

Risk thresholds for frontier AI

Leonie Koessler* Jonas Schuett Markus Anderljung

Centre for the Governance of AI

Abstract

Frontier artificial intelligence (AI) systems could pose increasing risks to public safety and security. But what level of risk is acceptable? One increasingly popular approach is to define *capability thresholds*, which describe AI capabilities beyond which an AI system is deemed to pose too much risk. A more direct approach is to define *risk thresholds* that simply state how much risk would be too much. For instance, they might state that the likelihood of cybercriminals using an AI system to cause X amount of economic damage must not increase by more than Y percentage points. The main upside of risk thresholds is that they are more principled than capability thresholds, but the main downside is that they are more difficult to evaluate reliably. For this reason, we currently recommend that companies (1) define risk thresholds to provide a principled foundation for their decision-making, (2) use these risk thresholds to help set capability thresholds, and then (3) primarily rely on capability thresholds to make their decisions. Regulators should also explore the area because, ultimately, they are the most legitimate actors to define risk thresholds. If AI risk estimates become more reliable, risk thresholds should arguably play an increasingly direct role in decision-making.




Figure 1: Risk thresholds can directly and indirectly inform high-stakes AI development and deployment decisions.

*Corresponding author: leonie.koessler@governance.ai.

Contents

1	Introduction	5
2	Risk thresholds and related concepts	6
2.1	Risk thresholds	7
2.2	Capability thresholds	8
2.3	Compute thresholds	9
3	How to use AI risk thresholds	10
3.1	Using risk thresholds to directly feed into decisions	10
3.2	Using risk thresholds to indirectly feed into decisions	11
4	The case for AI risk thresholds	12
4.1	Arguments for using risk thresholds	13
4.2	Arguments against using risk thresholds	14
4.3	Overall suggestions for using AI risk thresholds	14
5	How to define AI risk thresholds	15
5.1	Type of risk	16
5.2	Level of risk	17
6	Conclusion	18
	References	20

Executive summary

Frontier artificial intelligence (AI) systems could pose increasing risks to public safety and security (e.g. through cyberattacks on critical infrastructure, the acquisition of biological weapons, or loss of control over AI systems). These risks could largely stem from a small number of high-stakes development and deployment decisions made by frontier AI companies (e.g. whether to start a large training run or whether to release a model). When making such decisions, companies do not seem to use risk thresholds, i.e. limits for what likelihood and severity of harm they are willing to accept. Instead, where companies have defined thresholds for what AI systems are too risky to release, those thresholds have been defined in terms of model capabilities. This paper draws on other industries to discuss how to use risk thresholds for making high-stakes AI development and deployment decisions.

Risk thresholds serve a different function than capability thresholds and compute thresholds (Section 2).

- **Compute thresholds.** Compute thresholds are defined in terms of computational resources used to train a model (“training compute”). Training compute is a very imperfect proxy for risk, but can easily be measured and forecasted early on in the development process. Compute thresholds should thus be used as an initial filter to identify models that warrant further scrutiny, oversight, and precautionary safety measures.
- **Capability thresholds.** Model capabilities are a better proxy for risk than training compute and are easier to evaluate than risk. Capability thresholds may therefore serve as a key trigger for whether additional safety measures should be implemented before a high-stakes activity may go ahead.
- **Risk thresholds.** Risk estimates try to measure the level of risk directly, but they are still highly unreliable. In theory, risk thresholds are the ideal determinator for when additional safety measures are necessary. But in practice, risk thresholds cannot yet be relied upon for decision-making. More on the role they should play below.

In principle, there are two ways in which risk thresholds can be used: they can directly feed into high-stakes AI development and deployment decisions and they can indirectly feed into such decisions by helping set capability thresholds (Section 3). These two ways are illustrated in Figure 1.

- **Directly feeding into decisions.** Using risk thresholds to directly feed into high-stakes decisions is the most common use case for risk thresholds in other industries. Before making a high-stakes decision, many companies compare risk estimates to predefined risk thresholds. If the estimated level of risk is above the risk thresholds, companies implement additional safety measures and repeat the process. This process is similar to how some frontier AI companies evaluate model capabilities and compare them to predefined capability thresholds, but with a focus on risk rather than model capabilities. In this way, both risk thresholds and capability thresholds can directly feed into high-stakes decisions.
- **Indirectly feeding into decisions.** Using risk thresholds to indirectly feed into decisions is less common in other industries. One exception are U.S. nuclear regulators who use risk thresholds to determine adequate safety measures. In the context of frontier AI, capability thresholds and corresponding safety measures could be designed such that they would be estimated to keep risk below some risk thresholds. To that end, risk thresholds need to be defined. Next, risk models can be developed, i.e. mappings of pathways from risk factors to harm. These risk models can help identify the model capabilities at which risk would exceed the risk thresholds, and the safety measures that would keep risk below the risk thresholds. The identified model capabilities then serve as the capability thresholds that trigger the identified safety measures.

We argue that risk thresholds are a promising tool for frontier AI regulation (Section 4).

- **Arguments for using risk thresholds.** Risk thresholds may help align business conduct with societal concern; enable consistent allocation of safety resources; ensure risk estimation results are actually acted upon; prevent motivated reasoning regarding what level of risk is acceptable; and avoid locking in premature safety measures.

- **Arguments against using risk thresholds.** Risk thresholds rely on risk estimates but estimating risks from AI is extremely hard; AI is a dual-use, general-purpose technology; risk thresholds may create an incentive to produce artificially low risk estimates; and defining risk thresholds for AI involves handling thorny normative trade-offs.
- **How risk thresholds should be used.** Overall, we suggest that risk thresholds should be used to indirectly feed into decisions by helping set capability thresholds. Yet risk thresholds should only inform, but not determine, where to set capability thresholds: risk thresholds should not be the sole basis of a strict decision-rule. Other considerations should also be taken into account when setting capability thresholds. Further, we suggest that risk thresholds may be used to directly feed into decisions. However, again, risk thresholds should only inform decisions (e.g. as one of a number of considerations), and not determine decisions (e.g. as the sole basis for a strict decision-rule). If and when our ability to produce risk estimates improves, we can rely more on risk thresholds.

Finally, we propose a framework for how to define risk thresholds for frontier AI (Section 5).

Before regulators or companies can answer the question of what level of risk is acceptable, they need to decide which type of risk the threshold should refer to, that is, which risk scenarios are in scope. Next, when determining the acceptable level of risk, they need to handle three related normative trade-offs: (1) how to weigh potential harms and benefits, (2) to what extent should mitigation costs be taken into account, and (3) how to deal with uncertainty regarding all of the aforementioned.

We encourage frontier AI companies to start experimenting with risk thresholds today. Regulators should also explore the area because, ultimately, they are the most legitimate actors to define risk thresholds. To this end, we need a discussion about what level of risk we, as a society, are willing to accept.

1 Introduction

Frontier artificial intelligence (AI) systems¹ pose increasing risks² to public safety and security³ (Bengio et al., 2024; Hendrycks et al., 2023; Anderljung et al., 2023). For example, frontier AI systems may already increase cybercriminal productivity (Fang et al., 2024; Hazell, 2023; Lohn & Jackson, 2022; Mirsky et al., 2021), while future systems might increase the risk that terrorists will succeed in acquiring biological weapons (Boiko et al., 2023; Mouton et al., 2023; Sandbrink, 2023; Soice et al., 2023; Urbina et al., 2022). A more speculative concern is that, at some point, frontier AI systems might evade human control and cause large-scale harm on their own (Chan et al., 2023; Cohen et al., 2024; Hendrycks et al., 2023; Ngo et al., 2024).

These risks could largely stem from a small number of high-stakes development and deployment decisions made by frontier AI companies, such as whether to start a final large training run or whether to deploy a model, also referred to as “go/no-go decisions” (NIST, 2023). When making these decisions, companies necessarily accept some level of risk.⁴ For example, a company deploying a system could be accepting a 0.01% increase in the risk that a malicious actor will succeed in acquiring a biological weapon based on instructions from that system. Many frontier AI companies seem to consider potential harms and benefits to society in their decision-making (e.g. Anthropic, 2023; Google AI, 2018; Google DeepMind, 2024; Meta, 2023; Microsoft, 2024; OpenAI, 2023). However, companies do not appear to have clear limits for what likelihood and severity of harm they are willing to accept, so-called “risk thresholds”.⁵

At the 2024 AI Summit in South Korea, governments and companies both emphasized the importance of setting thresholds above which risk would be unacceptable. The Seoul Ministerial Statement includes the intention to “identify thresholds at which the level of risk posed by the design, development, deployment and use of frontier AI models or systems would be severe absent appropriate mitigations” (DSIT, 2024c). The Seoul Frontier AI Safety Commitments had 16 companies commit to “set out thresholds at which severe risks posed by a model or system, unless adequately mitigated, would be deemed intolerable”, while noting that “thresholds can be defined using model capabilities, estimates of risk [i.e. “risk thresholds”], implemented safeguards, deployment contexts and/or other relevant risk factors” (DSIT, 2024b). While the past year has seen frontier AI companies increasingly define thresholds in terms of model capabilities, it is unclear whether these thresholds keep risk to an acceptable level. This paper draws on other industries to discuss how regulators and companies should use risk thresholds for making high-stakes AI development and deployment decisions.

There is an extensive body of literature on risk thresholds in other industries. Technical standards provide high-level guidance on how to use risk thresholds for business risks (e.g. COSO, 2017; ISO, 2018; ISO & IEC, 2019). The scholarly literature provides more in-depth guidance for business and societal risks (e.g. Aven, 2012, 2015; Popov et al., 2021; Rausand & Haugen, 2020) and discusses

¹We define “frontier AI systems” as “highly capable general-purpose AI models or systems that can perform a wide variety of tasks and match or exceed the capabilities present in the most advanced models” (DSIT, 2024b). For example, this currently includes systems like GPT-4, Claude 3, and Gemini Ultra. Note that, in contrast to an earlier, otherwise identical definition (DSIT, 2023a), this definition has replaced “today’s most advanced models” with “the most advanced models”, which implies that the frontier changes as models become more capable. We also note that the term “frontier AI” has been accused of promoting a specific worldview (Helfrich, 2024).

²We define “risk” as the combination of likelihood and severity of harm (ISO & IEC, 2014). A recent trend in risk management uses a definition of risk that includes both negative impacts, i.e. harm, and positive impacts, i.e. benefits (COSO, 2017; ISO, 2018; NIST, 2023). However, in the context of risk thresholds, the understanding of risk typically only includes harm, whereas benefits come into play as the key consideration when choosing what level of risk is acceptable (see Section 5.2).

