UNIVERSITÀ DEGLI STUDI DI MILANO

UNIVERSITÀ DEGLI STUDI DI MILANO

CORSO DI DOTTORATO ed eventuale curriculum
The Human Mind and its Expllanations

DIPARTIMENTO DI AFFERENZA DEL CORSO
Dipartimento di Filosofia Piero Martinetti

TESI DI DOTTORATO DI RICERCA
The Other Option

M-FIL/02

DUCA GIOVANNI
R13776
ORCID n. 0009-0001-8054-8599

TUTOR: prof. HYKEL HOSNI

CO-TUTOR: prof. MARIO PIAZZA

COORDINATORE DEL DOTTORATO: prof. FRANCESCO GUALA

A.A. 2025/26

Contents

Introduction

You are in a supermarket, sweating cold, in front of a shelf with many, maybe endless bottles. You know you must leave with one bottle in your bag — but which one? A choice must be made. Like many everyday situations, this simple choice involves identifying, comparing, and ultimately selecting among a set of alternatives. When the number of possibilities is too large to consider individually, we naturally simplify the problem by grouping similar options into a smaller set of *focal alternatives*. These focal options share common, distinguishing features that allow us to focus on the most relevant differences and ignore minor variations.

A similar process occurs in more structured decision contexts. Consider a typical survey that asks respondents to rate their satisfaction with a product on a continuous scale between 0 and 1. Research in psychometrics (see [McK78]) has long shown that using a continuous rather than discrete scale offers no advantage in terms of the reliability or validity of the responses. With some variability due to the matter of the questions and to the metrics used to evaluate the answers, researchers have found that people are not able to reliably discriminate more than 12 categories on a scale. For instance, [Ram73] shows that the standard errors of measurement do not significantly decrease with more than nine categories. Whether in everyday life or in formal settings, people simplify complex choice spaces by clustering similar possibilities and focusing on a manageable set of alternatives. Once this reduced set of focal options is identified, a judgment still needs to be made—we must decide which option is *best* or *most accurate*.

Any sound judgment must rely on data which allow us to distinguish among alternatives. However, acquiring such data is not always straightforward. Psychological research has identified several systematic biases that limit our ability to seek or interpret new information. One notable example is the phenomenon of *pseudodiagnosticity*, first introduced by [Doh+79].

Even with a limited number of hypotheses, we tend to acquire more evidence about the options for which we already have information, exhibiting an achoring effect with respect to the already available information. In our supermarket example, pseudodiagnosticity would manifest as inspecting the familiar shampoo's ingredients to confirm it contains nothing harmful, while neglecting to examine competing brands that might be safer or more pleasant.

Another relevant phenomenon is illustrated by the well-known *Wason selection task* [GC82], which tests how individuals evaluate conditional statements. Although the task concerns hypothesis testing rather than information search, the results similarly show that people systematically fail to select the pieces of evidence that would most effectively test their hypotheses. Both pseudodiagnosticity and the Wason task reveal the limits of human reasoning when it comes to acquiring *informative* evidence.

In this dissertation, however, our focus is not on how agents *acquire* new information, but rather on how they *reason* when presented with a finite set of alternatives. Given such a set, how do agents determine which option they should regard as correct or best supported? To address this question, we build upon the formal framework developed by Hanti Lin and Kevin Kelly [LK12b]. As we just discussed, when dealing with a high, or even infinite, number of options, people focus on a finite partition of them so to simplify the decision task. For that reason, throughout this work we are going to assume a finite set of possibilities which are exhaustive and mutually exclude each other. We are also going to assume a subjective interpretation of probability. Following de Finetti, the probability of an option a is measured by the price that the agent is willing to pay for a bet that pays 1 if a is true and 0 otherwise. Under this subjective interpretation, probabilities are also called degrees of belief or credences. Then, the agent disposes of a finite set of alternative options and some degrees of belief over them. The other crucial component of the framework consists of the agent's qualitative beliefs. A is a belief if the agent believes that the best option is among the alternatives compatible with A.

1. Probabilities and Beliefs

A central philosophical question arises: what is the relationship between an agent's quantitative *degrees of belief* and their qualitative *beliefs*? At what point does confidence in a proposition become strong enough to constitute belief? Conversely, when does diminishing confidence lead to its abandonment? More generally, how do changes in probabilistic confidence translate into changes in categorical belief?

Such and many more similar queries have brought the attention of philosophers to the matter. One of the reasons for such interest stands in the following. Logicians and epistemologists have always been interested in how idealized agents reason. But reasoning is never done in isolation, it is always bound to an environment; and beliefs are the agent's best attempt to capture the surrounding environment. The study of agents' beliefs has followed two main approaches, one qualitative and another more quantitative. The former, often referred to as belief revision, considers beliefs in a binary manner: either you believe something or you don't. More importantly, it studies how the binary beliefs change when some new evidence is acquired. To do so, it relies more on logical, rule-based and set-theoretical formalisms. On the other hand, the more quantitative approach makes more extensive use of probability and utility functions in order to capture respectively how likely and how valuable the agent thinks the outcomes to be. This second approach has been assumed in the areas of Bayesian Epistemology and Decision Theory.

Due to the partial difference in the formalisms they use, the object they study and the philosophical interests of the researchers, these two approaches have developed in quite independent ways. Then, asking what is the relation between binary beliefs and credences also allows to bring together and reconcile two distinct traditions in the field of logic and epistemology. Therefore, such query is not just about distinct formalisms and ways of relating them; it also brings researchers to address more philosophical questions such as whether binary beliefs and credences refer to the same cognitive phenomenon in humans, or whether they represent two different kinds of beliefs, which obey to distinct laws. A natural reference here is to the general idea behind Thinking fast and slow by Daniel Kahneman: there both system one and system two allow us to reason and to get a representation of our surroundings; nevertheless, in a practical scenario their aims differ, and so does their utility. The same kind of distinction might be applied to the present notions of credences and beliefs.

It is no surprise then that the main solutions proposed so far differ not just at the formal level, but at the philosophical as well. In particular, the odds-threshold account by Hanti Lin and Kevin Kelly [LK12b] functionally relates a probability function with a set of beliefs; under such proposal, binary beliefs are determined by credences, and so presumably supervene on them. In contrast, Leitgeb's *stability* theory [Lei17] posits that while beliefs and credences must cohere, they are not reducible to one another. More recently, [Dor19] and [SF19] have developed arguments in support of a third position, the *Lockean thesis*, which proposes a threshold view linking belief and credence but with distinctive epistemic implications.

2. Preferences over Options

Once an agent has formulated a set of alternative hypotheses, the next step in reasoning is *comparison*: evaluating which options appear more plausible or desirable. Consider a physician who, after assessing a patient's symptoms, identifies a number of possible diseases. They then weigh these alternatives — excluding some as unlikely and favoring others as more probable — in order to make a diagnosis. This evaluative process reflects a *preference ordering* among options, grounded in the agent's degrees of belief.

The problem of establishing a preference between options from some real-valued function over them has been thoroughly studied in the past. Specifically, between the mid 50's and the mid 70's this problem was addressed in the context of decision theory and utility measures. Duncan Luce pioneered in this topic.[1] In their case, one utility u_1 is judged as *definitely greater* than u_2, in notation $u_1 \succ_\delta u_2$, if and only if u_1 is greater than u_2 by a sufficiently large δ, referred to as the *just noticeable difference* (j.n.d.). Formally, fixed some $\delta > 0$,

$$u_1 \succ_\delta u_2 \text{ if and only if } u_1 \geq u_2 + \delta. \tag{1}$$

[1] See [Luc56], [SS58], [SZ62], [Sco64], [Rob70], [Rob71], [Fis70] and [Fis73].

As a reason for their choice, the field of utility theory is interested in determining what is the smallest unit that a person can value, the j.n.d. In so doing, they determine a threshold below which all the utilities have practically zero value. In addition, they are able to establish a preference relation over the utilities as in 1 precisely in terms of the just noticeable difference. Later, in Section 4 we will get into the details of such definition of preference over utilities and we will present the main results concerning it.

While there are strong theoretical reasons to accept Luce's definition of *definitely greater*, in the present work we will opt for a different definition. Specifically, Luce assumes a notion of just noticeable difference which is constant, it does not change with the utilities that are compared. This is coherent with the role that the j.n.d. plays in Luce's study of the *definitely greater* kind of preferences. However, in most cases the j.n.d. is be constant but rather directly proportionate to the magnitudes of the options. We see this as a widespread phenomenon in human cognition. From monetary valuations to sensory perception, humans do not ground their preference on a unique, smallest quantity that they are able to detect. Rather, their ability to compare and order two measures depends on the magnitude of the measures involved, and the j.n.d. for each of those is directly proportionate to them. A clear example comes from psychometrics of perception and it is the so called *Weber-Fechner Law*. The law was first postulated in 1834 by the German physiologist Ernst Heinrich Weber. It describes the relation between the physical intensity of a stimulus S and the smallest detectable change ΔS in its intensity. Relevantly to the present case, ΔS is referred to as *just noticeable difference*. The law sets the ratio between ΔS and S equal to a constant K_W, called *Weber fraction*.

(Weber's Law)
$$K_W = \frac{\Delta S}{S}$$

The law establishes a constant growth of the smallest detectable change ΔS with respect to the increase of the initial stimulus S. The law was soon refined and it was given a stronger theoretical explanation by Weber's student, Gustav Theodor Fechner. The perceived change in a stimulus dp, also referred to as *Weber contrast*, is proportional to the ratio of the change in stimulus intensity dS to the original stimulus intensity S for some constant k.

$$dp = k\frac{dS}{S}$$

By integrating on both sides, we can explain the so called *Fechner's Law*, which states that the relation between the stimulus and its perception is logarithmic. Fixed some base stimulus S_0, the difference between the perceptions of S_0 and of any other stimulus S is expressed by the following:

$$p(S) - p(S_0) = k\int_{S_0}^{S} \frac{dS'}{S'} = k \ln \frac{S}{S_0}$$

(Fechner's Law) $$p = k \ln \frac{S}{S_0}$$

Weber-Fechner Law has found many fields of application. We already mentioned that originally the authors were interested in the ability to distinguish sensory stimuli. As another important application, Weber-Fechner Law strongly relates in economics to the *Law of Diminishing Marginal Utility*, which states that, as the consumption of a good increases, the satisfaction provided by consuming one unit of that good continuously decreases. In general, the quantitative approach by Weber and Fechner on sensory stimuli and their perception is strongly relevant to orthodox economics and to consumer studies, fields which are intrinsically related to theory of preference.

Despite such connection and the fact that Weber-Fechner Law was already established at the time of publication of [Luc56], for the reasons mentioned before Luce gives a fixed numerical interpretation to the just noticeable difference. On the contrary, Weber-Fechner Law constitutes a both empirical and theoretical ground to argue for a different numerical interpretation of *definitively-greater* preference relation which is not additive, but rather *multiplicative*. Under such new interpretation, the j.n.d. is not constant but directly proportionate to the magnitudes, as Weber's Law and Weber contrast suggests. As a result, we are going to assume the following definition of preference. Let Ω be a set of options and $u : \Omega \to \mathbb{R}_{>0}$ a positive real-valued function. Then we preference one option a to another b if and only if the value of a is at least t-times the value of b, where t is some real number greater than 1.

Formally, the preference relation $\succ_u$ over Ω is defined as follows:

$$a \succ_u b \text{ if and only if } u(a) \geq t \cdot u(b). \tag{2}$$

A preferential relation $\succ_{u,t}$ so defined will be called odds-threshold order. Since t is contextual threshold which in the present case remains constant, we will give for granted that an odds-threshold order depends on t as well. Thus, we will assume the simplified notation $\succ_u$ to refer to an odds-threshold order.

2.1. From preferences to beliefs. The notion of odds-threshold order will be applied in the following manner. Starting from a set of alternative options, there exists a probability function capturing the agent's degrees of belief. The next step for the agent is to compare the alternatives so to determine the optimal one. In order to do so, they compare the probabilities of the options. We are not always able to distinguish two real values and to determine which of the two is higher. But we prefer one option two another if the former is sufficiently more likely than the latter. Our presentation of the Weber-Fechner's law shows that it is a good choice to capture our preferences in terms of an odds-threshold order. Thus, given some contextually determined threshold $t > 1$, a probability distribution P over the set of alternatives $\Omega = \{\alpha_1, \dots, \alpha_n\}$ generates an odds-threshold order:

$$\alpha_i \succ_P \alpha_j \text{ if and only if } P(\alpha_i) \geq t \cdot P(\alpha_j) \text{ with } P(\alpha_j) > 0. \tag{3}$$

In such a way, our degrees of belief generate a preferential order $\succ_P$ over Ω. $\succ_P$ gives a qualitative representation of the order of the probabilities of the alternatives in Ω and it expresses our preferences over them.

We have now defined a sound way of establishing our preferences over the options from their probabilities. Once we have established a preference over the options, we can determine our beliefs, which options we believe to the true. Because our preferences give a graphic, qualitative representation of our degrees of belief, our beliefs consist of the maximal elements of the odds-threshold order generated by our degrees of belief. More precisely, our logically strongest belief is equivalent to the disjunction of such maximal elements. In such a way, the probability function representing our degrees of belief determines our preferences and our beliefs as well.

2.2. Plan of the work. So far we have given an introduction on the framework that the present work is going to assume and on the motivations behind it. We are going to exploit such relation established between degrees of belief and qualitative beliefs to study how we reason when comparing alternative options. The work is structured as follows.

Chapter 1 studies what are the properties that our preferences satisfy if they are captured by an odds-threshold order. Specifically, we define by a set of qualitative condition a new class of preferential orders, referred to as semi-orders. The first sections of the chapter are devoted to proving a representation theorem between the class of semi-orders and the class of odds-threshold orders (Theorem 3.4). The second half of the chapter takes into account the results by Luce et al. It is shown that the same class of semi-order can be represented by the preferential relations defined in terms of a fixed difference between the values assigned to the options as in 1.

Semi-orders captures the order-theoretic properties satisfied by odds-threshold orders. Such qualitative characterization allows to understand how our preferences evolve, and so also how our beliefs evolve, when new information is acquired. This is the main content of Chapter 2. There it is proven that each of the properties characterizing semi-orders corresponds to a logical property satisfied by the nonmonotonic consequence relations induced by the semi-orders. The main result of the chapter, Theorem 4.1, shows that, a nonmonotonic consequence relation can be generated by a semi-order if and only if it satisfies all the rule of a newly defined nonmonotonic system, labeled system **D**.

Finally, Chapter 3 combines the results previously obtained to show how our beliefs change when the probability function that generated them is updated in the light of new information. As a first case, it is proven in Theorem 2.1 that, when such probability function is updated by Bayesian conditioning, our beliefs change in such a way that they satisfy all the rules of system **D**. Secondly, we extend the same result to a generalization of Bayesian conditioning, namely Jeffrey conditioning. Theorem 4.6 shows that the impact of Jeffrey conditioning on our beliefs is compatible with system **D** as long as the agent believes the evidence to be true. Such restriction on Jeffrey conditioning will be referred to as *successful* Jeffrey conditioning. Finally, one main difference between Bayesian conditioning and Jeffrey conditioning is that the latter does not necessarily assign to the evidence maximal probability. As a result, we conclude the chapter by addressing the properties of iterated belief revision that are satisfied by our beliefs when the underlying probability function is updated by Jeffrey conditioning.

3. Language and Probability

Throughout the entirety of this work, we assume a fixed language $\ddagger\mathcal{L}$ generated from a finite set of n propositional variables $\{a_1, \ldots, a_n\}$ together with the classical logical connectives $\neg$, $\wedge$, $\vee$ and $\leftrightarrow$. Let $\mathcal{SL}$ be the set of all sentences that can be constructed from $\mathcal{SL}$ as usual. A world or atom α is a formula of the form $\pm a_1 \wedge \cdots \wedge \pm a_n$, where $+a_i := a_i$ and $-a_i := \neg a_i$. Ω is the set of worlds, with cardinality $N := 2^n$. The set of subsets of Ω is denoted by 2^Ω and it is the set of propositions. Given a sentence $A \in \mathcal{SL}$, we denote with $\widehat{A}$ the set of worlds $\alpha \in \Omega$ satisfying A, in notation $\alpha \Vdash A$.

For what concerns probability, a probability function P is any function $P : 2^\Omega \to [0, 1]$ such that

- $\displaystyle\sum_{\alpha\in\Omega} P(\{\alpha\}) = 1$;
- $P(\widehat{A} \cup \widehat{B}) = P(\widehat{A}) + P(\widehat{B})$ for every $\widehat{A}, \widehat{B} \in 2^\Omega$ such that $\widehat{A} \cap \widehat{B} = \emptyset$.

Finally, the probability of an arbitrary sentence $A \in \mathcal{SL}$ is defined as $P(A) := \sum_{\alpha\in\widehat{A}} P(\alpha)$.

CHAPTER 1

Preference and semi-orders

1. Introduction

Say that you prefer one option to another only if the former is sufficiently better in some respect than the latter. Specifically, say that you prefer one option to another if the value of the former is at least t times higher than the value of the latter, where t is some real value greater than 1. If you then systematically compare all pairs of options in such a manner, as a result you obtain a preferential order over the options, which will be called here an odds-threshold order. By the pair-wise comparison just discussed, a real-valued function Q over the options is able to generate an odds-threshold order, once the parameter t is contextually fixed. An odds-threshold order then captures the notion of strong preference induced by the function Q. By changing the function Q, the odds-threshold order induced by Q, written $\succ_Q$, might change as well.

The present chapter studies the class of odds-threshold orders. It does so by giving a qualitative characterization of the class of all odds-threshold orders that can be generated by any function Q. In other words, in the present chapter we give a set of order-theoretic conditions which do not rely on any real-valued function and which are satisfied by all and only the odds-threshold orders. Any preferential order which satisfies all such qualitative conditions is called semi-order. The representation Theorem 3.4 proves that the classes of semi-orders and odds-threshold orders are identical.

The next section formally introduces both the notions of odds-threshold order and semi-order. The first step to prove the equality between these two classes of orders is to simplify semi-orders so that a real-valued function can be defined over it. The notion of Rank-set (Definition 2.10) serves precisely the purpose of identiying those elements which have the same relations in the semi-order, and are so indistinguishable with respect to it; a rank-set then collapses such elements together so that the resulting semi-order is simplified. Then, Definition 3.1 defines a real-valued function over the semi-order which Theorem 3.4 proves to generate an odds-threshold order equal to the initial semi-order. Finally, Section 4 considers other important results from the literature concerning semi-orders. Specifically, such results proves that semi-orders can be defined also in terms of the difference between the values of the options: one option is preferred to another if and only if its valued is at least δ greater than the value of the latter.

2. Semi-orders: qualitative odds-threshold orders

The present section focuses on a specific class of preferential orders: semi-orders. Their main purpose is to give a qualitative characterization of the class of odds-threshold orders, here defined.

DEFINITION 2.1 (Odds-threshold order). *Let Ω be a finite non-empty set. Let $Q : \Omega \to \mathbb{R}_{\geq 0}$ be a positive real-valued function over Ω, and let $\Omega_0 \subseteq \Omega$ be the set of $a \in \Omega$ such that $Q(a) > 0$. Let t be a contextually determined real-valued threshold such that $t > 1$. Then, Q induces an odds-threshold order $\langle \Omega_0, \succ_Q \rangle$ as follows:*

$$(a, b) \in \succ_Q \iff Q(a) \geq t \cdot Q(b) \quad \textit{with} \quad Q(b) > 0. \tag{4}$$

For emphasis, in general we write $\alpha_i \succ \alpha_j$ when $(\alpha_i, \alpha_j) \in \succ$. Also, for notational convenience, $\alpha_i \sim \alpha_j$ is written when both $\alpha_i \not\succ \alpha_j$ and $\alpha_j \not\succ \alpha_i$.

In this section, two definitions of semi-orders are given in terms of two distinct order-theoretic properties. The first definition relies on forbidden configurations, that is, suborders which semi-orders are not allowed to have. The second definition instead establishes for semi-orders a correspondence between the *number* of relations of an element (its *rank*), with what elements it is related to. Proposition 2.9 will prove the two definitions of semi-orders to be equivalent. Both of the two will be essential to prove Theorem 3.4, i.e. to prove the correspondence between semi-orders and odds-threshold orders.

2.1. Defining a semi-order. The first order-theoretic properties stated in Definition 2.4, considers two *forbidden configurations*: it states that a semi-order is any preferential order which does not have any *2v2* or *3v1* suborders, depicted in Figure 1. In other words, semi-orders are preferential order with the *2v2* and the *3v1* orders as forbidden configurations. All the figures in the present work represent the preferential relation $\succ$ as an arrow $\longrightarrow$. At the same time, notice that throughout this whole work, transitivity is given for granted when a preferential order is depicted. For instance, in Figure 1b, it holds that $a \succ c$ but no direct arrow goes from a to c. As already mentioned, for notational convenience, it is assumed an indifference relation $a \sim b$, defined as $a \not\succ b$ *and* $b \not\succ a$. Coherently, the figures depict the indifference relation as the absence of any arrow between two nodes. This whole work will rely on semi-orders, and each of the two forbidden configurations will play a distinct role. As a consequence, it is useful to introduce them separately. In particular, the two notions of filtered and mezzanine orders are defined as those preferential orders rank with the *2v2* and the *3v1* as respective forbidden configurations. Semi-orders then become the class of filtered mezzanine orders.

(A) *2v2*.   (B) *3v1*.

FIGURE 1. The *2v2* and *3v1* order configurations.

DEFINITION 2.2 (Filtered order). *A strict partial order* $\succ$ *over the set* Ω *is filtered if* $\succ$ *has no 2v2 suborder.*

DEFINITION 2.3 (Mezzanine order). *A strict partial order* $\succ$ *over the set* Ω *is a mezzanine order if* $\succ$ *has no 3v1 suborder.*

Definition 2.4 (Semi-order). *An order $\succ$ over the set Ω is a semi-order if $\succ$ is a filtered mezzanine order.*

By excluding the two configurations *2v2* and *3v1*, it is possible to establish a *ranking* over the elements of the preferential order. When we restrict to finite semi-orders, such ranking can be made explicit by assigning to the elements of the semi-order an integer value, given by the *rank function*. As stated in Definition 2.5, for each element $a \in \Omega$, the rank of a, $r(a)$, amounts to the number of elements defeated by a, those b such that $a \succ b$, minus the number of elements defeating a, those b such that $b \succ a$.

Definition 2.5 (Rank). *Let $\succ$ be a strict partial order over some finite set Ω. The rank function $r : \Omega \to \mathbb{Z}$ is defined:*

$$r_{\succ}(a) := |Down_{\succ}(a)| - |Up_{\succ}(a)|$$

where $Down_{\succ}(a) = \{b \in \Omega : a \succ b\}$ and $Up_{\succ}(a) = \{b \in \Omega : b \succ a\}$, respectively, referred to as a Down-set and Up-set.

In general, for a strict partial order, such measure bears low significance in the sense that comparing the ranks of two elements a and b does not give any information about their $\succ$-relations. In order to see this, consider the example of a strict partial order but not a semi-order given in Figure 2. Notice that the one depicted is not a semi-order because $b \succ c$ and $f \succ g$ constitute a *2v2* suborder. Then consider the elements c and f, with ranks -1 and 0, respectively. However, the ranks do not give any indication of the relations of the two elements in the following sense: despite the fact that c has a lower rank than f, c is preferred to an element d, i.e. $c \succ d$, while f is incomparable to d, $c \sim d$. That c is preferred to d while f is not, would suggest that c has a higher position than f in the preferential order; this is in contrast, however, with their respective ranks, -1 and 0. Similarly, though f has a higher rank than c, f is defeated by e while c is not. From this example, one can conclude that, in general in a strict partial order, the comparison between the ranks of two elements is not informative with regard to the relation between the two elements.

