important then work execution or recovery capacity. Our scale is out of balance. Our bow tie is fat on one side and small on the other. We are amazing at prevention and almost capacity-less on the recovery side.

This imbalance between prevention capacity and recovery capacity is great when work is successful. It actually is even a bit "self-fulfilling" in that it looks like all the effort and resource we put into prevention actually paid off handsomely in the complete presence of production success and the complete absence of fatalities and serious events.

Think about our illustration case studies. Think about our poor movie crew, and think about serious events in your organization, and it is quite probable that you will find an imbalance between the capacities. The film crew seemed to put all their effort on shooting the scene on the railroad trestle. So much capacity was used to create this gorilla cinema that almost no capacity was spent on prevention capacity or recovery capacity. You could make a similar case on almost any case study you chose to use--the balance between prevention, execution, and recovery seems almost vital to doing stable and safe work.

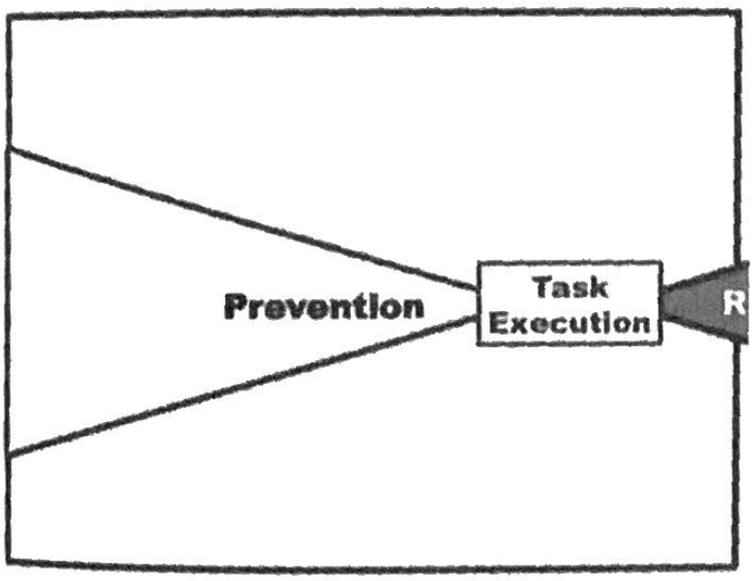

The best proof is when I observe high-performing teams, teams that are productive, efficient, stable, and safe. These high performance teams seem to naturally understand that you manage capacity in a balanced approach. When a tugboat goes in to the Harbor of Long Beach to provide ship assist, the crew manages prevention by ensuring the vessel is ready and prepped for the job. The crew of course executes the work in real time always detecting and correcting for the changing work and work hazards. Most impressively, a tugboat captain is always planning for where the tugboat will end up as the tugboat is executing work. Like a good billiards player, the tugboat captain must always position the

tugboat for the next task to be preformed while performing the current task.

The tugboat crew may not realize they have a balanced approach to managing safety capacity at all three levels of work resilience, but they absolutely do this work every time they perform their mission. This idea of multiple different safety capacities is not new; I am just not sure we have named it or thought of it as a way to address events. This became clear to me in the midst of all these fatalities peer assist visits I was doing. I was in a role where I was allowed to look at really safe and stable operations that had just suffered a catastrophic failure, and I began to notice a strong prevention culture but not a strong recovery culture.

I see the same examples with linemen working a storm response. These linemen manage capacity at all three levels because the work is so uncertain and unpredictable. Lineman always manage prevention before the go to the pole, they execute with capacity because of the hazards with which these guys and gals interface, and during a storm recovery-capacity is the entire reason they can do this high-risk work. Everything the lineman does during a storm is built upon the foundation that this task could screw up at any moment and when this task does screw up he or she wants to be where they have the most protection and the least exposure.

These and other examples of successful high-risk work is how a balanced approach, equal capacity at all three levels of work, shows itself in practice. I challenge you to not find examples of good teams managing multiple capacities when you look at your organization. We know this practice exists informally among seasoned experts who have done dangerous work for years. Could it be time to perhaps take a more formal approach to purposely creating this balance?

