Advances in Reliability, Safety and Security, Part 5 – Kołowrocki, Kosmowski (eds.)
© 2024 Polish Safety and Reliability Association, Gdynia, ISBN 978-83-68136-17-3 (printed), ISBN 978-83-68136-04-3 (electronic)

Advances in Reliability, Safety and Security

ESREL 2024
Monograph Book Series

Towards Ethical Risk Assessment

Vincent Philip Paglioni[a]

[a]*Risk, Reliability, and Resiliency Characterization (R3C) Lab, Department of Systems Eng., Colorado State University. Fort Collins, USA*

Abstract

Safeguarding the health and welfare of the public is the charge of all engineers, but there is perhaps no field where this commitment is more actionable than in risk assessment. Understanding and mitigating risks in complex engineering systems is critical to ensuring that the public is not exposed to undue risks. Part of risk assessment is, therefore, defending the common good – those aspects of life that are freely enjoyed by all in a community. However, this is where risk assessment could be more impactful. Despite the organizational and community benefits engendered by risk assessment, there is little focus on the distribution of benefits and negative consequences. As a result, large infrastructure projects that benefit a wide area tend to expose the local community to the majority of risk, creating a glaring discrepancy between the communities that benefit from large infrastructure projects, and those exposed to the risks. This is an ethical problem, especially considering that communities nearby to large infrastructure projects are often already underserved. This paper discusses the ethical considerations important to risk assessment and provides a preliminary view for incorporating ethics into the risk assessment process to protect the common good for *all* communities.

Keywords: risk assessment, ethics, risk management, community engagement

1. Introduction

The common good, as conceived by philosophers, is an ideal that is defended by societal rulers for the benefit of society itself. As a "rising tide raises all ships," the common good is just that – *common*, in the sense that it is for the benefit of *all* members of a society. As a result, the onus for defending the common good falls in principle on everyone, while in practice it is generally assumed that societal leaders, who are best positioned to do so, will safeguard this ideal to the benefit of those they lead, and themselves. As the stewards of complex engineering systems, upon which the global population relies for both the necessities of life and surplus benefits, engineers have a special role in safeguarding the common good, principally through the conduct of risk assessments.

Risk assessment is, namely, a three-part process of analysis, management and communication of risks in a system (Modarres & Groth, 2023). As part of this process, systems are studied for risks posed to the system, to system operators and stakeholders, and to the public and environment. Typically, risks are deemed "acceptable" or "unacceptable," and accommodated or mitigated as required through various risk management protocols. The outcomes of risk assessment are communicated to organizational stakeholders, regulatory officials, and the public. Risk assessments inform whether the system is ultimately "safe enough," what measures need to be taken to reach/maintain the desired level of safety, what the likelihood and consequences of risk-relevant events (externally- and internally-driven) might be. As a result, risk assessments essentially provide stakeholders with the understanding of the benefits and risks posed by a given system. However, risk assessments do not typically consider how the risks and benefits of the system are *distributed* across stakeholder groups (i.e., the public). As a result, there may be inherent inequities in the impact of systems on the public that are not accounted for in risk assessment. Risk assessments that do not account for the distributive effects of the system call into question whether risk assessment, as currently conceived and implemented, truly safeguards the common good and the public health and welfare.

This paper is not an attack on risk assessment as a concept, or even as a field. Rather, I attempt to show that, for all the improvements engendered since the inception of the field almost 100 years ago, there is still room for significant pragmatic and philosophic gains in addressing the ethical implications of risk assessment, particularly in its handling of distributive effects and public engagement. I will begin by briefly reviewing more closely the evolution of the common good as a philosophic ideal. In Section 3, I discuss the current implementations of risk assessment and their connections (or lack thereof) to ethical considerations regarding equity and fairness. Finally, I discuss an ethical risk assessment framework, based on contractualism, that could improve on the treatment of the distributive characteristics of risk, and conclude with my perspective on a path forward for the field.

2. The common good

The common good, broadly construed as those entities or ideals that are shared for the mutual benefit of all member of society (Hussain, 2018), is often viewed by philosophers as a critical aspect, if not the ultimate goal of, a flourishing society. While early conceptions of the common good were viewed as being protected by benevolent rulers or so-called "philosopher kings," modern interpretations view the protection of the common good as a societal responsibility. While incumbent upon all members of the society, there are certain members upon whom this responsibility is more pressing, particularly civil and technical stewards and leaders.

The common good, as conceived above, is analogous to the public health and welfare, and therefore imparted by complex, engineering systems. Such systems, including large infrastructure projects, health systems, water, energy, and transportation, are relied upon for common goods that enable survival and/or quality of life, e.g., clean water, electricity, or healthcare. As the designers, operators, and maintainers of these systems, engineers hold a special role in the protection of the common good. This principle is espoused in the NSPE Code of Ethics, which demands that engineers "… hold paramount the safety, health, and welfare of the public …" (Engineers, 2019), and is recognized as a critical aspect of risk assessment and reliability engineering.