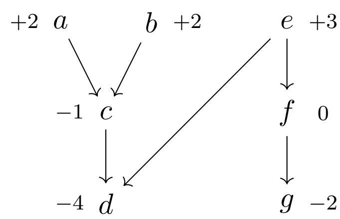

FIGURE 2. The rank function on a strict partial order which is not a semi-order.

On the contrary, this is the case when we restrict to semi-orders. As proven in Proposition 2.9, for a semi-order $\succ$, if $r_{\succ}(a) > r_{\succ}(b)$, then all the defeaters of a are defeaters of b ($Up_{\succ}(a) \subseteq Up_{\succ}(b)$), and all the elements defeated by b are defeated by a as well ($Down_{\succ}(a) \supseteq Down_{\succ}(b)$). Proposition 2.9 proves that such property is not just necessary but also sufficient for a preferential order to be a semi-order, once we restrict to finite cases.

The importance of the rank function $r_{\succ}$ stands in the fact that it allows to refine the original preferential relation $\succ$ so that a *ranking* can be established among all the elements, even among those elements which are incomparable with respect to $\succ$. While establishing such ranking, the rank function remains coherent with the preferential relation $\succ$ in the sense given by Proposition 2.8: if $a \succ b$, then $r_{\succ}(a) > r_{\succ}(b)$. Then, the rank function extracts some information which is only implicit in $\succ$, ranking also elements which are unrelated with respect to $\succ$. We will later see that the rank makes clear whether, even among $\succ$-incomparable elements, one needs to be assigned a real value which is higher than the one assigned to another, so that such real-valued assignment can represent the initial semi-order.

LEMMA 2.6. *Let $\succ$ be a strict partial order over Ω. For any two $a, b \in \Omega$, if $a \succ b$, then $Up_{\succ}(a) \subset Up_{\succ}(b)$ and $Down_{\succ}(a) \supset Down_{\succ}(b)$.*

PROOF. Take any two $a, b \in \Omega$ such that $a \succ b$. Since $\succ$ is a strict partial order, $\succ$ is transitive and irreflexive. For any $c \in \Omega$, if $c \succ a$, then by transitivity $c \succ b$; equivalently, $Up_{\succ}(a) \subseteq Up_{\succ}(b)$. In addition, $a \in Up_{\succ}(b)$ and $a \notin Up_{\succ}(a)$ by irriflexivity of $\succ$. Hence, $Up_{\succ}(a) \subset Up_{\succ}(b)$. Likewise, for any $c \in \Omega$, if $b \succ c$, then $a \succ c$; equivalently, $Down_{\succ}(a) \supseteq$

$Down_\succ(b)$. In addition, $b \in Down_\succ(a)$ and $b \notin Down_\succ(b)$. Hence, $Down_\succ(a) \supset Down_\succ(b)$. □

LEMMA 2.7. *Let $\succ$ be a strict partial order over Ω. For any two $a, b \in \Omega$,*

(1) if $Up_\succ(a) \subseteq Up_\succ(b)$ and $Down_\succ(a) \supseteq Down_\succ(b)$, then $r_\succ(a) \geq r_\succ(b)$;

(2) if $Up_\succ(a) \subset Up_\succ(b)$ and $Down_\succ(a) \supset Down_\succ(b)$, then $r_\succ(a) > r_\succ(b)$.

PROOF. Take any two a, $b \in \Omega$ such that $Up_\succ(a) \subset Up_\succ(b)$ and $Down_\succ(a) \supseteq Down_\succ(b)$. It follows that $|Up_\succ(a)| \leq |Up_\succ(b)|$ and $|Down_\succ(a)| \geq |Down_\succ(b)|$. From the two inequalities, it holds that $|Up_\succ(a)| - |Down_\succ(a)| \geq |Up_\succ(b)| - |Down_\succ(b)|$; by definition of rank, this is equivalent to $r_\succ(a) \geq r_\succ(b)$. The proof for the strict case is analogous. □

PROPOSITION 2.8. *Let $\succ$ be a strict partial order over Ω. For any $a, b \in \Omega$,*

$$\text{if } a \succ b, \text{ then } r_\succ(a) > r_\succ(b).$$

PROOF. The Proposition is a direct consequence of Lemmas 2.6 and 2.7. □

PROPOSITION 2.9. *Let $\succ$ be any strict partial order over Ω. The following properties are equivalent:*

(1) $\succ$ is a semi-order;

(2) for any $a, b \in \Omega$, $r_\succ(a) \geq r_\succ(b)$ (if and) only if $Down_\succ(a) \supseteq Down_\succ(b)$ and $Up_\succ(a) \subseteq Up_\succ(b)$.

PROOF.

(1) $\Rightarrow$ (2). The contrapositive is proven: *not* (2) $\Rightarrow$ *not* (1). Assume *not* (2). The right-to-left direction holds for every strict partial order as established by Lemma 2.7 (1). Thus, for the right-to-left direction of (2) *not* to hold, it has to be the case that for some $a, b \in \Omega$ such that $r_\succ(a) \geq r_\succ(b)$, either (A) $Down_\succ(a) \not\supseteq Down_\succ(b)$ or (B) $Up_\succ(a) \not\subseteq Up_\succ(b)$.

(A) $r_\succ(a) \geq r_\succ(b)$ and $Down_\succ(a) \not\supseteq Down_\succ(b)$.

$Down_\succ(a) \not\supseteq Down_\succ(b) \;\Leftrightarrow\; Down_\succ(b) \setminus Down_\succ(a) \neq \emptyset$;

Therefore, there exists $c \in \Omega$ s.t. $c \in Down_\succ(b)$ and $c \notin Down_\succ(a)$.

If $c \in Up_{\succ}(a)$, then $b \succ a$. This is in contradiction with $r_{\succ}(a) \geq r_{\succ}(b)$ by Proposition 2.8. Therefore, $c \notin Up_{\succ}(a)$. Since $c \notin Down_{\succ}(a)$ also holds, $c \sim a$ follows.[1]

Given $Down_{\succ}(a) \not\supseteq Down_{\succ}(b)$, it holds that either (A.1) $Down_{\succ}(a) \setminus Down_{\succ}(b) \neq \emptyset$, or (A.2) $Down_{\succ}(a) \subset Down_{\succ}(b)$.

(A.1) $Down_{\succ}(a) \setminus Down_{\succ}(b) \neq \emptyset$.

Therefore, there exists $d \in \Omega$ s.t. $d \in Down_{\succ}(a)$ and $d \notin Down_{\succ}(b)$.

In the following we prove that $a \succ d$ and $b \succ c$ are the only relations of the order, making it an instance on a *2v2* suborder. $c \sim a$ has already been proven. In order, we show that $d \sim b$, $c \sim d$ and $a \sim b$.

If $d \in Down_{\succ}(b)$, then $a \succ b$ follows from $d \in Down_{\succ}(a)$. Since $\succ$ is transitive, $Down_{\succ}(a) \supseteq Down_{\succ}(b)$. This is in plain contradiction with the assumption $Down_{\succ}(a) \not\supseteq Down_{\succ}(b)$. Therefore, $d \notin Up_{\succ}(b)$. As $d \notin Down_{\succ}(b)$ also holds, $d \sim b$ follows.

If $c \succ d$, given $b \succ c$, then by transitivity $b \succ d$. This contradicts $d \sim b$; therefore, $c \not\succ d$.

Likewise, $d \not\succ c$ follows from $a \succ d$ and $c \sim a$. Then, in short, $c \sim d$ holds.

Likewise, $a \not\succ b$ follows from $b \succ c$ and $c \sim a$.

Likewise, $b \not\succ a$ follows from $a \succ d$ and $d \sim b$. Then, in short, $a \sim b$ holds.

Therefore, $a \succ d$ and $b \succ c$ are the only relations of the order. It is thus obtained an instance of a *2v2* order. This contradicts (1); *not* (1) holds.

(A.2) $Down_{\succ}(a) \subset Down_{\succ}(b)$.

$r_{\succ}(a) \geq r_{\succ}(b) \Rightarrow |Up_{\succ}(b)| - |Up_{\succ}(a)| \geq |Down_{\succ}(b)| - |Down_{\succ}(a)|$.

$Down_{\succ}(b) \supset Down_{\succ}(a) \Rightarrow |Down_{\succ}(b)| - |Down_{\succ}(a)| > 0$.

$\Rightarrow |Up_{\succ}(b)| - |Up_{\succ}(a)| > 0$

$\Rightarrow$ there exists $d \in \Omega$ s.t. $d \in Up_{\succ}(b)$ and $d \notin Up_{\succ}(a)$.

---

[1] As a reminder, $a \sim b$ is written when both $a \not\succ b$ and $b \not\succ a$.

If $d \in Down_{\succ}(a)$, then the following chain is obtained: $a \succ d \succ b \succ c$. This contradicts the initial assumption $c \notin Down_{\succ}(a)$. Therefore $d \notin Down_{\succ}(a)$. Since $d \notin Up_{\succ}(a)$, $d \sim a$.

From $d \sim a$, $d \succ b \succ c$ and $c \sim a$, it follows that $b \sim a$. It is thus obtained an instance of a *3v1* order. This is the negation of (1); *not* (1) holds.

(B) $r_{\succ}(a) \geq r_{\succ}(b)$ and $Up_{\succ}(a) \nsubseteq Up_{\succ}(b)$.

As for case (A), the assumptions of case (B) lead to the existence of either a *2v2* or a *3v1* suborder, which negates (1). As the proof is analogous to case (A), it is left in the appendix.

We proved that assuming *not* (2) implies *not* (1). Therefore, (1) implies (2), as desired.

(2) $\Rightarrow$ (1). The contrapositive is proven: it is shown that the existence of any *2v2* or *3v1* suborder falsifies (2). We address the two cases separately.

2v2. There is a *2v2* suborder, say $\langle\{a, b, c, d\}, \{(a, c), (b, d)\}\rangle$. Take the two atoms a and b, and their ranks relative to the full order $\succ$. The following two cases are exhaustive: either $r_{\succ}(a) \geq r_{\succ}(b)$ or $r_{\succ}(b) \geq r_{\succ}(a)$. We address the two exhaustive cases separately.

- $r_{\succ}(a) \geq r_{\succ}(b)$. By premise, $b \succ d$ and $a \not\succ d$; hence, $Down_{\succ}(a) \nsupseteq Down_{\succ}(b)$, which contradicts (2).
- $r_{\succ}(b) \geq r_{\succ}(a)$. By premise, $a \succ c$ and $b \not\succ c$; hence, $Down_{\succ}(b) \nsupseteq Down_{\succ}(a)$, which contradicts (2).

Therefore, if there exists a *2v2* suborder, (2) is false.

3v1. There exists a *3v1* suborder, say $\langle\{a, b, c, d\}, \{(a, b), (b, c), (a, c)\}\rangle$. Take the two atoms b and d, and their ranks relative to the full order $\succ$. The following two cases are exhaustive: either $r_{\succ}(b) \geq r_{\succ}(d)$ or $r_{\succ}(d) \geq r_{\succ}(b)$. We address the two exhaustive cases separately.

- $r_{\succ}(b) \geq r_{\succ}(d)$. By premise, $a \succ b$ and $a \not\succ d$; hence, $Up_{\succ}(b) \nsubseteq Up_{\succ}(d)$, which contradicts (2).

- $r_\succ(d) \geq r_\succ(b)$. By premise, $c \succ b$ and $c \not\succ d$; hence, $Down_\succ(d) \not\supseteq Down_\succ(b)$, which contradicts (2).

Therefore, if there exists a *3v1* suborder, (2) is false.

Since both the *2v2* and the *3v1* cases negate (2), so does their disjunction: *not* (1) implies *not* (2), as desired. □

Proposition 2.9 proves the equivalence of the properties (1) and (2). If we restrict to finite orders, property (2) of Proposition 2.9 gives an equivalent characterization of semi-orders. Property (1) allows for a rank-independent and equivalent characterization of the class of semi-orders. So the notion of rank is clearly not required for the definition of such class. This is coherent with the fact that the rank function is entirely determined by the preferential order. Nevertheless, the notion of rank will play an essential role in defining a probability measure able to prove the Representation Theorem 3.4. From here on we will focus on semi-orders and assume by default this is the kind of strict partial orders that we are referring to. We will disambiguate when confusion might arise.

2.2. Rank sets. An important consequence of Proposition 2.9 is that two atoms have same rank if and only if they have the same relations, i.e. they have the same *Up* and *Down* sets. Thus, in the following, all rank-equivalent atoms, i.e. those with the same rank $r(\alpha_i)$, are grouped together into the set R_i, called *rank-set*.

Definition 2.10 (Rank set). *Let $\succ$ be a semi-order over a finite non-empty set Ω. Let $r_\succ$ be the rank function associated with $\succ$ and let $r_\succ[\Omega]$ be the image of Ω under $r_\succ$. For each $z \in r_\succ[\Omega]$, the rank-set R_z is the preimage of z under $r_\succ$:*

$$R_z = \{a \in \Omega : r_\succ(a) = z\}.$$

$\mathcal{R}_\succ$ is the set of all rank-sets of Ω. With an abuse of notation, the rank function applies to the rank-sets as follows:

$$r_\succ(R_z) := z$$

.

Proposition 2.9 states that all the elements of a rank-set are all and only those elements that share the same relations with each other. Proposition 2.11 also shows that, if two elements are in the same rank-set, then they must also be unrelated with respect to $\succ$. The converse instead does not hold: if two elements are unrelated, $a \sim b$, then it can be the case that $r_\succ(a) \neq r_\succ(b)$, as shown in Figure 3. Thus, these two facts show that in a semi-order for two elements to be incomparable it is not enough that they are unrelated, $a \sim b$, but they also need to have the same rank. However, once it is established that two elements have the same rank, grouping them together in the same rank-set will make easier proving the representation result, Theorem 3.4. Specifically, we will inductively define a real-valued function over the rank-sets (Definition 3.1) which we will prove to generate an odds-threshold order which is equal to the initial semi-order, as required to prove the representation theorem.

Proposition 2.11. *Let $\succ$ be a semi-order over Ω. For any $a, b \in \Omega$,*

(1) $r_\succ(a) = r_\succ(b)$ if and only if, $Up_\succ(a) = Up_\succ(b)$ and $Down_\succ(a) = Down_\succ(b)$;
(2) if $r_\succ(a) = r_\succ(b)$, then $a \sim b$.

Proof. By Proposition 2.9 (2), for any $a, b \in \Omega$ such that $r_\succ(a) = r_\succ(b)$, both the following cases apply:

$$r_\succ(a) \geq r_\succ(b) \;\Leftrightarrow\; Up_\succ(a) \subseteq Up_\succ(b) \text{ and } Down_\succ(a) \supseteq Down_\succ(b)$$

$$r_\succ(a) \leq r_\succ(b) \;\Leftrightarrow\; Up_\succ(a) \supseteq Up_\succ(b) \text{ and } Down_\succ(a) \subseteq Down_\succ(b)$$

Hence, $r_\succ(a) = r_\succ(b)$ holds if and only if both $Up_\succ(a) = Up_\succ(b)$ and $Down_\succ(a) = Down_\succ(b)$. Secondly, since $\succ$ is a strict partial order, $\succ$ is irreflexive, $a \notin Up_\succ(a)$ and $a \notin Down_\succ(a)$. Therefore, $a \notin Up_\succ(b)$ and $a \notin Down_\succ(b)$. So, a and b are unrelated with respect to $\succ$. □

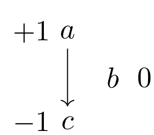

Figure 3. Two unrelated elements a and b have different ranks.

Notice that a semi-order $\succ$ entirely determines the rank function $r_\succ$ since $r_\succ$ just counts the relations of each $a \in \Omega$. In its own turn, the rank function $r_\succ$ entirely determines the Ω-partition $\mathcal{R}_\succ$. Therefore, in tern each semi-order induces a unique rank-based partition of Ω. Proposition 2.12 also shows that, since each rank-set is associated with an integer, its rank value, and no two rank sets share the same value, *the rank function completely orders* $\mathcal{R}_\succ$. For this reason, in the following we assume the notation of R_{min} to refer to the unique rank set with the lowest rank, R_{max} the one with the highest rank, R_{i+1} to refer to the rank set with the lowest rank such that $r_\succ(R_{i+1}) > r_\succ(R_i)$, and analogously R_{i-1} to refer to the rank set with the highest rank such that $r_\succ(R_i) > r_\succ(R_{i-1})$.

Proposition 2.12. *Let $\succ$ be a semi-order and let $r_\succ$ its associated rank function over Ω. Then, the set $\mathcal{R}_\succ$ of rank-sets induced by the rank function constitutes a partition of Ω. In addition, $r_\succ$ completely orders $\mathcal{R}_\succ$.*

Proof. For each $a \in \Omega$, the rank function counts the number of elements defeated by a and subtracts the number of defeaters of a. Thus, since Ω is assumed finite, $r_\succ$ is a function defined for all and only the elements of Ω, i.e. all and only $a \in \Omega$ receive a finite rank, and such rank is an integer number. Because each rank-set is the preimage of some $z \in r_\succ[\Omega]$, imagine of Ω under the function $r_\succ$, it follows that every $a \in \Omega$ belongs to one and only one rank-set, and there are no empty rank-sets. Hence, $\mathcal{R}_\succ$ is a partition of Ω.

Concerning the second part of the proposition, the image $r_\succ[\Omega]$ is composed of only integer numbers. The natural ordering over $r_\succ[\Omega]$ can be extended to the corresponding preimages, the rank-sets. Thus, the Ω-partition $\mathcal{R}_\succ$ is totally ordered by the rank function. $\square$

There is one final point to be noted about rank-sets. Proposition 2.11 shows that all the elements in a rank-set R_i share the same relations. As a consequence, the relations of a semi-order directly extend to the rank-sets. Two rank-sets are related, $R_i \succ R_j$, if and only if there is an element $a \in R_i$ that is preferred to some $b \in R_j$, if and only if all $a \in R_i$ are preferred to all $b \in R_j$. Hence, all the properties of Proposition 2.9 apply to the rank sets orders as well. The *Down* and *Up* sets are analogously defined:

$$Down_{\succ}(R_i) = \{R_j \in \mathcal{R} : R_i \succ R_j\} \text{ and } Up_{\succ}(R_i) = \{R_j \in \mathcal{R} : R_j \succ R_i\}.$$

Additionally, two distinct rank sets R_i and R_j cannot have the same relations:

$$Down_{\succ}(R_i) = Down_{\succ}(R_j) \text{ if and only if } Up_{\succ}(R_i) \neq Up_{\succ}(R_j).$$

2.3. Visualizing semi-orders. Semi-orders are preferential order with the *2v2* and *3v1* as forbidden configurations. Such negative characterization however does not facilitate coming up with examples of semi-orders as one would need to check for the absence of the *2v2* and *3v1* configurations. Figure 4 gives an example of what a semi-order looks like. In general, any semi-order can be included in an order of the kind depicted in Figure 4. Proposition 2.11 shows that in a semi-order two elements have the same rank if and only if they have the same relations. For this reason, in Figure 4 we collapse together all atoms with equal rank and we assume that each node corresponds to a rank set R_i, and is labeled by the corresponding rank value r_i. Coherently, there are no nodes with the same edges as there are no rank sets with the same relations. Finally, the nodes are placed in a way that the rank values associated to them increase as you move rightwards and upwards.

The same disposition is given in Figure 4b, where each node is labeled with a pair of natural number (f, r), the first one representing the *floor* and the second one the *room* associated with each node. The order given by the rank value corresponds to the order given by the pair (floor, room): given two distinct nodes, the one with the higher floor has higher rank, and if two nodes are in the same floor, the one with the higher room number has higher rank. Such way of visualizing semi-orders is what initially inspired the name *hotel orders*. That the rank sets are disposed in increasing order with respect to their rank is directly related to their relations: if in Figure 4 you move from one node R_i to its higher adjacent R_{i+1}, $Up_{\succ}(R_{i+1})$ has exactly one rank set less than $Up_{\succ}(R_i)$, as long as $Up_{\succ}(R_i)$ is not empty; likewise, if you move from one node R_i to its lower adjacent R_{i-1}, $Down_{\succ}(R_{i-1})$ has exactly one rank set less than $Down_{\succ}(R_i)$, as long as $Down_{\succ}(R_i)$ is not empty. Such regularly decreasing nature of both the *Up* and the *Down* sets is what makes the order in Figure 4 a peculiar kind of

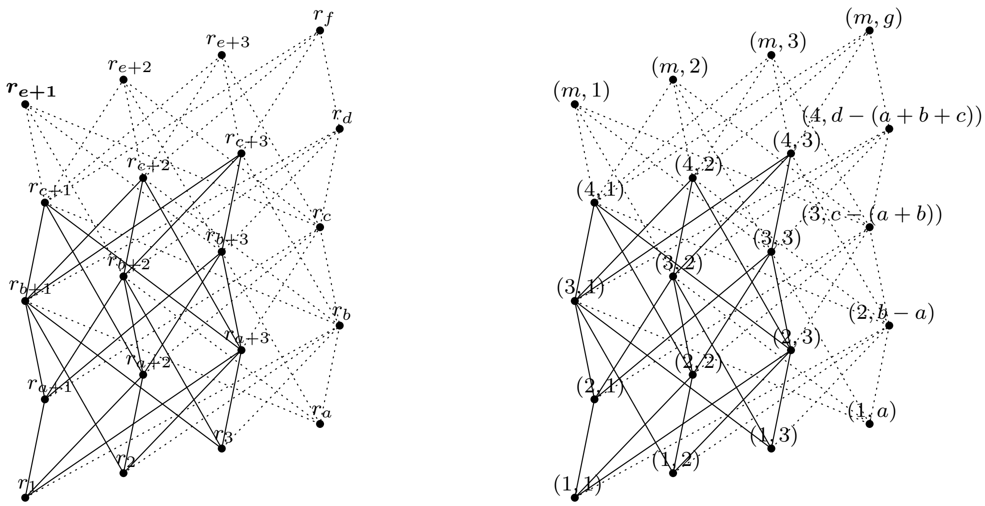

(A) Extended hotel order with increases rank values.

(B) Floors and rooms, (f, r), of an extended hotel order.

FIGURE 4

semi-order, and it is the reason why any semi-order can be included in it for some big enough $a, \ldots, f$ values.

One final note concerning Figure 4 and the intuition behind the rank function. As already mentioned before, our approach here is somewhat the opposite of the one taken by Luce. We started from probability distributions over atoms, and our intention was to characterize the class of orders that can be generated by comparing the probabilities of the atoms, by preferring an atoms to another if and only if it is at least t times more likely than the latter. In the process of giving a qualitative representation of the class of preferential orders so obtainable, the rank function has the role of ordering the atoms alongside their respective real-valued probabilities, while also taking into account the preferential relations, being a function of those. Thus, by focusing on the rank sets, and so collapsing together all those atoms with equal rank, we eliminate as much ambiguity as possible concerning which atoms are more likely than the others. In so doing, we get a total order over the rank sets which corresponds to the total order of the probabilities over the atoms that originated the initial semi-order.

