

Making systems

Volume 4

People

Richard Golding



Making systems

Volume 1—Fundamentals

Volume 2—Life cycles

Volume 3—Development

Volume 4—People

Volume 5—Management

Making systems

Volume 4

People

Richard Golding

March 4, 2026

Copyright ©2026 by Richard Golding

All rights reserved.

No portion of this work may be reproduced in any form without written permission from the publisher or author, except as permitted by U.S. copyright law.

Release: 0.4-review

Contents

XII	Team organization	I
56	Introduction	3
57	Structure	5
58	Communication	7
59	Agreement, disagreement, and bias	9
59.1	Agreement and disagreement	9
59.2	Biased work	9
XIII	Team members	II
60	Team roles	13
61	Skills	15
62	Leadership and culture	17
63	System-building ethics	19
63.1	Purpose	19
63.2	Definitions	20
63.3	Group versus personal ethics	22
63.4	Harms and benefits	23
63.5	Framework	27
63.6	After an incident	39
63.7	Case studies	43
63.8	Using the ethics framework	53
63.9	Commentary	56
	Bibliography	59

Part XII

Team organization

Introduction

- Basic collective action problem is related to the equilibrium between contribution and resulting common good
 - Depends on fungibility and homogeneity of contribution and good
- Collective good is homogeneous: a successful system
 - More specifically, good decisions, designs, and implementations that lead to the system as a whole working in a way that meets its purposes, which will in turn provide individual goods like pay, satisfaction, status
- The cost is heterogeneous: each person should be doing different things to contribute to the good; this is about sharing out work, with different people doing different tasks
 - Modeling contribution as effort (time spent) is not helpful, because if everyone does the same thing with their effort, the result will have zero value
- Perception of what one person should do (is incentivized to do, is assigned to do) does not necessarily result in an optimal common good
 - Value is related to satisfaction of the system structure constraints
 - Value of one contribution can be low if someone else does something not consistent with it; the same contribution can have high value if the other does something consistent
- Free rider problem (not contributing full share for the good received) increases with size of group
 - Detection of free rider becomes harder
 - Perception of value of contribution decreases because it is a smaller share of contribution, and share of common good is lower
 - Work (cost) in a team is transferable, so if one member tries to be a free rider, someone else will have to pick up the work. This creates an incentive for enforcement
- If contributions are not homogeneous and shared equitably, is there increased motivation to free ride? Is there greater motivation for the “greater” contributor to dictate?
- Common interest of a team: pay, satisfaction in work, product
 - Pay and satisfaction are private and excludable; product is neither

- Organization interest is the product, which may or may not be an interest of the members
- Within the team, common interest can be about ability to work on some part or make certain decisions
- Antagonistic interest in that people want more scope or to do less work
- A successful system product is a common good (public good); non-exclusivity of the result

Consider, for example, meetings that involve too many people, and accordingly cannot make decisions promptly or carefully. Everyone would like to have the meeting end quickly, but few if any will be willing to let their pet concern be dropped to make this possible. And though all of those participating presumably have an interest in reaching sound decisions, this all too often fails to happen. When the number of participants is large, the typical participant will know that his own efforts will probably not make much difference to the outcome, and that he will be affected by the meeting's decision in much the same way no matter how much or how little effort he puts into studying the issues. [...] The decisions of the meeting are thus public goods to the participants (and perhaps others), and the contribution that each participant will make toward achieving or improving these public goods will become smaller as the meeting becomes larger. It is for these reasons, among others, that organizations so often turn to the small group; committees, sub-committees, and small leadership groups are created, and once created they tend to play a crucial role. [Olson65, p. 53]

that needs careful design just as much as the system product

- This is the “machine” that, when run, builds the system product
- People form organization whether given or not
- Required in order to divide up work and to share things
- Know who is working on what, to avoid duplicated work and so direct people to the right person to talk to
- Know how to escalate issues when needed
- Know when a situation is exceptional and should be escalated, so that delegation can work
- Know that escalated issues and side-channel reporting will be dealt with
- Know how checks and approvals happen
- Ensure that no work or system parts fall through the cracks

Structure

- objectives
 - people know who they work with, and who has authority
 - promotes communication and interaction
 - allows for independence when needed
 - minimising overheads
 - clear ownership (responsibility for) artifacts or procedures
- patterns
 - organized around part of system
 - organized around discipline
 - matrixed to be both
 - hybrid—mostly around system, with consulting by disciplines
- documenting structure
 - org chart
 - roles, and who has them
 - teaching about the structure
- structure versus size, and reorganization
 - simple when small
 - formal structure when large
 - transition points between structures

Communication

What kinds of communication needed—normal circumstances

- questions while building something
- informing about part of system (e.g. a design decision)
- reviews and feedback
- approvals
- resource coordination (people, lab, budget)
- plan management (strategy)
- operations updates, questions, feedback

Off-nominal communications

- detected technical problem (immediately detectable)
- interpersonal problems
- anticipated or likely technical problems (forward looking, not immediate)

Tie to Section 22.4: defined responsibilities for communication and avoiding thoughtless broadcast meetings

Agreement, disagreement, and bias

59.1 Agreement and disagreement

- definitions of consensus, agreement, disagreement
 - when consensus is good, what benefits it has
 - when agreement, not consensus, is enough
 - when disagreement is okay
 - when disagreement is a problem
- Achieving agreement
- Detecting problemating disagreement
- Resolving disagreement

59.2 Biased work

- What bias is
 - Confirmation bias
 - Motivated reasoning
 - Echo chamber
 - Commons
- Utility of bias (when it's good)
- Problems with bias
- Addressing biases
 - Prevention
 - Detection
 - Correction

Part XIII

Team members

CHAPTER 60

Team roles

- the roles people fill in the team
- leadership
- coordination
- planning
- tasking
- specialists

Skills

- important skills and abilities
- listening (active)
- negotiation skills that allow for collaborative problem solving and for competitive negotiation
- ability to detect when one doesn't know enough or as much as someone else, and then focus on learning
- ability to follow a complex graph of things and find how they relate to each other

CHAPTER 62

Leadership and culture

- team leadership and culture
- need for responsibility

System-building ethics

63.1 Purpose

This chapter presents a framework for thinking about ethics in systems-building projects. This is not a general treatment of ethics; that is a much bigger subject. The aim is to present guidelines for people building a system and a theory that provides a framework to reason about situations that go beyond those guidelines.

The building of complex systems causes impacts on the world. The time and effort involved affect all those who do the work. Those who provide resources used to make the system do so trusting that there will be benefit from their investment. The system that is built and deployed will affect those who use it and beyond. It is inescapable that the acts of making a system have effects.

Ethics can be viewed as taking responsibility for the effects that one's actions cause, and making choices to result in the best possible effects.

Ethical responsibility is separate from legal responsibility. Legal responsibility is general closely circumscribed, covering only a small fraction of ethical responsibilities and even then limited to those responsibilities that affect the functioning of society at large. The vast majority of ethical responsibilities are matters between individuals or with limited effects, and are not encoded in law.

Most people have a general sense of morality and ethics. They have a sense that honesty and fairness matter, for example. They understand ideas of trust and contracts. They are aware of the maxim “first, do no harm” that is reputed to come from the Hippocratic Oath.

People working on systems projects need more of a guide than those basics. The basics are correct and a sound foundation—but building complex systems often leads to complex ethical situations. I have found many situations where it was not immediately obvious that there was a basic ethical matter to be addressed, and the people involved were busy trying to get work done and not thinking about potential ethical issues. In other situations a decision had to balance multiple separate effects; reducing one harm increased a different harm, for example.

In researching for this chapter, I looked for ethics guidelines for technical disciplines. The Hippocratic Oath in medicine is the best-known example. It is of ancient origin, and many modern doctors give an oath to follow updated versions of it. The IEEE publishes a code of ethics for its members [IEEE24]. It is about one page long,

and includes several useful principles. The INCOSE has a similar, and similarly brief, code for systems engineers [INCOSE24].

These codes all list useful principles, and they are worth reading and taking to heart, but they are not a sufficient guide to the complexities of actual systems-building ethics. They are at the level of the general understanding people generally have. Compare, for example, the Hippocratic Oath to the volumes of work on medical and bioethics. The Hippocratic Oath is valid, but it does not address the complexity arising from modern medical capabilities.

This chapter presents a framework for thinking about ethical matters that come up in a systems-building project. It does *not* attempt to provide a succinct guide, because there are plenty of those already available (as the examples cited above). Instead, this chapter defines a number of basic terms and some foundational axioms, based on ideas of *responsibility*, then discusses different aspects of those responsibilities.

63.2 Definitions

I use the word *ethics* in this book to mean “the principles of conduct governing an individual or a group” and “the discipline dealing with what is good and bad”.¹ I have avoided the word *morals*, which is often defined to be more or less equivalent to ethics, because it has connotations of religious or other fundamental sources of good and bad.

Ethics, in this account, is concerned with *harm* and *benefit* (rather than *good* and *bad*). Harm is some loss or damage that someone experiences. This idea is related to that of accident and loss in safety or security analysis (Chapter 46). Benefit is something that helps someone or that they value. What constitutes benefit or harm is not necessarily absolute; it can depend on whose point of view and other circumstances. Reduced harm can be treated as a benefit, and reduced benefit can be considered harm (for example, the idea of opportunity cost).

The objective of ethical behavior (and ethical decisions) is to maximize benefit and minimize harm. The first objective is to produce benefit, minimizing harm while doing so. It is not enough only to avoid harm; indeed, it is not possible to avoid harm by risk-averse inaction.

Harm and benefit are most often not quantifiable. Two different harms or benefits are usually not even commensurable: they cannot be directly compared against each other. If deciding yes produces some amount of benefit A and some of B, but deciding no produces different amounts of A and B, is yes or no the better decision? Such value judgments must be made when weighing alternative choices, and how the different benefits are weighted becomes a matter of policy.

¹Per Merriam-Webster Online Dictionary.