STOP Work

We have reached the end of the usefulness of STOP work. I am not sure

STOP work was ever effective, but it felt like the right thing to do, so we did the crap out of it. I would guess the data collected by observers when a worker was observed stopping a job before it failed and injured or killed someone is totally anecdotal and not scientific.

You know, of course, if your organization counts on STOP work as a prevention strategy that means that every other prevention strategy that your organization uses failed. Really using STOP work as an explanation for why a bad outcome happened is perhaps the weakest (and dumbest) excuse and organization could have for why a bad outcome happened.

I am not saying STOP work is wrong. It just strikes me as painfully simple and incredibly retrospective. We tend to use the STOP work criteria as a finding in an investigation and never a tool to execute safe and stable work. Asking workers to STOP a job before the job has bad consequences feels like management is creating a "back door" that removes all responsibility from the organization. When something fails, it would be easy to simply say, "The worker should have stopped the job." Management is relieved of the responsibility for creating work that can be done safely and the accountability for the accident goes directly on the workers who, not only failed to stop work, but also had an accident.

Throughout this entire discussion so far, we are building a case that prevention is not sufficient to manage serious outcome events like fatalities and serious events. We know this because organizations with amazingly formal and culturally rich prevention process still kill workers and have serious failures. These more catastrophic events seem to defy our prevention models at every turn. Try as we might, the events that cause fatalities happened in spite of all our efforts to prevent them. Prevention does not seem to be the most correct choice for an attempt to curtail serious accidents.

That said, at its best STOP work is a prevention tool. STOP work is not, nor ever has been nor ever will be a control. STOP work is supposed to happen before the undesired outcome or consequence happens. That is the entire point of STOP work. STOP work is supposed to stop the job before someone gets hurt. IF this use of STOP work is accurate, then STOP work authority is a prevention tool and not a control or a safeguard.

I really hate to admit this, but I had never thought about this nuanced view of STOP work until I was doing a learning team with a group of workers who had had a fatality. This crusty old mechanic looked me straight in the eye and told me, "STOP work is at best a prevention tool and not a good prevention tool at that. " That stalled the conversation for a bit. I nodded and acted like I knew that, but inside I was processing this information as fast as I could. That crusty mechanic is exactly right, he could not have been more right, and he should make you think a bit as well.

So what does all this mean? It seems a bit terrible to say this process is bad and offer no alternative. Let's establish that STOP work is ineffective and abusive towards your organization's workers. STOP work doesn't really prevent events, because if the workers knew the action they were about to take would cause harm to them or their coworkers, they would stop work whether they had the authority or not. STOP work is used after a consequence to ask workers why they failed to use their STOP work authority. STOP work dumbs down the complex and context rich conditions that create environments that foster accidents. They make an accident look like it would be as simple to avoid as saying the word, "stop."

Instead of asking worker to psychically identify the precise uncertainty that will lead to an accident, why not push the effort towards ensuring the right controls and safeguards are in place for the work to fail safely *before* the work starts. Perhaps STOP work should be called START work. Why not empower the workers to have the power to not start a job if the appropriate safeguards are not present or engaged.

Lets move from placing our operational emphasis on STOP work to placing renewed operational emphasis on START work. Every worker has the power to not start the task until the workers are assured the proper safeguards are in place to allow for the task to fail without consequence to the workers or the facility. Can we fail safely and gracefully? Does this feel more in line with the shift in thinking that we have been discussing throughout this book?

This START work idea removes the pressure on prevention and places this pressure on the recovery capacity that must exist to successfully execute potentially high-consequence work. The jury is out if this idea will work. START work still has a slight tinge of "weaponization" if used retrospectively after a failure, to be sure, but does directly address the shift toward creating a balance between prevention and recovery efforts while doing work.

All of these efforts indicate a new way of thinking, not necessarily a new way of doing work. Highly effective work groups have been managing safeguards for years. If you watch the successful execution of high-risk work by really experienced and skilled workers, you will see these workers address recoverability before these workers ever touch a tool.