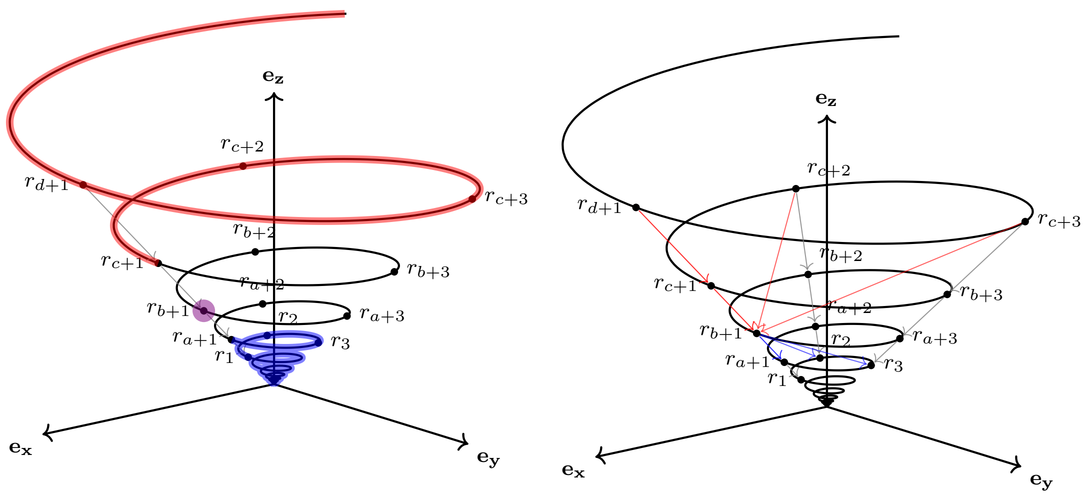

(A) In red the values t-times more likely and in blue t-times less likely than the value of reference r_{b+1}.

(B) The preference relations induced by comparing the real values.

FIGURE 5. Helix.

In order to appreciate this correspondence between the ordering of the rank sets given by their rank values, and the probabilities of the respective elements, notice that Figure 4 is the two-dimensional development of the helix in Figure 5. The helix is composed of the half-line of the positive real numbers, bent in such a way that at each upwards complete revolution of 2π the value has increased by a factor t, with $t > 1$. A helix with such growth rate takes the name of *logarithmic helix*. In Figure 5, we also assumed a constant growth rate on the z-coordinate, and a clockwise rotation of the spiral, given by the minus sign on the y-coordinate, tough none of them are necessary conditions for our representation. Thus, the helix is defined by the following parametric equations:

$$\text{(Logarithmic helix)} \qquad \begin{cases} x & = ae^{k\phi}\cos\phi \\ y & = -ae^{k\phi}\sin\phi \\ (z & = ae^{k\phi}) \end{cases}$$

We want the half-line of the positive real numbers to be bent in such a way that at every upwards complete revolution of 2π, the value increases by a factor of t. This property allows us to express the parameter k in terms of t. First, we need to obtain the formula for the arc length $L(\phi_1, \phi_2)$ between two angles ϕ_1 and ϕ_2 for such a helix. In order to do so, it is enough to consider the arc formula for a generic 3D function, Eq. 5.

$$L(\phi_1, \phi_2) = \int_{\phi_1}^{\phi_2} \sqrt{\left(\frac{dx}{d\phi}\right)^2 + \left(\frac{dy}{d\phi}\right)^2 + \left(\frac{dz}{d\phi}\right)^2} d\phi \tag{5}$$

In order to apply the equation, we determine the derivatives from the parametric equations of the logarithmic helix.

$$\begin{cases} \frac{dx}{d\phi} = ae^{k\phi}(k \sin\phi - \cos\phi) \\ \frac{dy}{d\phi} = -ae^{k\phi}(k \cos\phi + \sin\phi) \\ \frac{dz}{d\phi} = ake^{k\phi} \end{cases} \tag{6}$$

Finally, after substituting the derivatives from Eq. 6 into Eq. 5, by integration we obtain the Arc Length Equation.

$$L(\phi_1, \phi_2) = \frac{a\sqrt{2k^2+1}}{k}(e^{k\phi_2} - e^{k\phi_1}) \tag{Arc Length}$$

Notice that the parametric equations of the helix are expressed in terms of the parameter ϕ, corresponding to the angle of revolution, which ranges over all real values. The definition of the helix is such that it intersects the origin of the axes in Figure 5 only in the limit at $\phi = -\infty$. Nevertheless, the total length of the helix from $-\infty$ up to a certain angle ϕ converges to a finite value, which is given by the following:

$$L(-\infty, \phi) = \frac{a\sqrt{2k^2+1}}{k}e^{k\phi} \tag{7}$$

We can now state our condition on the length of the helix increasing by a factor of t at each 2π revolution:

$$L(-\infty, \phi + 2\pi) = t \cdot L(-\infty, \phi) \tag{8}$$

which results in the following relation between the constants k and t:

$$k = \frac{ln(t)}{2\pi} \tag{9}$$

Notice that the initial condition 8 refers to a generic angle ϕ, but ϕ is not mentioned in Equation 9. This follows from one of the properties of logarithmic spirals, that their radius grows with the angle ϕ at a constant rate, proportionate to k. At the same time, we also imposed that the length of the helix increases with ϕ at a constant rate. As a result, Equation 9 shows that k and t are proportionate as well.

Figure 4 gives an idea of the *regularity* in which the elements of a semi-order are disposed. In support of that, Figure 5 shows how it is possible to characterize both the *Up* and *Down* sets of a given element in terms of some real-values assigned to them. In the next section we will prove the representability of every semi-order in terms of real valued-functions over its elements. In other words, starting from a semi-order we inductively define a real-valued function which is able to generate the same order via the odds-threshold method.

3. The Q Function and the Representation Theorem

The present section focuses on showing that every semi-order is representable in terms of a real-valued function, and in terms of a probability function. Given a semi-order $\succ$ over a finite set $\Omega_0 \subseteq \Omega$ and the rank-sets partition of Ω_0, $\mathcal{R} = \{R_0, \ldots, R_m\}$, we first define a real-valued function $Q : \mathcal{R} \to \mathbb{R}_{>0}$ that recursively assigns to each R_i a strictly positive real value (Def. 3.1). The function Q is proven to be monotonic with respect to the rank values of the rank sets (Lemma 3.2); Q is also proven to match the preference relations given by the initial hotel order, in terms of a generic odds-threshold $t > 1$ (Lemma 3.3). The real-valued function which is able to generate the initial semi-order and so to verify Theorem 3.4, is obtained by extending Q from $\mathcal{R}$ to Ω by giving to the atoms the same Q value as the rank set which they belong to, i.e. $Q(a) := Q(R_i)$, if $a \in R_i$, and $Q(a) = 0$ if $a \in \Omega \setminus \Omega_0$. Finally, for the general purpose of the present work, we need to show that for any semi-order, there is a probability function P able to generate it via the odds-threshold method. This is simply obtained as a corollary of the representation theorem by normalizing the function Q over the set of atoms Ω as the normalization preserves the odds between the probabilities of two atoms.

Definition 3.1 defines the Q function, that we just introduced. Since Ω is finite, so is the partition of its subset Ω_0, $\mathcal{R}$. Hence, the definition of Q proceeds from bottom to top. It starts from R_{min}, the rank-set with the lowest rank, which is assigned a generic strictly positive real value q. Secondly, because the rank-sets are linearly ordered by the same rank function, the Q function is recursively defined for all the other rank-sets R_i. The Q has already been defined for R_{i-1}. If $Down_{\succ}(R_i) = Down_{\succ}(R_{i-1})$, then the two rank-sets defeats the same rank-sets. Nevertheless, the rank of R_i is greater than the rank of R_{i-1} thus it has to be the case that $Q(R_i) > Q(R_{i-1})$ in order for the Q function to represent $\succ$; thus, a small enough positive value ϵ is $Q(R_{i-1})$ to defined $Q(R_i)$. On the contrary, if $Down_{\succ}(R_i) \neq Down_{\succ}(R_{i-1})$, then R_{i-1} and R_i are not preferred to the same rank-sets, and since R_i has higher rank, R_i defeats some rank-set more than R_{i-1}. Thus, its Q value is set to t times the Q value of the rank set $R'_i \in Down_{\succ}(R_i)$ with the highest rank.

Definition 3.1. *Let Ω be a finite non-empty set. Let $\mathcal{R}$ be the rank sets partition of Ω induced by the semi-order $\succ$ over Ω, and let q and t be any two real values such that $q > 0$ and $t > 1$. The following defines the function $Q : \mathcal{R} \to \mathbb{R}_{>0}$ assigning to each rank set a strictly positive value.*

Input: $\Omega \neq \emptyset$, $t > 1$, $q > 0$

Output: $Q : \mathcal{R} \to \mathbb{R}_{>0}$

$i \leftarrow 0$

$Q(R_i) \leftarrow q$

$s \leftarrow 0$

$\epsilon_s \leftarrow \frac{(t-1)\cdot Q(R_{min})}{(|\Omega|)^s}$

while $R_i \neq R_{max}$ ***do***

    $i \leftarrow i + 1$

    if $Down_{\succ}(R_i) = Down_{\succ}(R_{i-1})$ ***then***

        $s \leftarrow s + 1$

$\epsilon_s \leftarrow \frac{(t-1) \cdot Q(R_{min})}{(|\Omega|)^s}$

$Q(R_i) \leftarrow Q(R_{i-1}) + \epsilon_s$

else

$R_i' \leftarrow$ *the rank set with the highest rank in* $Down_{\succ}(R_i)$

$Q(R_i) \leftarrow t \cdot Q(R_i')$

end if

end while

Lemma 3.2 proves that the Q function is monotonic with respect to the rank function, and viceversa. Finally, Lemma 3.3 shows that the Q function generates an odds-threshold order which is the same as the initial semi-order. The two lemmas ease the proof of Theorem 3.4 which follows. The theorem proves that a strict partial order is semi-order if and only if it can be generated by positive real-valued function Q over its elements. Finally, Corollary 3.5 shows that $\succ$ is semi-order if and only if it can be represented by a special kind of positive real-valued function, namely a probability function P.

LEMMA 3.2. *Let* $Q : \mathcal{R} \to \mathbb{R}_{>0}$ *be a real valued function defined by Def. 3.1 for some fixed value* t. *Then, for any two rank sets* $R_j, R_i \in \mathcal{R}$,

$$r(R_j) > r(R_j) \iff Q(R_j) > Q(R_i).$$

PROOF. First of all, notice that if there is only one rank set, the terms on the left and right of the logical equivalence ($\Leftrightarrow$) are never satisfied, making Lemma 3.2 vacuously true. Therefore, we can assume that there are at least two distinct rank sets, R_{min} and R_{max}. In such a case, since $\mathcal{R}$ is a partition of Ω, $|\mathcal{R}| \geq 2$ implies $|\Omega| \geq 2$. This will later come in handy.

($\Rightarrow$). An exhaustive proof of the monotony of the Q function with respect to the rank is given in the Appendix, Proof 7. We prove ($\Rightarrow$) by induction, whose steps rely on a proof by cases. As a result, the proof is quite tedious and is left in the Appendix. In the following, we provide several examples to show why Q is monotonous with respect to r ($\Rightarrow$).

Definition 3.1 starts by assigning R_{min} some strictly positive real value q, and, at each iteration of the **while**, it assigns a new value $Q(R_j)$ either by multiplying an already defined[2] value $Q(R_i)$ by some real $t > 1$ (**else**), or by adding to an already defined[2] value $Q(R_i)$ some strictly positive value ϵ_s (**if**). To see that $\epsilon_s > 0$, consider the following: it is assumed $Q(R_{min}) = q > 0$ and $t > 1$, from which $t - 1 > 0$. Also, being s a counter of how many times the ***if*** at line 7 was satisfied, $s \geq 0$. Finally, $|\Omega| \geq 2$. Therefore,

$$\epsilon_s = \frac{(t-1) \cdot Q(R_{min})}{(|\Omega|)^s} > 0.$$

Getting into the matter of ($\Rightarrow$), take any two rank sets R_i, $R_j \in \mathcal{R}$ such that $r(R_i) < r(R_j)$. By the definition of rank set, R_i and R_j are distinct. If $Down(R_i) = Down(R_j)$, then we are in a situation like the one depicted in Fig. 6. There, as in the following figures, we omit the rank sets that do not pertain to the case; notice however that, being the depicted nodes distinct rank sets, they must have different relations; hence, there must be additional rank sets to which their are related to, and the relations with which makes them distinct rank sets. Back to Figure 6, we have established that each ϵ_s value is strictly positive. It is immediate to see that $Q(R_i) < Q(R_i) + \sum_{s=p}^{r} \epsilon_s = Q(R_j)$, as required.

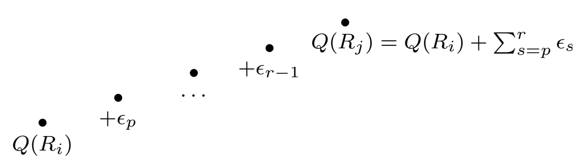

FIGURE 6. Rank-sets sharing the same *Down*-set.

[2]That $Q(R_i)$ has already been defined by Def. 3.1 is due to the following. For **if** at line 7 to apply, $R_i = R_{j-1}$, for which $Q(R_{j-1})$ was defined at the previous iteration of **while**. For **else** and line 11 to apply, $R_i \in Down_{\succ}(R_j)$, from which we know by Prop. 2.8 that $r(R_j) > r(R_i)$; since at each of its iterations, **while** moves form one rank set R_k to its successor R_{k+1}, and there are no two rank sets with the same rank by Def. 2.10, once **while** defines $Q(R_j)$, $Q(R_i)$ has already been defined for all the rank sets R_i with lower rank than R_j.

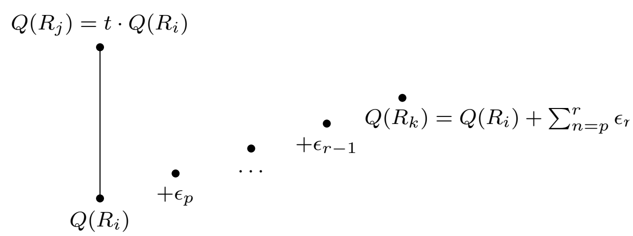

FIGURE 7. R_j has a distinct *Down*-set from its predecessor (case 1).

If $Down(R_i) \neq Down(R_j)$ instead, by the properties of semi-orders it follows from $r(R_i) < r(R_j)$ that $Down(R_i) \subset Down(R_j)$. Consider Figure 7, where $Q(R_j)$ is defined by **else** of Def. 3.1 as $Q(R_j) = t \cdot Q(R_i)$. $Q(R_i) < t \cdot Q(R_i) = Q(R_j)$ holds as required. More importantly, consider all the rank sets R_k between R_i and R_j, whose Q value is defined by **if** of Def. 3.1 as

$$Q(R_k) = Q(R_i) + \sum_{s=p}^{r} \epsilon_s,$$

with $r \geq p \geq 1$. That is the case if and only if $Down(R_k) = Down(R_i)$. Since $\succ$ is strict partial order, it holds that $R_i \not\succ R_i$, i.e. $R_i \notin Down(R_i)$; then $R_i \notin Down(R_k)$. At the same time $R_i \in Down(R_j)$. This is enough to establish that $r(R_j) > r(R_k)$. In the following, we show that, regardless of how big r is, it is true that $Q(R_j) > Q(R_k)$, as required. In other words, it does not matter how many times **if** is applied and how many ϵ's are added to $Q(R_i)$ because $t \cdot Q(R_i)$ is never reached nor exceeded. This aspect, on the one hand, contributes to the monotonicity of the Q function with respect to the rank ($\Rightarrow$); on the other hand, it plays a pivotal role for establishing that the Q function generates an odds-based order which is exactly equal to the initial hotel order (see Lemma 3.3).

$$
\begin{aligned}
Q(R_k) = Q(R_i) + \sum_{s=p}^{r} \epsilon_s &= Q(R_i) + (t-1) \cdot Q(R_{min}) \cdot \sum_{s=p}^{r} \left(\frac{1}{|\Omega|}\right)^s \\
&< Q(R_i) + (t-1) \cdot Q(R_{min}) \cdot \sum_{s=p}^{\infty} \left(\frac{1}{|\Omega|}\right)^s \\
&= Q(R_i) + (t-1) \cdot Q(R_{min}) \cdot \left(\frac{1}{|\Omega|}\right)^{p-1} \cdot \frac{1}{|\Omega|-1} \\
&< Q(R_i) + (t-1) \cdot Q(R_{min}) \\
&< Q(R_i) + (t-1) \cdot Q(R_i) = t \cdot Q(R_i) = Q(R_j)
\end{aligned}
$$

The last and more general case to be considered is when $Down(R_i) \subset Down(R_j)$ holds as before, but $R_j \not\succ R_i$. Figure 8 exemplifies such situation; for simplicity, it is additionally assumed that both $Q(R_i)$ and $Q(R_j)$ are defined by **else** of Def. 3.1. As for the previous case, since $R'_j \notin Down(R_i)$ and $R'_j \in Down(R_j)$, it can be established that $r(R_i) < r(R_j)$ and that $r(R_{j-1}) < r(R_j)$. Additionally, since $R'_i, R'_j \in Down(R_j)$, $Q(R_j) = t \cdot Q(R'_j)$, and by Def. 3.1 R'_j is the rank set with the highest rank in $Down(R_j)$, it follows that $r(R'_i) < r(R'_j)$. As shown in Proof 7, it is possible to establish that, when $r(R'_i) < r(R'_j)$, $Q(R'_i)$ and $Q(R'_j)$ differ of at least a value ϵ_o, with $o < p$; for simplicity, assume that $Q(R'_j) - Q(R'_i) = \epsilon_o$. Below, it is shown that, regardless of how high r is, $Q(R_{j-1}) < Q(R_j)$ holds, as required. This proves the following: even when two rank sets R_i and R_j are defined by **else** of Def. 3.1 as t times $Q(R'_i)$ and $Q(R'_j)$ respectively, where $Q(R'_i)$ and $Q(R'_j)$ differ only the smallest amount ϵ_o, it does not matter how many ϵ values are added to $Q(R_i)$; $Q(R_j)$ is never reached nor exceeded. This is due to Def. 3.1 diminishing the ϵ value at each iteration of **if**, and to the properties of finite geometrical series with ratio lower than 1 that is derived from that.

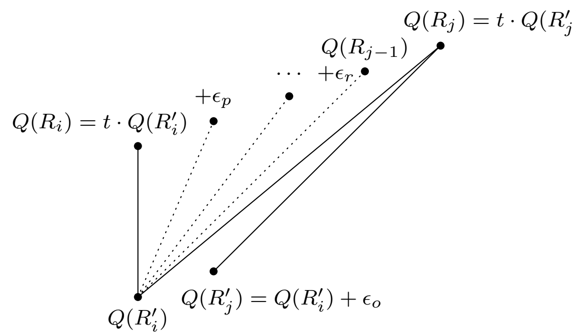

FIGURE 8. R_j has a distinct *Down*-set from its predecessor (case 2).

$$\begin{aligned}
Q(R_{j-1}) &= Q(R_i) + \sum_{s=p}^{r} \epsilon_s = t \cdot Q(R'_i) + \sum_{s=p}^{r} \epsilon_s \\
&< t \cdot Q(R'_i) + \sum_{s=p}^{\infty} \epsilon_s = t \cdot Q(R'_i) + (t-1) \cdot Q(R_{min}) \cdot \sum_{s=p}^{\infty} \left(\frac{1}{|\Omega|}\right)^s \\
&< t \cdot Q(R'_i) + (t-1) \cdot Q(R_{min}) \cdot \left(\frac{1}{|\Omega|}\right)^{p-1} \\
&\leq t \cdot Q(R'_i) + (t-1) \cdot Q(R_{min}) \cdot \left(\frac{1}{|\Omega|}\right)^{o} = t \cdot Q(R'_i) + \epsilon_o \\
&< t \cdot Q(R'_i) + t \cdot \epsilon_o = t \cdot Q(R'_j) = Q(R_j)
\end{aligned}$$

The three cases that we addressed exemplify the possible cases where there are two rank sets R_j and R_i such that $r(R_j) > r(R_i)$. For all three possible cases, we showed that any Q function defined by Def. 3.1 is such that $Q(R_j) > Q(R_i)$. Therefore, ($\Rightarrow$) holds.

($\Leftarrow$). It has to be proven that, for any two rank sets $R_i, R_j \in \mathcal{R}$, if $Q(R_i) > Q(R_j)$, then $r(R_i) > r(R_j)$.

The contrapositive is proven: for any two rank sets $R_i, R_j \in \mathcal{R}$, if $r(R_i) \leq r(R_j)$, then $Q(R_i) \leq Q(R_j)$. We proceed by cases:

- $r(R_i) = r(R_j)$. By the definition of rank set, there are no two distinct rank sets with the same associated rank; therefore $R_i = R_j$. Then, $Q(R_i) = Q(R_j)$, from which $Q(R_i) \leq Q(R_j)$, as desired.
- $r(R_i) < r(R_j)$. By the just proven ($\Rightarrow$) of the present Lemma, it follows that $Q(R_i) < Q(R_j)$, from which $Q(R_i) \leq Q(R_j)$ holds.

□

LEMMA 3.3. *Let $Q : \mathcal{R} \rightarrow \mathbb{R}_{>0}$ be a real valued function defined by Def. 3.1 for some fixed value t. Then, for any two rank sets $R_i, R_j \in \mathcal{R}$,*

$$R_i \succ R_j \iff Q(R_i) \geq t \cdot Q(R_j).$$

PROOF.
($\Rightarrow$). It has to be proven that for any two rank sets $R_i, R_j \in \mathcal{R}$, if $R_i \succ R_j$, then $Q(R_i) \geq t \cdot Q(R_j)$.

$R_i \succ R_j$ is equivalent to $R_j \in Down(R_i)$. Then, take the rank set R_k with the lowest rank such that $Down(R_k) = Down(R_i)$. By the definition of R_k, it holds that $r(R_i) \geq r(R_k)$; by Lemma 3.2, it follows

$$Q(R_i) \geq Q(R_k).$$

Additionally, $R_j \in Down(R_k)$ implies that $Down(R_k) \neq \emptyset$; then, given the definition of R_k, by property (2) of Prop. 2.9, it holds that $Down(R_{k-1}) \subset Down(R_k)$. Therefore, $Q(R_k)$ is defined by **else** of Def. 3.1 as:

$$Q(R_k) = t \cdot Q(R'_k),$$

where R'_k is the rank set with the highest rank in $Down(R_k)$. Finally, it is known that $R_j \in Down(R_k)$ and that R'_k is the rank set with the highest rank in $Down(R_k)$; it follows that $r(R'_k) \geq r(R_j)$. It follows by Lemma 3.2 that $Q(R'_k) \geq Q(R_j)$. The desired result holds:

$$Q(R_i) \geq Q(R_k) = t \cdot Q(R'_k) \geq t \cdot Q(R_j).$$

($\Leftarrow$). It has to be proven that for any two rank sets $R_i, R_j \in \mathcal{R}$, if $Q(R_i) \geq t \cdot Q(R_j)$, then $R_i \succ R_j$. The contrapositive is proven: for any two rank sets $R_i, R_j \in \mathcal{R}$, if $R_i \not\succ R_j$, then $Q(R_i) < t \cdot Q(R_j)$.