This account of ethics is organized around *responsibility*.² Each person or group has responsibilities for particular actions and decisions, and is therefore responsible for the effects of those actions. Responsibility implies an expectation that the person will consider the ethical implications of an action or decision, and make the decision that maximizes benefit and minimizes harm. Of course the outcomes will happen in the future and are not fully predictable at the time the decision is made, and so the person making the decision can only do as well as the information they have.

Culpability is the condition that comes from making a decision or taking an action that does not meet this ethical bar. It is defined as “meriting condemnation or blame” for the action.³ Ideas of culpability are generally associated with guilt, punishment, or reparation. I do not treat those questions here. As with safety or security matters (Chapter 46), I am more concerned with avoiding poor ethical decisions by providing people with the tools to understand when the decisions are to be made and how to reason through the choices. After a poor ethical decision has been made, my focus is understanding why and how the decision happened and learning how to avoid such mistakes in the future. (Punishment or retribution may have a role in building incentives to avoid poor decisions.)

Responsibility is associated with two other concepts: *authority* and *power*.

Authority is the right, according to social and organizational rules of behavior, to make particular decisions or to take certain actions. A person is expected not to perform actions for which they do not have authority. Responsibility must be matched by authority, and vice versa; one cannot have a responsibility for some outcome of an action without authority to control that action. At the same time, one cannot reasonably claim responsibility outside their scope of authority.

Power, on the other hand, is the ability to cause actions regardless of whether one has the authority for those actions. Power can follow from authority, where having some authority leads to having the power to do the corresponding action, but in an organization with problems someone can have authority without having the power to perform the actions. Power is a matter of interpersonal relationships outside the formal organizational structures like authority. Someone can gain power, for example, by friendships with others, gaining respect and trust from others in a group, or by coercing others. Someone can lose power by losing respect or trust.

Power and authority are not ethical matters *per se*. Rather, how they work in a group can help or hinder ethical behavior. One can have power but refrain from using it outside authority. When someone with authority starts to make ethically poor decisions, someone else with power can act to correct the situation. Problems most often occur when someone with power takes unethical actions, and no one else

²This approach to basing ethics on responsibility is inspired by Drucker [Drucker93, Chapter 25] and Nollkaemper [Nollkaemper20], among others.

³Per Merriam-Webster Online Dictionary.

has authority to deal with and correct those actions. A disjunction between power and authority can also hinder the way a team builds and maintains trust in each other and its organization, which can in turn lead to unethical actions—but the power and authority in themselves are not the problem.

Finally, all these definitions are related to actions taken by or attributes of groups and individual people. Following the practice of security work, I will refer to them as *principals*.

Principals can be *individuals* or *groups*. Only individuals can make decisions or take actions. However, groups have joint responsibilities and take collective actions. These joint responsibilities are emergent properties of the collection of individuals that make up the group. The responsibilities and authorities of groups can be a complex subject, which I will discuss below; it leads to ideas such as collective culpability, which is sometimes used to justify inappropriate collective punishment.

63.3 Group versus personal ethics

A group, such as a project or team, should behave ethically in aggregate. People treat an organization or a project as a principal that has ethical responsibilities. A project is responsible for being honest about the system it has built; for treating other stakeholders honestly and with respect; for how it treats its team members; and for using resources invested in it to good effect. An organization is responsible for only taking on projects that have more benefit than harm in the world.

An organization can behave unethically even if all the individuals in it behave ethically. For example, it can cause harm when no one has the responsibility to check accuracy of some design or to consider the harms a decision may cause.

Individual persons should behave ethically in specific. Each person is responsible for the effects of their own actions. They have a responsibility to treat others fairly and honestly, to meet commitments they have made, and in general cause benefit and avoid harm.

An individual can behave unethically because of ignorance, because of carelessness, as well as out of malice. For example, a team member can put another in potential physical danger because they did not know the safety protocols they are to follow, or because they forgot to do some essential safety step.

An individual, especially in leadership, can behave unethically because of conflicts of interest. This occurs, for example, when someone expects to receive a personal benefit from making a particular decision that does not benefit the project, or a customer, or other stakeholder. Potential personal benefit can complicate someone's ability to make ethical decisions; it can lead to motivated reasoning to try to find rea-

sons why they should decide so that they see the personal benefit. Each team member must remain vigilant and exercise greater care when they may have such a conflict.

While a project or organization can have collective ethical responsibilities, only the individuals that make it up can take decisions and actions that cause benefit or harm. Individual ethical behavior is necessary to collective ethical behavior, but not sufficient. The project's structure determines how the collective responsibilities flow down to individual team members. In doing so the project sets the environment for individual ethical behavior. The project must also ensure that collective responsibilities are fully covered by individual responsibilities. Individuals are thus responsible for some share of the project's ethical behavior, and for reporting, protesting, refusing, or whistle-blowing when there are violations.

63.4 Harms and benefits

The first question to ask about harms or benefits is: to whom?

The list of project stakeholders in Section 16.2 is one place to start: customer, team, organization, funders, regulators.

A system-building project affects more than that, though. In general a project has effects on society and the world. A system affects the customer, and also the customer's customers. It uses resources in the world and produces waste. The resources invested in building a system and then in deploying and using it could be used for other purposes, and so the project affects a potentially broad part of society. A project will be in competition with others and thus affect them: for investment, for customers, for resources.

Next, harms and benefit can have their effects at different times. Some effects are clear right away: when a system helps a customer (benefit) or when someone is injured (harm). Other effects occur later, like pollution that was not understood when the system was built and deployed (harm) or when the system creates an opportunity later for follow-on business (benefit).

Finally, what are some examples of benefits and harms that one should consider in a system-building project? While many people will immediately think of "harm" as meaning physical injury, most harms are actually economic or other less material harms. Many of these harms occur when the system does not provide the benefit that a better choice would have: resource misallocation or lost opportunities.

The lists below are illustrative, not complete. They are meant to inspire thought.

Customer. The customer acquires and uses the system that the project is building. They depend on the system to help them with whatever their purpose is—the customer they intend to serve, or the service they intend to provide. They expect that

the system will meet that need, as well as being reliable and safe, and other concerns (Section 16.2.1).

Responsibility to the customer is limited by the lifetime of the system, not by the lifetime of the project. There are many examples of companies addressing their ethical responsibilities only as long as the company exists or offers the system product, while potential harms from system defects can last for decades. While legal responsibility is usually more limited, and certainly ends when the corporate entity ceases to exist, ethical responsibility remains and goes beyond what is compelled by law.

Harms:

- Building a system that is not what they want or need, especially if this happens because of conflict of interest.
- Failing to deliver a system that has been promised or paid for.
- Persuading the customer to misallocate their own resources by acquiring and deploying the system instead of something else.
- Providing a system that has errors that increase the customer's cost of use or requires more effort than it should have.
- Building a system that cannot evolve to meet changing customer needs.
- Building a system that exposes them to danger or insecurity.
- Leaking customer confidential information.
- Treatment of the customer with disrespect.

Benefits:

- Deploying a system that improves the customer's ability to provide service to their customers.
- Deploying a system that decreases the customer's costs.
- Building a system that improves the customer's ability to control or manage their own work.
- Improving the customer's safety or security.

Customer's customer. The customer of the system customer receives the effects of using the system.

Harms:

- Getting a product or service that costs more than it needed to, or provides less benefit than it could have, because of the system provided to the customer.
- Exposing customers to injury or insecurity by the use of the system.
- Exposing customers to injury or insecurity through the products or results of the customer using the system.

Benefits:

- Getting better product or service than they otherwise could have because of the system provided to the customer.

- Improving their safety or security.

Team. The team is the group of people who make the system (Section 16.2.2). They must work together to get the system built. They do the work in part for the compensation received and for the satisfaction in achievement.

Harms:

- Exposure to physical or mental danger; exposure to security breaches.
- Interference with a team member's ability to work well.
- Undercompensation.
- Social drama and alienation.

Benefits:

- Compensation for work done.
- Satisfaction in achievement and improved social status.
- Social bonding with others.

Organization. The organization hosts the system-building project (Section 16.2.3). It provides a legal entity for the project (and likely others).

Harms:

- Interference with the organization's ability operate, such as interfering with its ability to deliver on strategy.
- Leaking organization's confidential information, especially when that impairs the organization's function.
- Reputational harm that impairs the organization's ability to attract staff or customers.
- Legal liability for illegal or unethical activities, or from harm caused by the system or by the building of the system.
- Inability to attract future capital investment.

Benefits:

- Profit that improves the ability to invest in the future of this project or new projects
- Reputational improvement that helps the organization attract staff and customers

Funder. A funder provides investment to pay for building the system (Section 16.2.4). They do this in the expectation of receiving benefits in the future. Those benefits typically include either returned capital to invest in other projects, some external benefit (such as meeting public policy), or some combination of both.

Harms:

- Misuse or waste of the provided capital investment.

- Achieving less than the expected return on investment or external benefit.
- Liability for unethical or illegal behavior by the project, whether legal liability or reputational harm that impairs its ability to raise capital and invest.

Benefits:

- Receiving at least the expected return on investment.
- Efficient use of the invested resources.
- Getting the expected external benefit.
- Reputational benefit that increases its ability to attract capital and invest.

Regulator. A regulator has the mission to oversee and potentially limit the activities of a project in order to achieve some public aim that may not be the project's own objective, or that involves collective action that the project would not do on its own (Section 16.2.5). Safety drives many regulators today; the regulators have been brought into existence to protect public safety after some organizations have injured the public. Other regulators push for standardization among many organizations, which benefits them all as long as they all conform—a classic collective action problem.

Harms:

- Subverting or interfering with their mission to achieve public policy, such as safety.

Benefits:

- Achieving intended public policy aims.

Society. Every project is embedded in society at large. The people on its team are part of the society. The work uses resources that could be used for different work and thus affect society differently. The purpose of the system is to have effects on society.

Harms:

- Exposing the public to danger or insecurity.
- Breaking laws.
- Causing societal strife and decreasing social well-being, through the use of the system or from technology developed in building the system.
- Increases in geopolitical tension.
- Using resources for the project that could be used to better societal effect doing something else.