Clearly, risk assessment is used to safeguard the public health and welfare, and therefore the common good. In most industries, from consumer products to large infrastructure projects, risk assessment is used to ensure that the end-user, public, and environment are not exposed to unnecessary risk. However, where risk assessment fails is in the accounting of who is shouldering the risk versus reaping the benefits. This is particularly salient in large infrastructure projects which, while bringing benefits to thousands or millions of people, can reinforce existing inequities through the risks imposed on a relatively smaller cohort of the public. For example, coal-fired power plants have been systematically sited closer to red-lined "hazardous" neighborhoods, compounding the risk burden faced by already-disadvantaged communities (Cushing et al., 2023). Similar patterns have been observed with freeway siting, where the benefits of increased transportation capacity were wrought at the expense of increasing risk to underserved communities (Coombs, 2022).

It is undeniable that these large infrastructure projects did, or were at least designed to, enhance the common good through electricity and transportation. However, it is also undeniable that these projects, despite their risk assessment, have greatly exacerbated the risks faced by the local communities. As a result, it is difficult to state that such projects represent an ethical implementation of risk assessment, which would have identified the adverse effects to these communities and developed mechanisms to mitigate or compensate for the increased risk. The question of the worth of such projects, not only in monetary or technical metrics, but in the human cost, is rightly asked.

Infrastructure and engineering projects are required to maintain the common good, and must be placed somewhere, keeping in mind operability, accessibility, and cost; therefore, it is almost inevitable that some community will have to accept greater risk to accommodate distributed benefits. Which communities, then, can be engaged and incentivized to shoulder higher risk, if the benefits are so widely distributed? This idea is encapsulated in the so-called "Not In My Backyard" or NIMBY opposition to infrastructure projects, a relatively common sentiment in well-off communities (Balintec, 2023). The idea of the NIMBY opposition is that infrastructure (even "social good" infrastructure such as affordable housing) is acceptable only if the project is located elsewhere. This type of grassroots organizing has succeeded in delaying, cancelling, or moving infrastructure projects in well-off communities (Balintec, 2023). Interestingly, the risk assessment is somewhat disconnected from the opposition, which is reinforced by the public perception of the risks, which are typically markedly different from the expert assessment of the risk (Slovic et al., 2000a, 2000b).

In the face of these obstacles, risk assessment (as currently implemented) is a poor solution. Decades of what to the public feels like disaster after disaster have left them untrusting of government- and/or industry-driven risk assessments in favor of their own intuitive thinking. The public perception of risk is driven by qualitative aspects of both the technology and risk that are not accounted for in a risk assessment, such as the "dread" factor, feeling of controllability, and the distributed nature of the risks and benefits (Slovic et al., 2000a). Clearly, the public's perception of dread or controllability should not be used as a determinant in the assessment of risk, as such characteristics have no bearing on the actual risk. However, the distribution of risks and benefits does affect the actual risk faced by a community, and this necessitates accounting within the risk assessment.

3. Risk assessment and ethics

Risk assessment as a discipline was developed during and immediately after World War II, principally for military and aerospace applications. The field has consistently evolved to understand systems of increasing complexity, and entered the civil arena with the rise of commercial air travel and nuclear power plants. We will herein briefly review the fundamentals of risk assessment, in addition to the relevant methods. A thorough understanding of risk assessment can be found through any number of textbooks, articles, and lectures.

3.1. The fundamentals of risk

As a construct, risk is composed of three main elements: 1) scenario, 2) probability/frequency/likelihood, and 3) consequences. The scenario defines the events and conditions that lead to the exposure of a hazard in a system; this occurrence has a probability or likelihood which characterizes the uncertainty of the event's occurrence. Finally, the exposure of a hazard is associated with some negative consequence(s), the avoidance of which is of course preferred.

Colloquially, risk entails a similar, but distinct, meaning. To many people, risk is more strongly associated with a negative consequence than probability or likelihood. This colloquial conception of risk, as contrasted with the technical definition, is in part responsible for the differences in risk perception between experts and the general public. The below discussions on each aspect of the risk triplet are broad-based overviews rather than detailed investigations. It should be further noted that, for complex systems such as the coal-fired power plant discussed in the examples, the scenario, frequency, and consequences are developed with increasingly complex tools including Event Tree Analysis and Fault Tree Analysis, which together form the basis of *probabilistic risk assessment* (PRA). For the purposes of the discussion herein, a thorough understanding of PRA is not necessary. Those readers interested in a more in-depth discussion of PRA and risk assessment tools are referred to (Modarres, 2006; Modarres & Groth, 2023).

