Hubbard Decision Research

Every Organization's Biggest Risk: Broken Risk Management

Measure What Matters — Make Better Decisions

Introduction

One of an organization's top priorities should be to make better decisions by identifying risks and determining how they can be mitigated. Enterprise risk management is how an organization accomplishes that priority. But as this white paper will show, popular risk management practices in use today actually represent the biggest risk to an organization: using broken methodologies that don't adequately protect the firm and create a false sense of security.

The purpose of this white paper is to:

1. Identify key problems risk managers face;

2. Explain how risk management is currently done;

3. Point out critical flaws with most risk management methods in use today; and

4. Demonstrate a better, more scientific way to manage risk.

The outcome from reading this white paper should be a better understanding of the benefits of using rigorous quantitative and scientifically validated risk management methods - and the risk of using anything else.

Contents

Why Risk Management is Failing

Enterprise risk management has become a crucial part of an organization's decision-making process due to the diverse nature of threats and the impact they can have on the organization. Sometimes identifying risks is easy; other times they're difficult to detect and assess. Sometimes a mitigation strategy is simple; other times figuring out what to do about a risk can be challenging and complex.

Most organizations large enough to have a COO, CFO or CIO have implemented risk management in some form or are busy attempting to do so. The methods used today, however, are not making organizations safer from threats. They are either doing nothing at all, or are luring decision-makers into a false sense of security - while exposing companies to the potential for devastating losses.

In his book *The Failure of Risk Management: Why It's Broken and How to Fix It*, Doug Hubbard argues that most risk management methodologies in use today are little better than astrology. Using these methodologies ironically imposes a risk even though they're supposedly designed to *reduce* risk.

Three Reasons RM Has Failed

Risk management has failed today due to one of three reasons:

1. Organizations do not measure and validate methods in whole or in part;

2. Organizations use components that are known to not work; and

3. Organizations don't use components that are known to work.

In other words, "best practices" put into place today are based on methods that, at the end of the day, amount to little more than guesswork at best. And yet, a company's future and livelihood are based on such best practices.

The failure of an organization's risk management processes places an extreme burden on the people specifically tasked with helping protect against risk: the company's risk management team. Risk managers have a difficult job as it is; it becomes even more difficult when they don't have the tools and methods to do their job the right way.

As the next chapter shows, risk managers face problems that are largely created by the very systems put in place to help them solve the problem of risk itself.

Major Problems Risk Managers Face

Risk managers have difficult jobs. Whether you're a chief risk officer, another C-suite executive tasked with risk mitigation and compliance, a risk manager or analyst, or some other decision-maker, your job is to identify threats and figure out ways to avoid them. In other words, you're trying to make good decisions in the face of uncertainty.

In his book, Doug Hubbard identifies **seven major problems** that risk managers face today.

1. **Confusion regarding the concept of risk.** Among different specialties in risk management, analysts and managers are using the word risk to mean some very different things.

2. **Completely avoidable human errors in subjective judgments of risk.** Most of the methods of risk assessment must rely on at least some subjective inputs by human experts, but, without certain precautions, human experts make surprisingly consistent types of errors in judgment about uncertainty and risk. Although research shows that there are methods that can correct for certain systemic errors that people make, very few do so and the net result is an almost universal understatement of risk.

3. **Entirely ineffectual but popular soft scoring and "heat map" methods.** The numerous arbitrary rules and values created in scoring methods not only fail to consider the problems with subjective risks, they introduce errors of their own and may actually make decisions worse. There is no large, important decision that would not be better served with some other analysis approach.

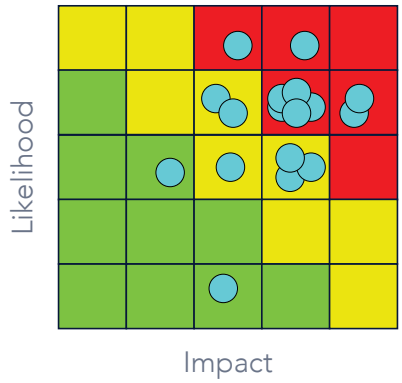

Figure 1: Example of Soft-Scoring Heat Map

4. **Misconceptions that block the use of better, existing methods.** Even some experienced risk analysts defend the use of ineffectual methods by arguing that better, more sophisticated methods will not work. But each of these arguments is based on fundamental fallacies about the nature of quantitative risk analysis.

