

An Ontological Analysis of Risk in Basic Formal Ontology

Federico Donato^{1,*}, Adrien Barton²

¹*Department of Philosophy, University at Buffalo, Buffalo NY, USA*

²*IRIT, CNRS, Université de Toulouse, Toulouse, France*

Abstract

The paper explores the nature of risk, providing a characterization using the categories of the Basic Formal Ontology (BFO). It argues that the category Risk is a subclass of BFO:Role, contrasting it with a similar view classifying Risk as a subclass of BFO:Disposition. This modeling choice is applied on one example of risk, which represents objects, processes (both physical and mental) and their interrelations, then generalizing from the instances in the example to obtain an overall analysis of risk, making explicit what are the sufficient conditions for being a risk. Plausible necessary conditions are also mentioned for future work.

Keywords

Ontology, Risk, BFO, Role, Disposition

1. Risk Databases and Ontologies

1.1. The need for a risk ontology

Risks are ubiquitous. Your likelihood of receiving a loan and the interest rates associated with the loan are directly correlated to the risk the bank sees in losing its money if you fail to repay it. Companies in the stock market are evaluated by brokers and investors based on their risk profile. Policymakers ponder whether a hydrogeological plant will put the environment at risk and how this will negatively affect their net political campaign. Physicians recommend cost-benefit analyses to their patients for surgical operations that could potentially extend their lives, but also pose risks. Risks can be severe or minor and present in the population according to various statistical distributions. Risky situations, such as in a terrorist emergency, can arguably exist even if the terrorist attack is just suspected but never occurs.

These plausible natural language descriptions of events related to risk can be found in terminologies, contracts, and many other documents. Moreover, in today's information-driven society, information about risk is stored and analyzed using databases and risk analysis equations. The databases focusing on this specific domain are called risk registers, and they can track all kinds of risks, connecting them to people, processes, economical assets, medical status, locations and times [1]. However, the conceptual foundations of risk lack clarity and there are ongoing discussions in both the natural language and database communities regarding risk [2]. Despite the prevalence of risk discussions, there is no universally accepted definition among all scholars or institutions, which hinders effective action to prevent or mitigate risk. The confusion is exacerbated by the numerous complex vocabulary terms and database fields that seem to be linked to or depend on risk, such as "risk propagation", "risk prevention", "risk aggregation", "vulnerability" and "cybersecurity", among many others [3] [4]. In fact, precise knowledge of what is at risk, what causes risk, and how many risks there are is essential for an intervention that is timely and proportional to the prospected damage.

Questions that still pertain to the contemporary theory of risk include: What is the nature of risks (material entities, beliefs, special kinds of probabilities)? Is the risk of losing the reputation of my company identical to the opportunity of a rival company to gain reputation? Does a risk that never

Proceedings of the Twelfth Semantic Technology for Intelligence, Defense, and Security (STIDS) Conference, October 22-23, 2024, Woodbridge VA, USA

*Corresponding author.

✉ fdonato@buffalo.edu (F. Donato); adrien.barton@irit.fr (A. Barton)

🆔 0009-0001-6600-240X (F. Donato); 0000-0001-5500-6539 (A. Barton)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

realizes actually exist? Is risk objective or subjective, absolute or relative? As a preliminary to address these questions, we will propose an explicit and logically-based definition of risk based on top-level ontology categories.

2. Definitions of risk and their elements

Ontologies are information artifacts that specify meanings of terms and capture our knowledge of the world, categorize the relevant entities of a domain in hierarchies, leverage logical deduction, and enable interoperability between diverse databases [5]. We align our framework with the Basic Formal Ontology [6], a widely used top-level ontology that provides common ground for building domain ontologies and establishing a common semantics across many domains, from biology to jurisprudence [7]. To our knowledge, only two other risk ontologies have been aligned with top-level ones, indicating that the research field is still relatively underdeveloped [8] [9]. This paper will argue step by step for the following sufficient condition on being a risk:

If: (i) r is a role &
(ii) r is externally grounded in an aversion that an agent or group of agents have towards the possible realizations of the role,
then r is a risk.

In order to understand risk, we need to define the terms that are involved in the condition, namely something that is risky or the source of risk, something that is put at stake or at risk, something that is valuable or desirable, and some process that might involve uncertainty, as risk is commonly understood [10]. The implication of values, interests or desires in the definition of risk might raise questions on the subjective or objective nature of risks [11] [12]. We will consider the interplay between subjectivity and objectivity, especially with the relation of external grounding.