3.1.1. Scenario

The scenario, as the name suggests, is the set of conditions and events that lead to the exposure of some hazard in the system. This is roughly equivalent to the idea of an *accident sequence* in the parlance of nuclear probabilistic risk assessment (PRA). As an example, consider a coal-fired power plant as a system. There are many hazards associated with this system, many dependent on the scenario or use-case (e.g., normal operations, shutdown operations, maintenance, etc.). Therefore, the first step in assessing the risk associated with this system is to understand the different scenarios, which allows analysts to not only think more critically about all phases of a system lifecycle and use, but also structures and limits the analysis to focus on a single scenario at a time.

In this example, let's consider that the scenario is a coal power plant undergoing normal operations. The hazards in this scenario include, but are not limited to, unexpected shutdown, scrubber failure (e.g., release of unfiltered exhaust gases), or fly ash release. For now, we will focus on the hazard of unexpected shutdown. The scenario influences the probability of the hazard occurring. For instance, unexpected shutdown cannot occur if the power plant is already shut down for maintenance, scrubber failure can only occur if the power plant is operating. Rigorously defining the scenario surrounding different risks is critical to arriving at an appropriate understanding of the probability and consequences.

3.1.2. Frequency

The frequency, which is sometimes used interchangeably with *probability* or *likelihood*, provides a quantitative metric of the chance that the hazard will be exposed *given the conditions as defined in the scenario*. This is an important consideration – probabilities in risk assessment should not be treated as if the events occur in a vacuum. As discussed above, the scenario can heavily influence the frequency of the hazard exposure, and this is true for most scenarios. The frequency associated with a risk in one scenario will not necessarily hold as the scenario changes. Returning to the example of the coal power plant, we could postulate that the frequency of an unexpected shutdown in "ideal" conditions might be 1 in 10,000 operating hours, while in sub-optimal conditions (e.g., differing fuel purity, operating on backup systems, etc.) this frequency may be much higher, maybe 1 in 5,000 operating hours. Note that these frequencies are examples only and not based on actual data, but the point remains the same. The definition of the scenario should be used to guide assessments of frequency. Historical data, test results, and observations are generally used to identify the frequency of failure for various hardware components under various conditions.

3.1.3. Consequence

The consequence associated with a given risk is often the most visible aspect of the risk. Whereas the frequency of risks are not always observable, the consequences usually are. For example, an unexpected shutdown of a coal-fired power plant is a highly visible event, as the operators will be immediately notified that the power plant is not generating electricity, the grid regulators will have to balance the grid, and there may even be brown or blackouts in surrounding areas. Consequences are typically qualitatively rated on a scale from, e.g., "minor" to "severe," where each classification is defined for the system at hand. For example, an unexpected shutdown might be a "minor" consequence if deaths/injuries are the focus of consequences; alternatively the same consequence could be "major" depending on the scope and purpose of the analysis. Generally, a risk assessment does not discount any consequence based on severity, but only once it is proven to fall below a threshold of likelihood (e.g., 1×10^{-7}, which is 1 in 10,000,000 or lower). Thus, every consequence identified in a risk assessment is ***a priori*** important, until it is proven to be of negligible likelihood.

3.2. The ethics of risk assessment

Contrasted to the complex and time-consuming process that characterizes many of the steps in a typical risk assessment, relatively little work is devoted to understanding the *distribution* of consequences and what it means for local communities that will likely bear the brunt of any system failure. Admittedly, this is perhaps more of a fault with risk *management*, rather than risk analysis, but risk management only functions insofar as the information provided from risk analysts allows. In the absence of concern generated by the risk analysis for the distribution of consequences, there is likely to be little discussion of the same by risk management.

What does this mean in reality? This means that, although the risks associated with most complex systems are exceedingly small, when they do occur, they can disproportionately affect a localized area. A laundry list of examples is illustrative to this point. The Seveso Disaster (1976), Lake Peigneur Drilling Disaster (1980), Bhopal Disaster (1984), AZF Fertilizer Factory Explosion (2001), and Deepwater Horizon (2010) all had catastrophic *localized* effects despite the risk assessment process inherent in such complex systems. In some of these cases, including Seveso (Eskenazi et al., 2018), Bhopal (Dhara et al., 2002), and Deepwater Horizon ((D'Andrea & Reddy, 2018), the affected local communities are still grappling with long-term physical and mental health effects.

The question arises: is risk assessment ethically sound if there is little (or, often, no) consideration for *who* is affected by the consequences, as compared to who reaps the benefits? I argue that, as currently conceived and applied, risk assessment is *not* ethically sound in this regard. While clearly technically robust and useful, risk assessment is inadequate for addressing the concerns of local populations. This issue is compounded when risk analyses and risk management decisions are made in the absence of input and buy-in from the very people who will bear the brunt of the consequences.