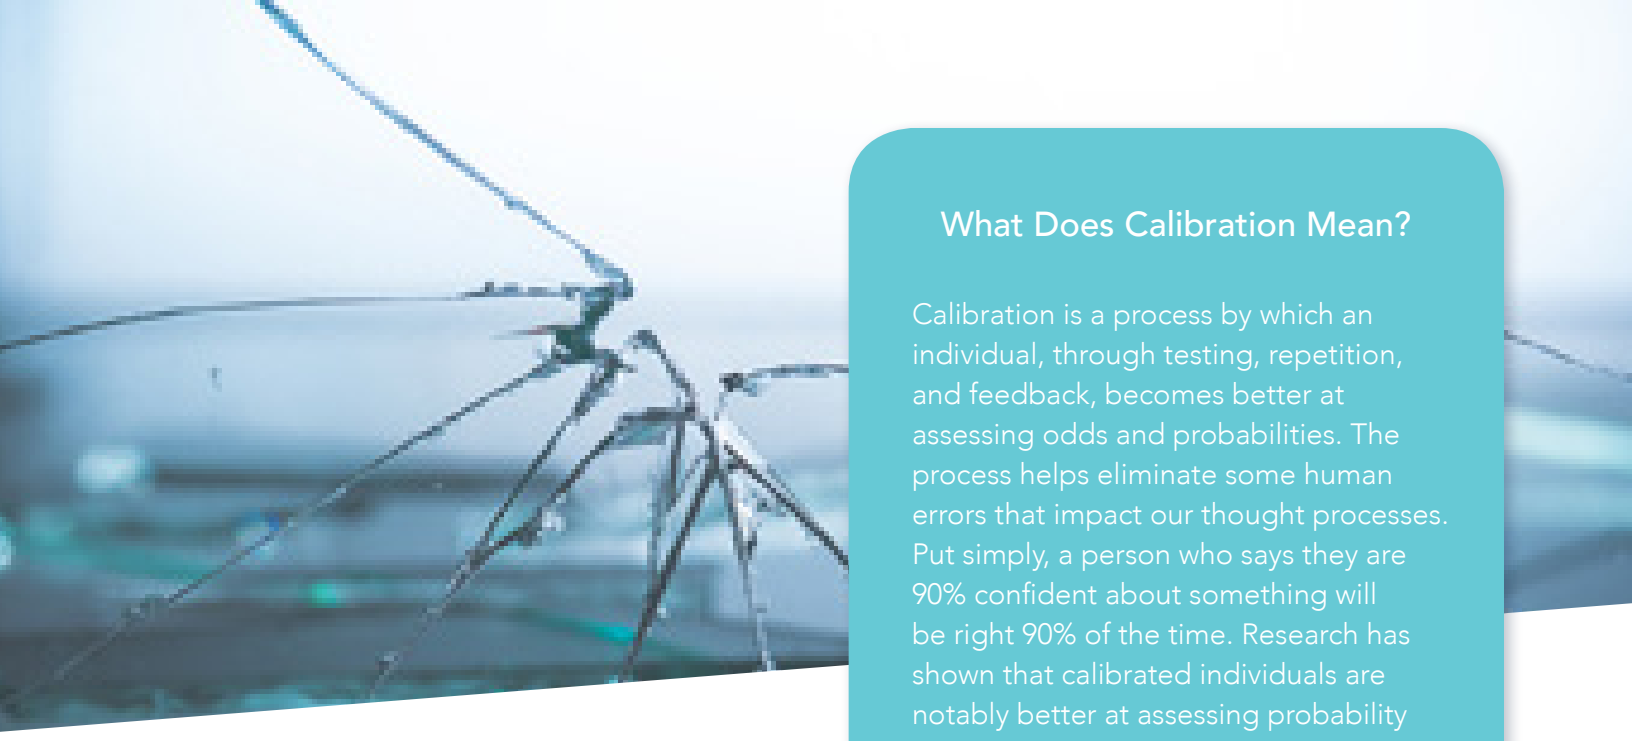

What Does Calibration Mean?

Calibration is a process by which an individual, through testing, repetition, and feedback, becomes better at assessing odds and probabilities. The process helps eliminate some human errors that impact our thought processes. Put simply, a person who says they are 90% confident about something will be right 90% of the time. Research has shown that calibrated individuals are notably better at assessing probability than uncalibrated individuals.