Benefits:

- Improving social well-being.
- Improving safety or security.
- Providing capital to invest in other socially-beneficial work.

World. Every project is also part of the world: it exists in space and time and has effects on the land, water, air, space, or whatever other environment in which it exists.

Harms:

- Increase in pollution, including contamination and debris.
- Consumption of limited natural resources.

Benefits:

- Reduction in previous pollution.
- Improved habitability of a place.

Fiduciary duty. This duty is a legal concept for situations where one person places trust in another to act on their behalf. This situation creates an ethical responsibility for the trustee:

A fiduciary's duty requires "[n]ot honesty alone, but the punctilio of an honor the most sensitive", with "undivided and undiluted loyalty to those whose interests the fiduciary is to protect".

—Capital Heat Inc v. Michael Blatner Family Trust, 2009

[New York09]

Fiduciary relationships impose more stringent requirements than those between, for example, a project and its customer, which is legally an arms-length relationship. Any relationship must be examined to determine what level of trust and duty is involved. Most responsibilities in a system-building project will involve lesser levels of responsibility, but those involving money and other resources (managing project funds, for example) likely do involve fiduciary duty.

63.5 Framework

I treat ethics as a matter of personal behavior and responsibility to maximize benefit and minimize harm. While a group may have collective responsibility, only the individual people ("natural persons") in the group have the ability to make decisions and thus be responsible for the decisions' effects.

That said, personal behavior occurs in the context of the project and the team, and the ethics of the group must be considered. The project and team have identities. The project can collectively have a contract or agreement with another person or organizations. The team's social environment and culture place constraints on how its members behave. The project can control resources, such as finance, space, equipment, or intellectual property, that affects how each person can behave.

There are three classes of personal behavior to consider:

1. A person's behavior that affects the project's collective ethical behavior;
2. A person's ethical behavior with respect to the team and its members; and
3. A person's ethical behavior toward themselves.

63.5.1 Collective ethical behavior

Collective activities, such as groups, projects, and organizations, are commonly treated as if they can take aggregate actions and make aggregate commitments. In many countries a group can be formalized as a partnership or corporation that gives the group a degree of legal personhood: a legal identity and the ability to make contracts, for example. People buy things from these corporations; if the thing they buy has a flaw, they expect the corporation to make it right. Similarly, unincorporated groups are often viewed as a collective by those outside them. Such a group can have a reputation, and people adjust their expectations of individual members of the group based on the group's reputation. Non-members expect they can communicate with "the group", perhaps by communicating with a few of its members.

These expectations mean that ideas of ethical behavior by a group are valid in practice.

The philosophical basis for collective ethics is subject to debate. It leads to difficult ideas about collective culpability, which has led in turn to punishment of individuals who fall into a class of people deemed to constitute the collective, but who are not in fact responsible. There are many ways to define a group (a collective class), and short of those ways defined in law for corporations and partnerships, most of those ways are *ad hoc*. As one example, consider the ideas of shared responsibility of states in international law [Nollkaemper20].

For system-building projects, however, the situation is simpler. The membership of the project—the team—is clearly defined for the most part. At the fringes, part-time consultants or *ad hoc* advisors have a slightly ambiguous role but they can be treated, inclusively, as team members for the work they do related to the project. There are largely defined procedures for people to join and leave the team, and a team member's responsibility largely ends when they leave the team.

The behavior of the team is an emergent property of the behavior of its members. This implies that any ethical responsibility of the team as a whole devolves in some way to ethical responsibilities on its members.

The methods for designing safety, security, and other emergent properties (Chapter 46) apply to reasoning about how collective responsibility is related to individual responsibility, and to designing the team's organization and work patterns to achieve desired team ethics objectives. That method begins with naming the harms to be avoided and the benefits to achieve. Avoiding the harms and achieving the benefits are the collective ethical responsibilities for the project. The lists above can serve as

a starting point for harms and benefits to consider. A real project must work out specific variations.

These collective ethical responsibilities are only abstract objectives because the team *as a whole* does not take decisions or actions. These objectives must be mapped onto individual responsibilities, which I address in the next section.

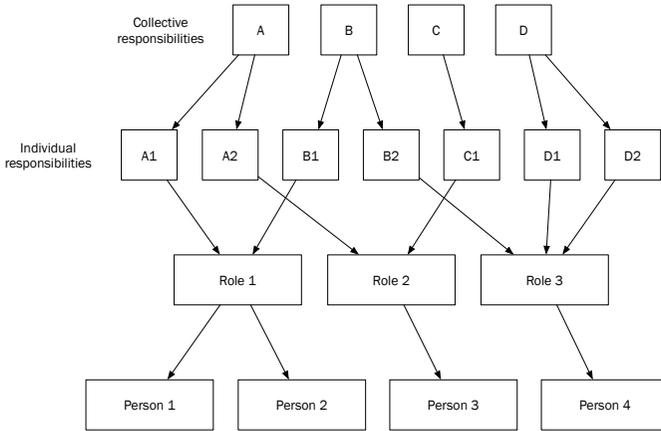
After identifying the harms and benefits, the next step is identifying ways that they could happen. For situations that could lead to harm, the following step is to work out ways that they can be completely eliminated, or ways to reduce the likelihood or severity if they do happen. In situations where harms can happen, one also works out ways to detect that they have happened and how to repair the harm. For situations that can lead to benefit, the following step is to work out ways to make those situations likely.

63.5.2 From collective responsibility to individual responsibility

Collective responsibilities are met or missed based on the behavior of the people in the team—that is, meeting collective responsibilities is an emergent behavior arising from the behaviors of everyone in the team.

Each person has a responsibility to behave ethically in their work, but the aggregate ethical behavior of all the people on the team does not necessarily yield ethical behavior as a collective. For example, the aggregation of each person saying true things might yield a misleading aggregate if there is some information for which no one is responsible and thus for which no one says anything. And separately, I believe there are systemic problems where each person does their part ethically and well, but the procedures they all share end up causing harm. All the individual behaviors must be coordinated in ways that make the aggregate behavior ethical.

The collective responsibilities, then, must be mapped to individual responsibilities. The result is that each person in the team has individual responsibility for some portion of the team's collective responsibilities.



The simplest way to map collective to individual responsibilities is simply to say that every team member shares responsibility for each and every collective responsibility. This can work for a very small team, perhaps up to three or four people, where everyone does all the kinds of work in the project equally. For any larger team, however, this approach fails. Different team members fill different roles and have different skills. Different people may have different authority or power within the team. When people do different work, they cannot all be responsible in the same way for group responsibilities.

When too many people are all responsible for something, each person can expect that someone else will take some action, and that they can get on with their own more urgent-seeming issues. In the end, no one takes the action and the shared responsibility is not met. This is a classic collective action problem.

The mapping is *designed*; it does not happen by chance. This mapping is based on the roles that each person fills, with the attendant authority, and the power that each person has to perform the duties of that authority. The mapping might be implicit in the definition of all the roles, but a project cannot have confidence that the mapping is sufficient without actually analyzing the mapping.

For example, one or two people might have responsibility for tracking the project's finances and the authority to approve or reject expenditures. In filling that role, those people meet part of the project's collective responsibility to prevent misuse or waste of funder's investments, and part of the responsibility to avoid harming a customer or team member by running out of funds.

When two or more people have overlapping responsibilities, or they have related responsibilities, there is the possibility that they will take inconsistent actions. One accounting person might check and approve a purchase request, while another might in parallel check and deny it. Two people might see that some action is needed, but they take actions that interfere with each other. This is equivalent to problems in

safety systems where two control systems can both affect the same controlled components.

Those individual roles are designed and assigned in ways that lead to the desired emergent behaviors. In a properly designed project, every collective responsibility is mapped to a set of personal responsibilities, with no collective responsibility left even partially unmet. This means that all the needed roles are defined and that they are filled by people who are able to do the work associated with the role.

A project's leadership has a responsibility for designing and maintaining this part of the team's organization, just as they have responsibility for designing other ways that the team is organized (Chapter 57). This design begins when the team is first put together. The leadership is responsible for monitoring whether the team's organization is doing a good job of leading to ethical behavior—which will involve defining ways to observe or measure how the team is doing. The leadership is also responsible for making changes when the team's organization is no longer suitable, either because the team has changed in size or complexity, or because the previous design isn't working as intended.

63.5.3 Robust collective responsibility

In practice, no team is entirely perfect. The members will occasionally misunderstand their responsibility, or forget, or lack key information needed to take ethical action. Most teams, at some point, have people who take unethical actions maliciously.

The organization of ethical responsibilities throughout the team must therefore be designed to handle lapses and ongoing problems. The team must be structured so that ethical lapses are made unlikely. The lapses that do occur must be detected and rectified.

There are ways the team's organization can help.

1. Provide clear definitions of ethical responsibilities to everyone on the team, so that they know their responsibilities and the standards of behavior expected of them. Providing definitions and regular training helps avoid people not understanding their responsibilities.
2. Build checking and correction into tasks. This allows quick catching of potential problems and defines who has the responsibility for rectifying problems that do occur. It also normalizes checking, so that people accept it rather than reacting as if the checks mean lack of confidence. Checking the ethical aspects of decisions and actions is just as important as checking the technical content of designs and implementations (Section 8.2.6).
3. Incentivize ethical behavior and disincentivize unethical behavior. This reduces the likelihood of someone choosing to perform an unethical behavior.

4. Build a safety net to detect and handle ethical lapses that happen despite the other methods. This includes clear and credible ways to report and address potential lapses.

Incentives. Managing the incentives for behaving ethically is the primary mechanism for guiding people toward making good decisions. People behave according to their understanding of their interests. Some interests derive from internal sources: their own sense of self-worth deriving from morality, for example. Other interests are external: for what will the person be rewarded or punished? The most effective external incentives come from a team's culture. A team that has a culture promoting ethical behavior will create a sense in each person that their social standing within the team is based in part on how they behave.