3.2.1. The role of the public in risk assessment

Before discussing the role of various ethical frameworks for mitigating the concerns just raised, it is prudent to address another, related question: what should be the role of the public in risk assessment? This is a question of norms, i.e., deliberating *how* risk assessment should be performed. The question of public involvement in risk assessment is informative for understanding how ethics and risk assessment interact.

There are, broadly, two broad camps in the normative question of the role of the "general public" within risk assessment. The first, which has been called "technocratic" (Sunstein, 2002) or even "paternalistic" (Kusch, 2007), are those who argue that risk analysts, by virtue of their expertise, should be insulated from public debate on risk-relevant projects. This view argues that the general public is ill-suited for performing a risk assessment, and thus experts should be making risk management decisions, guided by cost-benefit (or risk-benefit) analysis and insulated from the fears and opinions of the public (Sunstein, 2002). The main argument proffered in this approach is that the public is generally in error when disagreements with experts arrive (Sunstein, 2002). While well-founded on the recognition that lay people are generally worse at estimating risk metrics (principally, annual fatalities) than experts, this thesis ignores the qualitative characteristics that non-experts deem important for risk assessment, and that experts are subject to the same cognitive biases as non-experts (Slovic, 2000b; Slovic et al., 2000b). Furthermore, a system that seeks to insulate risk decision making from public discourse (Kusch, 2007) would rob individuals of agency in serious matters relative to their welfare. Finally, as Jasanoff notes, if risk perception is so influenced by social and cultural factors, the strict delimitation of factors included versus excluded in formal risk assessment should be carefully considered (Jasanoff, 1998).

The alternative view, which might be labelled "public-driven" or "participatory" (Kusch, 2007), argues that risks are largely social constructs, hence the social and cultural factors that affect perception (Jasanoff, 1998), and therefore the public should play an *outsized* role in the analysis of risks. This view embraces the public as stakeholders in risks that will affect their welfare. Instead of expert-driven risk assessment with no public discourse, the deliberative democratic approach of public-driven risk assessment views the people as the best perceivers and guardians of their common good. However, a non-expert public would quickly be out of their depth in assessing risks for complex engineering projects. While perhaps more ethically sound in regards to public agency than the expert-driven view, it can be reasonably argued that a public-driven risk assessment would not be capable of meaningfully assessing complex risks, even if it is able to achieve community consensus regarding risk management decisions.

Thus, both normative views on risk assessment, presented above in their respective extremes, are flawed, but in different respects. The expert-driven approach, while likely to be more technically sound, fails to provide meaningful agency to the public, rendering them essentially captive to the technocratic system, albeit with a thin assurance that the experts have their best interests in mind. The public-driven approach restores agency to the public at the expense of producing potentially flawed risk analyses. Neither extreme seems befitting for questions as important as those tackled by risk assessment. The resolution, in my view, must lie somewhere in between the extrema, wherein an expert-driven risk assessment is buoyed by extensive public input and information sharing. In such a system, the public would have greater agency than is currently afforded, while allowing experts to handle the technical and complex work of assessing the risks. Thus, such a system could provide for an ethically sound assessment and management of risks, guided by the very people facing those risks, who may have critical insights not considered by risk experts. Local insight may be critical for effectively managing risks without destroying the community; this was the case when attempting to manage the aftermath(s) of the incidents at Sellafield, which alienated the local community because experts would not engage with them (Wynne, 1992). An ethical risk assessment process should thus combine the technical expertise of experts with the appreciation for distributive and qualitative aspects of risk and local insights from the public (Peschard et al., 2023).

3.2.2. The common good in risk assessment

The preceding section presented two broad normative views on risk assessment as related to the public. However, the focus of this paper is on the use of risk assessment to safeguard the common good. The question remains, then, as to the role of the common good in risk assessment. If we assume a modern conception of the common good, as those conditions mutually advantageous to life and shared by all members of a community, then risk assessment should be focused on safeguarding the common good from large system failures.

To a certain extent, this is true: risk assessment determines the consequences to the system itself and surrounding areas/populations, and prioritizes those risks for mitigation and management. As a result, it could be argued that the common good is being safeguarded by risk analysts. Risk assessment, by nature, seeks to minimize the negative consequences of system failure. This means that the system will maintain operation, providing critical functions (e.g., water, power, etc.) as reliably as possible. These functions are necessary to life and/or welfare, and thus could be arguably considered common goods. Under this argument, risk assessment safeguards the common good by maximizing production/service and minimizing (where possible) failures, and therefore the associated consequences.

However, this argument does not consider the common good that is not *produced* by the system in question, but rather *extant independent of the system*. Large infrastructure projects, which are typically the most visible applications of risk assessment, exist within, and often use as input, other entities classified as "common good." The air, land, flora/fauna, and water around infrastructure projects are all common goods, particularly the air and water. When risk assessment is performed, the consequences to the surrounding environment are considered as any other consequence, and (if applicable) these risks are mitigated through risk management. However, there is no special consideration or prioritization specifically for this "localized" common good. As a result, risk assessment *de facto* prioritizes the needs of the majority, those to whom the service/function is provided, over the rights of the few – those participating in the common good surrounding the system and who will bear the brunt of failure consequences.