5. **Recurring errors in even the most sophisticated models.** Even when risk managers adopt more quantitative approaches, they do not attempt to measure the reliability of their models. Some analysts assume that their models take on a level of authority and "truth" that is never justified. Half-understood models are misapplied in a variety of situations.

6. **Institutional factors.** Unnecessary isolation of risk analysts from each other - both within the same organization and among organizations - means that important shared risks and relationships will be ignored in overspecialized models.

7. **Unproductive incentive structures.** The methods will not matter much if the incentives to make better decisions and manage risks are not improved. Minimizing risk is not a factor in most executive bonus calculations. Human experts are not incentivized to give reliable forecasts and there is little incentive to verify old forecasts against observations. A key motivator is compliance and the use of so-called best practices. If a ship is sinking, at least the captain can point out that he followed established procedures.

Each of these contributes to even more uncertainty - or, worse, a false sense of confidence that the risk management methodologies used by an organization are actually working when, in fact, they're not.

One common theme: most risk management methods today are **unscientific** and based on **subjective human judgment without calibration*** and **self-assessments** that are **prone to avoidable errors.** Risk managers rely too much on expert intuition - or "gut feel" - which has shown, by research, to be inferior to statistical models.

This isn't to say that subjectivity is always bad. Indeed, even an uncalibrated subjective assessment can be better than a flawed system that introduces more error into the decision-making process. Subjectivity is bad whenever there aren't processes in place to correct for common, known errors in subjective estimates.

Risk Management Methods Commonly Used & Why They Don't Work

In the risk management industry, there is a wide array of methodologies and best practices in use. Some of them are proprietary; some have been codified by national and international institutions. Some use only one technique; others are a combination of techniques.

These methods most commonly include:

- **Expert intuition.** Risks are assessed by experts who use experience as a foundation for their judgments.
- **Expert audits.** More in-depth, but still based on subjective expert assessments.
- **Simple stratification and scoring methods.** These include heat maps, risk maps, risk matrices, risk scores, and other soft methods based on rating threats on a spectrum (e.g. high-low, most severe-least severe). These are quantified in that they have numbers attached to them, but they aren't scientific; they're based on expert judgment.
- **Traditional financial analysis.** Originally created for the finance industry, this method incorporates some practices that attempt to be quantitative, such as "discount rates" and cost/benefit analyses.
- **Calculus of preferences.** Methods such as multi-attribute utility theory (MAUT), multi-criteria decision making (MCDM), and analytic hierarchy process (AHP). These are more structured than simple weighted scores, but are still based on subjective judgment rather than empirical evidence.
- **Probabilistic models.** More rigorous and quantitative than other methods, and are used in finance, engineering, insurance, and other fields. These models can still be misapplied, but are an improvement over other methods.

One common denominator across these methods is a reliance on uncalibrated subjective human judgment and error. Research has shown, though, that expert judgment is prone to more error than statistical probabilistic modeling. Daniel Kahneman and Amos Tversky, psychologists who pioneered several areas of research in the field of judgment and decision-making (Kahneman received the Nobel Prize for Economics for part of this work), found that a wide array of factors and biases contribute to either overestimating or underestimating risk and probability.

> Companies are making very critical decisions with major potential impacts...based on little or no scientific basis - and they believe they are making good decisions without any evidence to prove it.

These factors include the facts that experience is a non-random and non-scientific sample of lifetime events based on selective memory, and that judgments based on experience can be full of logical error and often lack reliable feedback (which, when it comes from humans, is subject to the same limitations). Also, experience is often inconsistently applied.

The result: humans - even experts in their field - are naturally bad at assessing the probabilities of events.

What's more is that we are very bad at determining if our uncalibrated subjective, non-scientific, non-empirical risk management methods actually work.

As he relates in *Failure of Risk Management*, Doug Hubbard routinely asks rooms of risk managers and experts a series of questions answered by raising their hands. He starts by asking if they have a defined approach to managing risks. Most raise their hands. He then asks if they measure risks. Many lower their hands. He follows up by asking if they use probabilities in their measurements, and even more lower their hands. Finally, he asks if these measurements of probabilities and losses are in any way based on statistical analysis or methods used in actuarial science.