³Note that for the purposes of this paper, we focus on risks to individuals, groups, and society as a whole, i.e. societal risks (e.g. fatalities, economic damage, and societal disruption). This means we ignore risks to the company itself, i.e. business risks (e.g. financial risks, legal risks, and reputational risks). We also focus on risks to public safety and security, but the tools we discuss can likely be applied to many other types of societal harm, too.

⁴We define “level of risk” as the combined measure of the likelihood and severity of harm.

⁵Google AI (2018) commits that it “will not design or deploy AI (...) applications [that] cause or are likely to cause overall harm. Where there is a material risk of harm, we will proceed only where we believe that the benefits substantially outweigh the risks, and will incorporate appropriate safety constraints.” This can be understood as a risk threshold, albeit a very vague one. Much depends on how Google operationalizes this risk threshold.

various issues with risk thresholds, including substantial uncertainties in risk estimates (e.g. Fischhoff et al., 1984; Klinke & Renn, 2002; Starr, 1969). Regulators in many safety-critical industries mandate or recommend specific risk thresholds, such as in the nuclear (ANVS, 2020; IAEA, 2005; NRC, 1983), maritime (IMO, 2018), aviation (EUROCONTROL, 2001; FAA, 1988; ICAO, 2018), and space industries (ESA, 2023; FAA, 2016). In addition to a large corpus of industry-specific literature, many reports survey the use of risk thresholds across industries and jurisdictions (e.g. CCPS, 2009; Ehrhart et al., 2020; Flamberg et al., 2016; Linkov et al., 2011; Marhavidas & Koulouriotis, 2021).

By contrast, in the context of frontier AI development and deployment, regulators and scholars are only starting to discuss risk thresholds. The NIST AI Risk Management Framework recommends that companies define “risk tolerances” (NIST, 2023), but does not provide much guidance for how to define or use them. DSIT’s policy paper Emerging Processes for Frontier AI Safety recommends that companies use risk thresholds in responsible capability scaling (DSIT, 2023b), but it only provides high-level guidance. Further, the forthcoming EU AI Act mandates that risk management measures for general-purpose AI models with systemic risk “shall be proportionate to the risks [and] take into consideration their severity and probability” (Article 56(2)(d)). This could be ensured by using risk thresholds.⁶ There is only tangential scholarly treatment of AI risk thresholds (Clymer et al., 2024). Taken together, there is a clear need for more concrete guidance on how to use risk thresholds in the context of frontier AI. This paper aims to help fill this gap.

The paper proceeds as follows. First, we introduce the concept of risk thresholds as a specific type of risk acceptance criteria and differentiate it from the related concepts emerging in the frontier AI context: capability thresholds and compute thresholds (Section 2). We then outline how risk thresholds can be used to directly and indirectly feed into high-stakes AI development and deployment decisions (Section 3). Next, we argue that risk thresholds should only be used to inform, but not determine, high-stakes decisions, unless risk estimates become more reliable (Section 4). We also highlight key considerations and provide initial guidance for defining AI risk thresholds (Section 5). We conclude with a summary of our main contributions and suggestions for further research (Section 6).

2 Risk thresholds and related concepts

In frontier AI regulation, different thresholds are currently emerging: risk thresholds, capability thresholds, and compute thresholds. These thresholds are predefined values above which additional safety measures are deemed necessary. The thresholds differ regarding the *metric* in terms of which they are defined (risk, model capabilities, and training compute), and the *function* they serve in frontier AI regulation (we discuss this for each threshold below). For example, Anthropic’s Responsible Scaling Policy maps specific model capabilities to specific safety measures (Anthropic, 2023), whereas the EU AI Act classifies general-purpose AI models trained on more than 10^{25} floating-point operations as posing systemic risk (Article 51(2)) and imposes more stringent requirements on their providers (Article 55(1)).

In this paper, we are most interested in thresholds that can be used in high-stakes decision-making to determine whether the risk from the development and deployment of a frontier AI system is acceptable. This includes both risk thresholds and capability thresholds, though we will focus on risk thresholds based on risk estimates (see DSIT, 2024b).

In the remainder of this section, we first conceptualize risk thresholds as risk acceptance criteria and outline how they are used in other industries to directly and indirectly feed into high-stakes decisions (Section 2.1). We then argue that capability thresholds can also be considered risk acceptance criteria that may serve as a key trigger for when to implement additional safety measures in the frontier AI context (Section 2.2). Finally, we assert that compute thresholds should not be considered risk acceptance criteria but only serve as an initial filter to identify models of potential concern (Section 2.3).

⁶Similarly, risk management measures for high-risk AI systems “shall be such that the relevant residual risk associated with each hazard, as well as the overall residual risk of the high-risk AI systems is judged to be acceptable” (Article 9(5)). For related discussions, see (Fraser & Bello y Villarino, 2023; Laux et al., 2024), and (Schuett, 2023). Moreover, the EU AI Act puts AI systems into risk categories, the boundaries of which have been referred to as risk thresholds (Novelli et al., 2024), although they are not defined in terms of likelihood and severity of harm, and therefore do not qualify as risk thresholds according to our definition.

2.1 Risk thresholds

Risk thresholds are limits to what level of estimated risk is acceptable (Aven, 2015). Thus, they are also referred to as “risk limits” (ISO & IEC, 2019), “tolerability limits” (Aven, 2015), or “risk tolerances” (NIST, 2023). In the context of business risks, risk thresholds are also sometimes referred to as companies’ “risk appetite” (COSO, 2017; ISO & IEC, 2019). Risk thresholds vary across different industries and jurisdictions (Ehrhart et al., 2020; Flamberg et al., 2016; Linkov et al., 2011). For example, in the U.S. aviation industry, the probability of “failure conditions which would prevent continued safe flight and landing” should not exceed 1×10^{-9} (one in a billion) per flight-hour (FAA, 1988). As another example, in the UK nuclear industry, the risk of death of a member of the public is “unacceptable” if it is above 1×10^{-4} per plant-year and “broadly acceptable” if it is below 1×10^{-6} per plant-year (ONR, 2020).

Risk thresholds can be understood as a particular type of “risk acceptance criteria”, i.e. criteria that establish the conditions under which risk is acceptable to an organization (e.g. a regulator or a company). Therefore, risk acceptance criteria are also referred to as “risk evaluation criteria”, “decision criteria for risk management decision making”, or simply “risk criteria” (Aven, 2012, 2015, 2016; ISO, 2018; ISO & IEC, 2019; Morgan & Henrion, 1990).⁷ Risk acceptance criteria beyond risk thresholds can take many forms. For example, risk may be acceptable if it is “as low as reasonably practicable” (“ALARP”), if the “best available technology” (“BAT”) is used, or if the affected individuals have given consent (Klinke & Renn, 2002; Morgan & Henrion, 1990; Vanem, 2012). Compared to other types of risk acceptance criteria, risk thresholds are more often quantitative, although they can also be qualitative (e.g. “only proceed if risk is deemed low”). However, in the regulatory context, qualitative risk thresholds appear to be very uncommon.

We highlight that choosing a type of risk acceptance criteria may reflect a particular ethical viewpoint. Although this viewpoint can significantly affect which risks are deemed acceptable, it is rarely made explicit. Common ethical principles that may underlie different types of risk acceptance criteria include principles of utility, fairness, and human rights (Morgan & Henrion, 1990; Vanem, 2012). Risk thresholds may draw most strongly on the principle of utility, because they focus on potential harms and benefits, outcomes rather than processes, and general welfare rather than individual liberties.⁸ However, other principles can be taken into account via the design of the risk thresholds (see Morgan & Henrion, 1990; Vanem, 2012). For example, U.S. oil and gas facilities have to observe stricter risk thresholds regarding particularly vulnerable groups in places such as schools, hospitals, and prisons (NFPA, 2023). Furthermore, participatory elements can be included when setting risk thresholds, for instance, through public consultations (e.g. NRC, 1983). Finally, we do not argue that risk thresholds should be the only risk acceptance criteria in the frontier AI context.

Risk thresholds are defined in terms of likelihood and severity of harm. Likelihood scales refer to the probability of events, which can be estimated using historical data, models, or expert judgment, among other things. Severity scales refer to the magnitude or degree of some type of harm, such as fatalities, injuries, or economic damage. They can also be defined in terms of potentially harmful events, such as a successful cyberattack, the acquisition of a biological weapon, or the creation of a deepfake.⁹ Both likelihood and severity scales can be quantitative (i.e. numeric values, e.g. probabilities or numbers of fatalities), semi-quantitative (i.e. ranges of numeric values, e.g. 1 – 5% or 10, 000-100, 000 fatalities), or qualitative (i.e. categories based on non-numeric values, e.g. “likely” or “severe”) (ISO & IEC, 2019). Risk thresholds consist of a single pair of likelihood and severity values (e.g. an expected value) or several pairs of likelihood and severity values (e.g. a probability distribution). The latter seems to be much more common in the regulatory context, at least for fatalities (e.g. EUROCONTROL, 2001; HSE, 2001; NRC, 1983). Quantitative risk thresholds can be visualized in graphs (e.g. F/N diagrams with fatalities N on the x -axis and frequencies F on the

⁷Note that concepts and terminology vary among sources or are simply unclear. Some authors seem to equate risk thresholds with risk acceptance criteria (e.g. Linkov et al., 2011), whereas other authors seem to understand risk thresholds as quantitative risk acceptance criteria (e.g. (Flamberg et al., 2016)). For most authors, it simply remains unclear how they conceptualize the relationship between risk thresholds and risk acceptance criteria.

⁸On the relationship between utility and rights, see e.g. (Hart, 2017).

⁹In this paper, for simplicity, we focus on events that are intrinsic harms. On the one hand, the likelihood of potentially harmful events will usually be easier to estimate than the likelihood of intrinsic harms. On the other hand, the question of at what level to set the threshold is even more complicated for potentially harmful events than it already is for intrinsic harms (Section 5.2).




Figure 2: F/N-diagram (quantitative) and risk matrix (semi-quantitative / qualitative) (ISO & IEC, 2019)

y-axis), whereas semi-quantitative and qualitative risk thresholds can be visualized in risk matrices (Figure 2).

Risk thresholds can feed into high-stakes decisions in two ways: *directly* and *indirectly*. First, when companies make high-stakes decisions, risk thresholds can be used to help decide whether an activity may go ahead (ISO, 2018). In this way, risk thresholds *directly* feed into high-stakes decisions. This is the most common way in which other industries use risk thresholds. Second, instead of using risk thresholds on a case-by-case basis, risk thresholds can also be used to help specify which safety measures need to be implemented under which circumstances. In this way, risk thresholds *indirectly* feed into high-stakes decisions. In the U.S. nuclear industry, “safety goals (...) are to be used (...) in making regulatory judgments on the need of proposing and backfitting new generic requirements on nuclear power plant licensees” (NRC, 2021). Similarly, Anthropic evaluates for “capability improvements (...) [that] would significantly increase the risk (...) past an unacceptable threshold” to decide when additional safety measures are necessary (Anthropic, 2023). We elaborate on how to use risk thresholds in the frontier AI context in Section 3.