$R_i \not\succ R_j$, i.e. $R_j \notin Down(R_i)$. If $r(R_i) \leq r(R_j)$, then by Lemma 3.2 $Q(R_i) \leq Q(R_j)$, and with $t > 1$, $Q(R_i) < t \cdot Q(R_j)$, as desired. Otherwise, $r(R_i) > r(R_j)$. It immediately follows that there are at least two rank-sets, and because $\mathcal{R}$ is a partition of $\Omega_0 \subseteq \Omega$, we get that $|\Omega| \geq 2$.

Assume that $Down(R_i) = \emptyset$. Property (2) of Proposition 2.9 implies that every rank-set between R_{min} and R_i has empty $Down$-set. Thus, up to the moment of definition of $Q(R_i)$ Def. 3.1 has applied only the **if** clause, and $Q(R_i)$ is defined by the pth iteration of the **if** clause as: $Q(R_i) = Q(R_{min}) + \sum_{n=1}^{p} \epsilon_n,\ p \geq 1$. Turning to $Q(R_j)$, since $r(R_j) \geq r(R_{min})$, by Lemma 3.2 it follows that $Q(R_j) \geq Q(R_{min})$. Finally, given $|\Omega| \geq 2$ and $p \geq 1$, the following holds:

$$\sum_{n=1}^{p} \left(\frac{1}{|\Omega|}\right)^n = \frac{1 - \left(\frac{1}{|\Omega|}\right)^p}{|\Omega| - 1} < 1.$$

Given $t > 1$, we can combine the three inequalities just obtained so to prove $Q(R_i) < t \cdot Q(R_j)$, as desired:

$$\begin{aligned} Q(R_i) = Q(R_{min}) + \sum_{n=1}^{p} \epsilon_n &= Q(R_{min}) + (t-1) \cdot Q(R_{min}) \cdot \sum_{n=1}^{p} \left(\frac{1}{|\Omega|}\right)^n \\ &= Q(R_{min}) \left(1 + (t-1) \frac{1 - \left(\frac{1}{|\Omega|}\right)^p}{|\Omega| - 1}\right) < t \cdot Q(R_{min}) \leq t \cdot Q(R_j). \end{aligned}$$

After proving that $Q(R_i) < t \cdot Q(R_j)$ under the condition that $Down(R_i) = \emptyset$, we can turn to the case where $Down(R_i) \neq \emptyset$.

If $Down(R_i) \supset Down(R_{i-1})$, then $Q(R_i)$ is defined by the **else** clause of Def. 3.1 as $Q(R_i) = t \cdot Q(R_i')$ where R_i' is the rank-set with the highest rank in $Down(R_i)$. By assumption $R_i \not\succ R_j$, while $R_i \succ R_i'$; hence, $r(R_j) > r(R_i')$, from which by Lemma 3.2 $Q(R_j) > Q(R_i')$. Then, the desired result holds: $Q(R_i) = t \cdot Q(R_i') < t \cdot Q(R_j)$.

Otherwise, $Down(R_i) = Down(R_{i-1}) \neq \emptyset$. We need to go back to the last iteration of Def. 3.1 when **else** was applied. To do so, take the rank set R_k with the lowest rank such that

$Down(R_i) = Down(R_k)$. By its definition, $Down(R_k) \supseteq Down(R_{k-1})$, which guarantees that $Q(R_k)$ is defined by **else** of Def. 3.1 as $Q(R_k) = t \cdot Q(R'_k)$ where R'_k is the rank-set with the highest rank in $Down(R_k)$. By their definition, we also know that $r(R_i) > r(R_{i-1}) \geq r(R_{K+1}) \geq r(R_k)$ and that $Down(R_i) = Down(R_{i-1}) = Down(R_{k+1}) = Down(R_k)$. Notice that R_{k+1} might coincide with R_{i-1}; it is nevertheless true that $Q(R_{k+1})$ is defined by the qth iteration of **if** of Def. 3.1, for some $q \geq 0$, and because R_i and R_{k+1} have same $Down$-sets, $Q(R_i)$ is defined by the later rth iteration of **if** of Def. 3.1, with $r > q$. Therefore, $Q(R_i)$ can be defined in terms of $Q(R_k)$, and in terms of $Q(R'_k)$ at its turn, as follows:

$$Q(R_i) = Q(R_k) + \sum_{n=q}^{r} \epsilon_n = t \cdot Q(R'_k) + \sum_{n=q}^{r} \epsilon_n.$$

We still need to prove that $Q(R_i) < t \cdot Q(R_j)$. To get there, consider that $Down(R_j) = Down(R_k)$ by assumption, and because $\succ$ is a strict partial order $R_j \notin Down(R_j)$, it follows that $R_j \notin Down(R_k)$. On the contrary, by assumption $R'_k \notin Down(R_k)$. Therefore, by (2) of Prop. 2.9, $r(R_j) > r(R'_k)$. Then, $r(R_j) \geq r(R'_{K+1})$, with the two possibly coinciding, from which by Lemma 3.2 $Q(R_j) \geq Q(R'_{k+1})$. The proof of the previous Lemma, Lemma 3.2 actually proved that for every R'_k, $Q(R'_{k+1}) - Q(R'_k) \geq \epsilon_p$, where ϵ_p is the ϵ value at the time of definition of $Q(R'_k)$ with $p \geq 0$. Then, it is also true that $Q(R_j) - Q(R'_k) \geq \epsilon_p$, equivalent to $Q(R_j) \geq Q(R'_k) + \epsilon_p$. We previously assumed $Q(R_{k+1})$ to be defined by the qth iteration of **if** of Def. 3.1, with $q \geq 0$. Since $R'_{k+1} \in Down(R_{k+1})$, it follows that $r(R_{k+1}) > r(R_k) \geq r(R'_{k+1})$, with the latter two rank-sets possibly coinciding. Hence, $p < q$; since both are natural numbers, this is equivalent to $p \geq q-1$. Finally, we exploit the property of the geometrical series for which, given $q, r \in \mathbb{N}$ such that $1 \leq q \leq r$, and $|\Omega| \geq 3$:

$$\sum_{n=q}^{r} \left(\frac{1}{|\Omega|}\right)^n < \sum_{n=q}^{\infty} \left(\frac{1}{|\Omega|}\right)^n = \left(\frac{1}{|\Omega|}\right)^q \cdot \sum_{n=0}^{\infty} \left(\frac{1}{|\Omega|}\right)^n = \left(\frac{1}{|\Omega|}\right)^q \cdot \left(\frac{1}{1 - \frac{1}{|\Omega|}}\right)$$
$$= \left(\frac{1}{|\Omega|}\right)^q \cdot \left(\frac{|\Omega|}{|\Omega| - 1}\right) = \left(\frac{1}{|\Omega|}\right)^{q-1} \cdot \left(\frac{1}{|\Omega| - 1}\right) < \left(\frac{1}{|\Omega|}\right)^{q-1}.$$

Given p, q, $r \in \mathbb{N}$ such that $p < q \leq r$, and t, $Q(R_{min}) \in \mathbb{R}$ such that $t > 1$ and $Q(R_{min}) > 0$, we can obtain the desired result:

$$\begin{aligned} Q(R_i) &= t \cdot Q(R'_k) + \sum_{n=q}^{r} \epsilon_n = t \cdot Q(R'_k) + (t-1) \cdot Q(R_{min}) \cdot \sum_{n=q}^{r} \left(\frac{1}{|\Omega|}\right)^n \\ &< t \cdot Q(R'_k) + (t-1) \cdot Q(R_{min}) \cdot \left(\frac{1}{|\Omega|}\right)^{q-1} \\ &\leq t \cdot Q(R'_k) + (t-1) \cdot Q(R_{min}) \cdot \left(\frac{1}{|\Omega|}\right)^{p} = t \cdot Q(R'_k) + \epsilon_p \\ &= t \cdot \left(Q(R'_k) + \frac{\epsilon_p}{t}\right) < t \cdot \big(Q(R'_k) + \epsilon_p\big) \leq t \cdot Q(R_j). \end{aligned}$$

□

THEOREM 3.4 (Representation theorem for semi-orders). *Let Ω be a finite non-empty set. $\succ$ is a semi-order over Ω if and only if, for any real odds threshold $t > 1$, there exists some positive real valued function $Q : \Omega \to \mathbb{R}_{>0}$ which generates an odds-threshold order $\succ_Q$ such that $\succ_Q = \succ$.*

PROOF. To prove the right-to-left direction, we have to show that for any t strictly greater than 1, and any a positive function over Ω, Q, the odds-threshold order $\succ_Q$ is a semi-order. The result is obtained by showing that $\succ_Q$ has no *2v2* or *3v1* suborder.

If $|\Omega| < 4$, the property is vacuously satisfied by the definition of *2v2* and *3v1* order. Then assume $|\Omega| \geq 4$ and take any suborder $\succ'_Q$ of four atoms. The proof proceeds ad absurdum. Then assume that there exists a *2v2* or *3v1* suborder $\succ'_Q$. By cases:

2v2. $\succ'_Q = \langle \{a, \dots, d\}, \{(a, c), (b, d)\} \rangle$. Since the preference relation $\succ_Q$ is defined by the odds threshold method, the following inequalities hold:

$$Q(a) \geq t \cdot Q(c) > 0 \ \text{ and } \ Q(b) \geq t \cdot Q(d) > 0.$$

Compare $Q(a)$ and $Q(b)$: either $Q(a) \geq Q(b)$ or $Q(a) < Q(b)$.
If $Q(a) \geq Q(b)$, then $Q(a) \geq t \cdot Q(d)$, from which $a \succ_Q d$. This contradicts the initial definition of $\succ'_Q$.

Otherwise, $Q(a) < Q(b)$; then $Q(b) \geq t \cdot Q(c)$, from which $b \succ_Q c$. This contradicts as well the initial definition of $\succ'_Q$.

3v1. $\succ'_Q = \langle\{\alpha_1, \ldots, \alpha_4\}, \{(\alpha_1, \alpha_2), (\alpha_2, \alpha_3), (\alpha_1, \alpha_3)\}\rangle$. Since the preference relation $\succ_P$ is defined by the odds threshold method, the following inequalities hold:

$$P(\alpha_1) \geq t \cdot P(\alpha_2) \geq t^2 \cdot P(\alpha_3) > 0.$$

Compare $P(\alpha_2)$ and $P(\alpha_4)$: either $P(\alpha_2) \geq P(\alpha_4)$ or $P(\alpha_2) < P(\alpha_4)$.
If $P(\alpha_2) \geq P(\alpha_4)$, then $P(\alpha_1) \geq t \cdot P(\alpha_4)$, from which $\alpha_1 \succ_P \alpha_4$. This contradicts the initial definition of $\succ'_P$.
Otherwise, $P(\alpha_2) < P(\alpha_4)$; then $P(\alpha_1) \geq t \cdot P(\alpha_4)$, from which $\alpha_1 \succ_P \alpha_4$. This contradicts the initial definition of $\succ'_P$.

Therefore, any order generated by some probability distribution P via an odds threshold method, for any real $t > 1$, has no *2v2* or *3v1* suborder.

To prove the left-to-right direction, we show that for any semi-order $\succ$, there exists a real-valued function Q over Ω such that $\succ_Q = \succ$. Let $\succ$ be a semi-order over some finite non-empty Ω and choose any real value $t > 1$. Definition 3.1 defines a positive real valued function over the rank set partition of Ω, $\mathcal{R}$. Then Q is naturally extended over Ω by giving the same $Q(R_i)$ value to all the elements in R_i: $Q(a) = Q(R_i)$ if $a \in R_i$. As $\mathcal{R}$ is a partition of Ω, Q is so defined for all elements in Ω. It is left to show that the Q function generates an odds-threshold order equal to the initial semi-order. Take any two elements $a, b \in \Omega$ such that $a \in R_i$ and $b \in R_j$. By the definition of rank sets and by Lemma 3.3, we get:

$$a \succ b \;\Leftrightarrow\; R_i \succ R_j \;\Leftrightarrow\; Q(R_i) \geq t \cdot Q(R_j) \;\Leftrightarrow\; Q(a) \geq t \cdot Q(b) \;\Leftrightarrow\; a \succ_Q b$$

□

Corollary 3.5. *Let Ω be a finite non-empty set of atoms. $\succ$ is a semi-order over $\Omega_0 \subseteq \Omega$ if and only if, for any real odds threshold $t > 1$, there exists some probability function $P : 2^\Omega \to \mathbb{R}$ which generates an odds-threshold order $\succ_P$ such that $\succ_P = \succ$.*

Proof. For the right-to-left direction, choose a probability function P which assigns strictly positive values to all the atoms in Ω. As a consequence of Theorem 3.4, it holds that P induces an odds-threshold order $\succ_P$ which is a semi-order.

For the left-to right direction, assume $\succ$ to be a semi-order over $\Omega_0 \subseteq \Omega$. Consider the rank sets partition $\mathcal{R}$ of Ω_0 uniquely generated by $\succ$. Let $Q : \mathcal{R} \to \mathbb{R}$ be any real valued function over $\mathcal{R}$ defined as in Def. 3.1. First, since $\mathcal{R}$ is partition of $\Omega_0 \subseteq \Omega$, Q is extended over Ω by giving the same real value $Q(R_i)$ to all atoms $\alpha_i \in R_i$ and by assigning 0 to all atoms in Ω/Ω_0.

$$Q(\alpha_i) = \begin{cases} Q(R_j) & \text{if } \alpha_i \in R_j \\ 0 & \text{if } \alpha_i \in \Omega/\Omega_0 \end{cases}$$

Finally, to obtain a probability function, such real values are normalized by dividing them by their sum $Q(\Omega)$:

$$Q(\Omega) := \sum_{\alpha_i \in \Omega} Q(\alpha_i) = \sum_{\alpha_i \in \Omega_0} Q(\alpha_i) + \sum_{\alpha_i \in \Omega/\Omega_0} Q(\alpha_i) = \sum_{R_i \in \mathcal{R}} |R_i| \cdot Q(R_i).$$

We need to show that $Q(\Omega) \neq 0$. Since $\Omega_0 \neq \emptyset$ and $\mathcal{R}$ is a partition of Ω_0, there exists at least one rank set R_i in $\mathcal{R}$, and $|R_i| \geq 1$. Additionally, Definition 3.1 assigns $Q(R_{min})$ some strictly positive real value; it is also true that for any rank set $R_j \in \mathcal{R}$, $r(R_j) \geq r(R_{min})$. Therefore, by Lemma 3.2, it follows that for any rank set $R_j \in \mathcal{R}$, $Q(R_i) \geq Q(R_{min}) > 0$. Hence, since there exists at least one rank set R_i in $\mathcal{R}$, and $|R_i| \geq 1$, and $Q(R_j) > 0$ for any rank set $R_j \in \mathcal{R}$, it follows that $Q(\Omega) > 0$, as desired.

The function $P : 2^\Omega \to [0, 1]$ is now defined by letting:

$$P(\alpha_i) := \frac{Q(\alpha_i)}{Q(\Omega)};$$
$$P(S) := \sum_{\alpha_i \in S} P(\alpha_i), \text{ with } S \in 2^\Omega.$$

As usual, for notational simplicity, $P(\alpha_i)$ is written instead of $P(\{\alpha_i\})$. We first show that P is a probability distribution over Ω:

- $P(\alpha_i) \geq 0$, for all $\alpha_i \in \Omega$.

  It has already been established that $Q(\Omega) > 0$ and $Q(R_i) > 0$ for all rank set $R_i \in \mathcal{R}$.

Hence, for each atom $\alpha_i \in \Omega_0$, $P(\alpha_i) = \frac{Q(R_i)}{Q(\Omega)} > 0$. Additionally, for any $\alpha_i \in \Omega/\Omega_0$, $P(\alpha_i) = \frac{0}{Q(\Omega)} = 0$.

- $P(\Omega) = 1$.

  By the definition of Q and P, all atoms $\alpha_i \in R_i$ in a rank set R_i get the same real value, $Q(\alpha_i)/Q(\Omega)$. Since $\mathcal{R}$ is a partition, each atom belongs to one and one only rank set. Therefore, the following obtains:

$$P(\Omega) = \sum_{\alpha \in \Omega} P(\alpha) = \sum_{R_i \in \mathcal{R}} |R_i| \cdot \frac{Q(R_i)}{Q(\Omega)} = \frac{1}{Q(\Omega)} \cdot \sum_{R_i \in \mathcal{R}} |R_i| \cdot Q(R_i) = \frac{\sum_{R_i \in \mathcal{R}} |R_i| \cdot Q(R_i)}{\sum_{R_i \in \mathcal{R}} |R_i| \cdot Q(R_i)} = 1.$$

Finally, we show that P generates an odds threshold order which is equivalent to the original one. The initial order $\succ$ was defined over a subset Ω_0 of Ω. Coherently, the function P assigns probability zero to all $\alpha_i \in \Omega/\Omega_0$; since the definition of odds threshold order requires the probability of its elements to be strictly positive, $\succ_P$ is undefined for any $\alpha_i \in \Omega/\Omega_0$. On the contrary, for any $\alpha_i \in \Omega_0$, $P(\alpha_i) = \frac{Q(\alpha_i)}{Q(\Omega)} > 0$; therefore $\succ_P$ is defined over all v. Finally, for any two atoms $\alpha_i, \alpha_j \in \Omega_0$,

$$\alpha_i \succ \alpha_j \;\Leftrightarrow\; R_i \succ R_j \;\Leftrightarrow\; Q(R_i) \geq t \cdot Q(R_j) \;\Leftrightarrow\; \frac{Q(R_i)}{Q(\Omega)} \geq t \cdot \frac{Q(R_j)}{Q(\Omega)}$$
$$\Leftrightarrow\; P(\alpha_i) \geq t \cdot P(\alpha_j) \;\Leftrightarrow\; \alpha_i \succ_P \alpha_j.$$

□

4. A different representation of semi-orders

Theorem 3.4 that was proven in the previous section proves that $\succ$ is a semi-order is and only if it can be generated by positive real-valued function Q over its element via the odds-threshold method, which states that $a \succ_Q b$ if and only if $Q(a) \geq t \cdot Q(b)$, for some $t > 1$. Corollary 3.5 proves a stronger result by showing that $\succ$ is a semi-order if and only if it can be analogously generated by a special kind of positive real-valued functions, that is, a probability function P. In the present section, we consider another representation theorem for semi-orders, originally proven by Patrick Suppes and Joseph L. Zinnes in [SZ62]. Their representation result shows that all and only the semi-orders over Ω can be generated by a real-valued function u, a utility function, when it is assumed that u induces a preference relation over Ω in terms of a fixed difference as follows: for any $a, b \in \Omega$, $a \succ_Q b$ if and only if $Q(a) \geq Q(b) + \delta$, where δ is some positive real value. Their result relies on a definition of semi-orders first given by Duncan Luce in [Luc56]. Before presenting their result, we prove the Luce's definition of semi-order is equivalent to ours.

4.1. Luce's definition of semi-orders. Definition 4.1 gives the definition of semi-orders introduced by Luce. While our definition relies on forbidden configurations, Luce's definition of semi-orders captures a notion of preference which allows indifference to be intransitive. Indeed, in Definition 4.1 the indifference relation $\sim$ is only assumed to be reflexive and to satisfy axioms S3 and S4, which are compatible with intransitivity.

DEFINITION 4.1 (Semi-order by Luce). *Let $\succ$ and $\sim$ be two relations over Ω. Then $\langle \succ, \sim \rangle$ is a semi-order if for every $a, b, c, d \in \Omega$ the following axioms are satisfied:*

S1 exactly one of $a \succ b$, $b \succ a$ or $a \sim b$ obtains,

S2 $a \sim a$,

S3 $a \succ b$, $b \sim c$, $c \succ d$ imply $a \succ d$,

S4 $a \succ b$, $b \succ c$, $b \sim d$ imply not both $a \sim d$ and $c \sim d$.

Despite the differences between the two definitions of semi-orders, it is possible to establish a correspondence between axioms S3 and S4 respectively with the *2v2* and *3v1* forbidden configurations. As a reminder, in Definition 2.2 *filtered orders* are characterized as all strict partial orders for which the *2v2* is a forbidden configuration. Likewise, Definition 2.3 states that *mezzanine orders* have the *3v1* as forbidden configuration. In the following, Proposition 4.2 shows that an order is filtered if and only if it satisfies axioms S1, S2 and S3. On the other hand, Proposition 4.3 shows that if an order is mezzanine, then it satisfies axioms S1, S2 and S4, but the converse does not hold. In order for that to be the case, axioms S1, S2 and S4 are insufficient; we additionally have to assume $\succ$ to be transitive. The reason for such asimmetry is that, as proven in Proposition 4.2, axiom S3 on its own implies transitivity, thus it does not have to be assumed for the case of filtered orders. Instead, axioms S1, S2 and S4 together do not imply transitivity, which has so to be assumed. Finally, Proposition 4.4 proves that the two definitions are equivalent.

PROPOSITION 4.2. *An order $\succ$ over Ω is a filtered order if and only if $\succ$ satisfies axioms S1, S2, S3.*

PROOF.
(If) The axioms in Def. 4.1 do not assume a, b, c, d to be distinct. Thus, assume for S3 b and c to coincide. S3 implies that, if $a \succ b$ and $b \succ d$, then $a \succ d$. Thus, from S3 we get that $\succ$ is a transitive relation. S2 holds that $a \sim a$, from which by S1 $a \not\succ a$, making $\succ$ irreflexive. Because $\succ$ is both transitive and irreflexive, $\succ$ is a strict partial order. Finally, S3 directly implies that $\succ$ has no *2v2* suborder.
(Only if) Being $\succ$ a strict partial order, S1 and S2 hold. By the premises of S3, a is distinct from all b, c, d; likewise, b and c from d. If b and c are the same, by transitivity $a \succ d$ holds as desired. Then, assume them distinct. By reduction, assume $a \not\succ d$, from which $a \sim d$ since $d \succ a$ would be in contradiction with the premises of S3. Since $c \succ d$, $a \not\succ c$ follows from $a \sim d$; and since $a \succ b$, it follows $b \not\succ d$. From the first two precises of S3, $a \succ b$ and $b \sim c$, we also know $c \not\succ a$, hence $a \sim c$; from the latter two, $d \not\succ b$, hence $b \sim d$. An instance of a *2v2* order has thus been obtained. Thus, axioms S3 follows from the definition of filtered order. □

PROPOSITION 4.3. *An order $\succ$ over Ω is a mezzanine order if and only if $\succ$ is transitive and satisfies axioms S1, S2, S4.*

PROOF.
(If) In the previous proposition, axioms S1 and S2 have been proven to imply $\succ$ to be irreflexive. As $\succ$ is also assumed to be transitive, $\succ$ is a strict partial order. Finally, assuming transitivity of $\succ$, S4 shows that, give a chain of three elements $a \succ b \succ c$, any forth element d has to be $\succ$-related to a or c. Thus, $\succ$ does not have any *3v1* suborder.
(Only if) Being $\succ$ a strict partial order, S1 and S2 hold and $\succ$ is transitive. Then, take $a, b, c, d \in \Omega$ such that they satisfy the premises of S4, namely $a \succ b$, $b \succ c$ and $b \sim d$. By

transitivity, also $a \succ c$ holds. By reduction, assume that both $a \sim d$ and $c \sim d$. An instance of a *3v1* suborder has already been obtained, in contradiction with the hypotheses. Thus, S4 follows from the definition of mezzanine orders. □

PROPOSITION 4.4. *An order $\succ$ over Ω is a semi-order if and only if it satisfies axioms S1-S4.*

PROOF. The proof is an immediate consequence of the proofs of Propositions 4.2 and 4.3. □

4.2. The second representation theorem: weak-orders instead of rank functions. In order to prove Theorem 3.4, we introduced the notion of *rank* by defining the rank function associated with a semi-order. The importance of the rank function is that it allows to establish a ranking even among elements which are unrelated with respect to $\sim$. So it refines the initial semi-order and allows to determine which elements are actually indistinguishable. By grouping together all the elements in Ω with the same rank into the rank-sets, the initial semi-order is simplified enough that it is possible to recursively define the Q function (Definition 3.1), and so to prove both Theorem 3.4 and Corollary 3.5.