To develop a comprehensive understanding of risk, we aim to create an Aristotelian definition, that categorizes the entity within a genus, namely a superclass, and using a differentia, which explains how it differs from other entities within the same genus, and clarifies its ontological relations with other entities [13]. Structuring a definition in an Aristotelian way helps translating it in a logical form, enabling automatic reasoning.

This short paper proposes an exploration in a new direction that the ontology literature has not considered so far. The paper is structured as follows: the third section contrasts two conceptions of risks as BFO:*Disposition* and BFO:*Role*; the fourth section considers the subjective side of risk, involving a mental attitude of an agent and the target of this mental attitude; the conclusion provides a way of categorizing risk with sufficient conditions, also listing further works to expand the ontology and limitations of the present ontological model.

3. Risks as Realizable Entities

Having introduced a general picture of what risk plausibly is and some open questions about it, we now turn to a unified definition of risk. In this section, we propose a view of risk as a realizable and externally grounded property. In this formulation, externality is understood as involving two or more mereologically distinct entities [14], while grounding is a relation between entities in which one exists because of the other [6, p. 100] [15]. We will examine two views that meet those requirements: risk as a disposition and risk as a role, ultimately favoring the latter. Both views on risk are developed in the BFO framework, which is briefly explained in the following.

BFO is a top-level ontology [6] that identifies two very general classes of entities, continuants and occurrents. Continuants can endure in time, existing at different times and changing through them, whereas occurrents are temporally extended and have temporal parts, but cannot change in time. Examples of continuants include an orange and its mass of 200g, both of which exist for a certain period

of time. In contrast, occurrents include the ripening process of the orange and the time span during which it occurs. Continuants encompass various entities, including material entities and properties, named “specifically dependent continuants” (SDC) in BFO. SDCs include qualities, such as the orange’s mass, as well as latent properties that can be realized in a process. Those latent properties, called “realizable entities”, include dispositions, that depend solely on the physical make-up of the entity. For instance, an orange has the disposition to dry out under the hot sun, even when it is fully hydrated and no drying process is occurring. This disposition exists independently of the orange’s current state. They also include properties that do not depend only on the physical make-up of the entity, like the role of this orange of being the winner in a fruit competition, after being chosen by a jury. More specifically, BFO defines dispositions and roles as follows [15]:

d is a disposition $=_{def}$ (i) d is a realizable entity & (ii) d ’s bearer is some material entity & (iii) d is such that, if it ceases to exist, then its bearer is physically changed & (iv) d ’s realization occurs because this bearer is in some special physical circumstances & (v) this realization occurs in virtue of the bearer’s physical make-up.

r is a role $=_{def}$ (i) r is a realizable entity & (ii) r exists because there is some single bearer that is in some special physical, social, or institutional set of circumstances in which this bearer does not have to be & (iii) r is not such that, if it ceases to exist, then the physical make-up of the bearer is thereby changed.

To illustrate the analysis of risks within the BFO framework, consider the following scenario: Paul is walking under a ceiling (named $ceiling_0$) that is unstable and might fall over him. Paul is a BFO: *Object* and $ceiling_0$ a BFO: *Material entity*, which bears a disposition $instability_0$ that might be realized in a process of collapsing (in which case it will be called $collapsing_0$), which has as participant $ceiling_0$. If it is undesirable for Paul that $ceiling_0$ would collapse on him, then a corresponding risk arises when he is in the room. This scenario raises a fundamental question: to which ontological category do individual risks belong? We will examine two competing views: the first [16] posits that risks are dispositions, while this paper will propose an alternative view according to which they are roles. We will present both views in order to compare their strengths and weaknesses.

3.1. First Alternative: Risks as Dispositions

Grenier and Barton [16] argue that risks are dispositions whose realizations are undesirable for certain agents. For every agent x_i , they introduce the class of dispositions whose possible realizations are undesirable for x_i , named here $Risk^d\text{-for-}x_i$. Then, they introduce a class (written here $Risk^d$) which is the union class of the classes $Risk^d\text{-for-}x_i$ for all i (thus, $Risk^d\text{-for-}x_i$ is a subclass of $Risk^d$ which is a subclass of BFO: *Disposition*; note that since $Risk^d\text{-for-}x_i$ explicitly mentions an individual x_i , it is arguably not a *bona fide* universal but a defined class or collection of particulars, as introduced by Smith and Ceusters [17]).