2.2 Capability thresholds

For risks to public safety and security, model capabilities can be considered a key risk factor and even an imperfect proxy for risk. Fundamentally, risks from frontier AI stem from the capabilities a model possesses, because many of these capabilities are dual-use: they can be used for good or for evil (Anderljung & Hazell, 2023; Bommasani et al., 2021; Shevlane & Dafoe, 2020). For example, a model that can be used by scientists to help develop new pharmaceuticals might also be used by terrorists to help develop new toxins (Urbina et al., 2022). Thus, model capabilities can be considered a key risk factor. They can even be considered a proxy for risk, a claim that has been made explicitly by some (e.g. Sastry et al., 2024) and to some extent implicitly relied upon by others (Anderljung et al., 2023; Shevlane et al., 2023; OpenAI, 2023). But factors other than model capabilities are crucial for risk too, such as the number, capacity, and willingness of malicious actors to use the model or the level of societal preparedness (Anderljung et al., 2023; Bernardi et al., 2024; Kapoor et al., 2024). Nevertheless, a model’s capabilities are easier to evaluate than its risk (Section 4.2), making them a useful metric for frontier AI regulation.

Already, frontier AI companies increasingly rely on capability thresholds to make high-stakes development and deployment decisions. Capability thresholds are predefined model capabilities at which additional safety measures are deemed necessary. Three frontier AI companies have published policies that define capability thresholds and when additional safety measures should be implemented




Figure 3: Different metrics and the relationships between them

before these capability thresholds are crossed. This includes Anthropic’s Responsible Scaling Policy (Anthropic, 2023), OpenAI’s Preparedness Framework (OpenAI, 2023), and Google DeepMind’s Frontier Safety Framework (Google DeepMind, 2024). These policies focus on chemical, biological, radiological, and nuclear (CBRN); cyber; persuasion; autonomy; and some other capabilities, measured by so-called “model evaluations” (Shevlane et al., 2023; Phuong et al., 2024). Regulators have yet to make use of capability thresholds, but some of them already seem to be thinking along these lines (DSIT, 2023b).

While concepts in the frontier AI context are still evolving, capability thresholds can be considered their own type of risk acceptance criteria. Capability thresholds essentially define conditions under which a risky activity may go ahead, namely if a model’s capabilities are below the threshold or if they are above the threshold but adequate safety measures have been implemented. In this way, capability thresholds can be considered a type of risk acceptance criteria that is distinct from risk thresholds. However, concepts in the frontier AI context are still evolving. In particular, not all model evaluations only measure inherent model capabilities; they may also include assessments of how users or even society as a whole interact with models (DSIT, 2024a; Patwardhan et al., 2024; Solaiman et al., 2023). Moreover, risk thresholds can be used to help with setting capability thresholds, a process that blurs the lines between risk thresholds and capability thresholds.¹⁰ We get back to this in Section 3.2.

2.3 Compute thresholds

Under the current deep learning paradigm, the amount of computational resources used to train a model (“training compute”) can be considered a very imperfect proxy for a model’s risk. Empirically, recent advances in model capabilities to a large extent stem from increasing amounts of computational resources being used to train the model, a phenomenon also referred to as “scaling laws” (Sutton, 2019; Kaplan et al., 2020; Hernandez et al., 2021; Hoffmann et al., 2022). While how long scaling laws will hold is somewhat contentious (Lohn & Musser, 2022; Villalobos et al., 2022), training compute can be, at least currently, considered a proxy for a model’s capabilities and thereby also the model’s risk (Anderljung et al., 2023; Sastry et al., 2024). However, training compute is even further removed from risk than model capabilities – they already are an imperfect proxy for risk – meaning that training compute is only a very imperfect proxy for risk. Still, training compute is relatively easy to measure, making it a useful metric to build on in frontier AI regulation (Anderljung et al., 2023; Heim & Koessler, forthcoming; Pistillo et al., forthcoming; Sastry et al., 2024). We show the relationship between the three metrics in Figure 3.

Indeed, regulators in the U.S. and the EU already make use of *compute thresholds* to identify models that might be of concern and require increased scrutiny, oversight, and precautionary security measures. The U.S. Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence imposes requirements on companies developing and deploying models above a training compute threshold of 10^{26} operations to notify the government before development; report on ownership and possession of model weights and measures taken to secure them; and report on the results of red-teaming tests and measures taken based on them (Section 4.2(i)). Setting a lower threshold while imposing more extensive requirements, the EU AI Act uses a compute threshold of 10^{25} floating-point operations to identify “general-purpose AI models” that may pose “systemic risk” (Article 51(2)), and requires providers of such models to conduct model evaluations; assess and

¹⁰Indeed, OpenAI refers to its capability thresholds as “risk thresholds” (OpenAI, 2023) – presumably because it aims for its capability thresholds to keep risk at an acceptable level. However, OpenAI does not define its thresholds in terms of likelihood and severity of harm, but model capabilities. Therefore, according to our definitions, these thresholds are capability thresholds, and not risk thresholds.

Compute thresholds	Capability thresholds	Risk thresholds
Initial filter for further scrutiny, oversight, and precautionary safety measures (e.g. security)	Key trigger for when additional safety measures are necessary, including fast responses (e.g. pausing)	Ideal, though immature, determinant for when additional safety measures are necessary; risk thresholds can <i>directly</i> feed into high-stakes decisions and <i>indirectly</i> feed into high-stakes decisions by helping set capability thresholds

Table 1: Different thresholds and their functions in frontier AI regulation

mitigate systemic risks; track, document, and report serious incidents; and ensure an adequate level of cybersecurity for the model and its physical infrastructure (Article 55(1)).

Compute thresholds should not be used as risk acceptance criteria to directly feed into high-stakes decisions. Because compute thresholds are such an imperfect proxy for risk, they should not be used to define conditions under which a risky activity may go ahead. Instead, compute thresholds may serve as an initial filter for further scrutiny, oversight, and some precautionary safety measures (Heim & Koessler, forthcoming; Pistillo et al., forthcoming). Capability thresholds and risk thresholds can then be used to help decide whether high-stakes decisions may go ahead and under which circumstances additional safety measures are warranted. The U.S. Executive Order on AI and the EU AI Act laudably use compute thresholds mainly in this way. Overall, risk thresholds, capability thresholds, and compute thresholds are not substitutes for each other; each has a distinct function in frontier AI regulation (Table 1)²

3 How to use AI risk thresholds

In this section, we discuss two ways in which risk thresholds can be used: to directly feed into high-stakes AI development and deployment decisions (Section 3.1) and to indirectly feed into decisions by helping set capability thresholds (Section 3.2). These two use cases are illustrated in Figure 1.

3.1 Using risk thresholds to directly feed into decisions

Using risk thresholds to directly feed into high-stakes decisions is the most common use case for risk thresholds in other industries (Section 2.1). In the standard risk management process, organizations estimate the level of risk (“risk analysis”) and compare the results to predefined risk thresholds (“risk evaluation”). If the estimated level of risk is above the risk threshold, companies need to implement additional safety measures and repeat the process (ISO, 2018). This process is similar to how some frontier AI companies evaluate model capabilities and compare them to predefined capability thresholds (Section 2.2), but with a focus on risk rather than model capabilities. In this way, both risk thresholds and capability thresholds can directly feed into high-stakes decisions (Figure 1).

If the estimated level of risk exceeds a risk threshold, the company needs to implement additional safety measures. In the simplest version of risk thresholds, the company can freely choose among safety measures as long as it brings the level of risk below the risk threshold before proceeding. But risk thresholds can also require companies to take more specific safety measures. For example, companies may take measures to reduce risk or uncertainty about risk as much as possible, or they may notify some internal or external stakeholder. In the case of frontier AI, risk thresholds could also require companies to notify the board of directors or the competent regulator (who may be allowed to veto the decision or be required to give permission to go ahead), or to conduct an extra suite of in-depth model evaluations (which may have to include external parties). Using safety measures other than a clear “no-go” also provides a way for risk thresholds to inform, but not determine, high-stakes decisions. We will get back to this in Section 4.3.




Figure 4: The ALARP framework (Melchers, 2001)

It is possible to set a single risk threshold or multiple risk thresholds that trigger different safety measures. The simplest approach is to set a single risk threshold that distinguishes two risk tiers. If this risk threshold is crossed, the activity in question may not go ahead unless either the risk has been reduced or specific safety measures have been taken. But it is also possible to set multiple risk thresholds at different levels of risk to distinguish between more than two risk tiers. For example, two risk thresholds could distinguish between risk being unacceptable under any circumstances, risk being acceptable if specific safety measures have been taken (e.g. risk has been reduced as much as possible, specific information has been gathered, or specific internal or external actors have been notified), and risk being acceptable without further safety measures. Stacking multiple risk thresholds in this way allows predefining more fine-grained decision rules for different levels of risk compared to a single risk threshold.

For example, many industries stack two risk thresholds: one threshold above which risk is unacceptable and another one above which risk must be “as low as reasonably practicable” (“ALARP”), also sometimes referred to as “as low as reasonably achievable” (“ALARA”) (Linkov et al., 2011; Melchers, 2001). The ALARP framework, as illustrated in Figure 4, originated in the UK Health and Safety at Work etc. Act 1974 and has since been used in various industries worldwide, including the UK and U.S. nuclear industry (HSE, 1992; NRC, 2016), the U.S. aerospace industry (Dezfuli et al., 2015), and the international maritime industry (IMO, 2018). Typically, risk is considered ALARP if the costs of further risk reduction measures are “grossly disproportionate” to their benefits (HSE, 1992). The ALARP principle may ensure continuous risk reduction efforts (Aven, 2015). However, it has also been criticized for being vague, leading to harmful risk aversion, and stifling innovation (Melchers, 2001; Oakley & Harrison, 2020). Theoretically, any other types of risk acceptance criteria could be applied at the risk tier in the middle, allowing risk thresholds to be combined with the other types of risk acceptance criteria mentioned in Section 2.1.

3.2 Using risk thresholds to indirectly feed into decisions

Risk thresholds can also be used to indirectly feed into decisions, such as whether to deploy a model, by helping set capability thresholds. U.S. nuclear regulators use risk thresholds to help determine adequate safety measures (Section 2.1). In the context of frontier AI, capability thresholds are emerging as a key trigger for safety measures (Section 2.2). Capability thresholds and corresponding safety measures could be designed such that they would be estimated to keep risk below some risk threshold. In this way, as capability thresholds directly feed into high-stakes decisions, risk thresholds indirectly feed into decisions (Figure 1).