In the case of Luce a similar approach is taken. In order to prove the Representation Theorem 4.9, it is always required to determine what are the elements that are indistinguishable with respect to the semi-order. This allows the semi-order to be simplified and a real-valued function to be recursively defined on it, as we did in Definition 3.1. In their case however, instead of using a rank function which ranks the elements by giving them some value, the authors use the qualitative analogous of a ranking, namely a weak-order.

DEFINITION 4.5 (Weak order). *Let $>$ and $\approx$ be two relations over the set Ω. The pair $\langle >, \approx \rangle$ is a weak order if the preference relation $>$ and the indifference relation $\approx$ satisfy conditions W1-W3:*

W1 for every $a, b \in \Omega$, exactly one of $a > b$, $b > a$, or $a \approx b$ obtains;

W2 $\approx$ is an equivalence relation;

W3 $>$ is transitive.

Weak-orders capture the notion of a ranking over a set. In Definition 2.5 the rank function is generated by some preferential order. Likewise, in Definition 4.6 Luce defines the relation $\langle >, \approx \rangle$, induced by a preferential order $\succ$. Theorem 4.7 proves that, when $\succ$ is a semi-order, then the relation $\langle >, \approx \rangle$ given by Definition 4.5 is a weak-order. This shows that Definition 4.6 correctly captures the ranking over Ω induced by a semi-order. Despite the fact that the two formalisms of the rank function and the weak-order play an analogous role, the qualitative one has the advantage that it does not rely on counting the relations to establish the ranking over Ω, which the rank function does. As a result, when we consider cases where Ω is infinite, the ranking function can be undefined and so inapplicable, while the qualitative notion of weak-order still applies. Such result allows the proof of Theorem 4.9.

DEFINITION 4.6 ([Luc56] Def. 1). *The relation $\langle >, \approx \rangle$ induced on Ω by a given relation $\langle \succ, \sim \rangle$ on Ω is defined as follows. Let $a > b$ if and only if one of the following holds:*

(1) $a \succ b$,

(2) $a \sim b$ and there exists $c \in \Omega$ such that $a \sim c$ and $c \succ b$,

(3) $a \sim b$ and there exists $d \in \Omega$ such that $a \succ d$ and $d \sim b$.

If neither $a > b$ not $b > a$, then $a \approx b$.

THEOREM 4.7 ([Luc56] Thm. 1). *$\succ$ is a semi-order if and only if $\succ$ is transitive and $\langle >, \approx \rangle$ is a weak order.*

Theorem 4.7 shows that the definition of semi-order is adequate to capture preferences with an intransitive indifference relation, which was Luce's interest. It does so by proving that any semi-order can be related to a weak-order representing the appropriate corresponding ranking of its elements. After showing such correspondence, the final step is to show any semi-order can be represented by some utility function, i.e. some real-valued function over Ω. On that note, Luce introduces a *just noticeable difference function* (j.n.d.) which varies with the individual elements of the set Ω: as the name suggests, such function captures for each magnitude the interval inside of which another magnitude would be judged as indistinguishable from the former. Luce does indeed prove that for any semi-order, there exists a utility function over

Ω and a j.n.d. function over Ω able to represent the semi-order. However, the j.n.d. is a function of the elements of Ω, thus the j.n.d. depends on the single values assigned by the utility function. Thus, Luce fails to give a *fixed numerical interpretation* of the relations $\succ$ and $\sim$ for all elements of Ω.

The existence of such fixed numerical interpretation is proven in *Basic measurement theory* by Patrick Suppes and Joseph L. Zinnes [SZ62]. Such numerical interpretation is captured by the notion of *numerical semi-order*, formally stated in Definition 4.8. Finally, Theorem 4.11 denotes with $\Omega/\approx$ the quotient of Ω with respect to the relation $\approx$.

DEFINITION 4.8 ([SZ62]). *The binary system $\langle R, \succ_\delta \rangle$ is a numerical semi-order if and only if R is a set of real numbers, $\delta \in \mathbb{R}_{>0}$ and the relation $\succ_\delta$ is the binary relation satisfying the following property: for all $x, y \in R$, $x \succ_\delta y$ if and only if $x > y + \delta$.*

THEOREM 4.9 ([SZ62] Thm. 11). *Let $\succ$ be a semi-order over Ω and let $\Omega/\approx$ be a finite set. Then $\succ$ over $\Omega/\approx$ is isomorphic to some numerical semi-order.*

The proof of Theorem 4.9 shows something slightly more specific, which is the statement given in Theorem 4.10. In his *Measurement Structures and Linear Inequalities*, Dana Scott discusses this more emphatic, though equivalent, version of the result ([Sco64] Problem 1 and Theorem 2.1) with 1 instead of δ as value for the j.n.d.

THEOREM 4.10 ([Fis70] Thm. 3). *If the order $\succ$ over Ω is a semi-order and $\Omega/\approx$ is finite, then there is a real-valued function u on Ω such that, for all a and b in Ω,*

$$a \succ b \text{ if and only if } u(a) > u(b) + 1.$$

The two theorems 4.9 and 4.10 prove the existence of a fixed numerical interpretation of the preferential relation $\succ$. Not just that, they show that the preferences expressed by a semi-order can be represented numerically if and only if there exists a real value δ which defines the agent's *just noticeable difference* and the agent prefers an option to another if and only if the value assigned to the former is at least δ-more intense than the value assigned to the latter.

4.3. A representation theorem for filtered orders. We conclude the section by presenting one final result. In [Fis70] Peter C. Fishburn proves another result which is relevant to the present work. This new theorem focuses on a generalization of semi-orders, namely filtered orders. The result shows that such generalization does not come without a prize when it comes to real-valued representability. Even though it is still possible to numerically represent an interval order, the preferential relation looses the fixed numerical interpretation which is possible for the case of semi-orders. Instead, we go back to the kind of representation initially given by Luce for semi-orders: the preference relation is still defined in terms of difference between the real values assigned to the options, but such difference is not generally fixed and it is a function of the elements of the order.

THEOREM 4.11 ([Fis70] Thm. 4). *If $\succ$ over Ω is an interval order and $\Omega/\approx$ is countable, then there are real-valued functions u and ρ with $\rho(a) > 0$ for each $a \in \Omega$ such that, for all $a, b \in \Omega$,*

$$a \succ b \quad \textit{if and only if} \quad u(a) > u(b) + \rho(b)$$

5. Conclusion

The aim of the present chapter was to give a qualitative characterization of the preferential orders that can be generated by a probability function via the odds-threshold method. Theorem 3.4 proved that such characterization is possible, and that an order can be defined via the odds-threshold method if and only if it is what is called a semi-order, viz. any strict partial order with the *2v2* and *3v1* as forbidden configurations. So, we are now able to qualitative characterize what our preferences look like when we prefer one elementary outcome to another if and only if the probability of former is sufficiently greater than the probability of the latter.

In our case *sufficiently greater* was interpreted in a multiplicative sense as t times more likely. This is not the only option. In the literature, Suppes and Zinnes [SZ62] proved another representation theorem of this kind for semi-orders, in which however the *sufficiently greater* was interpreted in an additive fashion as greater at least of some fixed quantity δ. The two results shed light on what are the properties of our preferences, especially because it is often the case in everyday scenarios that we choose one option to another only if the evaluation that we make of the former is sufficiently greater than the latter.

Giving a qualitative characterization of the class of preferential orders obtained by the odds-threshold method is essential for what comes in the next chapter. We will study how our beliefs change after learning some new evidence. In particular, we will assume that we believe true one and only world among the maximal elements of our preferential order, and we will also assume that, when new evidence is learnt, we eliminate all the alternatives which are inconsistent with it while keeping our preferences unchanged. Under such conditions, each of the two forbidden configurations which characterize semi-orders will result is a corresponding rule of belief revision satisfied by our beliefs. The resulting belief revision system satisfied by our belief is logically strong. This indicates that, even though comparing alternatives can be costly, it allows to generate believes which are very robust to new evidence. All of this will completely rely on the results that were obtained in the present chapter.

CHAPTER 2

Reasoning with semi-orders

1. Introduction

The previous chapter gave a qualitative characterization of odds-threshold orders, viz. the preferential orders that can be obtained by comparing the probabilities assigned to the atoms that compose the order, and preferring one atom to another if and only if the former is at least t-times more likely than the latter. Theorem 3.4 proves the equivalence between the class of odds-threshold orders and semi-orders. These are the strict partial orders which have two forbidden configurations, the *2v2* and the *3v1*, depicted in Figure 1.

(A) *2v2*. (B) *3v1*.

FIGURE 1. The *2v2* and *3v1* order configurations.

In this chapter, we shift our attention towards the logical significance of such orders. Specifically, in the following we will consider semi-orders over sets of alternative and exclusive hypotheses, or worlds. We will study preferential orders without taking into account any real-valued function that generates the order. Nevertheless, in accordance with the representation theorem given in the previous chapter, the preference relation $\alpha_i \succ \alpha_j$ of a semi-order is to be interpreted here as α_i *is a hypothesis sufficiently more plausible than* α_j. For each order $\succ$, the agent believes that the actual hypothesis is among the $\succ$-maximal worlds. As a result, the strongest proposition that the agent believes in is the disjunction of the $\succ$-maximal elements.

In addition, in this chapter we will assume that, when a new piece of information A is learned, the agent only eliminates from the preferential order $\succ$ all the worlds that are inconsistent with A while keeping the rest of $\succ$ unchanged. In other words, we will assume that, when some evidence A is learned, the preferential order is updated by quotient $\succ$ over A, in symbols $\succ |A$. Once the preferential order is updated, the new strongest logical proposition the agent believes in is the disjunction of the new $\succ |A$-maximal element. This mechanism of revision allows each semi-order $\succ$ to define a nonmonotonic consequence relation $A \mid\sim_{\succ} B$, which states that B *is believed to be true given that* A *is true*. Let us formally define $\mid\sim_{\succ}$.

Let $\mathcal{SL}$ refer to the set of well-formed formulas over a set of propositional variables, using the classical propositional connectives $\neg, \wedge, \vee, \rightarrow$ and $\leftrightarrow$. As for the semantics of such language, each world α_i corresponds to one complete assignment of truth values to the propositional variables. Assuming a finite number of propositional variables, and so a finite language $\mathcal{SL}$, we get a finite number of worlds. A world satisfies a formula $\alpha_i \Vdash A$ as usual: $\alpha_i \Vdash A \vee B$ if and only $\alpha_i \Vdash A$ or $\alpha_i \Vdash B$, and $\alpha_i \Vdash \neg A$ if and only if *not* $\alpha_i \Vdash A$. We also assume the notation $\widehat{A}$ to refer to the set of worlds that satisfy A. Finally, $Cn(A)$ denotes the set of formulas that classically follow from A. Fixed a language $\mathcal{SL}$, a preferential order $\succ$ over the set of worlds Ω defines a nonmonotonic consequence relation $\mid\sim_{\succ}$ as follows: a formula B is a nonmonotonic consequence of A, i.e. $a \mid\sim_{\succ} B$, if all the maximal worlds of $\succ |A$ satisfy B, that is, if all the worlds in $max(\succ |A)$ are included in $\widehat{B}$, the set of worlds satisfying B.

DEFINITION 1.1. *Let $\mathcal{SL}$ a finite language, let Ω be the associated set of worlds and $\succ$ a preferential order over Ω. The nonmonotonic consequence relation $\mid\sim_{\succ}$ is defined as follows:*

$$A \mid\sim_{\succ} B \quad \textit{if and only if} \quad max(\succ |A) \subseteq \widehat{B}.$$

$C_{\succ}(A)$ denotes the set of all formulas B that preferentially follow from A, i.e. $A \mid\sim_{\succ} B$.

By relying on the just presented framework, in this chapter we will consider what logical properties are satisfied by the class of all nonmonotonic consequence relations that can be defined by semi-orders. Having the *2v2* and having the *3v1* as forbidden configuration are independent order-theoretic properties; thus, we will treat them separately. For the *2v2* forbidden configuration, it was already proven in [Fre93] Theorem 2.4 a correspondence between filtered orders and the nonmonotonic consequence relations satisfying system **P** from [KLM90] together with the rule of *Disjunctive Rationality*. For the *3v1* forbidden configuration instead, as a new result, Theorem 3.11 will prove a correspondence between mezzanine orders and the new rule *Plausible Monotonicity*. The two result will be combined to obtained in Theorem 4.6 a correspondence between semi-orders and the nonmonotonic system **D**, equivalent to system **P** together with both Disjunctive Rationality and Plausible Monotonicity.

2. Representing nonmonotonic consequence relations

The present chapter strongly relies on the construction and the results first proposed by Michael Freund in his 1993 paper *Injective Models and Disjunctive Relations* [Fre93]. The paper is a continuation of the work on nonmonotonic consequence relations started by S. Kraus, D. Lehmann and M. Magidor [KLM90]. The main result of [Fre93] characterizes the kind of consequence relations, defined on a language $\mathcal{SL}$, that can be defined by an order on the set of worlds associated with $\mathcal{SL}$. Here we only consider cases with a finite number of worlds so we can also assume a finite propositional language $\mathcal{SL}$. Under such condition, [Fre93] show that a nonmonotonic consequence relation $\mid\sim$ can be represented by a preferential model if and only if it satisfies the following condition $\star$:

$$C(A \vee B) \subseteq Cn(C(A) \cup C(B)) \qquad (\star)$$

Such result is formally stated in the following:

THEOREM 2.1 ([Fre93] Thm. 4.13). *Let $\mathcal{SL}$ a finite language and $\mid\sim$ a preferential inference relation on $\mathcal{SL}$. Then, $\mid\sim$ is represented by an injective model if and only if $C(A \vee B) \subseteq Cn(C(A) \cup C(B))$ for any pair of formulas A and B.*

One could think of an *injective model* as just an order on the set of worlds associated with $\mathcal{SL}$. To see why such models are called *injective*, we introduce a bit more background. Back to [KLM90], the models that are used to define a preferential consequence relation $\mid\sim$

are triples of the following kind, $\langle S, \succ, l \rangle$, referred to as a *preferential model*, where S is a set of *states*, $\succ$ is a preferential order over S, and $l : S \to \Omega$ is a *labelling function* that assigns to each state a world associated with the language $\mathcal{SL}$. In [KLM90] the authors prove (see here Theorem 2.2) a one to one correspondence between preferential models and the preferential consequence relations satisfying the rules of the nonmonotonic system **P**. Here we give the canonical definition of **P**: the six rules allow to derive all the valid rules of **P**.

Reflexivity Axiom (Ref) $\quad A \mid\sim A,$

Logical Left Equivalence (LLE) $\quad \dfrac{\models A \leftrightarrow B \quad A \mid\sim C}{B \mid\sim C}$

Right Weakening (RW) $\quad \dfrac{\models B \to C \quad A \mid\sim B}{A \mid\sim C}$

Cautious Monotonicity (CMon) $\quad \dfrac{A \mid\sim B \quad A \mid\sim C}{A \wedge B \mid\sim C}$

Or $\quad \dfrac{A \mid\sim C \quad B \mid\sim C}{A \vee B \mid\sim C}$

And $\quad \dfrac{A \mid\sim B \quad A \mid\sim C}{A \mid\sim B \wedge C}$

THEOREM 2.2 (Representation Theorem for Preferential Relations – [KLM90] Thm. 5.18). *A consequence relation is a preferential consequence relation iff it is defined by some preferential model.*

Yet, even under the restriction on the labeling function to assign to each state a single world, one world can be assigned to multiple states. Thus, it is not always possible to translate a strict partial order over the set of states into a strict partial order over the set of worlds, since irriflexivity might fail. A trivial example is given in Figure 2. In order for that to be the

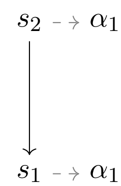

FIGURE 2. A strict partial order over the set of states allows a world to be preferred to itself.

case, the labeling function l has to be injective, that is, each world can at most be assigned to a single state. Such property gives the name to *injective* models. Under such condition, the set of states constituting the model can always be taken to be a subset of Ω, and the order relation to be the image of $\succ$ by l. While this is not in general the case for any propositional language $\mathcal{SL}$, if we restrict to finite languages, injective models have the additional property that for any formula $A \in \mathcal{SL}$, the maximal worlds wrt $\succ$ in $\widehat{A}$ are exactly those that satisfy $C(A)$, the nonmonotonic consequences of A for that model. Such property characterizes in general what Freund calls *standard models*. In the case of a finite language then, injective and standard models are equivalent notion, while in general this is not the case – see [Fre93] p. 236 for an example.

2.1. Filtered models and Disjunctive Rationality. The restriction to injective models comes with a price. As mentioned before, not all preferential consequence relations can be represented by an injective model, that is, a model on worlds. Assuming again a finite language, only those consequence relations satisfying condition $\star$ are representable by an injective model. In such case, by combining the Theorems 2.1 and 2.2, we get that a nonmonotonic consequence relation satisfies the rules of system **P** together with $\star$ if and only if it is representable by a strict partial order over the set of ($\mid\!\sim$-consistent) worlds.

One relevant application of Theorem 2.1 concerns *disjunctive* consequence relations, viz. all $\mid\!\sim$ satisfying the rules of the system **P** together with *Disjunctive Rationality* (DRat). If a consequence relation satisfies Disjunctive Rationality, it automatically satisfies $\star$ – this is an immediate consequence of [Fre93] Theorem 4.11 and Lemma 5.3. As a result, we obtain Theorem 2.4.

DEFINITION 2.3. *A disjunctive consequence relation is a preferential consequence relation that satisfies Disjunctive Rationality:*

(DRat) $$\frac{A \vee B \mid\!\sim C \quad A \not\mid\!\sim C}{B \mid\!\sim C}$$

THEOREM 2.4 ([Fre93] Thm. 5.2). *A nonmonotonic consequence relation* $\mid\!\sim$ *is disjunctive if and only if it can be defined by a standard filtered order.*

In addition, precisely for the just mentioned reason, condition $\star$ does not have to be assumed for $\mid\!\sim$ in order for it to be represented by a *standard* filtered model, and a fortiori to be represented by a generic filtered model on a set of states. This is going to come up later again in the two representation theorems that are proven. Theorem 4.1 does not require $\mid\!\sim$ to satisfy $\star$ as this directly follows from Disjunctive Rationality. On the other hand, Theorem 3.11 has to do so since the new rule of inference *Plausible Monotonicity* (PlaMon) together with system **P** do not imply condition $\star$. In the context of the nonmonotonic system **P**, PlaMon and $\star$ are independent. Therefore, in order to obtain a correspondence between preferential consequence relations satisfying Plausible Monotonicity and the class of standard mezzanine models condition $\star$ has to be imposed on $\mid\!\sim$.

3. Mezzanine orders and Plausible Monotonicity

So far we have focused on filtered models and disjunctive consequence relations. As it was previously shown, for a model being filtered is equivalent to not having any suborder with the forbidden configuration *2v2*. However, the object of study of the present section are *semi-orders*. Thus, we are also interested in the additional forbidden configuration *3v1*, whose absence characterizes *mezzanine orders*. In the following, we give a logical characterization of the nonmonotonic consequence relations characterized by mezzanine orders. Such consequence relations are defined by the logical property *Plausible Monotonicity* together with the rules of the nonmonotonic system **P**.

DEFINITION 3.1 (Pla-monotonicity). *An pla-monotone consequence relation $\mid\!\sim$ is a preferential consequence relation that satisfies Plausible Monotonicity (Mon):*

(PlaMon) $$\frac{A \mid\!\sim C,\ A \not\mid\!\sim \neg B,\ A \wedge B \mid\!\sim D}{A \wedge B \wedge C \mid\!\sim D}$$

The name of the rule Plausible Monotonicity was inspired by [Spo88], where a notion of plausibility is formally defined. Without entering into the details of the definition, it suffices to know that in [Spo88] one formula C is more plausible than another formula B if and only if $\max\{r(\alpha_i) : \alpha_i \Vdash C\} > \max\{r(\alpha_i) : \alpha_i \Vdash B\}$, that is, the maximal rank of the worlds satisfying A is greater or equal than the maximal rank of the worlds satisfying B. Then, the first two conditions of PlaMon, $A \mid\!\sim C$ and $A \not\mid\!\sim \neg B$, guarantee that, in the context of A, C is at least as plausible as B, or equivalently, that B is not more plausible than C. To see this, we turn to the semantics. By the second condition, there exists at least one world among the $\succ$-maximal worlds satisfying A that also satisfies B, call it α_i. Because $A \mid\!\sim C$, we know that all the $\succ$-maximal worlds satisfying A also satisfy C; thus, also α_i satisfies C. If $r(\alpha_i)$ is the highest rank among the maximal worlds satisfying A, then B and C would be equally plausible. Otherwise, there exists a distinct world $\alpha_j \in \max(\succ |A \wedge C)$ such that $r(\alpha_j) > r(\alpha_i)$, making C more plausible than B in the context of C. Thus, the first two conditions imply that, give A, C is at least as plausible as B. Given such conditions, PlaMon states that, in the context of A, all the nonmonotonic consequences D of B must be preserved if we learn that C also holds. When satisfied, the rule PlaMon implies that the consequence relation is monotone to at least as plausible premises. More emphatically, the rule states that the defeasible consequences are stable to new premises that were initially at least as plausible.

3.1. Soundness. Plausible Monotonicity is a variation of the rule Rational Monotonicity (RMon), initially discussed in [KLM90].

(RMon) $$\frac{A \mid\!\sim C,\ A \not\mid\!\sim \neg B}{A \wedge B \mid\!\sim C}$$

Intuitively, Rational Monotonicity says that we should not change our beliefs when we learn something not completely unexpected, something consistent with our initial beliefs. In the

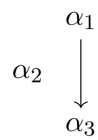

FIGURE 3. Ranked models and the *2v1* forbidden configuration.

context of system **P**, Rational Monotonicity both implies Disjunctive Rationality and Plausible Monotonicity, as shown below in Proposition 4.2. To prove however that Disjunctive Rationality and Plausible Monotonicity together are not enough for Rational Monotonicity to hold we rely on the following results. In [LM92], Theorem 3.12, the authors prove that a preferential consequence relation additionally satisfies Rational Monotonicity (rational) if and only if it is defined by some ranked model. In their case, ranked models are defined over sets of states. Nevertheless, [Fre93] Theorem 6.3 extends that result to show that every rational consequence relation can be defined by a strict partial order over sets of worlds.

DEFINITION 3.2 (Ranked model). *A ranked model is a preferential model* $\langle \Omega, \succ \rangle$ *for which the strict partial order* $\succ$ *is modular, viz.* $\succ$ *is such that for every* $\alpha_1, \alpha_2, \alpha_3 \in \Omega$*, if* $\alpha_2 \sim \alpha_3$ *and* $\alpha_1 \succ \alpha_3$*, then* $\alpha_1 \succ \alpha_2$*.*

THEOREM 3.3 ([Fre93] Thm. 6.3). *The standard model of a disjunctive relation* $\mid\!\sim$ *is ranked if and only if* $\mid\!\sim$ *is rational.*

One can think of ranked models as strict partial orders with the forbidden configuration *2v1*, shown in Figure 3. Any such suborder would directly falsify the condition given in Definition 3.2.