According to this view, a disposition d_0 can be an instance of $Risk^d\text{-for-}x_i$ at one time but not at another, depending on the agent x_i ’s attitude towards the possible realizations of d_0 at those times: more precisely, depending on whether x_i finds d_0 ’s possible realization undesirable or not. In the above mentioned example, when Paul leaves the room and becomes indifferent to the possible collapse of $ceiling_0$, $instability_0$ is no longer an instance of $Risk^d\text{-for-Paul}$, although it continues to exist and is still a disposition. This conception attempts to capture the dual nature of risk: the dispositional nature of an instance of $Risk^d$ captures the objective side of risk, whereas the fact that this disposition instantiates $Risk^d\text{-for-Paul}$ is based on Paul’s desires and reflects the subjective nature of risk, thus resolving one of the questions we posed in the introduction.

Let’s revisit two features of risks that must be taken into account: realizability and external grounding. In this view, risk is a disposition, and thus a realizable entity. A disposition is a risk (or not) based on external entities, namely the desires of some agent; thus, although its existence does not depend on such external entities ($instability_0$ exists independently of the mental attitudes of Paul), its risk

status (the fact that it instantiates $Risk^d$) does. To summarize, in this view, a disposition that is a risk is internally grounded in the sense that it exists independently of the external world (see [18] for more considerations on grounding); but its risk status (that is, the fact that it instantiates the class $Risk^d$) is externally grounded, as it is contingent upon external factors. To illustrate this distinction, consider a similar example. The rectangular shape of this table (a quality) is not externally grounded, but it could instantiate the class *Shape-seen-by-Paul* (which is, again, a defined class and not a *bona fide* universal) based on external features (namely Paul seeing this shape). Therefore, in this view, a disposition is not a risk essentially but accidentally: a disposition can become a risk, and a risk can lose its risk status and become a mere disposition, based on the desires of agents. Note also that in the dispositional approach, $collapsing_0$ is an instance of a more specific subclass than *BFO:Process*, in this case an instance of the defined class *Process-undesirable-for-Paul*. This instantiation is accidental, similarly to the accidental character of the instantiation of $Risk^d$ -for-Paul by $instability_0$. Fig. 1 summarizes this approach:

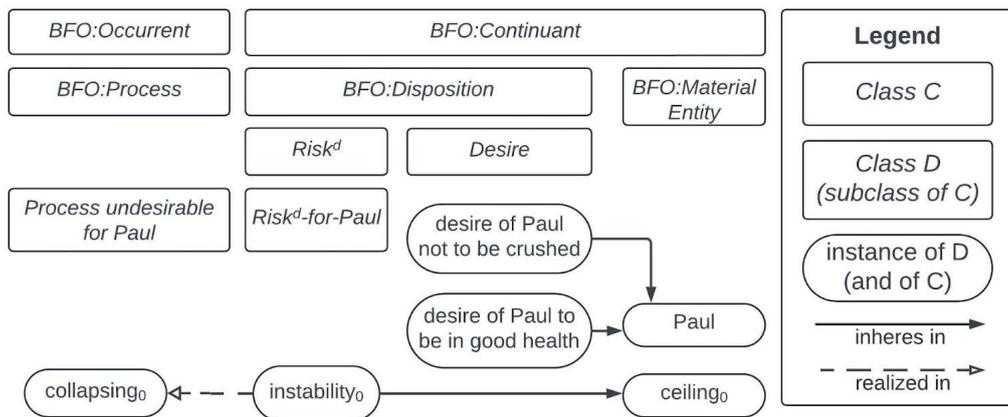


Figure 1: $Risk^d$ as a BFO:Disposition

3.2. Second Alternative: Risks as roles

Let us now contrast this dispositional approach with a new approach according to which risks are another kind of realizable entity, namely roles. To contrast the two views, a risk according to the dispositional approach will be called a “dispositional risk”, and a risk according to the role approach will be called a “risk role”.

To put it informally, roles are SDCs that existentially depend on the context in which they are present. Instead of having a disposition that starts and stops instantiating $Risk^d$ -for-Paul at different times, we can think of an instance of a class named $Risk^r$, itself a subclass of *BFO:Role*, which remains a risk throughout its existence. As in the dispositional approach, the role approach considers a risk as a *BFO:Realizable Entity*, since it can be activated and realized. It also agrees with the dispositional approach in considering the source of risk as $ceiling_0$, a material entity bearing properties relevant for the risky situation. In the ceiling example, we will say that $ceiling_0$ bears a risk role $hazard_0$ when Paul is located under it and wants to preserve his health.