A helpful tool when using risk thresholds to help set capability thresholds is “risk models” (see Google DeepMind, 2024), also referred to as “threat models” (Anthropic, 2023).¹¹ Risk models

¹¹The latter term is currently the most common, but may not be the most suitable. It stems from a security context and may thus lead to a narrow focus on security risks. Moreover, “threat modeling” encompasses more than outlining risk scenarios, in particular, prioritizing among safety measures (Shostack, 2014). In standard risk management, common terms to refer to risk models are “fault trees” and “event trees” (Barrett & Baum, 2017),

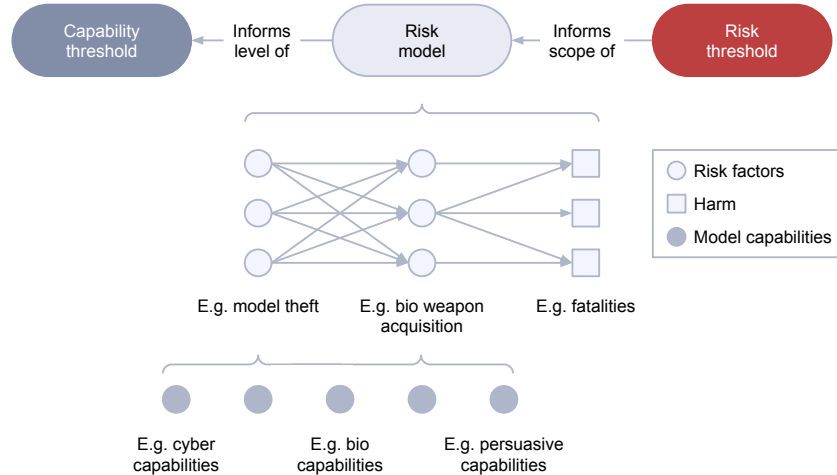


Figure 5: Risk thresholds, for example via risk models, can help set capability thresholds

outline the pathways from risk factors to harm, or “risk scenarios” (see OpenAI, 2023). They can be used to identify model capabilities that may cause large-scale harm and safety measures that may prevent such harm. For example, the capability to provide instructions for the acquisition of biological weapons (dark circle) may increase the risk of fatalities, economic damage, and societal disruption (squares). Risk models may help identify the level of bio capabilities at which the level of risk would exceed the risk threshold unless safety measures are implemented (Figure 5).

In more detail, risk thresholds can be used to set capability thresholds via risk models with the following steps. First, define risk thresholds. We provide some guidance for doing so in Section 5. Second, develop risk models. Ideally, risk models are comprehensive, meaning they contain all possible pathways from risk factors to harm. However, developing comprehensive risk models is generally extremely difficult (Shostack, 2014) and particularly so in the case of a general-purpose technology like AI (Section 4.2). Therefore, at least in the beginning, risk models may focus on a small number of key risk scenarios. Third, identify model capabilities that would lead to unacceptable risks as defined in the first step. This can draw on, for example, the risk models developed in the second step, data gathered about the occurrence of risk factors, near misses, and small-scale harm, as well as methods like trend extrapolation and sensitivity analysis (Frey & Patil, 2002).

Frontier AI companies setting capability thresholds already aim to identify the model capabilities that may lead to large-scale harm. In particular, companies increasingly engage in risk modeling (Anthropic, 2023; OpenAI, 2023; Google DeepMind, 2024). Yet, when doing so, companies currently seem to mostly focus on the *possibility* that model capabilities may cause large-scale harm rather than also considering the *likelihood* of this happening. Ignoring likelihood means ignoring a key component of risk and can lead to overly restrictive capability thresholds, because other factors may prevent harm from materializing, such as malicious actors not having access to the model or society ramping up its defenses. At least one company is planning on taking likelihood into account in the future (Anthropic, 2023).

4 The case for AI risk thresholds

In this section, we argue that risk thresholds are a promising tool for making high-stakes AI development and deployment decisions. Risk thresholds may help align business conduct with societal concern; enable consistent allocation of safety resources; ensure risk estimation results are actually acted upon; prevent motivated reasoning regarding what level of risk is acceptable; and avoid locking in premature safety measures (Section 4.1). We also discuss the most important objections to using AI risk thresholds and how they might be overcome. In particular, estimating risks from AI is extremely hard; AI is a dual-use, general-purpose technology; risk thresholds may create an incentive to produce

“attack trees” as a variation of fault trees for security risks (Salter et al., 1998; Schneier, 2011), and “causal maps” to depict non-linear relationships between risk factors (ISO & IEC, 2019; Koessler & Schuett, 2023).

artificially low risk estimates; and defining risk thresholds for AI involves handling thorny normative trade-offs (Section 4.2). Overall, we suggest that risk thresholds should be used indirectly inform high-stakes decisions by helping set, though not determine, capability thresholds. Further, we suggest that risk thresholds may be used to directly inform, though not determine, decisions. If and when our ability to produce risk estimates improves, we can rely more on risk thresholds (Section 4.3).

4.1 Arguments for using risk thresholds

First and foremost, risk thresholds are focused on potential harms to society and may thereby help align business conduct with societal concern. Risk thresholds directly pertain to externalities: the likelihood and severity of harm to individuals, groups, and society as a whole. In contrast to compute thresholds and capability thresholds, risk thresholds do not run into the issue of focusing on wrong proxies for risk, such as harmless models or capabilities. As a result, risk thresholds can help ensure companies only go ahead with risky activities if the risk is acceptable to society.

Second, risk thresholds can enable consistent allocation of safety resources. Risk thresholds can use the same units (e.g. expected number of fatalities or amount of economic damage in USD) for different risks (e.g. cyber and CBRN risks). As a result, risk thresholds can be set at the same level for different risks. If done well, this leads to a consistent allocation of safety resources. By contrast, capability thresholds (set without the help of risk thresholds) may inadvertently be set at different levels of risk for different model capabilities (e.g. autonomy and persuasion capabilities). This leads to an inconsistent allocation of safety resources. However, we note that this benefit can also be achieved by merely conducting risk estimates and consistently allocating safety resources based on their results, without also setting risk thresholds.

Third, in contrast to merely conducting risk estimates, risk thresholds can help ensure risk estimation results are actually acted upon. When using risk thresholds to directly feed into decisions, risk thresholds link the results of risk estimates to decisions (in the simplest version, go or no-go). When using risk thresholds to indirectly feed into decisions, risk estimation results are “enshrined” in capability thresholds, which in turn are integrated in decision rules (again, in the simplest version, go or no-go). In both ways, risk thresholds may help avoid situations where risk estimates are produced but not acted upon.

Fourth, risk thresholds may prevent companies from engaging in motivated reasoning regarding what level of risk is acceptable. Companies have strong incentives to argue that risk estimates are acceptable *in hindsight*. Risk thresholds can prevent this by determining criteria for what level of risk is acceptable *in advance*.¹² Yet, on the flipside, risk thresholds increase the incentive for companies to provide lower risk estimates in the first place. We discuss this concern below (Section 4.2).

Fifth, risk thresholds are future-proof and may help avoid locking in premature safety measures. Given that AI is still evolving (and rapidly so), regulators face the question of how prescriptive their requirements should be. Here, risk thresholds can provide a way out, as they do not require regulators to specify which safety measures companies need to implement. Instead, if regulators mandate risk thresholds to directly or indirectly feed into decisions (leaving it to companies to set capability thresholds), regulators put the burden on companies to find ways to reduce risk and can even continuously incentivize companies to innovate on safety measures, which may lower costs and result in more effective safety measures (Decker, 2018; Schuett, Anderljung, et al., forthcoming). On the other hand, if regulators mandate risk thresholds, they need a lot of effort and expertise to verify compliance (Decker, 2018; Schuett, Anderljung, et al., forthcoming). Regulators need to check whether the risk estimates that companies have produced are sound, which involves a case-by-case analysis of companies’ risk estimates. But regulators can require companies to provide them with detailed information about their risk estimates and the reasoning behind them, for example, through established tools like safety cases (Bishop & Bloomfield, 2000; Buhl et al., forthcoming; Kelly, 1998). Still, given the incentive for companies to provide low risk estimates (see previous paragraph), verifying compliance with risk thresholds may necessitate regulators to conduct their own risk estimates.

¹²Note that for risks where AI exacerbates a baseline risk, such as the current risk of cyberattacks on critical infrastructure, this baseline risk needs to be estimated before risk thresholds can be defined. However, even in these cases, risk thresholds are defined before the increase in risk caused by the AI development or deployment decision is estimated. This means that risk thresholds for any risk are defined before the increase in risk caused by AI is estimated.

4.2 Arguments against using risk thresholds

The key argument against using risk thresholds is that risk estimation is extremely hard for risks from frontier AI development and deployment. Using risk thresholds requires estimating the level of risk. In general, estimating risks from complex technological systems is hard (Apostolakis, 2004). This issue is aggravated in the case of frontier AI. There is little data from past incidents, meaning risk estimates mostly have to draw from modeling and expert judgment, which are less reliable. In general, risk estimation struggles with low-probability, high-impact events and “unknown unknowns”, which may be features of many risks from frontier AI. On top of that, understanding of how AI systems work and why they fail is poor, risk taxonomies and risk models are underdeveloped, and relevant information is split between companies and regulators – companies have knowledge of AI capabilities and usage, while regulators possess intelligence data, including about societal vulnerabilities and the number, capacities, and incentives of malicious actors. It might be possible to alleviate these issues, for instance, by improving risk estimation methodologies and gathering data about the occurrence of risk factors, near misses, and small-scale harm (Schuett, Baumoebl, et al., forthcoming). Nevertheless, the lack of reliable risk estimates currently is the main limitation of risk thresholds. The more strongly high-stakes decisions rely on risk thresholds, the more reliable these risk estimates should be.

Second, and relatedly, a common objection to using risk thresholds in frontier AI regulation is that foundation models, similar to electricity, are a dual-use or general-purpose technology that can be used in a tremendous number of ways and have a tremendous number of consequences that are both impossible to foresee and not the responsibility of frontier AI companies to prevent. This is a valid concern. However, this is a common issue in tort and criminal law, where mere causation is not enough (Wright, 1985). Likewise, in this context, this issue does not refute risk thresholds in general but means that regulators need to specify which effects are in scope (see also Section 5.1). Where to draw this line is a strategic decision that involves a variety of considerations (including economic, geopolitical, fairness, and safety considerations). Relevant qualitative criteria may be what type and amount of harm is at stake and whether intervention at later stages can be expected to be sufficiently effective (Anderljung & Hazell, 2023). Based on these criteria, imposing risk thresholds on frontier AI companies may be especially warranted for scenarios where single events cause large-scale harm and where no downstream developers are involved who could be held accountable instead.