Poor incentives have perverse outcomes and drive people to behave unethically, placing some other value above ethics. Those incentives are often about parts of people's work that is not in itself an ethical question. Consider, for example, what happened when Wells Fargo Bank created targets and incentives for its staff to meet high sales goals. Those targets are not, in themselves either ethical or unethical; setting goals is a standard practice used to promote business growth. In the Wells Fargo case, however, those incentives led to widespread fraud by the staff and to criminal investigation of the bank [DOJ20].

It is thus a collective ethical responsibility to implement incentives that drive people toward ethical behavior. This is especially the project leadership's responsibility. These people are also responsible for monitoring the effects of policies they set and changing them when they find that the policies are causing problems.

Safety net. The individual ethical responsibilities and the procedures that team members follow will not catch every lapse, since people sometimes simply forget and sometimes someone acts maliciously. Checking and incentives will reduce the rate at which these lapses happen, but not eliminate them.

The project's ethical structure must, therefore, include some kind of safety net to catch those situations that still happen.

These safety nets use on reporting and auditing to detect when problems occur. Reporting can be formal or informal.

A project should have a formal means for people to report problems they observe that aren't being addressed using normal channels. These problems can be ethical lapses that have occurred, or problems with the ethical structure of the team that are likely to lead to lapses. The project should have a clear method for reporting, and the team should have confidence that reports will be taken seriously and handled.

Informal reporting complements the formal channels. Anyone in the team should be confident that they can have a conversation with people who have authority to handle problems, and that if they convey their concern clearly those with authority will look into the problem.

Audits provide a different safety net. From time to time, people should independently check ethically-sensitive work. This is common in accounting, to check whether money and other resources have been properly accounted for and, in the process, detect problems.

Reporting can be misused. People might create false reports, abusing the system. It can interfere with team cohesion, leading to an environment of mutual suspicion. The best solution I know is for the project leadership to make it clear that reporting problems is a personal responsibility and it is valued by the leadership, or that taking action to correct a lapse is valued. Those with responsibility to act on reports must show to the team that they will act on reports, and that they will do it fairly and promptly. That is, it is up to the leadership to model the behavior they expect of others in order to create a social norm for the team.

The team leadership is responsible for designing the safety net and for ensuring that it works. The leadership monitors the reporting process and ensures that the roles for investigating and handling reports are staffed, and that those who fill the roles are given the support they need. They need clear authority and power to investigate and take action.

Everyone in the project is responsible for reporting problems they observe.

When the safety net fails. Sometimes the safety net will not work as planned—or a project may have failed to implement a reasonable safety net to catch ethical problems. A problem will occur and will not be addressed and repaired.

Team members have a responsibility to report and try to get a problem addressed, even if others fail to act on the report. While everyone has a responsibility to make a report, few people have the authority or power to respond to a problem. If someone tries to report a problem, they cannot be held culpable if others don't address it.

The most serious cases like this that I have encountered have happened when a project leader does not want to hear or believe that there is a problem. One project was organized as a large number of competing contractors. The project leadership expected that the contractors would work together to provide joint solutions. At the same time, their contracts did not allow collusion, and being in competition disincentivized collaboration. The project leadership was convinced that this would work, and did not believe reports that the approach was not working. In the end, the problem was not corrected and the project was canceled with few productive results.

This kind of situation creates an ethical problem for team members. In the example project, some of those who saw that the project's contracting structure was not working, and tried to bring attention to the problem. Others felt that they were obliged to follow the leaders, regardless of whether the leaders were right or not; they chose not to address the problem with project leadership. Those who choose to turn a blind eye to a problem are failing an important ethical responsibility. If they truly believe that a report will not be acted upon, then continuing to support a project is itself a further ethical lapse, and they have to decide whether to continue on the project or leave. Sometimes the potential benefit—to society or customers, perhaps—outweighs the harm caused by continuing to work on the project. Other times it does not. Each person who finds themselves in this situation must decide for themselves.

63.5.4 Ethical behavior toward team members

Each person in the team interacts with other team members, and so affects them. This creates a responsibility between team members for these effects. I discussed the model of teams and more of the responsibilities involved in Section 19.3.

The ethical responsibility is, as before, to provide benefit where possible and to avoid harm.

Beneficial behaviors toward others on the team include:

- Building and maintaining trust by following social norms, treating others honestly, and doing work well. This helps people work together (Section 19.2.3).
- Communicating clearly and respectfully about technical work, asking questions, raising issues, and reporting progress (Section 19.2.1). Clear communication is a responsibility of both the speaker and the listener; the listener should be confirming their understanding of what is being said and showing that they have heard the information.
- Acting to translate between people or groups that use different jargon and concepts, so that they communicate or negotiate accurately.
- Helping others improve their skills. This can come through mentoring or guidance. It can also happen through group participation in learning. Even just providing encouragement helps.
- Helping people understand when they have made a mistake or hold erroneous belief, especially if done in a kindly way.

On the other hand, there are harms to avoid:

- Breaking trust through dishonesty, failing to do agreed-upon work, doing work of poor quality, or talking maliciously about others.

- Making false claims about a skill, capability, or experience that one does not actually have, which leads to being unable to do needed work and causing loss of trust.
- Causing injury to others, or placing them in danger of injury. This includes mental and emotional injury as well as physical injury.
- Treating others disrespectfully.
- Interfering with someone else's ability to do their assigned tasks and meet their responsibilities. This includes making decisions that someone else has authority for.
- Failing to help someone in order to cause problems for them.
- Sharing malicious information, whether true or false, in ways that cause division within the team, make it harder for them to work together, and make it harder for people to repair harms that have occurred.
- Breaking legal workplace standards.

There is another harm that is often not understood as harmful: covering for someone to hide a problem. I have seen many people try to be what they consider nice to someone by cleaning up a problem the other person has made so that that person doesn't face consequences. This can be harmful to that person (and to the team) if it prevents the person who created the problem from learning how to do better, or if it prevents them from being reassigned to tasks they are better suited to doing. In the worst case, this leads them to repeat the problem because they have not recognized that they made mistakes. Helping someone is good, but it must be done in a way that it does not stop that person from growing.

Leadership responsibility. Project leadership has additional responsibility. Some of them are responsible for assigning tasks to team members. A team member's sense of achievement comes from being assigned work that they are able to do, that challenges them, and that they believe has purpose. Those who are responsible for assignment have, therefore, a responsibility to assign work that takes advantage of a person's strengths and provides opportunities for them to grow. Of course not every task will be challenging and fulfilling, but over all the task assignments should lead to the team member's sense of achievement.

The responsibility for task assignment comes with a responsibility not to demand more time and effort from people on the team than they can reasonably provide without harming themselves. Projects that work their people to burnout are behaving unethically.

Project leadership also has a responsibility to give people autonomy to do their work (see the sidebar in Section 19.2.4). This must be coupled with clear communi-

cation and providing ways for the assignee to ask questions, report progress, and get guidance.

Those with experience or power in the team have a responsibility to help other team members grow. They should be providing guidance or mentoring people with less experience.

Finally, some people in the team are responsible for selecting people to join the team. They are responsible for screening potential team members so that those who join will respect the team's social norms and behave ethically.

Problems. Sometimes someone on the team behaves unethically. They may be dishonest; they may disrupt team work; they may not act safely. The team's organization and procedures must be robust enough to handle these situations and restore the team to good working order. This is related to the principle in Section 8.4.5: that projects define how team members should communicate about exceptional situations. The discussion about the ethical safety net above applies.

Note that this kind of problem is especially serious when the person behaving unethically toward others is in project leadership.

It is reasonably likely that a team will include one or more people with a mild form of antisocial personality disorder, based on a cited prevalence of “between 0.2% and 3.3%” [DSM-5, pg. 661]. Hiring practices may screen some of these people out. Others may be aware of their behaviors and take mitigating actions. Nonetheless, teams should expect that sometimes a member will exhibit ongoing deceit, lack of empathy, and superficial charm. People who exhibit these behaviors can behave unethically toward others on the team and cause a breakdown of trust. (I have been part of one team where this happened.) It is thus a responsibility of a team's leadership to be alert to people who exhibit these behaviors and take steps to prevent them from harming the team.

Making ethical decisions depends on having accurate information. Withholding necessary information can lead to decisions that cause harm, and accurately sharing information reduces this risk. Shared false beliefs, however, are a more serious problem. I have seen many teams that begin to treat some false information as true, and that communication within the team, along with confirmation bias, provides social reinforcement for that shared false belief. This has led to adopting designs that do not meet requirements, or adopting requirements that do not reflect actual customer needs. Avoiding, detecting, and resolving this kind of self-reinforcing false belief is necessary to avoid unethical decisions (Section 59.2).

63.5.5 Ethical behavior toward self

The general ethical treatment of one's self is a complex topic, and here I focus only on those aspects that relate to one's work on a project. Even so, this account is surely incomplete.

There are a few principles that I use as guides.

1. Sustainable effort. One should work as hard as possible that allows one to continue for a long time. Each person on a team carries knowledge about the work they have done, and their value to the project generally increases as time goes on. If someone works too hard or too much, they can get to a point where they make mistakes at an increasing rate. Moreover, the work that someone does on a project will almost never be more than a small fraction of their life: they will likely work on many projects over time, and they will spend the majority of their time not working. Preserving one's ability to keep going matters both to that person and to those around them.
2. Refrain from self-harm. Few projects are worth injuring one's self. Work in ways that avoid physically dangerous situations, using safety equipment, following safety protocols. Avoiding emotional and mental self-harm is also important.
3. Achievement. Work should be about more than just getting paid for effort spent. It is a source of satisfaction and worth. It is also a source for social standing and connection. Ethical behavior works to avoid compromising these benefits.
4. Right job. When one is in a role that fits, it can lead to a sense of achievement. Sometimes, though, a role is not the right fit; the skills one has are not what the role needs. While this can lead to education and new skills, it can also lead to a sense of lack of achievement or disillusionment. Communicating with others about what one needs, and identifying when a role does not fit, can help avoid these problems. At some points, if the fit is bad enough, one should try to move to a different role or to a different project.
5. Self-respect. When one believes that they have had to compromise their basic moral principles, or have had to do work they consider to be unreasonably poor quality, it is hard to maintain self-respect. A sense that self-respect is threatened or compromised can be a signal that there are ethical problems with the work one is doing. Recognizing and acting on the situation can help resolve the problem.
6. Curiosity and education. For the kinds of work involved in building systems, regular intellectual stimulation is important for ongoing satisfaction. Learning new things and improving skills are also important for maintaining one's contribution to a project over the long term.