By the time this last question is asked, there will be very few hands left raised, if any.

The downside is obvious: companies are making very critical decisions with major potential impacts on processes that are prone to human error, based on little or no scientific basis, and subject to inconsistency - and they believe they are making good decisions without any evidence to prove it.

Doing nothing about risk management may not actually be the worst case. That flies against the firms who invoke the usual "At least we're doing *something*" defense of the risk management strategy they're following. What's worse than doing nothing would be an organization luring itself into a false sense of security - and wasting resources - by using soft scoring or unproven methods and believing in them.

> What's worse than doing nothing would be...using soft scoring or unproven methods and believing in them.

The Question Every Risk Manager Should Ask

When assessing the performance and effectiveness of your risk management process, it helps if you undertake a rigorous, critical examination of the process starting with one question: **How do I know my methods work?**

Before you answer, we need to clarify what this means. By "works" we mean a method, *measurably* reduces error in estimates, and improves average return on portfolios of decisions compared to expert intuition or an alternative method. Note that this is not the same as merely *perceived* benefits. If, for example, estimates of project cost overruns are improved, that should be objectively measurable by comparing original estimates to observed outcomes. Merely using a survey to ask managers their opinions about the benefits of a method will not suffice.

The reason we can't rely on the mere perception of effectiveness is that we are all susceptible to a kind of "analysis placebo effect." That is, research shows that we can increase our confidence at a task while not improving or even getting worse.

For example, it has been shown that using more data or more "rigor", even when there is no real measurable improvement, has increased confidence in estimating the outcomes of sporting events and portfolio returns.

Merely having a system also doesn't guarantee effectiveness or improvement. In one study in Harvard Business Review, the authors found that an analysis of over 200 popular management tools and processes had a surprising result: "Most of the management tools and techniques we studied had no direct causal relationship to superior business performance."

So, how can we measure *real* improvements? Ideally, there would be some big survey been conducted which tracked multiple organizations over a long period of time which showed that some methods are measurably outperforming others. Did 50 companies using one method over a 10-year period actually have fewer big loss events than another 50 companies using another method over the same period? Or were returns on portfolios of investments improved for the first group compared to the second group? Or were events at least predicted better?

Large scale research like that is rare. But there is a lot of research on individual *components* of methods, if not the entire methodology. Components include the elicitation of inputs, controls for various errors, use of historical data, specific quantitative procedures, and so on. What does the research say about each of the parts of your method? Also, is there research that shows that these components make estimates or outcomes worse?

As mentioned earlier, this research has already been done and the results are conclusive. So the only other question is why not get started on improvements now?

Finding a Better Way Forward for Risk Management

An honest assessment of risk management methods should reveal a better way forward: a need for more empirical data and scientific processes to analyze that data and reduce uncertainty through objective means.

Risk management, at the end of the day, is about reducing uncertainty. By using hard quantitative methods to assess and measure risk, better decisions regarding risks can be made due to less uncertainty and more (rightful) confidence.

As Doug Hubbard outlines, risk management in your corporation can be improved by:

1. **Adopting the language and philosophy of modeling uncertain systems.** This means getting away from speaking the language of soft scoring methods and non-scientific methods.
2. **Using calibrated probabilities to express uncertainties.** Calibration measurably improves an expert's ability to assess odds.
3. **Switching to probabilistic modeling methods immediately.** Probabilistic methods most notably include Monte Carlo simulations, but incorporate other means of analyzing data and statistically modeling outcomes.