Third, risk thresholds may create an incentive to produce artificially low risk estimates. If risk thresholds are used to directly feed into decisions, they establish a clear link between risk estimates and high-stakes decisions, making the implications of risk estimates immediately obvious. If risk thresholds are used to indirectly feed into decisions, the conditional risk estimates for different model capabilities have a less clear, but still perceivable impact on capability thresholds and thus decisions. Companies can take advantage of the uncertainty and subjectivity of the risk estimation process to produce the results they desire, with an added veneer of plausibility. To address this concern, regulators can verify companies’ risk estimates or mandate procedural requirements, such as that companies must involve more, diverse, and external assessors or that they break down risks into multiple events and ask assessors to estimate the risk of the individual events only. For example, instead of asking each assessor to estimate the increase in risk from biological attacks, companies could ask separate assessors to estimate the increases in risk regarding ideation, acquisition, magnification, formulation, and release of biological weapons (Patwardhan et al., 2024).

Fourth, it can be very difficult to decide what level of risk is acceptable (Aven, 2015). In particular, this decision involves making thorny normative judgments such as how much to value a human life (Reid, 2000; Vanem, 2012), how much to value future generations (Aven, 2012) or the environment (Vanem, 2012), and how cautious to be in the face of high uncertainty (Klinke & Renn, 2002). While making these decisions can be challenging for a single person, it will be even harder for different people or society as a whole to agree on the choice. Yet these decisions are currently being made implicitly through company development and deployment decisions. The fact that defining risk thresholds will be tough provides an argument for getting started sooner rather than later, such that important discussions and investigations can take place with sufficient time and rigor. We aim to help start this process by providing some guidance on how to define risk thresholds in Section 5.

4.3 Overall suggestions for using AI risk thresholds

Risk thresholds should be used to indirectly feed into high-stakes decisions, and may additionally be used to directly feed into such decisions. The respective benefits and limitations of risk thresholds

and capability thresholds mean that risk thresholds should complement, rather than replace, capability thresholds. Risk thresholds are directly focused on potential harms to society. However, they rely on risk estimates, which still have methodological limitations and involve substantial uncertainties, whereas capability thresholds rely on model evaluations, whose results are significantly less uncertain. Therefore, the key use case for risk thresholds should be to help set capability thresholds, ensuring that capability thresholds and corresponding safety measures, if followed, keep risk to an acceptable level. Additionally, using risk thresholds to directly feed into high-stakes decisions is helpful if capability thresholds miss the mark or become outdated. In conclusion, the two use cases of risk thresholds are not mutually exclusive but can make up for the limitations of the other. Hence, they should be applied in combination (Figure 1).

Nevertheless, as long as risk estimates are not reliable, risk thresholds in both of their use cases should not determine, but only inform, high-stakes decisions. The difficulty of producing reliable risk estimates is the strongest reason against using risk thresholds as the sole basis for a strict decision-rule for whether to go ahead or for where to set capability thresholds. It means that, in both cases, risk thresholds should currently only be used as one among a number of considerations. We provide some concrete examples for how risk thresholds can directly and indirectly inform, rather than determine, high-stakes decisions below. At the same time, to facilitate greater reliance on risk thresholds in the future, regulators and companies should invest in improving risk estimation methodology, gain experience in conducting risk estimates, investigate how much they can rely on them, and gather data about the occurrence of risk factors, near misses, and small-scale harm.

Concretely, when using risk thresholds to directly inform high-stakes decisions, they should be used among a number of other considerations, such as capability thresholds (which may or may not have been set with the help of risk thresholds). Moreover, many considerations unrelated to societal risk will come into play, including the company's appetite for business risks (e.g. liability risk or reputational risk) and various strategic considerations (e.g. whether a competitor is likely to release a similar model soon) (see ISO & IEC, 2019). Beyond that, could inform, rather than determine, decisions in that, if they are crossed, the board of directors or the competent regulator would have to be notified. The notified actor could potentially also be allowed to veto the decision or be required to give permission to proceed. Another option is that exceeding a risk threshold would trigger a requirement to conduct an extra suite of in-depth model evaluations, which may have to include third-party evaluators.

When using risk thresholds to indirectly inform decisions by helping set capability thresholds, other important considerations include expert judgment and safe design principles. Safe design involves, for instance, principles like redundancy, defense in depth, loose coupling of components to avoid cascading failures, separation of powers between decision-makers, and fail-safe design ensuring that systems fail gracefully (Dobbe, 2022; Leveson, 2016; Perrow, 1999; Reason, 1990). The company could also find that risk estimates suggest that one of its capability thresholds does not keep risk to an acceptable level, but nonetheless not change the capability threshold because the relevant model capabilities present substantial benefits or because they have good reasons to believe the risk estimates are unreliable. Moreover, the capability thresholds, set with the help of risk thresholds, may only inform, rather than determine, high-stakes decisions.

5 How to define AI risk thresholds

In this section, we propose a framework that consists of important considerations for defining AI risk thresholds. We did not find good general guidance for how to define risk thresholds. Thus, we conducted a non-systematic review of risk thresholds in various industries and jurisdictions, including aviation, nuclear, aerospace, maritime, and transportation and storage of hazardous materials, and tried to identify common ground on the most important considerations. Before regulators or companies can answer the question of what level of risk is acceptable, they need to decide which type of risk the threshold refers to (Section 5.1). Next, when determining the acceptable level of risk, they need to handle three difficult normative trade-offs: how to weigh potential harms and benefits, to what extent should mitigation costs be taken into account, and how to deal with uncertainty regarding all of the aforementioned (Section 5.2).

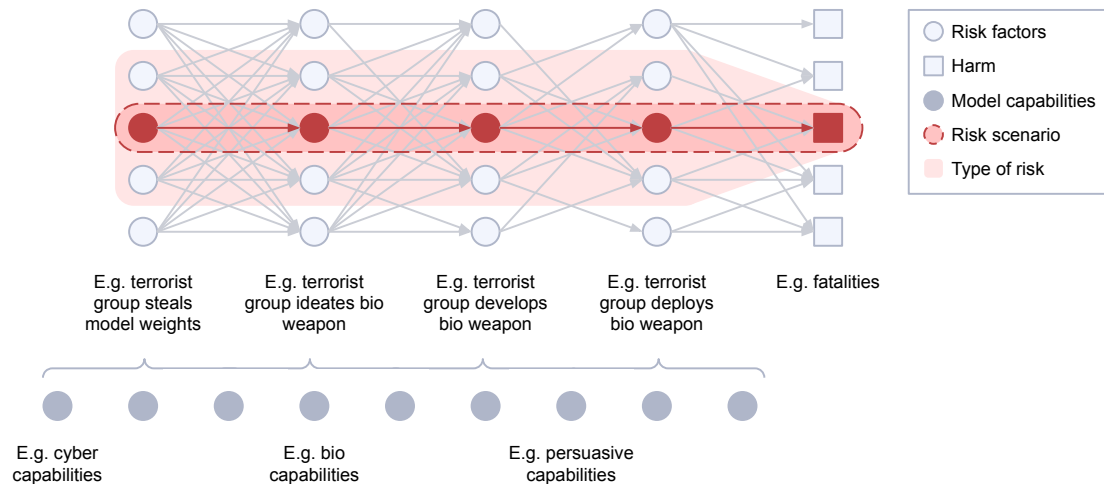


Figure 6: Representation of a linear risk model consisting of many risk scenarios

5.1 Type of risk

Every risk threshold is set for a specific “type of risk”. This term does not have a standard definition. In general, a type of risk seems to mean a group of risk scenarios that have similar impact, origin, or other characteristics, and may also be referred to as “area of risk” or “category of risk”. For example, common types of business risks include financial, legal, reputational, operational, and strategic risks. When it comes to risks from AI, types of risks could be distinguished by type of harm (e.g. fatalities, injuries, and economic damage)¹³ and potentially additionally by the domain or modality of occurrence of that harm (e.g. for fatalities, this could mean distinguishing between fatalities that stem from biological attacks, chemical attacks, and cyberattacks on critical infrastructure) (Figure 6).¹⁴

The choice of how many risk scenarios are in scope may affect where to set the risk threshold. All else equal, the fewer risk scenarios included, the lower, i.e. more strict, the risk thresholds should be because more fine-grained types of risks constitute a smaller fraction of the overall risk. For example, the U.S. aerospace industry used to have separate risk thresholds of 30×10^{-6} for the level of risk from each of the three main risk scenarios during rocket launch (explosive debris, toxic release, and blast overpressure). To simplify the licensing process, the industry switched to a risk threshold of 1×10^{-4} for the level of risk from all three risk scenarios combined (FAA, 2016). Note that the overall threshold remained about the same ($3 \times 30 \times 10^{-6} \approx 1 \times 10^{-4}$).

Many regulators also use separate risk thresholds for different numbers of fatalities. Two very common risk thresholds across many industries and jurisdictions are “individual risk thresholds” and “societal risk thresholds”. While definitions vary, individual risk thresholds usually refer to the risk of death of individuals, while societal risk thresholds refer to the risk of death of groups of people from a single event (e.g. HSE, 2001; IAEA, 2005; IMO, 2018). In short, regulators often use separate risk thresholds for risk scenarios where a single person dies and those where several people die.

The previous discussion can be considered to concern a risk threshold’s material scope – in addition, the temporal scope and territorial scope for harm to occur need to be defined. All else equal, the shorter the time period taken into account for harm to materialize, the lower, i.e. more strict, the risk

¹³Regulators and companies may want to begin with setting risk thresholds for types of harm that are relatively easy to measure (e.g. fatalities). However, types of harm that are harder to measure, such as discrimination, disinformation, or societal disruption, should not be neglected for this reason. But they require additional effort, because regulators and companies need to develop suitable metrics first. For a first effort in this regard, see Solaiman et al. (2023).

¹⁴Additionally distinguishing by domain or modality of occurrence of harm is not necessary but may allow setting more quantitative risk thresholds for types of risks where more data or better risk estimation methodologies exist. It may also make it easier to use risk thresholds to help with setting capability thresholds, which often focus on capabilities relevant for a particular domain (e.g. bio capabilities). For regulators, it can also make risk estimation easier, because information about different types of risks may be located within different government departments.

thresholds should be, because shorter time periods represent a smaller fraction of the overall risk. While longer time periods are more comprehensive, shorter time periods are easier to assess. For example, aviation has risk thresholds per flight-hour (ICAO, 2018), whereas the nuclear industry defines risk thresholds per reactor-year (IAEA, 2005). In the case of AI, a temporal scope of 12 months may align well with the yearly business cycle. However, developing biological weapons, for instance, may take several years and would not be in scope in this case. Similar considerations apply with regard to where the harm occurs. For example, in the U.S. nuclear industry, the individual risk threshold considers individuals within 1 mile of the power plant, whereas the societal risk threshold considers the population within 50 miles of the power plant (NRC, 1983). In the case of AI, the territorial scope may need to be unrestricted, because frontier AI companies provide their services globally, and harm may thus occur anywhere in the world. For instance, cyberattacks can target any system, especially if it is connected to the internet.