As in the first alternative, what makes $ceiling_0$ risky for Paul is external to $ceiling_0$. Additionally, $ceiling_0$ does not bear the risk role $hazard_0$ by necessity: it does so because of the context, which consists of Paul being in the vicinity of the ceiling, desiring not to be hurt, and of $ceiling_0$ being unstable. Expanding the scenario, if Vladimir, an enemy of Paul, wanted him crushed, then the instability of $ceiling_0$ would not be a risk for Vladimir, but an opportunity for him. So, $ceiling_0$ does not bear a risk on its own, independently of the rest of the world, but due to the mental attitudes of some agents. Similarly, in the dispositional approach, $instability_0$ is an instance of $Risk^d$ based on external features,

namely in virtue of the desires of Paul.

The dispositional approach and the role approach have similarities and differences in how they deal with this externally-grounded character of risk. In both approaches, as we saw, a material entity bears a dispositional risk or a risk role on the basis of external features. According to the dispositional approach, as explained earlier, a disposition can become or stop being a risk - similarly to a human who can continue to exist as he grows up even if he is not a child anymore. For example, $instability_0$ continues to exist even if it is not a risk anymore for anyone. According to the role approach, on the other hand, the risk cannot stop being a risk: if a situation is not risky anymore, the corresponding risk role disappears¹. For example, if Paul's desires change and he doesn't care anymore about being hurt, then the role $hazard_0$ disappears: it cannot exist without being a risk.

3.3. The realizations of the risk-related realizable entities

Note that the role approach does not deny the existence of dispositions such as $instability_0$: it just denies that this disposition is a risk and construes instead the risk as a role. Thus, the debate between both approaches can be seen as a question whether to map the natural language term "risk" with $Risk^d$ or $Risk^r$.

In this approach, the material object $ceiling_0$ bears two different realizable entities, the disposition $instability_0$ and the role $hazard_0$. Their realizations are related but different. If realized, $instability_0$ would bring about a BFO:Process of collapse in which the materials composing $ceiling_0$ lose cohesion and fall down under the action of gravity. In our scenario, this physical process is $collapsing_0$, which would happen independently of whether Paul is at risk or not. On the other hand, considering that the integrity of Paul, a normal human being, can be negatively impacted by the falling debris, and that he does not want such an outcome to happen, we identify the realization of $hazard_0$ (the risk) as a process of causing damage to Paul. This is a BFO:Process named here $damaging_0$ in which Paul's health is negatively affected. Note that $collapsing_0$ and $damaging_0$ happen approximately at the same time. If we just focus, to make things simple, on the ceiling debris that hit Paul and not those ten meters from him, the falling, physical impact, and biological damage all happen in the same approximate place and time, and with overlapping participants. This does not imply that those processes are identical though.

Fundamentally, $collapsing_0$ and $damaging_0$ can be analyzed as sums of spatial changes (change in spatial position of a material entity) and SDC changes respectively, where a SDC change is a BFO:Process that consists in a change of one or several SDCs (following Guarino, Baratella and Guizzardi [19], adapted to BFO by Toyoshima and Barton [20]). One can define $collapsing_0$ as the sum of processes of SDC changes (e.g. the structural cohesion of the ceiling that decreases) and processes of spatial changes (the spatial location of the ceiling debris which move down). $damaging_0$, on the other hand, might be defined as a change of SDCs characterizing Paul's health (without taking any specific position on the ontology of health here), which declines while he is hit by those bricks falling down². So, those two groups of quality changes co-occur in spatio-temporal regions that are close from each other but different. Therefore, the role approach admits an additional realizable entity compared to the dispositional approach, namely $hazard_0$; and $hazard_0$ and $instability_0$ are realized in different processes. This is another point of difference with the dispositional approach, which involved only one process, $collapsing_0$, and its classification in a specific subclass.

3.4. Merely possible processes

A central tenet of all discourses about risk is possibility. We have identified three processes so far: Paul actively entertaining a desiring process of not being hurt, $collapsing_0$, and $damaging_0$. However, even if none of these three processes would happen - in particular, if $instability_0$ and $hazard_0$ would have no realization, the whole situation is a risky one. How can we then express this unrealized possibility in the ontology?

¹Thus, one might say that an instance of $Risk^r$ is necessarily so, whereas an instance of $Risk^d$ is accidentally so.

²Thus, $collapsing_0$ and $damaging_0$ could be related to process profiles, which formed a subclass of BFO:Process until BFO2.0.

A first solution, following [21] would be to state that, if realized, $instability_0$ is realized in an instance of the class $ceiling_0_collapsing$ (which is, here too, a defined class and not a universal). This can be formalized with the axiom (O1):

(O1) $instability_0$ realized_in only $ceiling_0_collapsing$

However, a challenge with (O1) is that according to BFO, only classes that have an instance existing in our world (be it in the past, present, and/or future) can be introduced.