Later in Section 4.1 we will discuss in detail the relation of Rational Monotonicity with Disjunctive Rationality and Plausible Monotonicity. It is enough for the moment to notice that, were we to assume the *2v1* as a forbidden configuration, both the *2v2* and the *3v1* configurations would be forbidden as well. In light of Theorems 2.4 and 3.11 proven at the end of this section, this is expected since Rational Monotonicity implies both Disjunctive Rationality and Plausible Monotonicity. To see that the converse does not hold instead, it is enough to notice that there exists a preferential model validating Disjunctive Rationality and Plausible Monotonicity, but invalidating Rational Monotonicity (see Proposition 4.3); such countermodel would have no *2v2* or *3v1* suborders, but at least one *2v1* suborder. For this, we still need to prove the correspondence between mezzanine orders and Plausible Monotonicity. As a first step, Lemma 3.5 proves the soundness of Plausible Monotonicity with respect to mezzanine orders over sets of worlds. To ease the reading of the proof, notice that the two properties in Proposition 3.4 hold for all preferential orders; the proof is omitted.

PROPOSITION 3.4. *Let $\succ$ be a preferential order over a set Ω. Then, the two following properties hold. For any $\alpha \in \Omega$,*

(1) if $\alpha \in \max(\succ) \cap B$, then $\alpha \in \max(\succ |A \wedge B)$;
(2) if $\alpha \notin \widehat{A}$, then $\alpha \in \widehat{\neg A}$.

LEMMA 3.5 (Soundness). *Let $\succ$ be a mezzanine order over a finite set Ω. Then, $\succ$ defines a preferential consequence relation $\mid\sim$ which satisfies Plausible Monotonicity.*

PROOF. By contraposition, it has to be proven that in order to invalidate PlaMon a model needs to have a 3v1 suborder, shows in Figure 4. To invalidate Plausible Monotonicity, assume that the model satisfies all the premises but not the conclusion of the rule. By the negated conclusion of PlaMon, $max(\succ |A \wedge B \wedge C) \nsubseteq \widehat{D}$, from which there exists $\alpha_3 \in max(\succ |A \wedge B \wedge C) \cap \widehat{\neg D}$. By the third premise, $max(\succ |A \wedge B) \subseteq \widehat{D}$; therefore, $\alpha_3 \notin max(\succ |A \wedge B)$. Then there exists $\alpha_2 \in max(\succ |A \wedge B)$ s.t. $\alpha_2 \succ \alpha_3$. Additionally, since $\alpha_2 \succ \alpha_3$ and $\alpha_3 \in max(\succ |A \wedge B \wedge C)$, it must be the case that $\alpha_2 \in \widehat{\neg C}$. Since, by the first premise, $max(\succ |A) \subseteq \widehat{C}$ and $\alpha_2 \in max(\succ |A \wedge B) \cap \widehat{\neg C}$, it holds that $\alpha_2 \notin \max(\succ |A)$. Then, there exists $\alpha_1 \in max(\succ |A)$ s.t. $\alpha_1 \succ \alpha_2$. Additionally, since $\alpha_1 \succ \alpha_2$ and $\alpha_2 \in max(\succ |A \wedge B)$, then $\alpha_1 \notin max(\succ |A \wedge B)$; but, given $\alpha_1 \in max(\succ |A)$, it follows that $\alpha_1 \in max(\succ |A) \cap \widehat{\neg B}$. By the second premise, $max(\succ |A) \nsubseteq \widehat{\neg B}$. Hence, there exists α_4 s.t. $\alpha_4 \in max(\succ |A) \cap \widehat{B}$. Then, by the first premise it follows that $\alpha_4 \in max(\succ |A) \cap \widehat{B} \cap \widehat{C}$; therefore, $\alpha_4 \in max(\succ |A \wedge B)$ and $\alpha_4 \in max(\succ |A \wedge B \wedge C)$. Also, by the third premise and $\alpha_4 \in max(\succ |A \wedge B)$, it follows that $\alpha_4 \in max(\succ |A \wedge B) \cap \widehat{D}$. To recap then, $\alpha_1 \in \widehat{\neg B}$ and $\alpha_4 \in \widehat{B}$, from which

$$
\begin{array}{ccc}
\alpha_1 \Vdash A \wedge \neg B & & \\
\downarrow & & \\
\alpha_2 \Vdash A \wedge B \wedge \neg C \wedge D & \qquad & \alpha_4 : A \wedge B \wedge C \wedge \neg D \\
\downarrow & & \\
\alpha_3 \Vdash A \wedge B \wedge C \wedge \neg D & &
\end{array}
$$

FIGURE 4. The soundness of PlaMon for Mezzanine Models.

α_1 and α_4 are distinct; also, $\alpha_1, \alpha_4 \in max(\succ |A)$, from which $\alpha_1 \sim \alpha_4$. Likewise, $\alpha_2 \in \widehat{\neg C}$, $\alpha_4 \in \widehat{C}$ and $\alpha_2, \alpha_4 \in max(\succ |A \wedge B)$; hence, α_2 and α_4 are distinct and $\alpha_2 \sim \alpha_4$. Likewise again, $\alpha_3 \in \widehat{\neg D}$, $\alpha_4 \in \widehat{D}$ and $\alpha_3, \alpha_4 \in max(\succ |A \wedge B \wedge C)$; hence, α_3 and α_4 are distinct and $\alpha_3 \sim \alpha_4$. It was also established that $\alpha_1 \succ \alpha_2 \succ \alpha_3$. An instance of a 3v1 order has thus been established, as desired. □

3.2. Representing Plausible Monotonicity. We have just proven that every mezzanine order over a set of worlds induces a nonmonotonic consequence relation which satisfies PlaMon. The rest of the present section proves the converse, that to satisfy PlaMon a model has to be a mezzanine order. In order to do so, we are going to use a construction introduced by Freund in his [Fre93]. This construction allows the author to prove Theorems 2.1 and 2.4. The general idea is as follows: any consequence relation $\mathrel{|\!\sim}$ induces a transitive relation on the language $\mathcal{SL}$, which can in turn define an order on the set of worlds Ω. This relation captures the notion of normality induced by $\mathrel{|\!\sim}$ over the language and over the set of worlds. This strategy does not always generate a preferential model representing the *pla-monotone* $\mathrel{|\!\sim}$. It does so only for those $\mathrel{|\!\sim}$ that satisfy condition $\star$, the necessary and sufficient condition for their representability. As a first step, the notion of A^+ set is introduced.

DEFINITION 3.6 ($\mathrel{|\!\sim}$-consistency). *A formula $A \in \mathcal{SL}$ is $\mathrel{|\!\sim}$-consistent if and only if $A \not\mathrel{|\!\sim} \bot$.*

DEFINITION 3.7. *Let A be a $\mathrel{|\!\sim}$-consistent formula. Then, A^+ denotes the set of all $B \in \mathcal{SL}$ such that $A' \mathrel{|\!\sim} B$ for every classical consequence of A, $A' \in Cn(A)$.*

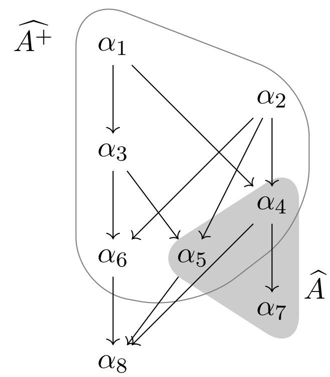

FIGURE 5. Example of an A^+ set in a preferential model with $\widehat{A} = \{\alpha_4, \alpha_5, \alpha_6\}$.

Definition 3.7 states that a propositional formula B belongs to the set A^+ if and only if B is a nonmonotonic consequence of all those formulas logically weaker than A. This property can be interpreted as follows: $B \in A^+$ if and only if we are already certain that B holds and no information weaker than A would change our mind. The set A^+ has properties very similar to $C(A)$, the set of nonmonotonic consequences of A. In a preferential model, one could think of the set A^+ not as the formulas satisfied by the most preferred worlds satisfying A, but rather as the formulas satisfied by the set of all those worlds α_i unrelated or preferred to the maximal worlds satisfying A. Figure 5 shows an example: given the set $\widehat{A} = \{\alpha_4, \alpha_5, \alpha_6\}$, the corresponding A^+ is the set of formulas satisfied by the worlds in the outlined area, i.e. the worlds from α_1 to α_6. The property satisfied by all the worlds in $\widehat{A^+}$ is that none of them is defeated by any of the maximal elements in $\widehat{A}$. In this way, one can see how the A^+ sets, syntactically defined by the initial consequence relation $\mid\!\sim$, allow to establish an order among the worlds in a model.

Definition 3.7 characterizes the A^+ sets straight from the $\mid\!\sim$ consequence relation. In [Fre93] however, to define A^+, the author first introduces an order $>_{\mid\sim}$ induced by $\mid\!\sim$ over the language $\mathcal{SL}$ as follows:

$$A >_{\mid\sim} B \quad \text{if and only if} \quad A \vee B \mid\!\sim \neg B$$

The relation states that A is more normal than B if and only if B is logically inconsistent with the nonmonotonic consequences of $A \vee B$. Semantically speaking, this is the case if and only if none of the most plausible $A \vee B$ worlds actually satisfy B. We do not show here that the relation $>_{\mid\sim}$ so defined is transitive. For the proof and for additional properties of $>_{\mid\sim}$, we refer to [Fre93] Lemmata 4.1 and 4.2. In the present setting, it suffices to know that every transitive relation $>_{\mid\sim}$ results in the set A^+ by the following definition. For any $A \in \mathcal{SL}$, A^+ is the set of formulas B whose negations are less normal than A:

$$B \in A^+ \quad \text{if and only if} \quad A >_{\mid\sim} \neg B$$

As one of the nice properties of the set A^+, D. Lehmann and M. Freund show ([Fre93] Lemma 4.4.) that this definition of A^+ is equivalent to the one given in Definition 3.7.

Moving on, Figure 6 considers three main cases, which will further clarify Definition 3.7 and its role in establishing the Representation Theorems 3.11 and 4.1. As was already mentioned, A^+ can be understood as the set of all those propositions that we already believe to be true, such that no information weaker than A will change our mind. To see this characterization from a semantic point of view in a preferential order, it suffices to consider three cases of new pieces of information that are not weaker than A. The first case is in Figure 6a: the weaker piece of evidence A', with $A \vDash A'$, is also satisfied by a world α_8 that is defeated by at least one of the worlds in $max(\succ |A)$, here both $\alpha_4 \succ \alpha_8$ and $\alpha_5 \succ \alpha_8$. In this case, this world will remain defeated for every piece of evidence weaker than A and, as a result, it will never satisfy all the formulas in A^+. Coherently, in Figure 6a the worlds α_8 and α_7 do not belong to the shaded circle. As a second case, consider Figure 6b and the worlds like α_6 that are unrelated to *all* the worlds in $\max(\succ |A)$, α_4, α_5. The reason why such kind of worlds satisfy the set of formulas A^+ is that there exists at least one A', for instance A s.t. $\widehat{A'} = \widehat{A} \cup \{\alpha_6\}$, such that, were we to update on A', then α_6 would be among the maximal elements, i.e. $max(\succ |A')$ and α_6 would so satisfy all the nonmonotonic consequences of A', $\alpha_6 \Vdash C(A')$. Since A^+ is the set of all formulas D such that $A' \mid\sim D$ for all $A \vDash A'$, we get that $A^+ \subseteq C(A')$; hence, $\alpha_6 \Vdash A^+$. Finally, as a third case, consider Figure 6c where $\widehat{A'} = \{\alpha_3, \alpha_4, \alpha_5, \alpha_7\}$, and α_3 is a world that is preferred to at least one world in $max(\succ |A)$, here both $\alpha_3 \succ \alpha_5$ and $\alpha_3 \succ \alpha_4$. Since in this case $\alpha_3 \in max(\succ |A')$, there is at least one A' such that $\alpha_3 \Vdash C(A')$. Hence, again by Definition 3.7, all the formulas in A^+ are satisfied by α_3.

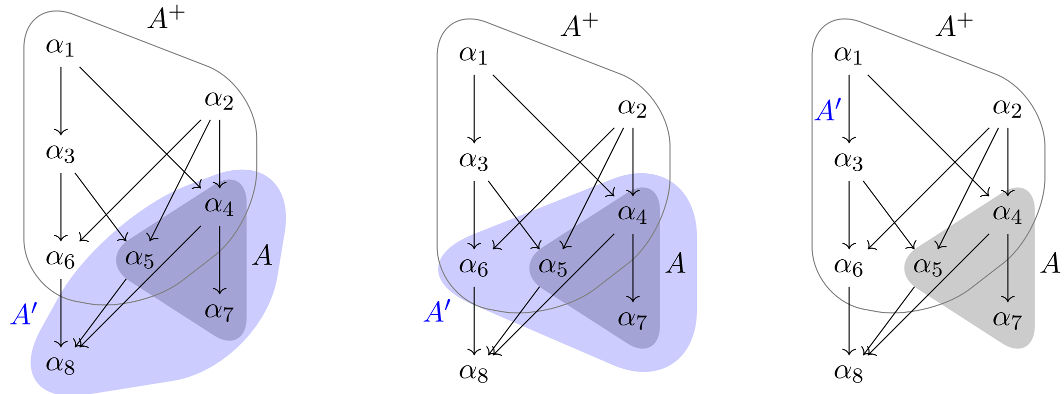

(A) Any world defeated by the ones in $max(\succ^A)$ does not satisfy A^+.

(B) The worlds similar to the ones in $\mathsf{B}(A)$ belong to $\bigcap A^+$.

(C) Example of an A^+ set in a preferential model with $A = \{\alpha_4, \alpha_5, \alpha_6\}$.

FIGURE 6

3.3. Inducing a preferential order from a consequence relation. The three cases that were just discussed and that are shown in Figure 6, are simple as for all of them A' is just A together with an additional world, of different kind for each of them. Yet, they cover all importance cases, and the more complex instances can be reduced to them. Now that the notion of A^+ set has been sufficiently elaborated, we can move on to consider why such sets are useful for proving the Representation Theorems 2.1, 2.4 already shown, together with 3.11 and 4.1 later. For all these cases, both a Soundness and a Completeness results between a class of consequence relations and a class of models have to be established. As is usually the case, proving the Soundness result is more straightforward, since we just need to prove that all the models under consideration induce a consequence relation which satisfies all the desired properties. The Completeness result, on the other hand, needs a bit more labor as we need to prove that for every consequence relation satisfying a set of desired properties, all the models representing such consequence relation are of the desired kind. In order to do so, we assume a strategy that generates a model corresponding to the initial $\mid\sim$. The A^+ sets that we have discussed serve the role of generating this model, inducing a preference relation over the set of worlds, starting from a syntactically defined nonmonotonic consequence relation $\mid\sim$. Specifically, $\mid\sim$ defines a preferential relation over Ω by the following relation:

(Def. $\succ_{\mid\sim}$) $\quad \alpha_i \succ_{\mid\sim} \alpha_j$ iff $\alpha_i \Vdash \neg A$ for all $\mid\sim$-consistent formulas A such that $\alpha_j \Vdash A^+$

In the model induced by $\mid\sim$, a world is preferred to another one if and only if the former falsifies every formula A for which the latter satisfies the set A^+. To get an idea why this is an adequate way of defining $\succ_{\mid\sim}$ starting from a nonmonotonic consequence relation $\mid\sim$, it is enough to consider two cases. First, consider a preferential order where two worlds are unrelated, like α_3 and α_4 in Figure 7. It is easy to see that under such premise there exists at least one formula – in the finite case, we can always take the disjunction of the maximal assignments assigned to the two worlds, here $\alpha_3 \vee \alpha_4$ – such that one of the two worlds satisfies A and the other A^+. Thus, such preferential order would be incompatible with a consequence relation where $\alpha_3 \Vdash \neg A$ for all $\mid\sim$-consistent formulas such that $\alpha_4 \mid\sim A^+$. Secondly, consider a preferential order with two worlds α_3 and α_5 such that $\alpha_3 \succ \alpha_5$ like in Figure 7b. Under such conditions, every formula A_i for which α_4 satisfies A_i^+, like the three examples in Figure 7b A_1, A_2, A_3, is such that A_i must exclude α_3 from the order, that is, $\alpha_3 \Vdash \neg A$. Thus, such preferential order would be incompatible with any nonomonotonic consequence relation for which there exists one formula A such that $\alpha_3 \Vdash A$ while $\alpha_5 \Vdash A^+$.

3.4. Representability of finite consequence relations. The two examples in Figure 7 are meant just to give a semantic intuition of why Def. $\succ_{\mid\sim}$ is a proper way to induce the preferential order $\succ_{\mid\sim}$ from $\mid\sim$. Indeed, we discussed before how a preferential consequence relation has to satisfy condition $\star$ in order to be represented by a preferential order over a set of worlds. Thus, for any of those $\mid\sim$ which do not satisfy condition $\star$, Def. $\succ_{\mid\sim}$ defines a relation over the set of worlds that is not a transitive irreflexive relation, in other words, it does not constitute a preferential model. This result is stated in Theorem 2.1.

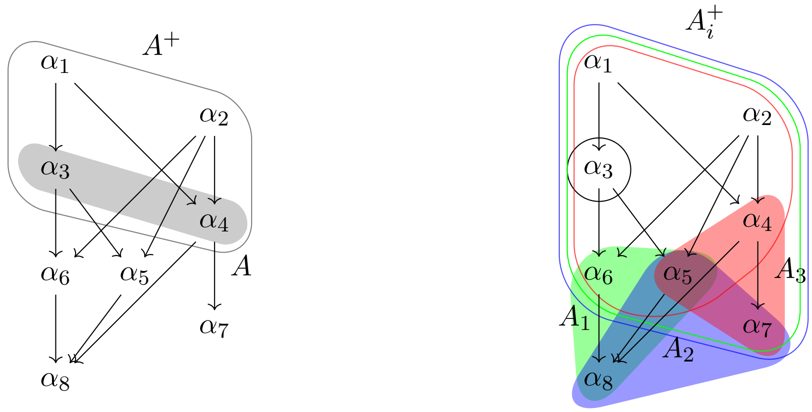

(A) An example of A that proves $\alpha_3 \not\mid\!\sim \alpha_4$.

(B) The worlds similar to the ones in $C(A)$ belong to A^+.

FIGURE 7

THEOREM 2.1 ([Fre93] Thm. 4.13). *Let $\mathcal{SL}$ be a finite language, $\mid\!\sim$ be a preferential consequence relation on $\mathcal{SL}$. Then $\mid\!\sim$ is represented by an injective model if and only if $C(A \vee B) \subseteq Cn(C(A) \vee C(B))$ for every pair of formulas A, B.*

We gave a thorough presentation of the notion of A^+ set because the proofs of Theorems 2.1 and 2.4 and of the following Theorems 3.11 and 4.1 rely on them in order to construct a preferential model representing a belief revision operator which satisfies the properties under consideration. In particular Lemmas 3.9 and 3.10 rely on the same construction, essential for proving the desired result. First, Lemma 3.8 states three properties of the A^+ sets, respectively corresponding to properties (5), (4) and (3) of [Fre93] Lemma 4.5. Properties (1) and (3) are quite straightforward. Remember that A^+ is the set of those beliefs which the agent is already certain of and which no information weaker than A would change their mind about. Hence, from $A \subseteq B$, we get (1) $A^+ \subseteq B^+$. And (3) directly follows. On the other hand, property (2) interestingly shows how the preferential consequences of A, $C(A)$, are embedded within A^+. Again, the set A^+ is defined by the preferential consequences of all those pieces of evidence A' weaker than A. For each of such A' we can apply a Reasoning by

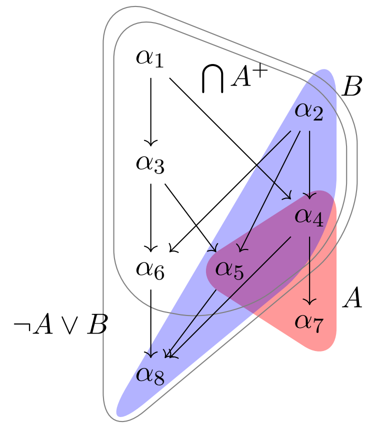

FIGURE 8. Example of an A^+ set in a preferential model with $A = \{\alpha_4, \alpha_5, \alpha_6\}$.

Cases on $A' \equiv A \vee (\neg A \wedge A')$ as follows. It holds that $A \mid\sim B$ by assumption. If A is the case, then B holds as well by $A \mid\sim B$, and by RW so does $A \to B$; on the other hand, if $\neg A \wedge A'$ holds, by Ref we get $\neg A \wedge A' \mid\sim \neg A$, and again by RW we obtain $A \to B$. Reasoning by Cases is a sound way of reasoning in the nonmonotonic system **P**. Then, as both cases conclude that $A \to B$ for every A' weaker than A, we know that $A \to B$ is in A^+.

Semantically, consider again our example in Figure 8. As shown before, A^+ is the set of maximal elements in $\succ |A$ together with all those worlds that are unrelated with all of the maximal elements in $\succ |A$ (like α_6) or preferred to at least one of them (like α_8). All worlds of the latter kind clearly falsify A, otherwise they would belong to the former kind. At the same time, we said that the other worlds satisfying A^+ are the maximal ones in $\succ^A$, and from $A \mid\sim B$ we know that they satisfy B. Therefore, we know that all the worlds satisfying A^+ belong to $\widehat{\neg A \vee B}$, that is, to $\widehat{A \to B}$; thus, $A \to B \in A^+$, as desired.

LEMMA 3.8. *Let A^+ be the set defined as in Definition 3.7. Then A^+ satisfies the following properties:*

(1) *if $\models A \to B$, then $A^+ \subseteq B^+$;*
(2) *$A \mid\sim B$ iff $A \to B \in A^+$;*

(3) $A^+ \cup B^+ \subseteq (A \vee B)^+$.

LEMMA 3.9. *A consequence relation* $\mid\!\sim$ *is U-monotone if and only if for any* $\mid\!\sim$*-consistent formulas* A, B *s.t.* $A \mid\!\sim C$ *and* $A \not\mid\!\sim \neg B$, $(A \wedge B)^+ = (A \wedge B \wedge C)^+$.

PROOF. That $(A \wedge B)^+ \supseteq (A \wedge B \wedge C)^+$ follows from $A \wedge B \wedge C \vDash A \wedge B$ by (1) of Lemma 3.8. To prove the other direction, assume that a formula $D \in (A \wedge B)^+$. This implies that $(A \wedge B) \vee \neg D \mid\!\sim D$, from which by LLE we get $(A \vee \neg D) \wedge (B \vee \neg D) \mid\!\sim D$.

By assumption, $A \mid\!\sim C$; by property (3) of Lemma 3.8, this holds if and only if $A \to C \in A^+$. By the definition of A^+ (Def. 3.7), it follows that $A \vee \neg D \mid\!\sim A \to C$. From $A^+ \supseteq (A \wedge B)^+$ and $D \in (A \wedge B)^+$, we get that $A \vee \neg D \mid\!\sim D$, from which $A \vee \neg D \mid\!\sim A$. By AND and RW, from the two $A \vee \neg D \mid\!\sim C$, and again by RW $A \vee \neg D \mid\!\sim C \vee \neg D$ follows.