Following these principles requires honest self-assessment. A false belief about one's capabilities leads to taking on tasks that one is not actually able to do. Overestimation of what one knows leads to poor decisions (see Section 8.2.7) and to failure to learn. Failing to pay attention to mental and physical fatigue leads to poor thinking and unsafe actions.

Some of these principles should not be taken too strictly. Maintaining a sustainable workload is necessary, but that does not mean never having a long day or week. Keeping one's energy and fatigue within reasonable bounds over the long term is what matters; a period of hard work needs to be balanced with some time for recovery. Similarly, no job is a perfect fit for skills and interests, but the work should fit the person well enough to lead to satisfaction. There will be periods when a position doesn't go so well, and one has to last through those; it is when a mismatch is going to continue without sufficient improvement that one should consider changing. In some (thankfully rare) projects, the importance of the benefit to others is so great that self-sacrifice is appropriate.

63.5.6 Some principles

The harms and benefits discussed in this section vary widely, but many of them can be boiled down to a few principles. These include:

- *Trust and honesty.* Customers, funders, regulators, and team members all want to know that they can trust that they are being presented with accurate and complete information, and that any information they give will be treated as correct.
- *Conflict of interest.* People and organizations want to believe that others are acting in their best interest, without hidden motivations to act otherwise. People don't want to have to question what other people do in their name.
- *Commitment and follow-through.* People want to believe that when someone make a commitment to do something, it will get done—or there will be honest notice that it will not be done as planned. People want to believe that they do not need to doubt what others say they will do.
- *Physical safety or security.* People want to be that if they use a system, they will not be in unreasonable or hidden danger. They do not want to have to wonder whether their information or property will be at risk using a system or working together. They expect that by using a system or working together the world and society will not be damaged.
- *Emotional safety and respect.* People expect that they will be treated reasonably and fairly, and that if they behave ethically and do good work they will be given respect and esteem from those they interact with. They expect not to be in danger of being falsely accused, whether directly or by insinuation.

63.6 After an incident

People will make mistakes. People will observe problems or lapses and be unable to take action to address the problem. A team will find itself with a malicious actor. What then?

The lapse must first be detected and repaired, if possible.

That is not the end, though. The next step is to work out why the lapse happened and to find ways to make it, or problems like it, unlikely to happen again. This can be treated as an incident analysis, and can use all the tools for that kind of work (see, e.g., Leveson [Leveson11, Chapter 11]).

63.6.1 Repairing harm

If there has been an ethical lapse, the situation must be repaired as much as possible. This means identifying the harms that have been done and finding ways to reverse those harms.

Repairing an ethical lapse requires making clear statements acknowledging that the lapse has occurred and show the plan for repairing it. These statements must be heard and understood by those directly affected by the lapse, and by those who have observed it or its effects. The statements should be public if the harms have been publicly visible; they should go to a smaller audience if the situation did not affect the public.

When one person has harmed another, we expect the one who has caused the harm to honestly take responsibility for what they have done as long as they are a competent adult. Indeed, one of the marks differentiating a child from an adult is that the adult is able to take this responsibility. Too often people fail to do this: they deflect responsibility to others, they deny that the harm happened or that it mattered, they attack those who try to hold them responsible. Sometimes they try to hide the harm that has been done, out of fear for the consequences they will suffer.

All these behaviors occur in system-building projects. Individuals are individuals, and some will have difficulty taking responsibility for the effects of their actions. Organizations and teams sometimes exhibit the same behaviors: they deny or avoid consequences; they try to shift responsibility; they try to minimize what happened. This has led to public cynicism about corporations running public relations campaigns to distract or deflect.

It takes a kind of courage to for an individual to face consequences of their actions and repair the harms. It can take collective courage for people in a team to do the same. Repairing harm causes a loss to the one who caused it: time, money, reputation. Accepting those losses is hard. Nonetheless, failing to take the steps to acknowledge and repair the harm caused is a serious moral failure, and the people or

organization involved deserves condemnation. It is a sign of lacking a basic capacity to function in a civilized society.

Repairing physical harm includes caring for those who have been injured so that they recover. Death and some injuries cannot be completely recovered and the repair must provide some kind of compensation. In the times when the *weregeld*⁴ applied, there was a code that defined what the compensation should be. In modern times there is no such simple guide, and so individuals have to work out for themselves what is right and just.

Economic harms are usually easier to repair. In many cases it can be repaired by paying money to compensate for losses. If the harm caused a company to shut down, or caused a permanent loss of capability to the economy, it is harder to see how to repair the harm.

Reputational harm requires making it clear to all who might have heard false information about the person or organization that has been harmed. Repairing this harm involves communicating with those who might have heard so that they know that the harm was done, that it was incorrect, and that they should restore their esteem of the one harmed.

Note that in this analysis of repairing harm I do not include punishment. Punishment in itself does not repair a harm. The threat of punishment for doing harm can create an incentive not to do harm, and punishment may be part of ensuring that such actions will not happen again in the future (as I discuss next). But punishment does not make things right, and must not be confused with repair.

63.6.2 Improving for the future

The other responsibility after an ethical lapse is to work out how to make it less likely that something similar will happen again. The analysis seeks to determine first what happened, and then what actions, decisions, or situations happened that led to the lapse. That understanding guides the choices for how to avoid problems in the future.

One goal of [the analysis technique] is to get away from assigning blame and instead to shift the focus to *why* the accident occurred and how to prevent similar losses in the future. To accomplish this goal, it is necessary to minimize hindsight bias and instead to determine why people behaved the way they did, given the information they had at the time.

—Leveson, Chapter 11 [Leveson11].

⁴*Weregeld* was a system of payments for the death or injury of a person in old Germanic societies, including Anglo-Saxon England. There was generally a schedule of the payment required to the injured person or their family, based on the social status and gender of the one injured, and the nature of the crime.

After the incident is understood, the team can make changes to its structure to avoid similar problems. The analysis may also reveal potential causes for ethical lapses that have not happened yet. The response is to change how the team works:

- Understanding about responsibilities or roles. Someone may have not understood about their responsibilities, or the responsibilities that others have. Education can address this.
- Procedures. There may have been insufficient checking done, or checking was not effective. Making the procedures more robust.
- The definitions of responsibility and authority of different roles. Perhaps there was a collective responsibility that did not map to individual responsibilities, or two people had conflicting responsibilities. Changing the responsibilities assigned to roles to make them complete or avoid conflict can help.
- Insufficient staffing. Some important role may be left unfilled, or the person in that role may be overworked so they cannot meet the responsibilities. Adjusting assignments or bringing in more resources can help.
- The power that some people have. Someone may have had responsibility but lacked the power to take action. Someone may have had power and took action, but did not have responsibility or authorization and thus created a problem. Educating people, redefining authorities, and adjusting power in the team can help resolve this.
- Role assignments. A role may be assigned to someone who lacks the skills to do the associated work ethically. The role can be assigned to someone else.
- Team membership. A person on the team may not be willing or able to behave in ways that consistently meet project ethical standards. That person should be removed from the team.

Again, note that the focus here is not on culpability, leading to blame or punishment. The focus is on making whatever changes are necessary to make it less likely that similar problems will happen again. In some cases the remedy will be to remove someone's authority or to remove them from the project—but this does not involve ideas of retribution or justice. It is solely in answer to the question: will the person in question improve their behavior, or will they repeat similar behaviors? If they will likely improve after being made aware of their behavior or after education, then they can continue in their role. If they are unlikely to, then they should not. In some cases, the person involved may be able to change their behavior, but trust within the team may be irreparably broken, which is effectively the same outcome as if they person will not change their behavior.

63.6.3 Who does the work

The responsibility for repairing harm and ensuring it does not happen again can fall on different people depending on what happened. If the harm was caused by one person or a small part of the team and its effects were visible only within the team, then the people who caused the harm are responsible for repair and changing their behavior. However, if the harm came from a systematic problem with the team, affected others outside the project, or was publicly visible, the project leadership must take the responsibility.

The work of repairing harm and changing behavior must be visible to those who were affected, and to those who had a part in causing the harm. This may include stakeholders like regulators, who represent others who have been affected (e.g. the public). While the objective is to include everyone involved, others need not necessarily be aware or involved. If two people within the team have a falling out and have to restore trust, the public does not need to know—and neither do people on the team who do not interact with those people and were not affected by their disagreement.

Some harms occur because of the behavior of many people in the team. Accident analysis has found that many accidents have multiple, systemic causes. Leveson [Leveson11, Section 2.2.4] discusses an analysis of the Union Carbide 1984 Bhopal incident, where the actual causes of the incident were a combination of many systemic failures, from failed or inoperative safety equipment, poor training and communication for workers involved, and corporate cost-cutting for several years before, and regulatory deficiencies. In incidents like this, addressing one specific person or problem will not avoid future harms; instead, it just shifts the potential causes around and gives a sense of complacency. Systemic problems are often less extreme. A corporate culture that gives power to some people beyond their actual authority can lead to small but serious breaches of trust between people on the team. The problem cannot be rectified just by taking power from one person; the breach ultimately happened because of a cultural problem that requires a change to the team's practices in general.

Some harms are visible outside the project. They may be harms to customers, to regulators, or to society. These are lapses of the collective ethical responsibility of the project, even if they stem from one person's behavior. These incidents require collective response to show that the project *as a whole* can be treated as responsible.

Responses to systemic and externally-visible problems require the involvement of a project's leadership because the responsibility is not limited to one or a few people. Analyzing what happened may require a team-wide perspective or a perspective that includes customers and others. Repairing the harm may involve the project as a whole, including making public statements.