This means that the class $ceiling_0_collapsing$ might not exist, as there might be no collapse of $ceiling_0$ in the history of our world, and thus should not be introduced.

A possible way to deal with this problem is to formalize an axiom (O2) slightly differently by introducing instead the universal $Collapsing$ ³.

(O2) $instability_0$ realized_in only ($Collapsing$ and has_participant value $ceiling_0$)

For sure, some material entities have collapsed in the history of our world; therefore, the universal $Collapsing$ exists according to BFO. Thus, the axiom O2 and the entities mentioned there can fit into BFO's metaphysical framework. This strategy can be straightforwardly adapted to the processes of Paul being damaged or Paul actively desiring not to be hurt.

4. Conative theories or risk

4.1. A knowledge agnostic approach

According to both the dispositional and the role approach, risks come into existence because some unwanted outcomes that would negatively impact an agent might occur. So far, we have modeled this condition with a process of desiring entertained by an agent, in our case Paul, or some other conative mental state. Paul might walk under the ceiling with or without knowledge of its instability, holding a desire of self-preservation in both cases. Many publications on risk analysis focus on what an agent knows or forecasts about possible situations and which actions he takes or avoids given their beliefs [22]. Some other studies consider whether some of these actions are rational or ethical [23]. Here, we only focus on the metaphysics of risk, and we argue that the knowledge of Paul is not relevant for determining what is a risk for him: when $ceiling_0$ is unstable and Paul stands below, the risk is there, independently of his knowledge of the situation. Similarly, the appropriate ethical response to a risk has no bearing on the metaphysics of risk. Therefore, epistemological or ethical features are largely irrelevant to our current investigation. Invoking desires allows us to have a belief-agnostic approach in modeling risk, which means in our scenario that the beliefs and knowledge of Paul are irrelevant to determine what is risky. Paul's beliefs and course of actions could determine whether he is a good risk assessor and his degree of risk aversion, but this is a matter for decision theory or psychology of decision making. Additionally, as already hinted in the section on merely possible processes, Paul could be ignorant about his own desires, by having a dispositional desire that remains unrealized in explicit mental processes. In this we follow a tradition analyzing many mental entities as dispositions, including beliefs, desires, and intentions [24]. The desire disposition can remain latent for some time and can be realized in mental desiring processes, typically conscious, but that might be unconscious [25]. If Paul is sleeping or comatose under $ceiling_0$, there is still a risk that it would collapse on him, even without a conscious process of desiring safety, and independently of his (dispositional) beliefs.

³The authors thank Barry Smith for suggesting this idea in a private discussion. Note that 'has_participant value $ceiling_0$ ' refers, in OWL Manchester syntax, to the class of entities that have as participant $ceiling_0$

4.2. Desires and aversions

We consider here desires as mental dispositions directed towards some objects. Moreover, given that risks concern scenarios in which some damage or problem might occur, it seems the desires relevant for risk scenarios are connected to something that might not happen. What seems easy from a linguistic point of view is more difficult from an ontological point of view. It is classically problematic for ontology to deal with negation or inexistence [26].

First of all, it is relevant to distinguish between desiring not- x and not desiring x . The first case concerns an active process of the mind wishing that a certain state of affairs would not obtain. In the second case, a mental process is lacking: even stones can be said to not desire something, whereas they certainly cannot desire not- x .

When risk roles are realized in processes, those processes are contrary to an agent's desires. But realist ontologies do not accept things such as negative processes or negative objects - or more generally negative entities (see also [27] and [28]). The negation is rather embedded in a peculiar mental entity opposite to desiring. As in the mental realm love is the opposite of hate and hope is the opposite of despair, we can consider the mental attitude of aversion, which is the opposite of desiring, and might satisfy statements such as: "if an agent desires x , then he does not avert x ". Similarly to desires, aversions are dispositions that can be realized in mental processes - in this case, processes of averting. This therefore dispenses us from introducing problematic entities as negations in our model, and only connects existing entities together. To be more precise, mental attitudes like aversions and desires are connected to what they represent by what is usually called a relation of "intentionality" or "aboutness" [29], which we will call here "is_directed_towards".