There is, therefore, a need to consider the ethical ramifications of risk assessment. Is it ethical to potentially jeopardize the common good for a minority, in order to serve the needs of many? Is there a way to provide special consideration for those directly impacted by system failures, that does not simply reverse the problem? These questions are normative in nature, asking what *ought* to be the case. Accordingly, the answers are normative rather than prescriptive – characterizing *how risk assessment should be performed*, rather than providing specific (and potentially limiting) instructions.

4. Contractualist ethics

The questions raised in the previous section are largely questions of values, that is questions of what is right or wrong (Scanlon, 2000), and particularly questions of controversial (i.e., reasonably argued against) non-epistemic (non-scientific) values (Hansson, 2007). Questions of controversial values (e.g., should nuclear power be banned?) are divisive because of the wide diversity of ultimate values people use to inform their beliefs – so divisive that it is unexpected to arrive at any consensus in values (Kusch, 2007). Risk assessment, particularly risk management, is a value-laden process (Hansson, 2007), and therefore has to grapple with how to assess and combine values in a decision space. With little hope of consensus on the values themselves, however, this is a difficult problem. Instead, it may be that developing *procedural principles*, that is normative rules defining *how* ethical risk assessment should be performed, is a pathway to overcoming this obstacle (Kusch, 2007).

Predicating risk management decisions on consequences alone, i.e., following a consequentialist framework like utilitarianism, can lead to violating the values held by some individuals and could permit morally reprehensible actions in a short-sighted attempt to bring about the maximal benefit (Cranor, 2007). Furthermore, risk assessments that prohibit certain outcomes may lead to the prohibition of large and generally beneficial infrastructure projects. On the other hand, if certain outcomes are *mandated*, some risks may be mismanaged in order to meet the required outcome. In either case, rooting our conception of ethical risk assessment in the consequence space is antithetical to the goal of improving the ethical basis of risk assessment. In a non-consequentialist framework, the rightness or wrongness of a risk management decision is determined by whether the decision is based on sound, deliberated, and agreed-upon principles. Contractualism provides a useful heuristic for identifying such principles based on whether they could not be reasonably rejected or could be justified to those most adversely affected (Cranor, 2007).

In our use, contractualism will refer to the specific ethical theory developed relatively recently (in the span of philosophical thought) by T. M. Scanlon, although it follows in a long tradition of social contract theories. In his theory of contractualism, acts are "wrong if and only if any principle that permitted it would be one that *could reasonably be rejected*" (Scanlon, 2000). In the context of risk assessment, actions or decisions are wrong if they are based on a principle that could be reasonably rejected, e.g., by the public or an analyst. By contrapositive, actions or decisions based on a principle that *could not be reasonably rejected* are therefore "right."

While a thorough discussion of Scanlon's contractualism is outside of the scope for this work, it is worth reviewing one of the central tenets of the theory, namely the conception and importance of *reasons* to ethics and morality. Critically, in Scanlon's formulation, judgments of right and wrong are actually judgments of what we have reason to do (Scanlon, 2000). In the context of risk assessment, this translates to judging risk-relevant decisions on the basis of the underlying reason. For example, if the decision to X is based on a reason Y, derived from a principle Z which cannot be reasonably rejected, then it is a "right" decision.

A strict contractualist frame of risk assessment might seem to be easily circumventable – all someone would have to do is reject the decision/action to hopelessly mire the process. However, a contractualist system would not necessarily open the *decision/action* to debate, but rather the *reason* for taking that decision/action. While someone may reject the decision to, for example, mandate a certain safety system because of its high cost, the underlying reason for the decision (e.g., the reduction in harm to the surrounding population) is not so easily cast aside. One would have to show that the rejection of the principle is *reasonable*, which would require disproving, or at least casting reasonable doubt on, the principle.

Clearly, a strict contractualist frame of risk assessment could easily go awry, as the principles underlying risk-relevant decisions and actions may not be so clear, and even clear principles might have good reason for rejection. It may seem paradoxical to rely on an ethical framework that will permit, if not demand, even more debate over already-controversial projects. However, people, as legitimate sources of moral input (Cranor, 2007), deserve agency in decisions that will affect their lives and potentially expose them to harm. Furthermore, large infrastructure projects have historically been sited in or near underserved communities, progressing a vicious cycle as incidents invariably occur and negatively impact the community. It is critical that such communities, indeed any community, be allowed some agency in the decision to place large engineered systems nearby. Contractualism provides the necessary structure to facilitate effective community dialogue over relevant risk issues.