I discussed the idea of an ethical safety net, and what happens when the safety net fails, in Section 63.5.3 above.

63.7 Case studies

The discussion of ethics so far has been mostly abstract. The following case studies illustrate what can happen in real situations.

63.7.1 Private benefit, public harm in UTM systems

The situation. I worked for a while on the UTM problem of providing air traffic management for unmanned aerial systems (UAS, also known as UAV). This system would work with the existing air traffic control (ATC) system that manages aircraft flying in controlled parts of the national airspace. Both UTM and the existing ATC have common goals:

- Maintaining the safety of flying aircraft, through avoiding collisions, weather problems, and restricted airspace;
- Maintain national security by tracking what objects are in the airspace, keeping them on safe trajectories, and detecting potential hostile objects;
- Maintaining the safety of people and property on the ground, by avoiding collisions in the air and collisions with the ground; and
- Maintaining fair access to airspace, so that any qualified operator has the opportunity to perform a flight, subject to meeting safety constraints.

In the existing ATC, this has led to a system whereby controlled flights file flight plans, follow standard routes (especially at low altitudes), and receive permission and instructions from air traffic controllers while in flight. The controllers use a number of tools to track where aircraft are and where they are going. All these aircraft are under the control of a human pilot (even if the pilot is remote). The ATC service is provided by an Air Navigation Service Provider (ANSP); most of these are government agencies either for a nation (the FAA in the United States) or a geographic region.

UTM is intended to add unmanned systems, both remotely piloted and autonomous, to the airspace while maintaining the safety and policy goals. While the ATC system evolved gradually over several decades, the UTM system is being developed deliberately.

There is a fundamental design question in how UTM systems will be organized and operated: organizationally centralized or distributed? One possibility is to structure them in much the same way as the existing ATC system. This would mean each region or nation would create and run a UTM system; it would be government-funded and under direct government control. Another possibility is to focus

on private service providers, who must follow government regulations but operate separately. There could be multiple private service providers operating in a region, each managing its own set of aircraft. These private service providers would then have to cooperate in real time to handle potential conflicts between aircraft managed by different providers.

There are advantages to each option. An organizationally centralized system can follow the existing organization structures, with added capabilities. Such a system can also be managed as one, so that all the parts can be built to work together correctly. A distributed organization of multiple service providers, on the other hand, can move the capital investment to design and deploy UTM to private sources, avoiding government financial allocation processes. A distributed organization can also provide competition that leads to innovation in UTM capabilities.

Ethical responsibilities. The choice between centralized and distributed organizations is an ethical question that is informed by technical considerations. It can be addressed by an analysis of the ethical implications of either choice. This starts with the claims of potential benefits and harms coming from making one choice or the other.

Most of the benefits that a UTM system can bring do not depend on which choice. UTM can enable a number of uses for UASes, and these can be expected to improve economic efficiency for delivery, infrastructure inspection, and local environment sensing, among others. These benefits accrue whether the UTM system is organized centrally or distributed.

There are some possible benefits that differ between the choices. A distributed organization might create a competitive marketplace for UTM services, which would lead to lower prices for users and greater innovation; that innovation could lead to efficiency for users or enable new uses for UASes. A centralized organization that is aligned with the existing ANSPs could be more effective at managing the combination of manned and unmanned aircraft in a combined way.

Next consider potential harms to be avoided.

- Delaying the development and deployment of UTM services delays when the benefits from using UASes begin to arrive.
- The UAS industry is still new. Building and deploying something too early, before the real needs are understood, could lock the industry into UTM systems that will not enable important uses for a long time.
- Failures to accurately coordinate air traffic between UASes or between UASes and manned aircraft can lead to injury and damage on the ground and in the air.

- Security failures can lead to national security incidents, where UASes are hijacked and used as weapons.
- The design of the UTM system can lead to anti-competitive behavior by UTM service providers or UAS operators, or impede fair access to airspace.

The choice of which kind of UTM system organization can affect each of these choices. A centralized organization that is part of a government agency can be slow to develop and deploy, leading to the first two harms. Distributed organizations have difficulty with coordination and evolving their systems. The security exposures are different between the two organization choices. Questions of anti-competitive behavior are handled differently when the system is centrally run compared to a distributed organization.

The next step is to consider how these harms could happen. What situations would lead to them?

Once ways that the harms can occur are understood, one can then investigate each situation. Which ones are likely and which are not? Are there technical or social ways to prevent some of them occurring? Are some situations possible with one design choice and not the other? Because these questions are being asked early in the process of working out a system concept, it won't be possible to answer them with precision but they can be answered well enough to guide the big decisions such as whether centralized or distributed organization is a better choice.

Committing to a design approach like this without knowing whether a basic ethical requirement can be met is unethical behavior. Those who have responsibility for making basic design choices therefore have an ethical responsibility to explore these questions honestly, completely, and carefully, and then to use what they learn to inform their choice.

The principals that decide on the basic organizational structure of UTM systems have the primary ethical responsibility; I will call these the government principals. In practice this includes the civil aviation authorities of various nations, along with the international organizations that coordinate their work and the research or advisory bodies that help them. In the United States, the FAA is the civil aviation authority that is authorized to make regulation, subject to legislative direction; the FAA is informed by research groups at NASA and Federally-funded research organizations.

The government principals do not do their work in a void. The organizations wanting to use UASes and those developing UTM systems pressure these principals to adopt favorable regulations. At the time that the civil aviation authorities and their advisors were working out how they would structure UTM, there were several companies working on UTM systems and lobbying NASA, CAAs, and other organizations.

These campaigns to influence the high-level structural decisions impose ethical responsibilities on both the governmental principals and those who are lobbying or proposing solutions. The primary responsibilities for the government principals remain making decisions about the overall structure and capability of a national and international UTM system that maximizes benefit and minimizes harm. That decision must consider all potential stakeholders, including the public and world; this means including government and social policy goals like fair access to airspace and public safety. The organizations developing UAS and UTM systems have responsibilities to help the government principals meet their aims, and a corresponding responsibility to avoid interfering with the government principals meeting their responsibilities.

In particular, organizations developing UTM systems have a responsibility to act in one of two ways: either they provide information to the government principals only about the needs they will have as a UTM developer and what they project UAS users will need, or present potential solutions that will further the government principals' responsibilities. Pragmatically, a UTM developer should anticipate that the CAAs will consider decisions that are not what the UTM developer would want.

Example harm: poor coordination between UASes. Avoiding collision or interference between aircraft is one of the most important objectives for a UTM system. This is exercised on different time scales. Long-term management involves deciding how the airspace will be organized, where flights will not be allowed, and where routes or corridors might be placed. "Strategic deconfliction" involves managing airspace and flights in the medium term by ensuring that flight plans for each aircraft will keep them separated. "Tactical deconfliction" is about the short term, giving directions to aircraft in flight as conditions change: for example, when one aircraft deviates from its planned path, or when weather changes and aircraft must follow different routes or be spaced further apart. There is a final time scale for the immediate direction when one or more aircraft are in immediate danger of collision. Long-term management makes strategic deconfliction easier; strategic makes the need for tactical deconfliction less frequent, and so on. Long-term management is generally handled by a CAA or an ANSP. The other time scales are handled differently using a centralized UTM organization or a distributed one.

Consider two UASes that will be wanted to fly paths that cross each other. The overall UTM system must ensure that the two will have flight paths that allow them to maintain separation. This could be done by moving the geographic path or the departure time so that the two will not get close. It could also be done by tracking the projected overall density of aircraft over time, so that if the two follow paths that

will get too close there is room for one or both to deviate a bit to stay apart.⁵ These are all uses for medium-term strategic deconfliction.

Or consider two UASes that are nearby but maintaining separation, when one of them has an emergency that requires deviating from their planned route to land immediately. This can create a situation that must be handled right away, giving the aircraft with a problem a path to an emergency landing and moving other aircraft to give it room. These are uses for strategic deconfliction or emergency response.

The question is then: what would lead a UTM system to fail to keep them separated? That is, what scenarios would cause harm by leading to a collision or a near-miss?

Several scenarios occur in any UTM structure: radio communication failures, for example, or failures of sensor systems that track the aircraft. Flaws in the UAS control systems can lead an aircraft not to follow instructions. These must be understood and mitigated, but they do not bear on the question of what UTM system structure to choose.

A centralized UTM structure can use the approaches used in the current ATC system, where the airspace is divided into regions with a controller for each region. This means that one entity (whether server or person) maintains situational awareness and has all the information needed to detect a potential loss of separation and determine appropriate recovery actions. Failures of a server can be mitigated by redundancy. Bugs in the detection and reaction algorithms can be mitigated by careful design and analysis.

A distributed UTM structure has many more ways it can fail to detect and resolve loss of separation. Consider if the two aircraft are being managed by different UTM providers. Some potential situations include:

- The two providers using incompatible separation detection and resolution algorithms. Both providers might tell their aircraft to do maneuvers that end up moving them closer together rather than apart. This is an example of conflicting control actions (see [Leveson11, Figure 8.6]).
- The two providers might not share complete or accurate information. This may lead one or both providers to make inaccurate decisions about whether they need to issue commands to maneuver, or what that maneuver should be.
- The providers might implement consensus mechanisms that have bugs that lead to inconsistent control decisions.

⁵If there are not too many aircraft in an area, there will be room for them to maneuver without in turn moving into a potential collision with another aircraft [Golding18]. The current ATC system for manned aircraft relies on maintaining a low density so that aircraft can almost always maneuver to maintain separation. Some usage scenarios for UASes show that traffic in urban areas may far exceed the densities that allow such independent maneuvering to maintain separation unless the airspace is more actively managed than is the case with the current manned ATC system.

To propose one approach or the other, an organization has a responsibility to determine whether there are technical or social mitigations for the potential hazardous scenarios. For a centralized structure, can redundancy address failure scenarios? For a distributed structure, can one ensure that different implementations are compatible? (The assurance problems for a distributed structure are in fact beyond the current state of the art at the time of writing.)