By assumption, $A \not\mid\!\sim \neg B$ holds; by CMon, it follows that $A \vee \neg D \not\mid\!\sim \neg B$, and by RW, $A \vee \neg D \not\mid\!\sim \neg(B \vee \neg D)$. We can now apply Plausible Monotonicity and obtain the following: $(A \vee \neg D) \wedge (B \vee \neg D) \wedge (C \vee \neg D) \mid\!\sim D$; by LLE, $(A \wedge B \wedge C) \vee \neg D \mid\!\sim D$. So, we get $D \in (A \wedge B \wedge C)^+$. The desired equality is thus obtained. □

LEMMA 3.10. *Let* $\mid\!\sim$ *be a U-monotone consequence relation and let* A *be a* $\mid\!\sim$*-consistent formula such that* $C \in A^+$ *and* $\neg B \notin A^+$. *Then* $B^+ = (B \vee \neg C)^+$.

PROOF. By the definition of $A^+ = \{D : A' \mid\!\sim D \text{ for all } A \vDash A'\}$, we can derive from the first premise $C \in A^+$ that $A \vee B \vee \neg C \mid\!\sim C$; by RW, it follows that $A \vee B \vee \neg C \mid\!\sim B \vee C$. Suppose by reductio that $\neg(B \vee \neg C) \in A^+$. By the definition of A^+ and RW, we know that A^+ is closed under supsets. Therefore, from $\neg(B \vee \neg C) \in A^+$ it follows that $\neg B \in A^+$, which contradicts the second hypothesis. Therefore, we know that $\neg(B \vee \neg C) \notin A^+$. Then $A \vee B \vee \neg C \not\mid\!\sim \neg(B \vee \neg C)$. We can now apply Lemma 3.9: $((A \vee B \vee \neg C) \wedge (B \vee \neg C)^+ = ((A \vee B \vee \neg C) \wedge (B \vee \neg C) \wedge (B \vee C))^+$. We simplify both sides and obtain the desired result: $(B \vee \neg C)^+ = B^+$. □

THEOREM 3.11. *A preferential consequence relation* $\mid\!\sim$ *satisfies* $\star$ *and PlaMon if and only if it can be defined by a mezzanine order.*

$$C(A \vee B) \subseteq Cn(C(A) \cup C(B)) \qquad (\star)$$

(PlaMon) $$\frac{A \mid\!\sim C,\ A \not\mid\!\sim \neg B,\ A \wedge B \mid\!\sim D}{A \wedge B \wedge C \mid\!\sim D}.$$

Proof. If (soundness): it was proven in Lemma 3.5.
Only if (completeness):
$\mid\!\sim$ is assumed to satisfy condition $\star$. Theorem 2.1 is proven precisely by using the construction given by Def. $\succ_{\mid\sim}$. Thus, it guarantees that the canonical model defined by Def. $\succ_{\mid\sim}$ from $\mid\!\sim$ constitutes a preferential model which represents $\mid\!\sim$. It is left to show that, if $\mid\!\sim$ satisfies Plausible Monotonicity, then the same canonical model generated by Def. $\succ_{\mid\sim}$ is a mezzanine order. So, we have to prove that for any four $\alpha_1, \alpha_2, \alpha_3, \alpha_4 \in \Omega$ such that $\alpha_1 \succ_{\mid\sim} \alpha_2 \succ_{\mid\sim} \alpha_3$, $\alpha_1 \sim_{\mid\sim} \alpha_4$ and $\alpha_2 \sim_{\mid\sim} \alpha_4$, it holds that $\alpha_4 \succ_{\mid\sim} \alpha_3$. As $\alpha_1 \sim_{\mid\sim} \alpha_4$ there exists $A \in \mathcal{SL}$ s.t. $\alpha_1 \Vdash A$ and $\alpha_4 \Vdash A^+$. Additionally, as $\alpha_1 \succ_{\mid\sim} \alpha_2$ we know that $\alpha_2 \nVdash A^+$; then $\alpha_2 \vDash \neg C$ for some $C \in A^+$. Suppose for reduction that $\alpha_4 \sim_{\mid\sim} \alpha_3$; there exists $B \in \mathcal{SL}$ s.t. $\alpha_3 \Vdash B^+$ and $\alpha_4 \Vdash B$. Since $\alpha_4 \Vdash A^+ \wedge B$, we know that $A^+ \nvDash \neg B$, from which $\neg B \notin A^+$. From $C \in A^+$ and $\neg B \notin A^+$ it follows by Lemma 3.10 that $B^+ = (B \vee \neg C)^+$. Then, $\alpha_3 \Vdash (B \vee \neg C)^+$; but $\alpha_2 \Vdash \neg C$, so $\alpha_2 \Vdash B \vee \neg C$. Therefore, we have that $\alpha_2 \not\succ_{\mid\sim} \alpha_3$, contradicting the hypothesis. □

4. Semi-orders and System D

Theorem 3.11 proves a correspondence between mezzanine orders, i.e. those models with the *3v1* orders as a forbidden configuration, and those preferential consequence relations $\mid\!\sim$ satisfying Plausible Monotonicity, together with the representability condition $\star$. This is the main new result given in the present section. However, our focus is not just on mezzanine orders, but on semiorders. Semiorders have two forbidden configurations, *3v1* and *2v2*. If having the *3v1* as a forbidden configuration results in Plausible Monotonicity being a valid rule of inference, we also know that having the *2v2* as a forbidden configuration results in *Disjunctive Rationality* being valid – see Theorem 2.4. In addition, notice that being a filtered order and being a mezzanine order are two independent notions. Trivially indeed, a *2v2* order such as the one in Figure 9a is a mezzanine order as it has no *3v1* suborders; likewise, the order in Figure 9b is a filtered order.

(A) *2v2*. (B) *3v1*.

FIGURE 9. The *2v2* and *3v1* order configurations.

As a consequence, Theorems 2.4 and 3.11 imply that the two rules of inference Disjunctive Rationality and Plausible Monotonicity are logically independent, just as their respective order-theoretic characterizations. Therefore, the correspondences DRat-filtered models, and the one PlaMon-mezzanine models can be combined in a modular fashion. Theorem 4.1 is the result of such process. It proves that a nonmonotonic consequence relation satisfies system **D**, viz. the set of inference rules of system **P** together with Disjunctive Rationality and Plausible Monotonicity, if and only if it can be defined by a filtered model. Notice that Theorem 4.1 does not mention condition $\star$. The reason for this is that Disjunctive Rationality already implies $\star$, making any preferential consequence relation defined over a finite language and satisfying Disjunctive Rationality representable by a finite preferential model. Since system **D** includes DRat, $\star$ is automatically satisfied by $\mathrel{|\!\sim}$ and can so be dropped.

THEOREM 4.1 (Representation Theorem for system **D**). *A nonmonotonic consequence relation* $\mathrel{|\!\sim}$ *satisfies system* **D** *if and only if it can be defined by a semiorder.*

4.1. System D and *Rational Monotonicity*. Theorem 4.1 proves and Table 1 shows that the consequence relations induced by the class of semiorders satisfy all the rules of system **P** together with *Disjunctive Rationality* and *Uniderminer Monotonicity*. Table 1 also shows how the consequence relation does not satisfy *Rational Monotonicity* (RMon).

(RMon) $$\frac{A \mathrel{|\!\sim} C,\ A \not\mathrel{|\!\sim} \neg B}{A \wedge B \mathrel{|\!\sim} C}$$

Systems			Rules	Conditions on $\mid\!\sim$	
			LLE	$A \mid\!\sim C, \models A \leftrightarrow B \;\Rightarrow\; B \mid\!\sim C$	
			RW	$A \mid\!\sim B, \models B \rightarrow C \;\Rightarrow\; A \mid\!\sim C$	
			Ref	$A \mid\!\sim A$	
		P	And	$A \mid\!\sim B,\ A \mid\!\sim C \;\Rightarrow\; A \mid\!\sim B \wedge C$	✓
			CMon	$A \mid\!\sim B,\ A \mid\!\sim C \;\Rightarrow\; A \wedge B \mid\!\sim C$	
			Or	$A \mid\!\sim C,\ B \mid\!\sim C \;\Rightarrow\; A \vee B \mid\!\sim C$	
	D		DRat	$A \vee B \mid\!\sim C,\ A \not\mid\!\sim C \;\Rightarrow\; B \mid\!\sim C$	
			PlaMon	$A \mid\!\sim C,\ A \not\mid\!\sim \neg B,\ A \wedge B \mid\!\sim D \;\Rightarrow\; A \wedge B \wedge C \mid\!\sim D$	
R			RMon	$A \mid\!\sim C,\ A \not\mid\!\sim \neg B \;\Rightarrow\; A \wedge B \mid\!\sim C$	×

TABLE 1. Rules of inference satisfied by the class of semiorders.

Such rule was first introduced by Krauss, Lehmann and Magidor in [KLM90].[1] It suggests that, were we to learn something that we considered possible, that was consistent with our prior beliefs, then we should keep all our priors beliefs, that is, none of the prior conclusions should be given up. In [LM92] the authors argue that "a reasonable nonmonotonic inference procedure should define a rational relation,"[2] i.e. in our terms, it should define a consequence relation satisfying Rational Monotonicity. Despite their effort, the rule is very strong, possibly too strong, and during the years its plausibility has strongly decreased. If taken into consideration together with the rule *Cautious Monotonicity*, we get that no piece of evidence which is expected (CMon) or even only consistent (RMon) with our prior beliefs, should change our

[1]See [KLM90] p. 197.

[2][LM92] p. 1.

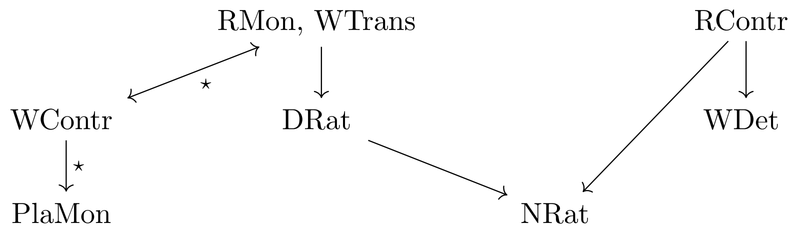

FIGURE 10. The relation between inference rules. ⋆ refers to properties that hold only on a finite language.

mind on any of our prior beliefs. In other words, we conclude that we are entitled to change our beliefs if and only if we learn some new fact that is inconsistent with our initial beliefs.

The strength of Rational Monotonicity can also be appreciated from the fact that it implies many rules of inference which have been studies in the nonmonotonic logics and belief revision literature. Figure 10 summarizes many of the results from [FLM91], Sections 2-3, [BMP97] Sections 3 and 7, and [LM92], Section 3. Notice on Figure 10 that Rational Monotonicity implies *Weak Contraposition* (WContr), *Negation Rationality* (NRat), and, more relevantly to the present case, both *Disjunctive Rationality* (DRat) and *Plausible Monotonicity* (PlaMon). In the following we give two proofs showing that Rational Monotonicity implies Disjunctive Rationality and Plausible Monotonicity.

(NRat) $$\frac{A \mid\!\sim C,\ A \wedge \neg B \mid\!\not\sim C}{A \wedge B \mid\!\sim C}$$

(WContr) $$\frac{A \wedge B \mid\!\sim C,\ A \mid\!\not\sim C}{A \wedge \neg C \mid\!\sim \neg B}$$

PROPOSITION 4.2. *In the context of system* $\mathbf{P}$*, RMon implies NRat, WContr, DRat and PlaMon.*

PROOF.

RMon implies NRat.

1.	$A \mid\sim C$	Assumption
2.	$A \wedge \neg B \not\mid\sim C$	Assumption
3.	$A \not\mid\sim \neg B$	1,2; CMon
4.	$A \wedge B \mid\sim C$	1,3; RMon

RMon implies WContr.

1.	$A \wedge B \mid\sim C$	Assumption
2.	$A \not\mid\sim C$	Assumption
3.	$A \mid\sim B \to C$	1; S
4.	$\vdash_{CPL} (B \to C) \to (\neg C \to \neg B)$	CPL
5.	$A \mid\sim \neg C \to \neg B$	3, 4; RW
6.	$A \wedge \neg C \mid\sim \neg C \to \neg B$	4, 2; RMon
7.	$A \wedge \neg C \mid\sim \neg C$	Ref, RW
8.	$A \wedge \neg C \mid\sim \neg C \wedge (\neg C \to \neg B)$	6, 7; And
9.	$A \wedge \neg C \mid\sim \neg B.$	7; RW

RMon implies DRat: by contraposition,

1.	$A \not\mid\sim C$	Assumption
2.	$B \not\mid\sim C$	Assumption
3.	$\vdash_{CPL} A \leftrightarrow (A \wedge (A \vee B))$	CPL
4.	$A \wedge (A \vee B) \not\mid\sim C$	3, 1; LLE
5.	Either $A \vee B \not\mid\sim \neg A$ or $A \vee B \mid\sim \neg A$	By cases
6.	$A \vee B \not\mid\sim \neg A$	Case 1
7.	$A \vee B \not\mid\sim C$	6, 4; RMon (contraposed)
8.	$A \vee B \mid\sim \neg A$	Case 2
9.	Either $A \vee B \mid\sim C$ or $A \vee B \not\mid\sim C$	By cases
10.	If $A \vee B \mid\sim C$, then $B \mid\sim C$	From **P**
11.	$A \vee B \not\mid\sim C$	

RMon implies PlaMon.

1.	$A \mid\sim C$	Assumption
2.	$A \not\mid\sim \neg B$	Assumption
3.	$A \wedge B \mid\sim D$	Assumption
4.	$A \wedge B \mid\sim C$	1,2; RMon
5.	$A \wedge B \wedge C \mid\sim D$	3,4; CMon

□

In [FLM91] Theorem 3.8 the authors proves that in the context of system **P**, Weak Contraposition and Disjunctive Rationality together imply Rational Monotonicity. As stated

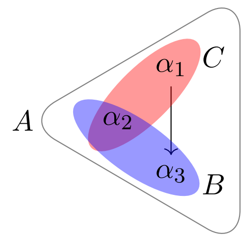

FIGURE 11. DRat and PlaMon do not imply RMon, a countermodel.

by the following proposition, the same does *not* hold if we assume Disjunctive Rationality and Plausible Monotonicity. If a consequence relation satisfies the rules of system **P** together with Disjunctive Rationality and Plausible Monotonicity, it does not necessarily satisfy Rational Monotonicity as well.

PROPOSITION 4.3. *In the context of system* **P**, *neither DRat, nor PlaMon, nor the two together imply RMon.*

In the light of Theorems 2.4 and 3.11, to prove Proposition 4.3 is enough to show that there exists a filtered mezzanine order which is not ranked. Figure 11 gives a countermodel. Clearly the model is a case of semiorder as it is constituted of only three worlds; therefore, no *2v2* nor *3v1* suborders are presents. Consequently, by Theorems 2.4 and 3.11 the model satisfies both Disjunctive Rationality and Plausible Monotonicity. At the same time however the maximal elements in $\widehat{A}$ satisfy C and at least one of them satisfies B, namely α_2, thus both $A \mid\sim C$ and $A \not\mid\sim \neg B$ hold; nevertheless, $A \wedge B \not\mid\sim C$ holds since not all the maximal elements in $\widehat{A \wedge B}$ satisfy C, specifically α_3 does not. Therefore, the model falsifies Rational Monotonicity while validating both Disjunctive Rationality and Plausible Monotonicity. Then it holds that the latter two rules do not imply Rational Monotonicity.

PROPOSITION 4.4. *In the context of system* **P**, *DRat implies NRat while PlaMon does not imply NRat.*

PROOF.

DRat implies NRat.

1.	$A \not\mid\!\sim C$	Assumption
2.	$A \wedge \neg B \not\mid\!\sim C$	Assumption
3.	$\vdash_{CPL} A \leftrightarrow ((A \wedge B) \vee (A \wedge \neg B))$	CPL
4.	$(A \wedge B) \vee (A \wedge \neg B) \mid\!\sim C$	2, 1; LLE
5.	$A \wedge B \mid\!\sim C$	4, 2; DRat

PlaMon does not imply NRat; in Figure 12 a countermodel is given where PlaMon holds while NRat is falsified. □

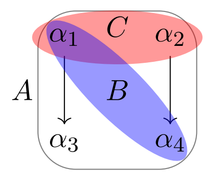

FIGURE 12. PlaMon does not imply NRat.

It was proven that Disjunctive Rationality implies Negation Rationality. As a consequence, any model that falsifies the latter also falsifies the former. Figure 12 constitutes an example where both Disjunctive Rationality and Negation Rationality are falsified. Consider Negation Rationality: all the elements of the order satisfy A, and it is also true that its maximal elements α_1, α_3 satisfy C; indeed, it holds that $A \mid\!\sim C$. However, when we restrict the order to the $A \wedge \neg B$ worlds, we get α_2 and α_3: both of them are maximal wrt to $\succ^{A \wedge \neg B}$ but α_3 does not satisfy C. So we get $A \wedge \neg B \not\mid\!\sim C$. Likewise, if we restrict to $A \wedge B$, we get the worlds α_1 and α_4, but α_4 is maximal and does not satisfy C; therefore, $A \wedge B \not\mid\!\sim C$. Negation Rationality is so falsified. We know already that Disjunctive Rationality and Plausible Monotonicity are logically independent by their order-theoretic representation. We also know that Negation Rationality follows from Disjunctive Rationality. Therefore, Negation Rationality cannot imply Plausible Monotonicity. At the same time, we just showed that Plausible Monotonicity does not imply Negation Rationality either. Hence, the two rules come out as logically independent, as shown in Figure 10.

Finally, we consider the relation between Plausible Monotonicity and Weak Contraposition. We already know that both are implied by Rational Monotonicity. In [FLM91] Theorem 3.8 the authors prove that Weak Contraposition and Disjunctive Rationality together imply Rational Monotonicity, while we showed before that even in the finite case Disjunctive Rationality and Plausible Monotonicity do not imply Rational Monotonicity. Therefore, it is clear that Plausible Monotonicity does not imply Weak Contraposition.

Theorem 4.5 ([FLM91] Thm. 3.10). *Let W be a preferential model in which, for every state, there exists a proposition true only in this state. If the preferential relation defined by W satisfies Weak Contraposition, then it satisfies Rational Monotonicity.*

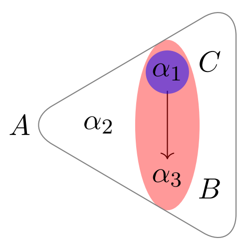

Figure 13. DRat and PlaMon do not imply RMon, a countermodel.

For what concerns the converse implication, that is, whether Weak Contraposition implies Plausible Monotonicity, the following result comes in handy since the present work assumes a finite language. Assuming a finite language means having a finite number of propositional variables $p_1, \ldots, p_n$. Each world corresponds to a unique assignment of truth values *true* or *false* to each propositional variable. Therefore, for each propositional variable p_i, each world satisfies either of the literals p_i or $\neg p_i$. As this truth assignment is unique for each world, there is a proposition corresponding to each world given by the maximal conjunction of literals satisfied by that world. Then, if we have a preferential model defined on state which is also assumed to be injective as in the present case – this is because here we consider strict partial orders over sets of worlds and not over sets of states – we get that here the initial condition of Theorem 4.5 is satisfied. Hence, given a finite language, Weak Contraposition implies Rational Monotonicity. To see this consider Figure 13: it shows that, assuming that each state uniquely satisfies a formula, any countermodel to Rational Monotonicity is a countermodel to Weak Rationality. The model represents a *2v1* order, which by Definition 3.2 and Theorem 3.3 implies that the model falsifies Rational Monotonicity. At the same time, it is true for the model that $A \wedge B \mathrel{|\!\sim} C$ and $A \mathrel{|\!\not\sim} C$, but $A \wedge \neg C \mathrel{|\!\not\sim} \neg B$, thus falsifying Weak Contraposition. It follows by transitivity that Weak Contraposition implies Plausible Monotonicity (and Disjunctive Rationality). Finally, we conclude by saying that we were not able to determine whether Weak Contraposition implies Plausible Monotonicity in general; the reason for this is that we are not aware of any order-theoretic characterization of the class of models satisfying Weak Contraposition. Further work needs to be done in such respect.

5. Conclusion

In the present section we studied the logical significance of representing our preferences over a set of alternative hypotheses in terms of semi-orders. Because semi-orders in our case are defined over sets of worlds, firstly we presented Theorem 2.1 by Freund which gives the condition $\star$ for the representability of a nonmonotonic consequence relation restricted to a finite language with respect to a finite preferential model over worlds. This means that not any nonmonotonic consequence relation, even if it is defined on a finite language, can be defined by a preferential model over worlds; that can be the case if and only if $\mathrel{|\!\sim}$ satisfies condition $\star$.

As an important application of such result, Freund proves that $\star$ is implied by the rule Disjunctive Rationality. Therefore, every nonmonotonic consequence relation satisfying DRat can also be represented by an order over worlds. In the same (Theorem 2.4 it is also proved that a preferential consequence relation satisfies Disjunctive Rationality if and only if it can be generated by a filtered order over a set of worlds, where a preferential consequence relation is any $\mathrel{|\!\sim}$ which satisfies the rules of system **P**. The main result of the present chapter is constituted by Theorem 3.11 which proves an analogous result but for Plausible Monotonicity and mezzanine orders. Specifically, Theorem 3.11 proves that a preferential consequence relation is defined by a mezzanine order over worlds if and only if it satisfies Plausible Monotonicity together with $\star$. Finally, we combined the two results to prove in Theorem 4.1 that a preferential consequence relation satisfies Disjunctive Rationality and Plausible Monotonicity if and only if it can be defined by a semi-order. In such case, since Disjunctive Rationality is assumed, condition $\star$ can be dropped as implied by the former. We were so able to give a sound and complete logical characterization of the nonmonotonic consequence relations defined by the class of semi-orders.

Such theorems will be the grounding work for what comes in the next chapter where we will generalize the kinds of update that our preferences might go through when a new piece of evidence is acquired. Those will not simply eliminate the worlds which are inconsistent with the evidence, but rather reshape the preferential model accordingly. Such generalization will allow to combine the theorems proved in this chapter with the representation results given in the first one. In particular, we will study how our preferences change when the underlying probability function that generates them is updated on some piece of evidence. We will consider two probabilistic updates, namely Bayesian conditioning and Jeffrey conditioning. The goal of the next chapter is to study what is the impact that these two probabilistic updates have on our preferences and so on our beliefs.

CHAPTER 3

Reasoning with uncertain evidence

1. Introduction

In the first chapter we showed the relation between preferences and semi-orders. In Theorem 3.4 we proved the correspondence between semi-orders and odds-threshold orders. Such result states that an order describing the agent's preferences can be generated by comparing the probabilities of the alternatives, and preferring one alternative over another if and only if the former is at least t-times more likely than the latter, if and only if it is such order is a semi-order, that is, it does not have any *2v2* or *3v1* suborders.

The second chapter focuses on the logical properties satisfied by the nonmonotonic consequence relations defined by the class of semi-orders. First of all, semi-orders are a special kind of strict partial orders. Therefore, the nonmonotonic consequence relations induced by such orders satisfy the nonmonotonic system **P**. In addition, because semi-orders have no *2v2* suborder, then such nonmonotonic consequence relations also satisfy the rule Disjunctive Rationality.