5. A contractualist risk assessment process

Assuming that a contractualist framework is adopted for ethical risk assessment, it is worth discussing what such a process would look like. Does the analysis of risks have to change, or just the management of risks? How can the community be effectively engaged without prohibitive schedule delays? These questions, and many others, will have to be answered before ethical risk assessment is truly achieved. However, the purpose of this paper is to provide a starting point for building ethical risk assessment; as such, these questions are briefly discussed with the hope that this will lead to a larger discussion about the specifics of ethical risk assessment.

The question of ethical risk assessment is not novel, but has hitherto been discussed chiefly in the domain of risk philosophy rather than within the risk assessment field itself, e.g., (Cranor, 2007; Hansson, 2007; Kusch, 2007). In this section, I present one possible framework for incorporating contractualist ethics into risk assessment. As mentioned, this is a preliminary framework that could serve as a starting point for developing ethical risk assessment. Doubtless, there are alternatives to each suggestion that will be more practicable in the real world. However, the suggestions outlined below would go a long way towards realizing an ethical risk assessment process.

5.1. A contractualist analysis of risks

Because contractualism deals largely with actions and their reasons and principles, the analysis of the risks themselves, meaning the process outlined in Section 3.1, largely does not have to change. The tools and techniques used by risk analysts – event trees, fault trees, Bayesian networks, etc. – remain the tools of the trade. However, risk analysis needs to provide the requisite information to inform risk management decisions, particularly regarding the distributed nature of risks and benefits. Risk management decisions that are made under incomplete information cannot be reasonably rejected unless the information to do so is available. What does change in risk analysis, therefore, is the computation of the risks to include a "consequence localization" metric that will indicate how consequences are distributed and which have the largest impacts on the community surrounding the project. Community, as an operative word, will need to be rigorously defined by regulatory agencies, so that similar projects are regulated in a consistent manner. For example, community might be defined for nuclear power plants as the population residing within the emergency planning zone (EPZ), or for transportation infrastructure community might mean those people relocated by, and/or remaining adjacent to, the project.

The appropriate measure by which to adjust the ultimate risk value (the product of consequence magnitude and frequency) should be left to a multidisciplinary team of experts from the relevant fields, including risk science, engineering, ethics, and community engagement to name a few. This type of multidisciplinary team is already a best practice in risk assessment and ensures that there is sufficient expertise to understand all parts of the project. The addition of ethicists and community organizers will ensure that the community itself has representation in the analysis of risks, in addition to the management and communication processes. This should be accompanied, where possible, by a transparent discussion of the processes used in the analysis of risks, so that the public can have more informed deliberation over the principles governing the ultimate risk management decisions. Thus, risk assessment is maintained as an expert-driven process with an informed public that can meaningfully contribute to public discourse over the resulting decisions. Build a risk assessment process that is more transparent to the public may provide the additional benefit of increasing public trust in industry and government decision makers, and the technology itself (Slovic, 2000a).

5.2. A contractualist management of risks

The analysis of risks does not change considerably under a contractualist risk assessment framework, except to both qualify and (ideally) quantify the localized consequences, with the goal of better informing risk management. Risk management should then proceed with public deliberation focusing on whether the principles upon which a decision is based are reasonable and whether the decision can be justified even to those most affected (Cranor, 2007). The public, informed by the transparent processes of the risk analysis and acting as necessary stakeholders in the project, can then work with the experts to evaluate the decisions and devise adjustments and/or alternatives that better address the specific risks and fears of the community. For example, a decision made to bisect a community with a new highway might be accompanied by a discussion of what incentive or reparation justifies this decision to those who would have to relocate. Hosting a large infrastructure project necessarily exposes the community to increased risks, and as a result they deserve more that more attention be paid to reducing those risks wherever possible. This is, of course, largely standard practice in current risk assessments – various methods including replacement, mitigation, and transferal are adopted to address risks found to be untenable.

The chief improvement in risk management is the inclusion of the public in deliberation over, principally, the soundness of the reasons for the decision, and what would be considered fair compensation in the event the risk is realized. This gives the public the agency over their lives demanded by their status as moral agents, while retaining the ultimate decision authority for the experts. This, in my opinion, combines the expertise available through a technocratic risk assessment with the participatory nature necessary for making risk management decisions regarding questions of value.

6. Discussion and conclusions

Clearly, the above contractualist framework for risk assessment is a high-level abstraction rather than a clearly defined process. If pursued through future research, this framework will become more concrete and workable, although care should be taken to preserve some of the abstract nature to allow for flexible application and further evolution. Just as it is unlikely to obtain consensus on questions of value, it is unlikely that a rigidly-constructed ethical risk assessment process will work in all, or even most, applications. Instead, consensus should be pursued on normative questions regarding the conduct of risk assessment, such as the timeline for engaging a community, the process for eliciting public discourse around risk management decisions, or the makeup of the risk assessment team. Classifying risk assessment as an "incompletely theorized agreement" allows for participants – the public and experts – to agree on outcomes without agreeing on underlying theory (Sunstein, 1995), and thus allows both parties to satisfice and compromise along predictable, but flexible lines. This flexibility, in both application and on underlying theory, will ensure that risk assessments are performed according to similar norms without sacrificing robustness or allowing either the public or experts to usurp control of the process.