There also need to be rules for what type of causation is in scope; for example, second-order effects may be excluded. At the very least, the model needs to be causal for the harm. Causation can be established via the “but-for test” from law (Hart & Honoré, 1985): “but for the model, would the harm have occurred?” But mere causation may not be sufficient for practical reasons, because it would include a tremendous number of cases where the activity contributes marginally to the occurrence of harm (see also Section 4.2). Therefore, for example, a risk threshold could focus on first-order effects, that is harms directly stemming from the AI development, deployment, or use. Based on this example definition, harm to users or harm caused by malicious actors would be in scope, whereas harm to workers that are displaced by AI systems would not be in scope. However, we highlight that the lines can be blurry, and clear rules need to be established.

Finally, for risks where AI exacerbates a baseline risk (e.g. cyberattacks) as opposed to creating a new risk (e.g. rogue AI scenarios), it will usually be preferable for risk thresholds to refer to the increase in risk caused by AI, i.e. “marginal risk”, rather than the total level of risk (Kapoor et al., 2024). Note that the increase in risk should still be expressed in absolute, not relative, terms: a 5% increase in deaths from heart attacks is far worse than a 5% increase in deaths from shark attacks. However, for many risks from AI it is unclear what should be the relevant baseline risk – the level of risk with or without current AI systems, and whether the former includes AI systems by the company itself or only AI systems by its competitors. It is also unclear whether and, if so, how risk estimates for risk thresholds should take into account the increase in risk caused by expected AI systems from competitors. If they did, that could create a situation where each frontier AI company behaves recklessly in part because it reasons that its competitors will behave recklessly. What constitutes the relevant baseline risk needs to be clearly defined.

5.2 Level of risk

There seem to be three main ways used to determine the acceptable level of risk: building on peoples’ revealed preferences, copying what other industries do, and doing cost-benefit analysis (Philipson, 1983; Reid, 2000). Some regulators have reviewed the level of risk that people accept through engaging in common activities like driving (e.g. HSE, 1992). Other regulators have reviewed the level of risk that society accepts in other industries with comparable benefits (e.g. NRC, 1983 – comparing nuclear to coal, “the competing form of generating electricity”).

Most regulators appear to have reviewed and copied the risk thresholds already used in other industries or jurisdictions. As a potential result, many regulators use the same individual risk threshold of 1×10^{-6} per fatality and year (e.g. HSE, 2001; IAEA, 2005; IMO, 2018). However, societal risk thresholds seem to vary more strongly. Among industries and jurisdictions, the acceptable risk threshold for 1,000 fatalities ranges from a likelihood of 1×10^{-6} to 1×10^{-11} per year (Ehrhart et al., 2020). This study finds, for example, that the UK Health and Safety Commission sets risk thresholds for the transport of dangerous substances, deeming 1×10^{-4} unacceptable and 1×10^{-6} acceptable. By contrast, the survey finds that the Swiss Federal Office for the Environment sets risk thresholds for fixed installations and tunnels, deeming 1×10^{-9} unacceptable and 1×10^{-11} acceptable.

Few regulators appear to have conducted systematic cost-benefit analysis to determine the acceptable level of risk. A notable exception seems to be the maritime industry (EMSA, 2015; IMO, 2018). Choosing the acceptable level of risk in a systematic way is extremely difficult (HSE, 1992). However, given that AI may not be comparable to any other industry in terms of the benefits it might generate,

this may be the necessary approach. In the following, we provide some initial guidance on the three key normative trade-offs that need to be handled in a systematic cost-benefit analysis: how to weigh potential harms and benefits, to what extent to take into account mitigation costs, and how to deal with large amounts of uncertainty.

The key question when determining the acceptable level of risk for an activity is how to weigh the many potential harms against the benefits that may come from it (Hubbard, 2020).¹⁵ Greater benefits can be accounted for by setting higher, i.e. less strict, thresholds.¹⁶ But can the benefits of scientific advances be weighed against the harms of discrimination or disinformation? This is extremely challenging (see Section 4.2). For example, regulators in the maritime industry have developed a target societal risk/benefit ratio, the amount of societal benefit necessary to outweigh the risk of a single fatality. They derived the target societal risk/benefit ratio from aviation – because aviation has “good statistical data” and an “excellent safety record” – estimating the benefits via company revenues. They then apply this target societal risk/benefit ratio to the maritime industry (EMSA, 2015; IMO, 2018). A key issue with using this approach for AI is that many societal benefits, such as fundamental scientific advances, may not be reflected in company revenue.

The second key trade-off when determining the acceptable level of risk is to what extent to take into account the costs of reducing risks, in terms of money, time, or effort. Greater mitigation costs can be accounted for by setting higher, i.e. less strict, thresholds. Alternatively, as discussed in Section 3.1, a common approach in other industries is to set two risk thresholds: one above which risk is unacceptable regardless of mitigation costs, and one above which risk must be “as low as reasonably practicable” or “ALARP”; that is, risk must be reduced until the costs would be “grossly disproportionate” to the benefits of further risk reduction (HSE, 1992). This means mitigation costs do not influence the acceptable level of risk, but above some level of risk they influence what must be done if the threshold is crossed.

The third key trade-off when determining the acceptable level of risk is how to set the expected ratio of false negatives to false positives. Estimates of harms, benefits, and mitigation costs will involve large amounts of uncertainty. An approach that is more risk tolerant and therefore more concerned about benefits and mitigation costs, i.e. false positives (either due to the risk threshold accidentally being set too low or the level of risk wrongly being estimated to be above the threshold), leads to higher, i.e. less strict, risk thresholds. The more the risk threshold should be risk averse and reflect concern about harms, i.e. false negatives (either due to the risk threshold accidentally being set too high or the level of risk wrongly being considered below the threshold), the lower, i.e. more strict, the risk thresholds should be to generate a “margin of safety”. It seems prudent to have a margin of safety that is larger the more consequential and irreversible the type of harm at stake (e.g. this applies more to fatalities than to injuries). Some regulators also choose to be more risk averse the larger the harm at stake. For example, the Dutch nuclear industry sets its societal risk threshold at a probability of $1 \times 10^{-5}/N^2$ for $10 \times N$ fatalities per year (ANVS, 2020). The division by N^2 instead of N means a steeper slope of the risk thresholds and reflects aversion to large accidents (the acceptable probability decreases exponentially instead of linearly with the number of fatalities increasing).

Generally, regulators have more legitimacy and better incentives to define socially desirable thresholds than companies do, in particular if companies’ activities may cause externalities to society (Abrahamsen & Aven, 2012). The public safety and security risks that may stem from frontier AI systems are such externalities. Therefore, ideally regulators, but at least companies, should define risk thresholds.

6 Conclusion

This paper has made four main contributions. First, we have clarified the concepts of risk thresholds, capability thresholds, and compute thresholds, arguing that they not only rely on different metrics, but should also serve different functions. Second, we have made the case that risk thresholds are a promising tool for frontier AI regulation to the extent that the reliability of risk estimates can be

¹⁵One often reads that benefits should not be taken into account above the unacceptable risk threshold (e.g. HSE, 2001). But that guidance seems to refer to the moment when the threshold is *used*. When the threshold is *defined*, benefits should always be taken into account.

¹⁶A key decision that needs to be made is which benefits are in scope – this raises parallel questions to which harms are in scope, so we refer to that discussion (Section 5.1).

improved. Third, we have argued that risk thresholds should be used to indirectly inform high-stakes decisions by helping set, but not determine, capability thresholds, and may also be used to directly inform, but not determine, high-stakes decisions. Fourth, we have developed initial guidance for defining AI risk thresholds.

Many questions around risk thresholds for high-stakes AI development and deployment decisions warrant much further research. We highlight some questions that seem especially important. Fundamentally, advancing the risk estimation methodology is of utmost importance if regulators and companies want to rely more on risk thresholds for high-stakes decisions. In that regard, developing risk taxonomies and risk models, gaining experience with risk estimation methods, as well as gathering data about the occurrence of risk factors, near misses, and small-scale harm may be among the most useful ways forward (see Schuett, Baumoehl, et al., forthcoming). The details of how to use risk thresholds to determine adequate capability thresholds and corresponding safety measures also need to be explored further. In this regard, companies and regulators may be able to learn from the U.S. nuclear industry (see NRC, 2021). Last but not least, the acceptable levels of risk for different types of risks need to be defined. To do so, regulators could conduct comparative studies of thresholds in other industries or systematic cost-benefit analyses.

Managing the risks from frontier AI systems is an important and urgent challenge. Frontier AI companies need to improve their risk management practices and should use risk thresholds to help set capability thresholds. Over time, high-stakes development and deployment decisions should be directly informed by risk thresholds set by regulators. To this end, we need a discussion about what level of risk we, as a society, are willing to accept.

Abbreviations

AI	Artificial intelligence
ALARA	As low as reasonably achievable
ALARP	As low as reasonably practicable
ANVS	Dutch Authority for Nuclear Safety and Radiation Protection
BEIS	UK Department for Business, Energy and Industrial Strategy
CCPS	Center for Chemical Process Safety
COSO	Committee of Sponsoring Organizations of the Treadway Commission
DSIT	UK Department for Science, Innovation and Technology
EMSA	European Maritime Safety Agency
ESA	European Space Agency
EUROCONTROL	European Organisation for the Safety of Air Navigation
FAA	U.S. Federal Aviation Administration
HSE	UK Health and Safety Executive
IAEA	International Atomic Energy Agency
ICAO	International Civil Aviation Organization
IEC	International Electrotechnical Commission
IMO	International Maritime Organization
ISO	International Organization for Standardization
NASA	U.S. National Aeronautics and Space Administration
NIST	U.S. National Institute of Standards and Technology
NFPA	U.S. National Fire Protection Association
NRC	U.S. Nuclear Regulatory Commission
ONR	UK Office for Nuclear Regulation

Acknowledgments

We are grateful for valuable input from the following individuals, listed in alphabetical order by surname: Jide Alaga, Bill Anderson-Samways, Anthony Barrett, Caroline Baumoehl, Samuel Bowman, Ben Bucknall, Marie Buhl, Francisco Carvalho, Alan Chan, Noemi Dreksler, Ben Garfinkel, James

Binns, John Halstead, Jonathan Happel, Lennart Heim, Samuel Hilton, Holden Karnofsky, Patrick Levermore, Eli Lifland, David Lindner, Sebastien Krier, Yannick Muehlhaeuser, Malcolm Murray, Aidan O’Gara, Cullen O’Keefe, Alex Rand, Luca Righetti, Josh Rosenberg, Gaurav Sett, Rohin Shah, Merlin Stein, Christopher Phenicie, Hjalmar Wijk, Zoe Williams, and Peter Wills. All views and remaining errors are our own.