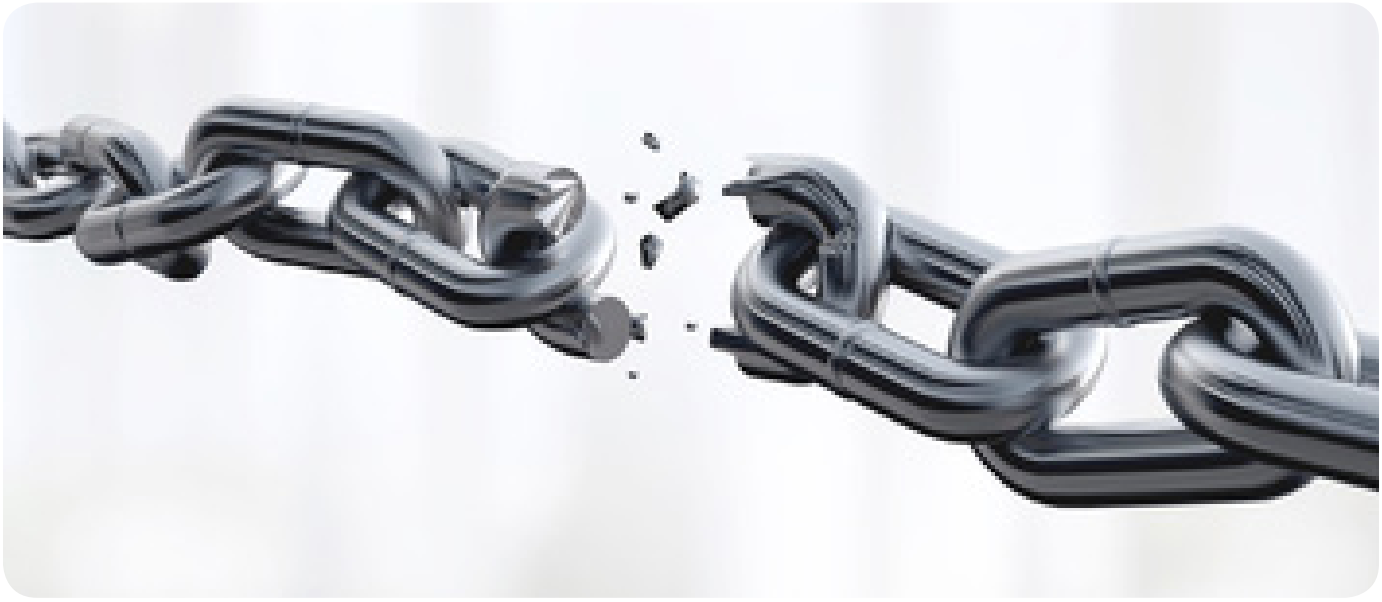

The Ideal Risk Management Process

Your risk management process should be described as follows:

"The firm builds quantitative models to run simulations of interconnected models of risk across the organization, any subjective estimates are from calibrated experts, additional empirical measurements are used where optimal, and risk tolerance is quantified. Always skeptical of any model, the modelers check against reality, are familiar with research about the validity of methods, and continuously improve the risk models."

Figure 2 below is an example of one output of a risk management process that works. Any decision-maker can look at this output and get actionable insight he or she can use to make a better decision - and any analyst, properly trained, can produce it.

What if we could measure risk more like an actuary —
"The probability of losing more than $10 million due to security incidents in 2016 is 16%"

What if we could prioritize security investments based on a "Return on Mitigation"?

	Expected Loss/Yr	Cost of Control	Control Effectiveness	Return on Control	Action
DB Access	$24.7M	$800K	95%	2832%	Mitigate
Physical Access	$2.5M	$300K	99%	727%	Mitigate
Data in Transit	$2.3M	$600K	95%	267%	Mitigate
Network Access Control	$2.3M	$400K	30%	74%	Mitigate
File Access	$969K	$600K	90%	45%	Monitor
Web Vulnerabilities	$409K	$800K	95%	51%	Track
System Configuration	$113K	$500K	100%	77%	Track

This means there is about a 40% chance of losing more than $10M in a year and about a 10% chance of losing more than $200M.