Because the choice of how to structure the UTM system affects the harms that can happen—and indeed the industry anticipates that some UAS accidents could kill or injure large numbers of people—making a careful decision based on this analysis is an ethical responsibility, not just a technical task.

Example harm: anti-competitive behavior. UTM service providers can attempt to act anti-competitively, leading to lower competition in some market. The lower competition is usually considered to cause harm by reducing product improvement or innovation or increasing prices. In some cases lower competition has led to lower product quality because there is nothing limiting the producer from reducing its costs. There are many kinds of anti-competitive behavior, including collusion between competitors (such as forming a cartel), pricing at lower than cost to drive out competitors (dumping), forcing customers to buy multiple unrelated products together (bundling), and creating exclusive contracts where a customer may only purchase from the company or a supplier may only supply to that company (exclusive dealing).

Some monopolistic practices are supported by governments because they yield benefits greater than any harms from the monopoly. Patents give the patent holder an exclusive right to an idea for a limited period in exchange for public disclosure of the idea. The argument is that time-limited exclusivity gives the inventor time to recover the cost of developing the idea, while ensuring that the idea is made public so that others can use it after the exclusivity term has ended, or improve on it. There are also natural monopolies, where there is no reasonable way for providers to compete. Water services in urban areas are one example; the cost of laying two sets of pipes throughout a region is greater than twice the cost of laying one set of pipes. Locks on a river are another example. The provision of public goods is a third example, discussed below.

Consider a geographic region where there is UTM service and UASes are operating. If the UTM service is implemented using a centralized structure, that service is a regional monopoly. As long as that regional monopoly is sanctioned, anti-competitive behavior is by definition not a concern. If it is implemented using a distributed structure, it *can* be competitive, and anti-competitive behaviors become a concern.

This is not a theoretical concern. At one public meeting, representatives of a company developing UTM services proposed that the company I was working for should join a cartel they were considering forming with a third company. This cartel would have provided superior service to the cartel's customers through enhanced coordination among the cartel members. This behavior was unethical, and I indicated that I would not support joining them. I believe that the people making the proposal did not actually understand they were planning a cartel; rather, they were driven by technical enthusiasm for providing better service and saw one way to do so. This example shows that unethical behavior does not always arise out of intent to do harm; it can also happen when people do not consider the ethical effects of what they are doing.

A different, theoretical, example might be if a UTM service provider had exclusive contracts with surveillance providers in a region. In that case, only that service provider would be able to sense where uncooperative aircraft were. Other service providers would not be able to provide a safe service in that area and thus be locked out.

As noted above, government principals have the primary ethical responsibilities related to anti-competitive behavior because they are the ones to decide how the system will be structured, and encode that into law or regulation. To meet this responsibility, they must first determine what harms they want to avoid, and then work out what behaviors can lead to those harms. They then must make decisions about structure and craft regulation that will mitigate or avoid those behaviors. The mitigations will include incentives for providers to behave well, and penalties for disallowed behavior.

Building a system of incentives and penalties requires the ability to monitor UTM provider behavior and enforce regulations. The monitoring must be accurate and complete, so that the regulator has a correct understanding of what the providers are doing. This can be done by the regulator itself, by the community of providers, or both. However, regulatory capture can render the first ineffective, and collusion can render the second ineffective.

In the end, a regulator cannot police every action by every participant at reasonable cost. Regulation only works if most providers follow the rules most of the time. The more the regulator has to monitor the details of what each provider does, the greater the cost to both the regulator and the provider. Greater monitoring also impedes the ability to innovate with new services or new ways of providing existing services.

Those organizations or projects developing distributed UTM systems thus have a complementary responsibility to avoid anti-competitive behavior. This responsibility includes working with regulators to define behaviors, monitoring, and incentives

in a way that will lead to the public benefit. This is quite different from ordinary lobbying, where a provider only considers what will benefit themselves. The organizations also have a responsibility to honestly implement and report monitoring information, and not alter or conceal information. This collective responsibility passes down to individuals throughout the provider’s organization.

A UTM system developer has an ethical responsibility to avoid *regulatory capture*. This is when a regulator or similar agency is co-opted to further the interests of the developer at the expense of others, including society or the world. Regulatory capture is an example of impeding the regulators from doing their mission to achieve public policy—in this case, subverting the regulator.

This implies that a system developer must take extreme care in their interactions with regulators and similar stakeholders. The developers have a legitimate interest in informing regulators about the developer’s needs, such as the kinds of customers they expect to support, the return on investment needed to get capital to develop a UTM system, or the results of safety and security analyses they have performed. They can suggest solutions or regulatory approaches, but if they do so they must disclose the full implications of those suggestions. They must be clear about the potential conflicts between their interests and the regulator’s interest.

For UTM system organization, several jurisdictions including the United States have implicitly adopted the distributed structure. Anecdotal reports are that this was done after potential UTM development organizations presented that approach as the preferred structure to government organizations. I have not been able to find any analysis from either the developer or government organizations that investigated the policy implications or even technical feasibility of this choice. If so, then this is an example of unethical regulatory capture.

Example harm: unfair airspace access. The general policy of many governments and their CAAs is to allow aircraft operators “fair access” to the airspace, subject to following safety regulations and subject to the capabilities of the aircraft. “Fairness” is a complex topic for which there is no single agreed-upon definition, but there has been preliminary work on defining it for UTM systems [Sachs20]. Informally, a UTM system would be unfair if it provides inequitable service to different UAS operators, where one might over some period have more flights approved or get more favorable reroutings during an in-flight conflict resolution. Unfairness might arise two ways: first, when regulations for conflict resolution are followed, but those regulations are flawed; or second, from one or more operators not following the regulations (cheating).

The short-term harm is that unfair access creates economic advantages and disadvantages for different UAS operators. The longer-term harm is that it creates distrust in the system, which in turn can lead to greater harms:

When participants feel that they are being treated unfairly, they are more likely to take action on their own to assure better outcomes for themselves. In aggregate, this results in a breakdown of cooperation more widely, and puts stress on the entire system. Fair allocation improves customer satisfaction; operators are more likely to participate in the decision-making process; trust grows; and enhanced understanding leads to sharing of improved data and user intent information. Such cooperation leads to efficiency improvements for all participants.

—Sachs *et al.* [Sachs20]

A UTM system design, therefore, must provide fairness. As with other system objectives, a design should start by defining the kinds of fairness that should be provided, and identify scenarios that can lead to unfairness. This then leads to designing mitigations that eliminate or reduce those situations.

The question of centralized versus distributed UTM structure affects how unfairness can occur. In a centralized system, unfair results can happen either because the central UTM services has problems, or because operators try to game the system (for example, by filing many unneeded flight requests to increase their apparent demand). For a distributed system, unfairness can come from flaws in negotiations between providers, misbehavior by UTM providers, and the same misbehaviors by UAS operators.

Before settling on a choice between a centralized or a distributed UTM structure, those responsible must look at the feasibility and cost of mitigation options for each of the scenarios that can lead to harms. A low-cost or high-value but infeasible solution should not be chosen. The option giving the best value-to-cost among feasible solutions should be chosen.

Anyone proposing a UTM system design, whether a regulator, a system developer, or a researcher, has a responsibility to consider and share how their design will address fairness.

I discussed designing systems to manage shared resources in Chapter 47. Airspace and UTM services are in most respects a common-pool resource, and the considerations for managing such resources apply [Ostrom08, p. 18]. I would expect that any discussion of a potential UTM approach would address:

- The definition of boundaries; for example, that UAS operators can be distinguished from each other;

- The involvement of system participants, such as how UTM providers and UAS operators have input into how airspace access is managed;
- The ability to accurately monitor behavior;
- The kinds of sanctions that are sure and also proportionate to the offense, showing how they create incentives for compliance; and
- Conflict-resolution mechanisms between users.

63.7.2 Pushing unwanted solutions in spacecraft projects

The situation. I discussed two spacecraft projects that involved, among other problems, participants pushing for the design to use certain existing components (Sections 4.1 and 4.4). In one project it was a spacecraft bus design and certain electronic components; in the other it was a key software component. In both projects, the proposed components did not contribute to the systems' objectives and would have resulted in systems that did not work. In both projects as well, those who were pushing for the use of those components were doing so because they had a vested interest in re-use and acceptance of their designs.

Ethical responsibilities. The primary ethical responsibility that parties in both projects failed to meet was to provide solutions to the customer that meet the customer's needs, and instead making false claims about the suitability of those components. There was a conflict of interest between those who were trying to provide the components and the customer, and the providers failed to consider the customer's needs first.

Harms. The harms that came in both projects were the public money was spent on designing and trying to build systems that would not work, and in both cases the projects were canceled. Not only was this a waste of investment, but it prevented the significant benefits that might have come had working systems been deployed.

63.7.3 Worker burnout

This final case study is about my own behavior on one particular project.

The situation. I was working in a leading engineer role on a multi-year spacecraft project. The project involved multiple companies, with teams all over the US. The role involved coordinating work by all those teams as well as leading the high-level design and systems engineering for the whole project. Some of the teams performed well, while other teams failed to deliver their parts of the work.

This led to constant travel, working too much, and sleeping too little. I was constantly fatigued and sometimes lost the thread of discussions. I was trying to keep the teams coordinated and keep the project moving forward, but was unable to.

Ethical responsibilities. Being in a lead role, many of the collective responsibilities of the project as a whole devolved to me, in full or in part. This included working with the customer to understand their needs, answer their questions, and deliver a system that met needs on time and budget. It included responsibilities to each of the organizations working on the project to understand their motivations and needs, apportion work amongst them effectively, and ensure that their teams were able to work together.

The responsibilities also included providing technical guidance and setting an example to the project team as a whole: providing a coherent and feasible system concept and high level design, and working with everyone during development to address problems that were found along the way.

Harms. The schedule meant that I did not take sufficient time to reflect on the project and to spend time with subject matter experts to understand well enough the details of some parts of the system. The result was that I missed some of the deeper problems in the design and had to spend more effort later trying to rein in teams that had gone off in an unhelpful direction. The fatigue meant that I sometimes made poor technical or managerial decisions. Sometimes I was too irritable, and that interfered with good relations among teams at different sites. In the long term, it also affected my health.