Declarations

One author’s spouse holds equity in a frontier AI company. Apart from that, the authors have no relevant financial or non-financial interests to disclose.

References

- Abrahamsen, E. B., & Aven, T. (2012). Why risk acceptance criteria need to be defined by the authorities and not the industry? *Reliability Engineering & System Safety*, *105*, 47–50. doi: 10.1016/j.ress.2011.11.004
- Anderljung, M., Barnhart, J., Korinek, A., Leung, J., O’Keefe, C., Whittlestone, J., ... Wolf, K. (2023). Frontier AI regulation: Managing emerging risks to public safety. *arXiv preprint arXiv:2307.03718*.
- Anderljung, M., & Hazell, J. (2023). Protecting society from AI misuse: When are restrictions on capabilities warranted? *arXiv preprint arXiv:2303.09377*.
- Anthropic. (2023). *Responsible Scaling Policy*. Retrieved from <https://www.anthropic.com/news/anthropics-responsible-scaling-policy>
- ANVS. (2020). *Guide on Level 3 PSA*. Retrieved from <https://english.autoriteitnvs.nl/documents/publication/2020/03/10/anvs-guide-on-level-3-psa>
- Apostolakis, G. E. (2004). How useful is quantitative risk assessment? *Risk Analysis*, *24*(3), 515–520. doi: 10.1111/j.0272-4332.2004.00455.x
- Aven, T. (2012). *Foundations of risk analysis*. Wiley. doi: 10.1002/9781119945482
- Aven, T. (2015). *Risk analysis*. Wiley. doi: 10.1002/9781119057819
- Aven, T. (2016). Risk assessment and risk management: Review of recent advances on their foundation. *European Journal of Operational Research*, *253*(1), 1–13. doi: 10.1016/j.ejor.2015.12.023
- Barrett, A. M., & Baum, S. D. (2017). A model of pathways to artificial superintelligence catastrophe for risk and decision analysis. *Journal of Experimental & Theoretical Artificial Intelligence*, *29*(2), 397–414. doi: 10.1080/0952813X.2016.1186228
- Bengio, Y., Hinton, G., Yao, A., Song, D., Abbeel, P., Darrell, T., ... Mindermann, S. (2024). Managing extreme AI risks amid rapid progress. *Science*. doi: 10.1126/science.adn0117
- Bernardi, J., Mukobi, G., Greaves, H., Heim, L., & Anderljung, M. (2024). Societal adaptation to advanced AI. *arXiv preprint arXiv:2405.10295*.
- Bishop, P., & Bloomfield, R. (2000). A methodology for safety case development. *Safety and Reliability*, *20*(1), 34–42. doi: 10.1080/09617353.2000.11690698
- Boiko, D. A., MacKnight, R., & Gomes, G. (2023). Emergent autonomous scientific research capabilities of large language models. *arXiv preprint arXiv:2304.05332*.
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., ... Liang, P. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Buhl, M., Sett, G., Koessler, L., & Schuett, J. (forthcoming). *Safety cases for frontier AI*.
- CCPS. (2009). *Guidelines for developing quantitative safety risk criteria*. Wiley. doi: 10.1002/9780470552940
- Chan, A., Salganik, R., Markelius, A., Pang, C., Rajkumar, N., Krasheninnikov, D., ... Maharaj, T. (2023). Harms from increasingly agentic algorithmic systems. In *ACM Conference on Fairness, Accountability, and Transparency* (pp. 651–666). doi: 10.1145/3593013.3594033
- Clymer, J., Gabrieli, N., Krueger, D., & Larsen, T. (2024). Safety cases: How to justify the safety of advanced AI systems. *arXiv preprint arXiv:2403.10462*.
- Cohen, M. K., Kolt, N., Bengio, Y., Hadfield, G. K., & Russell, S. (2024). Regulating advanced artificial agents. *Science*, *384*(6691), 36–38. doi: 10.1126/science.adl0625
- COSO. (2017). *Enterprise risk management: Integrating with strategy and performance*. Retrieved from <https://www.coso.org/guidance-erm>

- Decker, C. (2018). Goals-based and rules-based approaches to regulation. *BEIS*. Retrieved from <https://www.gov.uk/government/publications/regulation-goals-based-and-rules-based-approaches>
- Dezfuli, H., Benjamin, A., Everett, C., Feather, M., Rutledge, P., Sen, D., & Youngblood, R. (2015). NASA system safety handbook. Volume 2: System safety concepts, guidelines, and implementation examples. NASA. Retrieved from <https://ntrs.nasa.gov/citations/20150015500>
- Dobbe, R. I. J. (2022). System safety and artificial intelligence. In J. B. Bullock et al. (Eds.), *The Oxford handbook of AI governance* (pp. 441–458). Oxford University Press. doi: 10.1093/oxfordhb/9780197579329.013.67
- DSIT. (2023a). *AI Safety Summit: Introduction*. Retrieved from <https://www.gov.uk/government/publications/ai-safety-summit-introduction>
- DSIT. (2023b). *Emerging processes for frontier AI safety*. Retrieved from <https://www.gov.uk/government/publications/emerging-processes-for-frontier-ai-safety>
- DSIT. (2024a). *AI Safety Institute approach to evaluations*. Retrieved from <https://www.gov.uk/government/publications/ai-safety-institute-approach-to-evaluations>
- DSIT. (2024b). *Frontier AI Safety Commitments, AI Seoul Summit 2024*. Retrieved from <https://www.gov.uk/government/publications/frontier-ai-safety-commitments-ai-seoul-summit-2024>
- DSIT. (2024c). *Seoul Ministerial Statement for advancing AI safety, innovation and inclusivity: AI Seoul Summit 2024*. Retrieved from <https://www.gov.uk/government/publications/seoul-ministerial-statement-for-advancing-ai-safety-innovation-and-inclusivity-ai-seoul-summit-2024>
- Ehrhart, B. D., Brooks, D. M., Muna, A. B., & LaFleur, C. B. (2020). Evaluation of risk acceptance criteria for transporting hazardous materials. *Sandia National Laboratories*. doi: 10.2172/1602640
- EMSA. (2015). *Risk acceptance criteria and risk based damage stability, final report, part 1: Risk acceptance criteria*. Retrieved from <https://emsa.europa.eu/csn-menu/csn-background/items.html?cid=14&id=2419>
- ESA. (2023). *ESA space debris mitigation requirements*. Retrieved from <https://technology.esa.int/upload/media/ESA-Space-Debris-Mitigation-Requirements-ESSB-ST-U-007-Issue1.pdf>
- EUROCONTROL. (2001). *ESARR 4: Risk assessment and mitigation in ATM*. Retrieved from <https://www.eurocontrol.int/publication/esarr-4-risk-assessment-and-mitigation-atm>
- FAA. (1988). *Advisory Circular 25.1309-1A: System design and analysis*. Retrieved from https://www.faa.gov/regulations_policies/advisory_circulars/index.cfm/go/document.information/documentid/22680
- FAA. (2016). *Changing the collective risk limits for launches and reentries and clarifying the risk limit used to establish hazard areas for ships and aircraft*. Retrieved from <https://www.federalregister.gov/d/2016-17083>
- Fang, R., Bindu, R., Gupta, A., & Kang, D. (2024). LLM agents can autonomously exploit one-day vulnerabilities. *arXiv preprint arXiv:2404.08144*.
- Fischhoff, B., Lichtenstein, S., Slovic, P., Derby, S. L., & Keeney, R. (1984). *Acceptable risk*. Cambridge University Press.
- Flamberg, S., Rose, S., Kurth, B., & Sallaberry, C. (2016). Paper study on risk tolerance. *Kiefner*. Retrieved from <https://trid.trb.org/View/1469903>
- Fraser, H., & Bello y Villarino, J.-M. (2023). Acceptable risks in Europe’s proposed AI Act: Reasonableness and other principles for deciding how much risk management is enough. *European Journal of Risk Regulation*, 1–16. doi: 10.1017/err.2023.57
- Frey, H. C., & Patil, S. R. (2002). Identification and review of sensitivity analysis methods. *Risk Analysis*, 22(3), 553–578. doi: 10.1111/0272-4332.00039
- Google AI. (2018). *Our principles*. Retrieved from <https://ai.google/responsibility/principles>
- Google DeepMind. (2024). *Introducing the Frontier Safety Framework*. Retrieved from <https://deepmind.google/discover/blog/introducing-the-frontier-safety-framework>
- Hart, H. L. A. (2017). Between utility and rights. In L. Ten Chin (Ed.), *Theories of rights* (pp. 107–125). Routledge. doi: 10.4324/9781315236308-7