4.3. Intentionality towards Risk Realizations

Having tackled the negativity of certain mental conative entities like aversion doesn't settle the issue of which real-world entity does Paul avert - that is, towards which entity the mental attitude of aversion is directed. Entities that Paul might avert include ceiling_0 , its disposition instability_0 , the process collapsing_0 or the process damaging_0 . What Paul fundamentally averts is to be bodily hurt, so he would avert the process damaging_0 , which would realize the risk hazard_0 . However, his mental attitude of aversion would exist even if hazard_0 was not realized in the process damaging_0 . Exploiting a similar solution to (O2) that we have seen with merely possible processes, we can say that aversion_0 can only have an intentionality relation with something that is a damaging process, if any, and that this process has Paul as its passive participant, since he would take the damage⁴. In OWL:

(O3) aversion_0 is_directed_towards only (*Damaging* and has_participant value Paul)

In this way, the sentence (O3) can be added in an ontology and be true even if there were no specific instances of *Damaging* involving Paul at the time the mental attitude of aversion was held by Paul.

Given this mental aversion for a generic damaging process, we can actually model another situation, in which ceiling_0 has been heavily painted with a poisonous paint which could impact Paul's health if he passed close to it, given the paint can release volatile particles in the air. The same object, ceiling_0 , and the same latent disposition of aversion, can explain several risks in this way. In the scenario in which the ceiling is both unstable and poisonous, aversion_0 can be intentionally directed to both harmful realizations of damaging by collapse and damaging by poisoning, since averting bodily damage is generic enough to comprise them both. For sure, we could model another agent, more selective than Paul, who only averts being damaged by the collapse of something from above, a very specific process, and so the poisonous ceiling would not lead to a risk for this selective agent, since they would not care. The lesson to be learned is that the intentional direction of aversion determines what is a risk or not.

Instead of averting the realization of a risk role, consider if we had said that Paul averted a process of collapsing of ceiling_0 which is the realization of a disposition, not of a risk role: such a process

⁴The authors thank John Beverley for this suggestion about intentional objects in OWL following BFO.

could happen even when Paul was miles away from ceiling₀, with no possibility of being hurt by it. He might actually feel some schadenfreude and rejoice at the sight of a collapsing building, for all we know. Maybe Paul likes watching videos of collapsing buildings, so this kind of process is not averted in itself, but only because it brings about a process of damaging Paul's health. This is an important distinction between the realizations of instances of Risk^d and Risk^r, as well as a virtue of modeling risk as the latter, as it grasps what is really at the core of the conative states of an agent.

Risk roles are externally grounded properties in the sense that they exist in virtue of something else than the entity they inhere in. The risk hazard₀ is not there just because of ceiling₀, but because of some complex situation of an agent desiring something, which ontologically grounds it. With this in mind and having found all the suitable entities to express a situation of risk, we can finally clarify the “context” that we introduced earlier and understand better the entities and the connections present there. The objective side of risk that we have modeled first, with physical interactions between the parties involved, has now been articulated with the subjective side foreshadowed at the beginning. Paul's mental disposition aversion₀, which is intentionally directed towards a class of damaging processes, grounds the existence of the risk role hazard₀ in ceiling₀.

Which intentional objects are targeted by the mental attitude determines which and how many risks there are, so outside of our example, if Paul also cared for its clothes and not just his health, which could be ruined by the falling debris of ceiling₀, a second risk, let's say hazard₁, would pop out into existence. Such risk would be grounded in a new mental attitude such as aversion₁, that would be intentionally directed towards a class of processes of clothes being ruined.

5. Towards a definition of risk

So far, we have illustrated our framework with some specific instances in a very particular situation that enabled us to illustrate the agent-relativity of risk. We can now present the framework of risk as a role in Fig. 2. Note that we depicted the case in which ceiling₀ actually falls during the collapsing₀ process and hazard₀ is realized in damaging₀, but such processes might never occur. Moreover, to improve readability, the relations participates_in and is_directed_towards have been omitted.