Section 3.2.2. raised two questions regarding the ethics of risk assessment – namely whether it is ethical to jeopardize the common good for a minority to benefit the majority, and whether there are plausible mechanisms for fairly considering local consequences in a risk assessment. I would argue that the answer to the second question determines the answer to the first. If there are plausible mechanisms for considering (and mitigating or compensating for) local consequences in risk assessment, it may be ethically sound to jeopardize the common good for the minority to provide a majority with benefits, because that minority was consulted during the risk assessment. However, this only holds if those plausible mechanisms are developed and followed. The contractualist framework proposed in Section 5, while incomplete and abstract, provides the basis for developing such mechanisms and involving the public as moral agents in decisions that affect their lives.

Effectively accounting for localized consequences is no easy task, especially as it will involve calculating not only monetary damages, but also nonmonetary consequences to the environment, quality of life, and health of the population. Determining, for instance, the value of a human life is a difficult proposition for experts working with an abstract human, let alone when the public is called to weigh in on the value of their own lives. Determining the best methods by which to qualify, quantify, and address localized consequences, particularly nonmonetary consequences, should be an avenue of future research. Further research should also investigate the efficacy of various mechanisms for eliciting public input during the risk management phase. One method for doing this efficiently could be the creation of community advisory panels, where volunteers from the community can act as intermediaries between the technical experts and the community itself. Such an arrangement would allow the community to have actual representation in the risk assessment process, while distributing the labor-intensive process of collecting and implementing public opinion to ensure that the risk assessment remains tractable. An approximation of this community-facing architecture is the Diablo Canyon Independent Safety Committee (DCISC), which holds regular public meetings, but critically has neither community representation (by non-experts) nor authority over plant safety (Diablo Canyon Independent Safety Committee, 2024).

Establishing an ethical risk assessment framework is critical to protecting the common good, respecting communities, and, most importantly, curtailing the cycle of saddling underprivileged communities with unfair risks. The goal of creating equitable, ethical processes for infrastructure siting have been undertaken through programs like the U.S. Department of Energy's "Consent-Based Siting Process" (U.S. Department of Energy Office of Nuclear Energy, 2023) which aims to engage communities about their needs and concerns for siting a consolidated spent nuclear fuel storage facility. Additionally, through the Biden Administration's "Justice40"

Initiative, the government is attempting to ensure an equitable transition to a green economy in which no communities are left behind (Callahan et al., 2021).

These initiatives are some of the first in the U.S. to widely acknowledge the role of the community in deciding for itself on whether to shoulder the risk engendered by an infrastructure project. While there are important lessons to be gleaned from these projects, developing an ethical risk assessment process broadly applicable to infrastructure will require more research. The considerations outlined in this paper show that approaching risk through the lens of ethics is an important step in ensuring that risk assessments are performed equitably and communities are treated fairly. As we are able to develop an appropriate quantification for consequence localization and normative processes for community inclusion, risk assessment will approach a truly ethical state that democratizes risk.

Future work in this area should address critical standing questions related to creating ethical risk assessments, including investigating how to model and quantify consequence localization. Perhaps more importantly for communities, however, is developing strategies that can meaningfully engage communities on the subject of risk through education and outreach, and exploring how to best structure community advisory groups and their relationship to the risk assessment. Finally, while (thankfully) not widely applicable, there is a need to understand how risk experts can engage communities that have already experienced consequences of risky projects. For such communities, probability-driven appeals to system safety may be less effective, because their perception of the risk will likely be governed strongly by the availability heuristic (i.e., their perception is that probability is unity). The lion's share of the work remains, but the potential benefits afforded by ethical risk assessment are, simply, too good to pass up. This is an opportunity for industries, especially the nuclear industry, to work proactively to regain public trust and demonstrate their commitment to system safety and community involvement. In my estimation, the burgeoning "nuclear renaissance" hinges in no insignificant part on taking an ethical view of risk assessment.

References

Balintec, V. 2023. What is NIMBYism and how is it affecting how much housing is getting built? CBC News. https://www.cbc.ca/news/canada/nimbyism-explainer-1.6909852.

Callahan, C., Coffee, D., DeShazo, J.R. & Gonzalez, S.R. 2021. Making Justice40 a Reality for Frontline Communities: Lessons from State Approaches to Climate and Clean Energy Investments. UCLA Luskin Center for Innovation. https://innovation.luskin.ucla.edu/wp-content/uploads/2021/10/luskin-justice40-final-web-1.pdf.