- Hart, H. L. A., & Honoré, T. (1985). *Causation in the law*. Clarendon Press. doi: 10.1093/acprof:oso/9780198254744.001.0001
- Hazell, J. (2023). Spear phishing with large language models. *arXiv preprint arXiv:2305.06972*.
- Heim, L., & Koessler, L. (forthcoming). *Training compute thresholds: Features and functions in AI governance*.
- Helfrich, G. (2024). The harms of terminology: Why we should reject so-called “frontier AI”. *AI and Ethics*. doi: 10.1007/s43681-024-00438-1
- Hendrycks, D., Mazeika, M., & Woodside, T. (2023). An overview of catastrophic AI risks. *arXiv preprint arXiv:2306.12001*.
- Hernandez, D., Kaplan, J., Henighan, T., & McCandlish, S. (2021). Scaling laws for transfer. *arXiv preprint arXiv:2102.01293*.
- Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, E., . . . Sifre, L. (2022). Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.
- HSE. (1992). *The tolerability of risk from nuclear power stations*. Retrieved from <https://www.onr.org.uk/media/v1vi3v21/tolerability.pdf>
- HSE. (2001). *Reducing risks, protecting people*. Retrieved from <https://www.hse.gov.uk/enforce/assets/docs/r2p2.pdf>
- Hubbard, D. W. (2020). *The failure of risk management: Why it's broken and how to fix it* (2nd ed.). Wiley.
- IAEA. (2005). *Risk informed regulation of nuclear facilities: Overview of the current status*. Retrieved from <https://www.iaea.org/publications/7175/risk-informed-regulation-of-nuclear-facilities-overview-of-the-current-status>
- ICAO. (2018). *Doc 9859: Safety management manual* (4th ed.). Retrieved from <https://skybrary.aero/sites/default/files/bookshelf/5863.pdf>
- IMO. (2018). *Revised guidelines for formal safety assessment (FSA) for use in the IMO rule-making process*. Retrieved from <https://www.imo.org/en/OurWork/Safety/Pages/FormalSafetyAssessment.aspx>
- ISO. (2018). *ISO 31000:2018 Risk management*. Retrieved from <https://www.iso.org/iso-31000-risk-management.html>
- ISO & IEC. (2014). *ISO/IEC Guide 51:2014 Safety aspects — Guidelines for their inclusion in standards*. Retrieved from <https://www.iso.org/standard/53940.html>
- ISO & IEC. (2019). *ISO/IEC 31010:2019 Risk management — Risk assessment techniques*. Retrieved from <https://www.iso.org/standard/72140.html>
- Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., . . . Amodei, D. (2020). Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Kapoor, S., Bommasani, R., Klyman, K., Longpre, S., Ramaswami, A., Cihon, P., . . . Narayanan, A. (2024). On the societal impact of open foundation models. *arXiv preprint arXiv:2403.07918*.
- Kelly, T. P. (1998). *Arguing safety: A systematic approach to managing safety cases* (Unpublished doctoral dissertation). University of York.
- Klinke, A., & Renn, O. (2002). A new approach to risk evaluation and management: Risk-based, precaution-based, and discourse-based strategies. *Risk Analysis*, 22(6), 1071–1094. doi: 10.1111/1539-6924.00274
- Koessler, L., & Schuett, J. (2023). Risk assessment at AGI companies: A review of popular risk assessment techniques from other safety-critical industries. *arXiv preprint arXiv:2307.08823*.
- Laux, J., Wachter, S., & Mittelstadt, B. (2024). Trustworthy artificial intelligence and the European Union AI Act: On the conflation of trustworthiness and acceptability of risk. *Regulation & Governance*, 18(1), 3–32. doi: 10.1111/rego.12512
- Leveson, N. G. (2016). *Engineering a safer world: Systems thinking applied to safety*. MIT Press. doi: 10.7551/mitpress/8179.001.0001
- Linkov, I., Bates, M., Loney, D., Sparrevik, M., & Bridges, T. (2011). Risk management practices. In I. Linkov & T. Bridges (Eds.), *Climate: Global change and local adaptation* (pp. 133–155). Springer. doi: 10.1007/978-94-007-1770-1_8
- Lohn, A., & Jackson, K. (2022). Will AI make cyber swords or shields? *Center for Security and Emerging Technology*. doi: 10.51593/2022ca002
- Lohn, A., & Musser, M. (2022). AI and compute: How much longer can computing power drive artificial intelligence progress? *Center for Security and Emerging Technology*. doi: 10.51593/2021CA009
- Marhavilas, P. K., & Koulouriotis, D. E. (2021). Risk-acceptance criteria in occupational health and safety risk-assessment: The state-of-the-art through a systematic literature review. *Safety*, 7(4).

- doi: 10.3390/safety7040077
- Melchers, R. E. (2001). On the ALARP approach to risk management. *Reliability Engineering & System Safety*, 71(2), 201–208. doi: 10.1016/S0951-8320(00)00096-X
- Meta. (2023). *Meta's first responsible business practices report*. Retrieved from <https://about.fb.com/news/2023/07/metas-first-responsible-business-practices-report>
- Microsoft. (2024). *Responsible AI transparency report*. Retrieved from <https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RW115B0>
- Mirsky, Y., Demontis, A., Kotak, J., Shankar, R., Gelei, D., Yang, L., ... Biggio, B. (2021). The threat of offensive AI to organizations. *arXiv preprint arXiv:2106.15764*.
- Morgan, M. G., & Henrion, M. (1990). *Uncertainty: A guide to dealing with uncertainty in quantitative risk and policy analysis*. Cambridge University Press. doi: 10.1017/CBO9780511840609
- Mouton, C. A., Lucas, C., & Guest, E. (2023). The operational risks of AI in large-scale biological attacks: A red-team approach. *RAND*. doi: 10.7249/RRA2977-1
- NFPA. (2023). *NFPA 59A: Standard for the production, storage, and handling of liquefied natural gas (LNG)*. Retrieved from <https://www.nfpa.org/codes-and-standards/nfpa-59a-standard-development/59a>
- Ngo, R., Chan, L., & Mindermann, S. (2024). The alignment problem from a deep learning perspective. *arXiv preprint arXiv:2209.00626*.
- NIST. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. doi: 10.6028/NIST.AI.100-1
- Novelli, C., Casolari, F., Rotolo, A., Taddeo, M., & Floridi, L. (2024). AI risk assessment: A scenario-based, proportional methodology for the AI Act. *Digital Society*, 3(1), 1–13. doi: 10.1007/s44206-024-00095-1
- NRC. (1983). *Safety goals for nuclear power plant operation*. Retrieved from <https://www.nrc.gov/docs/ML0717/ML071770230.pdf>
- NRC. (2016). *Regulatory Guide 8.10: Operating philosophy for maintaining occupational and public radiation exposures as low as is reasonably achievable*. Retrieved from <https://www.nrc.gov/docs/ML1610/ML16105A136.pdf>
- NRC. (2021). *History of the NRC's risk-informed regulatory programs*. Retrieved from <https://www.nrc.gov/about-nrc/regulatory/risk-informed/history.html>
- Oakley, P. A., & Harrison, D. E. (2020). Death of the ALARA radiation protection principle as used in the medical sector. *Dose Response*, 18(2). doi: 10.1177/1559325820921641
- ONR. (2020). *Safety assessment principles (SAPs)*. Retrieved from <https://www.onr.org.uk/publications/regulatory-guidance/regulatory-assessment-and-permissioning/safety-assessment-principles-saps>
- OpenAI. (2023). *Preparedness Framework (Beta)*. Retrieved from <https://cdn.openai.com/openai-preparedness-framework-beta.pdf>
- Patwardhan, T., Liu, K., Markov, T., Chowdhury, N., Leet, D., Cone, N., ... Madry, A. (2024). Building an early warning system for LLM-aided biological threat creation. *OpenAI*. Retrieved from <https://openai.com/research/building-an-early-warning-system-for-llm-aided-biological-threat-creation>
- Perrow, C. (1999). *Normal accidents: Living with high risk technologies*. Princeton University Press.
- Philipson, L. L. (1983). Risk acceptance criteria and their development. *Journal of Medical Systems*, 7(5), 437–456. doi: 10.1007/BF00995743
- Phuong, M., Aitchison, M., Catt, E., Cogan, S., Kaskasoli, A., Krakovna, V., ... Shevlane, T. (2024). Evaluating frontier models for dangerous capabilities. *arxiv preprint arXiv:2403.13793*.
- Pistillo, M., Van Arsdale, S., Heim, L., & Winter, C. (forthcoming). The role of compute thresholds for AI governance. *George Washington Journal of Law & Technology*.
- Popov, G., Lyon, B. K., & Hollcroft, B. (2021). *Risk assessment: A practical guide to assessing operational risks*. Wiley. doi: 10.1002/9781119798323
- Rausand, M., & Haugen, S. (2020). *Risk assessment: Theory, methods, and applications*. Wiley. doi: 10.1002/9781119377351
- Reason, J. (1990). *Human error*. Cambridge University Press. doi: 10.1017/CBO9781139062367
- Reid, S. G. (2000). Acceptable risk criteria. *Progress in Structural Engineering and Materials*, 2(2), 254–262. doi: 10.1002/1528-2716(200004/06)2:2<254::AID-PSE30>3.0.CO;2-K
- Salter, C., Saydjari, O. S., Schneier, B., & Wallner, J. (1998). Toward a secure system engineering methodology. In *Workshop on New Security Paradigms* (pp. 2–10). doi: 10.1145/310889.310900

- Sandbrink, J. B. (2023). Artificial intelligence and biological misuse: Differentiating risks of language models and biological design tools. *arXiv preprint arXiv:2306.13952*.
- Sastry, G., Heim, L., Belfield, H., Anderljung, M., Brundage, M., Hazell, J., ... Coyle, D. (2024). Computing power and the governance of artificial intelligence. *arXiv preprint arXiv:2402.08797*.
- Schneier, B. (2011). *Secrets and lies: Digital security in a networked world*. Wiley. doi: 10.1002/9781119183631
- Schuett, J. (2023). Risk management in the Artificial Intelligence Act. *European Journal of Risk Regulation*, 1–19. doi: 10.1017/err.2023.1
- Schuett, J., Anderljung, M., Koessler, L., Carlier, A., & Garfinkel, B. (forthcoming). From principles to rules: A regulatory approach for frontier AI. In P. Hacker, A. Engel, S. Hammer, & B. Mittelstadt (Eds.), *Oxford handbook on the foundations and regulation of generative AI*. Oxford University Press.
- Schuett, J., Baumoechl, C., Murray, M., & Koessler, L. (forthcoming). *How to estimate the impact and likelihood of risks from AI*.
- Shevlane, T., & Dafoe, A. (2020). The offense-defense balance of scientific knowledge: Does publishing AI research reduce misuse? In *AAAI/ACM Conference on AI, Ethics, and Society* (pp. 173–179). doi: 10.1145/3375627.3375815
- Shevlane, T., Farquhar, S., Garfinkel, B., Phuong, M., Whittlestone, J., Leung, J., ... Dafoe, A. (2023). Model evaluation for extreme risks. *arXiv preprint arXiv:2305.15324*.
- Shostack, A. (2014). *Threat modeling: Designing for security*. Wiley.
- Soice, E. H., Rocha, R., Cordova, K., Specter, M., & Esvelt, K. M. (2023). Can large language models democratize access to dual-use biotechnology? *arXiv preprint arXiv:2306.03809*.
- Solaiman, I., Talat, Z., Agnew, W., Ahmad, L., Baker, D., Blodgett, S. L., ... Vassilev, A. (2023). Evaluating the social impact of generative AI systems in systems and society. *arXiv preprint arXiv:2306.05949*.
- Starr, C. (1969). Social benefit versus technological risk: What is our society willing to pay for safety? *Science*, 165(3899), 1232–1238. doi: 10.1126/science.165.3899.1232
- Sutton, R. (2019). *The bitter lesson*. Retrieved from <http://www.incompleteideas.net/IncIdeas/BitterLesson.html>
- Urbina, F., Lentzos, F., Invernizzi, C., & Ekins, S. (2022). Dual use of artificial intelligence-powered drug discovery. *Nature Machine Intelligence*, 4(3), 189–191. doi: 10.1038/s42256-022-00465-9
- Vanem, E. (2012). Ethics and fundamental principles of risk acceptance criteria. *Safety Science*, 50(4), 958–967. doi: 10.1016/j.ssci.2011.12.030
- Villalobos, P., Sevilla, J., Heim, L., Besiroglu, T., Hobbhahn, M., & Ho, A. (2022). Will we run out of data? Limits of LLM scaling based on human-generated data. *arXiv preprint arXiv:2211.04325*.
- Wright, R. W. (1985). Causation in tort law. *California Law Review*, 73(6), 1735–1828.