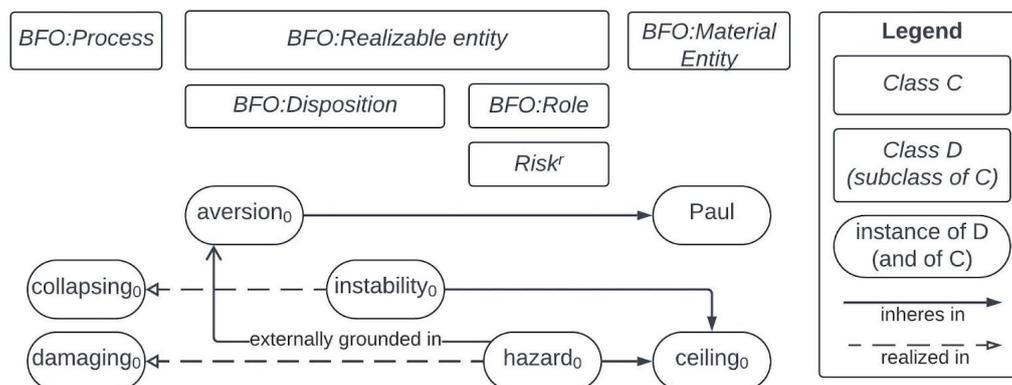


Figure 2: Risk as a BFO:Role

Now, we abstract away from this particular scenario and its instances to address risk more generally. The bearer of a risk might be a BFO:Object (think about a knife), or other kinds of BFO:Material entity, such as a body of water (e.g. in a flood) or instances of more specific BFO categories (consider the risk that a flock of birds, which is a BFO:Object aggregate, would attack a human). It is also important to consider that bearers of risk must be causally effective in producing the unwanted consequence an agent would avert, which supports the choice of a BFO:Material entity as the bearer of a risk, and that such bearer must also have some specific BFO:Disposition. Finally, the entity that averts the realization

of the risk could be an agent, like in our example, or a group of agents, like in the case of an entire population threatened by war (we do not discuss here whether such a group of agents would form an agent at another level).

Now, are these entities and relations enough to define risk? Arguably, some elements in our definition of risk might be too specific and unnecessary. Indeed, some risks might relate to teleological aspects that are more general than agents aversion - consider e.g. the risk for a plant to die when set aflame: burning would be the physical process involved, which is bad only within the perspective of the goal of the plant to stay alive. But it is debatable whether plants can desire or avert anything. Therefore, the conative approach might have to be generalized to also comprise this plant example.

Another element hinting towards a more general definition of risk, that we cannot elaborate here, is the possibility for non-material entities to bear risks. It is still to be determined if plausible risky situations like a financial loss or a cyber-attack only involve material entities, or if the source of damage could have a different nature.

Given these considerations, we propose the following sufficient (but not necessary) conditions for risk:

If: (i) r is a role &

(ii) r is externally grounded in an aversion that an agent or group of agents have towards the possible realizations of the role,

then r is a risk.

In case risk will be discovered to be only agentive and the risk bearer always material, then the condition above would actually be necessary and sufficient and constitute a definition. This shall be investigated in future work.

6. Conclusion and future work

To conclude, we have explored a novel way of understanding risk as a BFO:*Role*, after considering a similar approach which categorized it as a BFO:*Disposition*. In order to complete our modeling, we made use of realizable entities in BFO, opening the way for a future ontology of risks as roles. We proposed a sufficient condition for risks: future work should inquire whether this is also a necessary condition, or how it could be refined into a necessary and sufficient condition. A limitation of this exploratory investigation was its general character, and the lack of competency questions from domain experts to evaluate the model. Future work should also clarify the nature of a damaging process, what entities can be damaged and what is the output of damaging, since it might involve more considerations on negations, absences and realizable entities. Our risk ontology will need to be connected with risk probabilities (as started in [9] and [30]) and other relevant quantities: we merely covered qualitative possibilities so far, but the quantification of risks often involves probability measures and severity degrees. The ontology of risk propagation should also be investigated, where material entities can acquire and lose risk roles in a chain of causal processes. Furthermore, it will be important to evaluate this ontology of risk with real data and show how it can structure the information present in risk registers for better clarity and interoperability.

Declaration on Generative AI

No use of AI was made in the writing of this paper.

References

- [1] M. C. Leva, N. Balfe, B. McAleer, M. Rocke, Risk registers: Structuring data collection to develop risk intelligence, *Safety Science* 100 (2017) 143–156. doi:10.1016/j.ssci.2017.05.009.