In the end the project was canceled. It was not only from my actions, and I do not know even now if there was anything I could have done to change that outcome. But not keeping my own workload at a sustainable level meant that both I and others suffered.

63.8 Using the ethics framework

The framework in this chapter provides tools for thinking about the ethical aspects of system-building projects. It does not provide a complete and prescriptive set of rules; any such rules I could think of would depend on the specific project. Any nontrivial project will have unexpected situations happen that no one has prepared for, including anticipating their ethical consequences.

Any specific project, however, needs its ethical standards. The rules need to be specific and actionable, meaning that people understand correctly when the rules apply and what they mean. Those rules will never be complete, because strange new

events will happen; they should be complemented by principles that give guidance when the unexpected happens.

The ethical rules and principles are realized in multiple ways:

- Written down as part of the project's objectives;
- Recorded in policy documents, providing a kind of legal code for behavior within the project;
- Reflected in the team's structure, and how collective responsibilities map to the responsibilities of individual roles (Section 63.5.2);
- In the team's culture: its habits of communication and the interpersonal relationships among team members; and
- Modeled in the behavior of leaders in the team.

As I have argued with other aspects of a team's organization and culture, a project that functions ethically does so in part because its structure and culture are planned. The project's ethical objectives should be thought out and written down from the beginning. They do not need to be lengthy, but they need make it clear what will be expected of everyone on the project. Those objectives then should be translated into specific responsibilities for specific roles. Who has what responsibility will evolve, as unanticipated situations arise, as regulations change, or as the team finds that some team structure is not working.

All the written words will end up meaning little, however, if people do not follow basic ethical behaviors. I worked with a project that adopted the principle "safety, first and always"—and then did not follow that up by including safety in their engineering practices. It is the responsibility of those who lead, because of their authority or because of their social position in the team, to provide an example of ethical behavior at all times.

It is also an ethical responsibility for everyone on a project, including those who lead, to watch for ethical problems, accept concerns from others about potential ethical lapses, and act on them. Everyone on the team must have confidence that when they see something wrong, it will be dealt with. They must also have confidence that if they *do* something wrong it will also be dealt with.

Organizing a team to behave ethically does not add anything fundamentally new to how teams work. It involves including ethics considerations in a team's objectives and its structure. It involves addressing ethical problems in communication, in many of the same ways that people might address a technical problem. It requires each person to think about the ethical aspects of their work, just as they should be thinking about how their work affects the safety or security of the system being built.

63.8.1 When incidents happen

When an ethical lapse is likely about to happen or has happened, people have a responsibility to notice and act. This requires self-awareness, to be able to see that one's own decisions or actions have ethical consequences. It also requires awareness of what is going on around one, and to be able to think about the ethics of what the team is doing.

Responding to a lapse or potential lapse can come on a graduated scale depending on who is causing the lapse and who will be affected. It is usually best to start small and early, in order to correct a situation with the least disturbance possible.

The first step is to recognize when one is making a decision that can have ethical consequences themselves. This requires having an understanding of the ethical objectives a project aims to meet, and the time to think through a decision before committing to a choice. In this case, working out ethical questions only concerns the person making the decision (though they might well choose to document their thinking as part of the rationale for a design choice, for example). The person making the decision might choose to talk it over with someone else they trust in order to get perspective.

The next step is when someone notices a problematic decision being made by someone they work with. In that case, they should bring up their concern with their colleague. They may find out that their colleague has made a good decision, and they did not understand the rationale. If they still believe that their colleague is making a poor decision, and try to persuade them about the problem.⁶

When problems are larger—affecting several people or the whole team, or without a clear decision-maker to discuss with—the problem can be brought up in a group, or reported upward to someone with responsibility to address team-wide ethical problems. It might even go as far as invoking a formal ethics reporting mechanism (which every project should have).

If a problem has effects outside the team, such as involving a customer, funder, or regulator, then the project must respond collectively. When someone on the team detects that there is an ethical problem, they should report it to someone who has the responsibility for collective response (Section 63.5.2). When someone outside the project detects a problem, there should be a clear way for them to report the problem to someone responsible in the team.

This hierarchy of responses, from one-to-one up to collective response, contains the activities for handling an ethical lapse to those directly involved.

⁶I recommend Johnson's *Scaling People*, Chapter 2, Section 1, for some helpful advice on how to talk about difficult matters, including ethical ones [Johnson22].

But what does one do if reporting a problem or trying to persuade someone there is a problem does not work? What if the harms caused, potential or actual, are not addressed?

When this happens, one step is to escalate the issue upward to someone who will respond. This can mean going as far as making a whistleblower report to regulators if the problem is serious enough for their attention.

If, after escalation, the problem remains, then one has to decide what to do. There are no good choices. One option is to live with the situation, documenting what one has done to try to alert and respond. Another is to live with it, and try to work to mitigate the problem. The other option is to leave the group or project, so as not to be a part of the problem.

63.9 Commentary

Work on ethics often runs into complex situations, where ideas of rightness and wrongness are not clear. Because of this, one needs to maintain some humility to the process of working through ethical questions.

Most of the time one must make decisions without having all the information one might want, or with some information that may or may not be accurate. The decisions will often involve predictions about future consequences of different choices, but the future cannot be fully known. As a result, it is impossible to make perfect ethical decisions all the time. They must be made on the best information available, using as much care as can be in the time available.

A person who has to make a decision should not be held culpable for a bad choice if they made the best decision they could given the information and resources available. They might be blame-worthy if they did not seek out good information, or if they did not reason through the problem carefully—but not just because their choice had a poor outcome, as long as they did an honest job of investigating the situation.

Ethical lapses can also come from behaviors like holding inflexible opinions. Making ethical decisions must be based on truth, and one must be willing and able to check whether the beliefs they hold accord with the truth and discard those beliefs if they do not.

Dealing with ethical questions quite often challenges people's beliefs and their sense of self-worth. People often react as if they are being attacked when the ethics of something they are doing is called into question. This means that one must take care in how ethical questions are raised and discussed. Smugness and self-righteousness are likely to cause a defensive reaction in the person being questioned. While that may feel justified, it is often not particularly useful. Taking care about how to talk about such matters can help actually work through an issue.

Finally, ethical behavior requires courage. In many cases, making an ethical choice means refraining from doing something that seems attractive—boosting profit, hiding a problem, making someone feel good by over-promising. It takes a kind of courage to make the choice that won't cause harm to others. Once an ethical lapse has happened, it also takes courage to acknowledge that a mistake was made and take the steps to repair the harms caused. Repairing harms can be costly to a project in time or money. It can be equally costly to admit to those who have been harmed, or were potentially harmed, that a mistake has been made. In the medium to long run, though, it is only by being honest that a project will get better, and it is only by doing so that team members will develop trust in each other and the project.

Bibliography

- [DOJ20] Investigations into sales practices involving the opening of millions of accounts without customer authorization.
Office of Public Affairs, United States Department of Justice, Press release 20-219, February 2020.
<https://www.justice.gov/opa/pr/wells-fargo-agrees-pay-3-billion-resolve-criminal-and-civil-investigations-sales-practices>.
- [DSM-5] *Diagnostic and Statistical Manual of Mental Disorders*.
American Psychiatric Association, Washington, DC, Fifth ed., 2013.
- [Drucker93] Peter F. Drucker.
Management: Tasks, Responsibilities, Practices.
Harper Business, New York, NY, 1993.
- [Golding18] Richard Golding.
Metrics to characterize dense airspace traffic.
Altiscope Project, A³ by Airbus, Report TR-4, 7 June 2018.
https://www.chrysaetos.org/papers/TR-004_Metrics_to_characterize_dense_airspace_traffic.pdf.
- [IEEE24] Institute of Electrical and Electronics Engineers.
Section 7.8: IEEE Code of Ethics.
In *IEEE Policies*.
<https://www.ieee.org/content/dam/ieee-org/ieee/web/org/about/corporate/ieee-policies.pdf>, accessed 8 December 2024.
- [INCOSE24] International Council on Systems Engineering.
Code of Ethics.
<https://www.incose.org/about-incose/code-of-ethics>, accessed 8 December 2024.
- [Johnson22] Clair Hughes Johnson.
Scaling People: Tactics for Management and Company Building.
Stripe Press, South San Francisco, California, 2022.

- [Leveson11] Nancy G. Leveson.
Engineering a safer world: systems thinking applied to safety.
Engineering Systems. MIT Press, Cambridge, Massachusetts, 2011.
- [New York09] Capital Heat Inc v. Michael Blatner Family Trust.
Supreme Court, Appellate Division, Fourth Department, New
York, 10 August 2009.
<https://caselaw.findlaw.com/court/ny-supreme-court-appellate-division/1136465.html>, accessed 13 December 2024.
- [Nollkaemper20] André Nollkaemper, Jean d’Aspremont, Christiane Ahlborn, Berenice Boutin, Nataša Nedeski, and Ilias Plakokefalos.
Guiding principles on shared responsibility in international law.
European Journal of International Law, vol. 31, no. 1, August 2020,
pp. 15–72.
<https://doi.org/10.1093/ejil/chaa017>.
- [Olson65] Mancur Olson.
The Logic of Collective Action: Public Goods and the Theory of Groups.
Harvard Economic Studies. Harvard University Press, Cambridge,
Massachusetts, 1965.
- [Ostrom08] Elinor Ostrom.
The challenge of common-pool resources.
Environment: Science and Policy for Sustainable Development, vol.
50, no. 4, 2008, pp. 8–21.
doi:10.3200/ENVT.50.4.8-21.
- [Sachs20] Peter Sachs, Antony Evans, Maxim Egorov, Robert Hoffman, and Bert Hackney.
Evaluating fairness in UTM architecture and operations.
Airbus UTM, Report TR-010, February 2020.
https://storage.googleapis.com/blueprint/UTM_Fairness_Tech_Report-v1.1.pdf.