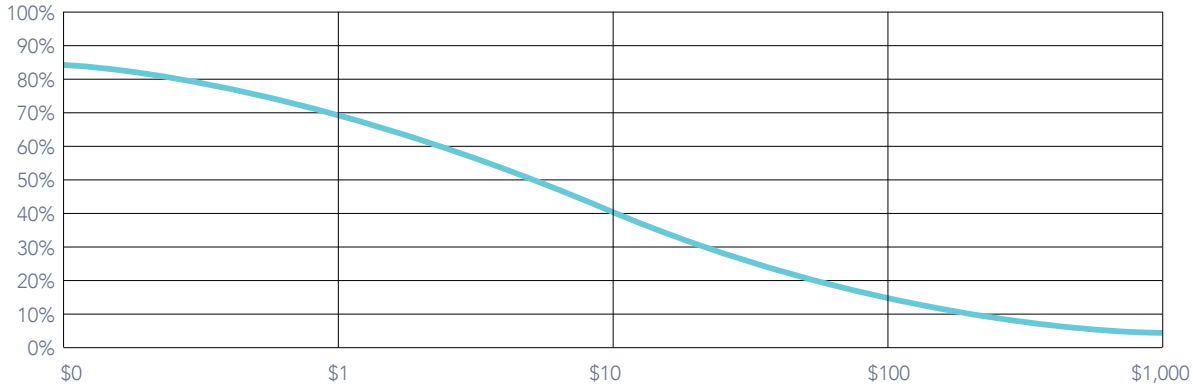

Figure 2: Loss Exceedance Curve for Risk Mitigation (from How to Measure Anything in Cybersecurity Risk)

Getting to this outcome involves following the four core principles that have helped organizations as diverse as Fortune 500 corporations, small businesses, nonprofits, and the military reduce risk, improve ROI, and achieve other critical goes:

Define the Decision(s):

Identify the real decision at the outset. Is the dilemma whether to simply approve a project or how to conduct a project given a vast combination of alternatives? Or is the decision a matter of when a given initiative should be approved?

Model What We Know Now:

Cost estimates, market forecasts, project risks, and other variables in a typical big investment decision are almost never known exactly. Usually, the uncertainty about some variables, especially long term forecasts, can seem extreme. But even extremely uncertain variables can be assessed. Methods have shown to work even when organizations have very little historical data, complex problems, and measurements that seem almost impossible.

Measure What Matters:

Much of measuring risk involves picking the right things to measure. Not all variables in a decision are worth measuring and those worth measuring are often a surprise to the decision makers. In fact, most managers measure exactly the wrong things – that is, the most uncertain variables tend to be ignored while the variables that usually receive a lot of attention actually have less bearing on the decision. Every variable in a model should have an "information value" that allows identification of high value variables in a decision.

Make Better Decisions:

This can include decisions that are more than just "accept/reject" choices but possible combinations of several choices. The final outcomes must consider the risk preferences of decision makers. The output of the decision model, updated with economically justified measurements, is compared to the risk appetite of the organization.

This process can guide the creation of a model that delivers that critical phrase - actionable insight - to the decision-maker. Figure 3 is an example of how a typical model accomplishes that in a way that any decision-maker can understand:

Figure 3: ROI Summary in a Typical Model

With effort, your organization's risk management culture can move away from the subjective methods it may currently favor and toward a more rigorous one in which decision-makers act as scientists seeking to make decisions based on actionable data and less uncertainty.

No organization can afford to have an error-prone risk management system. The costs of implementing a system that works pales in comparison to the costs an organization can - and will - face without one.

Following the principles discussed in this paper will hopefully give you a better idea of what risks are out there, how they can impact your organization, and - most importantly - how you can mitigate or avoid them and thereby avoid serious loss.

Hubbard Decision Research

About Doug Hubbard

Douglas Hubbard is the inventor of the Applied Information Economics (AIE) method and founder of Hubbard Decision Research (HDR). He is the author of *How to Measure Anything: Finding the Value of Intangibles in Business, The Failure of Risk Management: Why It's Broken and How to Fix It, Pulse: The New Science of Harnessing Internet Buzz to Track Threats and Opportunities* and his latest book, *How to Measure Anything in Cybersecurity Risk* (Wiley, 2016). He has sold over 100,000 copies of his books in eight different languages. Two of his books are required reading for the Society of Actuaries exam prep. In addition to his books, Mr. Hubbard has been published in several periodicals including *Nature, The IBM Journal of Research and Development, OR/MS Today, Analytics, CIO, Information Week, and Architecture Boston.*

About Hubbard Decision Research

Hubbard Decision Research (HDR) is a risk management consulting firm that applies quantitative analysis methods to the most difficult measurements and challenging decisions across many industries and professions. Using Applied Information Economics, HDR has developed quantitative analysis solutions to information technology investments, military logistics, entertainment media, major policy decisions, and business operations, for clients ranging from small businesses to Fortune 500 companies. More information can be found at hubbardresearch.com.

For more information on the training webinars and seminars we offer in project management, cybersecurity, risk management, organizational transformation, and other areas, visit our training page.

CONTACT US

For a consultation on the methods used by HDR to use scientific quantitative methods to identify, assess, measure, and mitigate risk, contact the team.