- [2] M. Andretta, Some considerations on the definition of risk based on concepts of systems theory and probability, *Risk Analysis* 34 (2014) 1184–1195. doi:10.1111/risa.12092.
- [3] P. J. Blokland, G. L. Reniers, The concepts of risk, safety, and security, in: *Safety and Reliability – Safe Societies in a Changing World*, 2020, pp. 9–16. doi:10.1007/978-3-030-47229-0_2.
- [4] M. Leitch, Iso 31000:2009—the new international standard on risk management, *Risk Analysis* 30 (2010) 887–892. doi:10.1111/j.1539-6924.2010.01397.x.
- [5] N. Guarino, Understanding, building and using ontologies, *International Journal of Human-Computer Studies* 46 (1997) 293–310. doi:10.1006/ijhc.1996.0091.
- [6] R. Arp, B. Smith, A. D. Spear, *Building Ontologies with Basic Formal Ontology*, MIT Press, Cambridge, MA, 2015.
- [7] A. Iliadis, The tower of babel problem: making data make sense with basic formal ontology, *Online Information Review* 43 (2019) 1021–1045. doi:10.1108/OIR-07-2018-0210.
- [8] T. P. Sales, F. Baião, G. Guizzardi, J. P. A. Almeida, N. Guarino, J. Mylopoulos, The common ontology of value and risk, in: *Conceptual Modeling (ER 2018)*, 2018, pp. 121–135. doi:10.1007/978-3-030-00847-5_11.
- [9] A. Barton, L. Jansen, A. Rosier, J.-F. Ethier, What is a risk? a formal representation of risk of stroke for people with atrial fibrillation, in: *Proceedings of ICBO 2017*, 2018, pp. 1–6. URL: <http://ceur-ws.org/Vol-2137/>.
- [10] J. H. Kim, D. J. Park, A study on the review of risk concepts, *Journal of the Korean Society of Safety* (2013) 90–96. doi:10.14346/JKOSOS.2013.28.6.090.
- [11] Society for Risk Analysis, *Society for risk analysis glossary*, 2018. URL: <https://backend.orbit.dtu.dk/ws/portalfiles/portal/377037938/Society>.
- [12] S. O. Hansson, Risk, in: E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*, 2014. URL: <https://plato.stanford.edu/archives/spr2014/entries/risk/>.
- [13] B. Smith, Introduction to the logic of definitions, in: *International Workshop on Definitions in Ontologies*, 2013, pp. 1–2.
- [14] A. Varzi, Mereology, in: *The Stanford Encyclopedia of Philosophy*, 2019. URL: <https://plato.stanford.edu/archives/spr2019/entries/mereology/>.
- [15] M. Rabenberg, et al., Grounding realizable entities, *arXiv* (2024). doi:10.48550/arXiv.2405.00197.
- [16] O. Grenier, A. Barton, Une ontologie dispositionnelle du risque, *Lato Sensu* 8 (2021) 58–69. doi:10.20416/lrsrps.v8i2.6.
- [17] B. Smith, W. Ceusters, Ontological realism, *Applied Ontology* 5 (2010) 139–188. doi:10.3233/AO-2010-0079.
- [18] F. Toyoshima, A. Barton, L. Jansen, J.-F. Ethier, Towards a unified dispositional framework for realizable entities, in: *FOIS 2021*, 2021.
- [19] N. Guarino, R. Baratella, G. Guizzardi, Events, their names, and their synchronic structure, *Applied Ontology* 17 (2022) 249–283. doi:10.3233/AO-220261.
- [20] F. Toyoshima, A. Barton, Two approaches to the identity of processes in bfo, in: *FOIS 2020*, 2020. URL: <https://hal.science/hal-04301542>.
- [21] J. Röhl, L. Jansen, Representing dispositions, *Journal of Biomedical Semantics* 2 (2011). doi:10.1186/2041-1480-2-S4-S4.
- [22] S. Roeser, *Handbook of Risk Theory*, Springer, 2012.
- [23] S. O. Hansson, *The Ethics of Risk*, Springer, 2013.
- [24] F. Toyoshima, A. Barton, O. Grenier, Foundations for an ontology of belief, desire and intention, in: *FOIS 2020*, 2020. doi:10.3233/FAIA200667.
- [25] J. Hastings, B. Smith, W. Ceusters, M. Jensen, K. Mulligan, Representing mental functioning, in: *ICBO 2012*, 2012.
- [26] W. Ceusters, P. Elkin, B. Smith, Negative findings in electronic health records, *International Journal of Medical Informatics* 76 (2007) S326–S333. doi:10.1016/j.ijmedinf.2007.02.003.
- [27] R. Demos, A discussion of a certain type of negative proposition, *Mind* 26 (1917) 188–196.
- [28] S. Mumford, Negation and denial, in: *Cambridge Handbook of the Philosophy of Language*,

Cambridge University Press, 2021.

- [29] P. Jacob, Intentionality, in: The Stanford Encyclopedia of Philosophy, 2023. URL: <https://plato.stanford.edu/archives/spr2023/entries/intentionality/>.
- [30] A. Barton, A. Burgun, R. Duvaferrier, Probability assignments to dispositions in ontologies, in: FOIS 2012, 2012, pp. 3–14.