Coombs, K. 2022). Highways and Environmental Justice. Environmental Law Education Center. https://elecenter.com/1394/highways-and-environmental-justice/.

Cranor, C. 2007. Toward a Non-Consequentialist Approach to Acceptable Risks. In T. Lewens (Ed.), Risk: Philosophical Perspectives (1st ed., pp. 36–53). Taylor & Francis Group: Routledge.

Cushing, L.J., Li, S., Steiger, B.B., & Casey, J.A. 2023. Historical red-lining is associated with fossil fuel power plant siting and present-day inequalities in air pollutant emissions. Nature Energy, 8(1), Article 1. https://doi.org/10.1038/s41560-022-01162-y.

D'Andrea, M. A. & Reddy, G.K. 2018. The Development of Long-Term Adverse Health Effects in Oil Spill Cleanup Workers of the Deepwater Horizon Offshore Drilling Rig Disaster. Frontiers in Public Health, 6. https://www.frontiersin.org/journals/public-health/articles/10.3389/fpubh.2018.00117.

Dhara, V.R., Dhara, R., Acquilla, S.D. & Cullinan, P. 2002. Personal exposure and long-term health effects in survivors of the union carbide disaster at bhopal. Environmental Health Perspectives, 110(5), 487–500. https://doi.org/10.1289/ehp.02110487.

Diablo Canyon Independent Safety Committee. 2024. Diablo Canyon Independent Safety Committee (DCISC). Diablo Canyon Independent Safety Committee. https://www.dcisc.org.

Engineers, N.S. of P. 2019. NSPE Code of Ethics for engineers (Vol. 1102). National Society of Professional Engineers. https://doi.org/10.1007/BF03223280.

Eskenazi, B., Warner, M., Brambilla, P., Signorini, S., Ames, J. & Mocarelli, P. 2018. The Seveso accident: A look at 40 years of health research and beyond. Environment International, 121(Pt 1), 71–84. https://doi.org/10.1016/j.envint.2018.08.051.

Hansson, S.O. 2007. Risk and Ethics. In T. Lewens (Ed.), Risk: Philosophical Perspectives (1st ed., pp. 21–35). Taylor & Francis Group: Routledge.

Hussain, W. 2018. The Common Good. In The Stanford Encyclopedia of Philosophy (Spring 2018). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/entries/common-good/.

Jasanoff, S. 1998. The political science of risk perception. Reliability Engineering & System Safety, 59(1), 91–99. https://doi.org/10.1016/S0951-8320(97)00129-4.

Kusch, M. 2007. Towards a Political Philosophy of Risk. In T. Lewens (Ed.), Risk: Philosophical Perspectives (1st ed., pp. 131–155). Taylor & Francis Group: Routledge.

Modarres, M. 2006. Risk Analysis in Engineering: Techniques, Tools, and Trends (1st ed.). CRC Press Taylor & Francis Group.

Modarres, M., & Groth, K.M. 2023. Reliability and Risk Analysis (2nd ed.). CRC Press, Taylor & Francis Group.

Peschard, I., Benetreau-Dupin, Y. & Wessels, C. 2023. Philosophy and Science of Risk: An Introduction (1st ed.). Taylor & Francis Group: Routledge.

Scanlon, T. 2000. What We Owe to Each Other (1st ed.). Harvard University Press.

Slovic, P. 2000a. Perceived Risk, Trust and Democracy. In The Perception of Risk (pp. 316–326). Earthscan Ltd.

Slovic, P. 2000b. Perception of Risk. In The Perception of Risk (pp. 220–231). Earthscan Ltd.

Slovic, P., Fischhoff, B. & Lichtenstein, S. 2000a. Facts and Fears: Understanding Perceived Risk. In The Perception of Risk (pp. 137–153). Earthscan Ltd.
Slovic, P., Fischhoff, B., & Lichtenstein, S. 2000b. Rating the Risks. In The Perception of Risk (pp. 104–120). Earthscan Ltd.
Sunstein, C. R. 1995. Incompletely Theorized Agreements. Harvard Law Review, 108(7), 1733–1772. JSTOR. https://doi.org/10.2307/1341816.
Sunstein, C. R. 2002. Risk and Reason: Safety, Law, and the Environment (1st ed.). Cambridge University Press.
U.S. Department of Energy Office of Nuclear Energy. 2023. Consent-based Siting Process for Federal Consolidated Interim Storage of Spent Nuclear Fuel. U.S. Department of Energy Office of Nuclear Energy. https://www.energy.gov/sites/default/files/2023-05/Consent-Based%20Siting%20Process%20Report-0424%203.pdf.
Wynne, B. 1992. Misunderstood misunderstanding: Social identities and public uptake of science. Public Understanding of Science, 1(3), 281–304. https://doi.org/10.1088/0963-6625/1/3/004.