(Disjunctive Rationality) $$\frac{A \vee B \mid\!\sim C,\ B \not\mid\!\sim C}{A \mid\!\sim C}$$

The real novelty of the last chapter was proving an analogous result for the *3v1* forbidden configuration. By using the same construction introduced by Freund, we proved that a strict partial order is a mezzanine order, viz. it has no *3v1* suborder, if and only if it defines a nonmonotonic consequence relation which satisfies **P** and the new rule Plausible Monotonicity:

(Plausible Monotonicity) $$\frac{A \mid\!\sim C,\ A \not\mid\!\sim \neg B,\ A \wedge B \mid\!\sim D}{A \wedge B \wedge C \mid\!\sim D}$$

The two results combined allowed us to capture which inference rules are valid if our preferences over the alternatives are captured by semi-orders and if, when some new evidence is acquired, we eliminate all the alternatives which are inconsistent with it.

In this chapter we will combine all the previous results so to study what it is like to reason by comparing alternatives. We will do so by relying on a framework introduced by Hanti Lin and Kevin Kelly in [LK12b]. Such framework links the agent's beliefs with the probability distribution over the set of alternatives which captures the agent's subjective evaluation of uncertainty over such alternatives, often referred to as the agent's *degrees of belief* or *credences*. The relation which the framework establishes between beliefs and probabilities takes the name of *odds-threshold method*. All of its components are already familiar to us. The starting block is the probability function capturing the agent's degrees of belief. As a first step, the agent establishes a preference over the alternatives by comparing their probabilities; the result is an odds-threshold order, which was defined in the first chapter. Secondly, the agent beliefs that the actual alternative is one among the maximal elements of the odds-threshold order. Therefore, the agent strongest belief is constituted by the disjunction of the odds-threshold order maximal elements.

By relying on the odds-threshold method as a way to relate probabilities and beliefs, in the present chapter we will study how updating the underlying probability distribution impacts the beliefs generated by it. As a first case, we will consider Bayesian conditioning as probabilistic update. As a first result, we will prove that if we update our degrees of belief by Bayesian conditioning, then our qualitative beliefs are revised in such a way that all and only the rule of the nonmonotonic system **D**, viz. **P** together with Disjunctive Rationality and Plausible Monotonicity, are satisfied. Then, we show that the same result generalizes to Jeffrey conditioning. In particular, we show that, as long as the evidence gets assigned a high enough probability that the agent ends up believing it to be true – this stronger version of Jeffrey conditioning is referred to in the literature as *Successful* Jeffrey conditioning – then the agent's beliefs are stable enough that their update satisfies the logical properties of system **D**. Such results conclude the section on simple belief revision.

Jeffrey conditioning has an additional advantage with respect to Bayesian conditioning. Since it does not necessarily assign to the evidence E maximal probability, it is possible that later pieces of evidence disconfirm E, decreasing its probability. As an extreme case, we are also able to accommodate pieces of evidence which might even contradict the initial evidence. As a result, relying on Jeffrey conditioning to update our credences does not rule out any of the alternatives as long as the parametric posterior probability e is not assigned any of the two extreme values 0 and 1. Hence, assuming Jeffrey conditioning allows us to consider sequences of evidence and iterated revisions. In response to that, the final results of the present chapter address what are the logical properties of iterated belief revision that are satisfied, were our beliefs induced via the odds-threshold method by some probability distribution updated by Jeffrey conditioning.

1.1. Notation. In the present chapter we will assume a simplified notation. Instead of considering nonmonotonic consequence relations, we will rely on the *belief revision operators* $\mathsf{B}(A)$. The intuition behind such operators is that $\mathsf{B}(A)$ is the strongest belief held by the agent after learning A. We formally define it as follows.

Definition 1.1. *Let $\mathcal{SL}$ be a finite propositional language and Ω the associated set of worlds. Let $\succ$ be a preferential order over Ω. For any $A \in \mathcal{SL}$, the belief revision operator $\mathsf{B} : \mathcal{SL} \to \mathcal{SL}$ is defined as follows:*

$$\mathsf{B}_{\succ}(A) \equiv \bigvee \{max(\succ |A)\}\,.$$

Let $C(A)$ denote the set of all $C \in \mathcal{SL}$ such that $\mathsf{B}(A) \vDash C$.

$\mathsf{B}_{\succ}(A)$ is equivalent to the disjunction of the most preferred worlds in $\succ |A$. This is just a different notation from the usual nonmonotonic consequence relations $\mid\!\sim_{\succ}$ in the style of Kraus, Lehmann and Magidor [KLM90]. As a reminder, the latter are defined as follows:

$$A \mid\!\sim_{\succ} C \;\text{ if and only if }\; max(\succ |A) \subseteq \widehat{C}.$$

The equivalence between a nonmonotonic consequence relation and a belief revision operator is proven in Proposition 1.2.

PROPOSITION 1.2. *Let $\succ$ be a strict partial order over a set of worlds Ω. Let $\mathsf{B}_{\succ}$ be the belief revision operator and $\mid\!\sim_{\succ}$ the nonmonotonic consequence relation defined by $\succ$. The following equivalence holds:*

$$\mathsf{B}_{\succ}(A) \vDash C \iff A \mid\!\sim_{\succ} C.$$

PROOF. $\mathsf{B}_{\succ}(A) \vDash C$ is equivalent to $\bigvee\{max(\succ |A)\} \vDash C$. Since $\vDash$ is a classical propositional entailment relation, the latter holds if and only if, for each $\alpha_i \in max(\succ |A)$, $\alpha_i \Vdash C$. This can be restated $max(\succ |A) \subseteq \widehat{C}$, which is equivalent by definition to $A \mid\!\sim_{\succ} C$. □

$\mathsf{B}(A)$ classically implies all and only the nonmonotonic consequences of A. Hence, $\mathsf{B}(A)$ allows to have all the nonmonotonic consequence of A encoded into one single proposition. This condensed notation will be convenient in the present chapter where we will constantly relate belief revision operators with probabilistic updates. Sticking to nonmonotonic consequence relations would have made the notation more convoluted and less straightforward.

2. Bayesian conditioning and the original tracking problem

One of the aims of the authors in [LK12b] is to solve what they call *the tracking problem*, depicted in Figure 1. Everything starts with a probability distribution which captures the agent's subjective estimate of uncertainty over a set of alternatives Ω. Given such a probability function and some new piece of evidence with a non-zero probability, the agent updates their own credences by Bayesian conditioning P on E so to obtain the posterior probability distribution P_E. At the same time, that initial probability function P allows the agent to acquire their own qualitative beliefs B_P, intended as the propositions that the agent believes to be true. More precisely, the framework assumes that the agent's beliefs result from P, are functionally determined by P. The same holds for the belief B_{P_E}, generated by the posterior probability $P(\cdot|E)$. Lin and Kelly ask whether it is possible to define a method to update the initial beliefs so that the revised beliefs are equivalent to $\mathsf{B}_{P|E}$.

To make the problem more clear, let us introduce a little bit of formalism. Let $\mathcal{P}_\Omega$ be the set of probability distributions over the propositional language $\mathcal{SL}$. Then, call an *acceptance rule* any mapping from $\mathcal{P}_\Omega \to \mathcal{SL}$, which takes a probability function P and gives a formula

$\mathsf{B}_P(\top)$ corresponding to the conjunction of all beliefs generated by P. With a slight abuse of notation, the $\top$ in input to $\mathsf{B}_P(\top)$ is only meant to indicate the initial strongest belief held by the agent, which has not been revised for any evidence yet. Secondly, call a belief revision method any mapping $\mathcal{SL} \to \mathcal{SL}$ which takes $\mathsf{B}_P(\top)$ and some logically consistent formula E, and gives the revised strongest belief held by the agent after learning E, namely $\mathsf{B}_P(E)$. Then, the initial tracking problem asks whether there exists a pair constituted of one acceptance rule and one belief revision method which are able to satisfy the condition on the bottom-right of Figure 1, that is, to track Bayesian conditioning. To put it in a different way, the tracking problem wants to define a belief revision method which commutes with Bayesian conditioning, modulo the same acceptance rule.

$$\begin{array}{ccccc} P(\cdot) & \xrightarrow{\textit{Bayesian conditioning}} & & & P(\cdot|E) \\ \Big\downarrow{\scriptstyle a.r.} & & & & \Big\downarrow{\scriptstyle a.r.} \\ \mathsf{B}_P(\top) & \xrightarrow{\textit{Belief revision}} & \mathsf{B}_P(E) & \equiv & \mathsf{B}_{P|E}(\top) \end{array}$$

FIGURE 1. Lin and Kelly's Tracking Problem

The solution to the tracking problem proposed by Lin and Kelly starts from noticing that Bayesian conditioning preserves the ratio between the probabilities of those worlds that satisfy the evidence: for any $\alpha_i, \alpha_j \Vdash E$, it holds that $P(\alpha_i) = P(\alpha_i \wedge E)$, from which the following obtains:

$$\frac{P(\alpha_i)}{P(\alpha_j)} = \frac{P(\alpha_i) \cdot P(E)}{P(\alpha_j) \cdot P(E)} = \frac{P(\alpha_i \wedge E)}{P(E)} \cdot \frac{P(E)}{P(\alpha_j \wedge E)} = \frac{P(\alpha_i|E)}{P(\alpha_j|E)}$$

In order to exploit this property of Bayesian conditioning, following Isaac Levi [Lev96], a probability distribution P induces an odds-threshold order $\succ_P$ over $\Omega_0 \subseteq \Omega$ for some $t > 1$:

$$\alpha_i \succ_P \alpha_j \;\Leftrightarrow\; P(\alpha_i) \geq t \cdot P(\alpha_j), \text{ with } P(\alpha_j) > 0.$$

The definition assumes that the $\succ_P$ is defined for all and only the worlds with strictly positive probability, Ω_0. Once the odds-threshold order is defined, the agent's strongest belief $\mathsf{B}_P(\top)$ is defined as the disjunction of the $\succ_P$-maximal elements:

$$\mathsf{B}_P(\top) \equiv \bigvee max(\succ_P).$$

$\mathsf{B}_P(E)$ also serves as a belief revision operator. The initial belief $\mathsf{B}_P(\top)$ has $\top$ because it has been updated on no evidence yet. The same acceptance rule which from P results in $\mathsf{B}_P(\top)$ is applied to the posterior probability distribution P_E obtained from Bayesian condition P on the proposition E. Such acceptance rule allows Lin and Kelly to give in [LK12b] Theorem 2 a solution to the tracking problem, depicted in Figure 2.

The theorem states that Bayesian conditioning on the underlying probability function P has the same impact on $\succ_P$ as eliminating all the worlds inconsistent with E: $\succ_P |E =\succ_{P|E}$. This is because the acceptance rule relies on odds-threshold orders, whose relations are stable to Bayesian conditioning in the following sense. For every two $\alpha_i, \alpha_j \in \widehat{E}$,

$$\alpha_i \succ_P \alpha_j \ \Leftrightarrow\ \frac{P(\alpha_i)}{P(\alpha_j)} \geq t \ \Leftrightarrow\ \frac{P(\alpha_i|E)}{P(\alpha_j|E)} \geq t \ \Leftrightarrow\ \alpha_i \succ_{P|E} \alpha_j$$

On the contrary, Bayesian conditioning assigns zero probability to all $\alpha_j \vDash \neg E$. Thus the odds-threshold order $\succ_{P|E}$ will not be defined on any of such worlds. By these two fact, the correspondence between the two orders $\succ_P |E =\succ_{P|E}$ is obtained. It immediately follows that the beliefs generated respectively by $\succ_P |E$ and $\succ_{P|E}$ are equivalent as well, since the maximal elements of the two orders are the same: $\mathsf{B}_{\succ_P}(E) \equiv \mathsf{B}_{\succ_{P|E}}(\top)$.

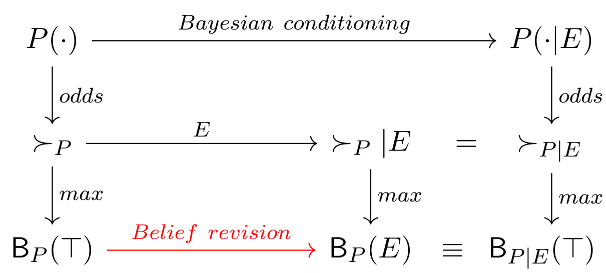

FIGURE 2. [LK12b] Theorem 2.

Notice however that, in order to solve the tracking problem, the authors provide a belief revision method which does not rely on the initial beliefs $\mathsf{B}_P(\top)$, but rather on the initial preferential order $\succ_P$. Indeed, there is no belief revision method which can both rely solely on the initial beliefs $\mathsf{B}_P(\top)$ and also track Bayesian conditioning. Figure 3 gives an example of why this is the case. Figure 3 shows the convex set of probability distributions over the

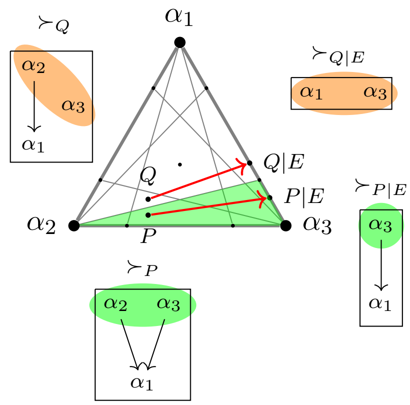

FIGURE 3. The belief revision method requires the preferential order to be specified.

set of three worlds, α_1, α_2 and α_3. Two initial probability distributions, P and Q, generate two different preferential orders, $\succ_P$ and $\succ_Q$, but the same initial belief, $\mathsf{B}_P(\top) \equiv \mathsf{B}_Q(\top)$. However, when P and Q are updated on the piece of evidence $E \equiv \alpha_1 \vee \alpha_3$, the resulting preferential orders $\succ_{P|E}$ and $\succ_{Q|E}$ end up being different one from the other, and even more importantly, their maximal elements differ and so do the beliefs that they generate. Indeed, $\mathsf{B}_{Q|E} \equiv \alpha_1 \vee \alpha_3$ while $\mathsf{B}_{P|E} \equiv \alpha_3$. The two cases start from the same beliefs, $\mathsf{B}_P(\top) \equiv \mathsf{B}_Q(\top)$, but end up generating different beliefs when the probabilities are updated, $\mathsf{B}_{P|E}(\top) \not\equiv \mathsf{B}_{Q|E}(\top)$. This example shows how a belief revision method cannot solely rely on the initial beliefs while tracking Bayesian conditioning. In order for the belief revision method to track Bayesian conditioning, additional information is required. Theorem 2 in [LK12b] proves that the preferential order defined by the odds-threshold order gives sufficient information to guarantee the tracking of Bayesian conditioning, and so it can guarantee the same final beliefs.

It seems then that in order to have a belief revision method which tracks Bayesian conditioning solely relying on the initial beliefs is not enough. Thus, the answer that Lin and Kelly give to the tracking problem is only partially positive: from their result it seems that there is no belief revision method which is able to entirely determine the beliefs after all possible updates. This is the reason why in Figure 1 the *Belief revision* arrow going from $\mathsf{B}_P(\top)$ to $\mathsf{B}_P(E)$ is colored red. The only way to determine $\mathsf{B}_P(E)$ from $\mathsf{B}_P(\top)$ is, loosely speaking, to go back to $\succ_P$ and update this order on E so to obtain $\succ_P |E$, and so $\mathsf{B}_P(E)$.

It is clear that there is no belief revision method relying solely on the initial beliefs that is capable of properly tracking Bayesian conditioning. Nevertheless, relying on the results previously obtained in this work, we are able to state the logical properties satisfied by our beliefs when their update is induced by Bayesian conditioning. More formally, in Theorem 2.1 we prove that a belief revision operator can be defined by Bayesian conditioning a probability distribution if and only if it satisfies system **D**.

THEOREM 2.1. *A belief revision operator* B *commutes with Bayesian conditioning for some probability function* P*, viz.* $\mathsf{B}(E) \equiv \mathsf{B}_{P|E}(\top)$*, if and only if* B *satisfies system* **D**.

PROOF. (Only if). It is enough to prove that the belief revision operator induced by a probability function P and Bayesian conditioning satisfies system **D**. By assumption, $\mathsf{B}_{P|E}(\top)$ is defined by the odds-threshold method. Thus, take the probability function P, whose odd-threshold order $\succ_P$ is by Corollary 3.5 a semi-order. By [LK12b] Theorem 2 we have $\succ_P |E = \succ_{P|E}$. Theorem 4.1 proves that, since $\succ_P$ is a semi-order, the belief revision operator $\mathsf{B}_{\succ_P|E}(\top) \equiv \bigvee\{max(\succ_P |E)\}$ satisfies system **D**. Because the two orders are equivalent, also the belief revision operator induced by Bayesian conditioning $\mathsf{B}_{P|E}(\top) \equiv \bigvee\{max(\succ_{P|E})\}$ satisfies system **D**.

(If). Let B be a belief revision operator satisfying system **D**. By Theorem 4.1, B can be defined by a semi-order $\succ_{\mathsf{B}}$ as $\mathsf{B}(E) \equiv \bigvee max(\succ_{\mathsf{B}} |E)$. By Corollary 3.5, $\succ_{\mathsf{B}}$ is a semi-order if and only if it can be represented by a probability function P, and by [LK12b] for such P, it holds that $\succ_{\mathsf{B}} |E = \succ_{P|E}$, from which $max(\succ_{\mathsf{B}} |E) = max(\succ_{P|E})$. Therefore, for each belief revision operator B satisfying system **D** there exists a probability function P such that $\mathsf{B}(E) \equiv \bigvee max(\succ_{\mathsf{B}} |E) \equiv \bigvee max(\succ_{P|E}) \equiv \mathsf{B}_{P|E}(\top)$. □

Systems			Rules	Conditions on B	B
			LLE	$\mathsf{B}(A) \vDash C, \vDash A \leftrightarrow B \;\Rightarrow\; \mathsf{B}(B) \vDash C$	
			RW	$\mathsf{B}(A) \vDash B, \vDash B \to C \;\Rightarrow\; \mathsf{B}(A) \vDash C$	
			Ref	$\mathsf{B}(A) \vDash A$	
		P	And	$\mathsf{B}(A) \vDash B,\, \mathsf{B}(A) \vDash C \;\Rightarrow\; \mathsf{B}(A) \vDash B \wedge C$	✓
			CMon	$\mathsf{B}(A) \vDash B,\, \mathsf{B}(A) \vDash C \;\Rightarrow\; \mathsf{B}(A \wedge B) \vDash C$	
			Or	$\mathsf{B}(A) \vDash C,\, \mathsf{B}(B) \vDash C \;\Rightarrow\; \mathsf{B}(A \vee B) \vDash C$	
	D		DRat	$\mathsf{B}(A \vee B) \vDash C,\, \mathsf{B}(A) \nvDash C \;\Rightarrow\; \mathsf{B}(B) \vDash C$	
			PlaMon	$\mathsf{B}(A) \vDash C,\; \mathsf{B}(A) \nvDash \neg B,\; \mathsf{B}(A \wedge B) \vDash D \;\Rightarrow\; \mathsf{B}(A \wedge B \wedge C) \vDash D$	
R			RMon	$\mathsf{B}(A) \vDash C,\, \mathsf{B}(A) \nvDash \neg B \;\Rightarrow\; \mathsf{B}(A \wedge B) \vDash C$	×

Table 1. Belief revision induced by Bayesian conditioning.

3. Belief change and Jeffrey conditioning

Theorem 2.1 shows that a belief revision operator tracks Bayesian conditioning if and only if it satisfies the rules of system **D**. In the present section we are going to extend such result to a generalization of Bayesian conditioning, namely, Jeffrey conditioning (JC). Here we will assume the following simplified version of Jeffrey conditioning, which always takes a two elements partition $\{E, \neg E\}$.

(Jeffrey conditioning) $$P_E^e(A) := e \cdot P(A|E) + (1-e) \cdot P(A|\neg E)$$

Almost the entirety of the fields of belief revision and nonomonotonic logics is interested in a notion of belief or consequence which is not certain. In the present case as well, beliefs are intended to be our best approximation of the true hypothesis. Nevertheless, we are not fully certain that our beliefs are true. It is reasonable then to have a notion of evidence which is not as restrictive as the one implied by Bayesian conditioning, which updates our credences only by assigning the evidence probability 1 so making it certain. On the contrary, we would like to also consider a weaker notion of evidence, which is reliable enough for us to update our credences and beliefs on it without assuming it to be certain.

In this section, we take the definition of JC that we have just given and we consider what are the properties that are always satisfied by the change that JC induces on our qualitative beliefs. Differently from Bayesian conditioning, Jeffrey conditioning does not assign maximal probability to E. Even more, it does not make any assumption on the evidence posterior probability $P_E(E) = e$. As a result, for every initial probability distribution P, every instance of Jeffrey conditioning requires us to specify not just the evidence E, but also the parameter e which quantifies the amount of uncertainty that we attribute to E after learning the new information about E. Since $e \in [0, 1]$, it might even be the case that JC lowers the probability of E, which has to be interpreted in the sense that the information learned about E has disconfirmed E, so lowering its probability. In other terms, this general definition of JC does not make any assumption on whether the evidence confirms or disconfirms E. It just assumes that the new information concerns E and that the probabilities of all the worlds or alternative which are respectively compatible or incompatible with E change proportionally. Formally, assuming $P(\alpha_j) \neq 0$, for all $\alpha_i, \alpha_j \Vdash E$ and $e \in (0, 1]$,

$$\frac{P^e_E(\alpha_i)}{P^e_E(\alpha_j)} = \frac{e \cdot P(\alpha_i|E) + (1-e) \cdot P(\alpha_i|\neg E)}{e \cdot P(\alpha_j|E) + (1-e) \cdot P(\alpha_j|\neg E)} = \frac{e \cdot P(\alpha_i \wedge E)}{e \cdot P(\alpha_j \wedge E)} = \frac{P(\alpha_i)}{P(\alpha_j)}.$$

This property is also true for Bayesian conditioning. However, for Jeffrey conditioning it holds as well for all $\alpha_i, \alpha_j \Vdash \neg E$ and $e \in [0, 1)$,

$$\frac{P^e_E(\alpha_i)}{P^e_E(\alpha_j)} = \frac{e \cdot P(\alpha_i|E) + (1-e) \cdot P(\alpha_i|\neg E)}{e \cdot P(\alpha_j|E) + (1-e) \cdot P(\alpha_j|\neg E)} = \frac{(1-e) \cdot P(\alpha_i \wedge \neg E)}{(1-e) \cdot P(\alpha_j \wedge \neg E)} = \frac{P(\alpha_i)}{P(\alpha_j)}.$$

If, on the one hand, JC preserves the odds between atoms respectively satisfying E and $\neg E$, on the other hand, the newly acquired information might confirm E, or just as well disconfirm E, to the two extreme cases of JC assuming E to be certain for $e = 1$ and to be impo ssible for $e = 0$. Indeed, the impact of JC on our beliefs has so much variability that most of the properties of belief revision are not in general satisfied. Most notably, *Reflexivity* fails as for $e = 0$, the agent not just does not believe in E to be true, but they are certain that E is false. As a result, following [CR19], we introduce a notion of qualitative update weaker than belief revision, named *belief change*. Along with the same strategy that we adopted to study the impact of Bayesian conditioning on our beliefs, in this section, firstly we qualitatively characterize the set of preferential orders that can be obtained from all the instances of Jeffrey conditioning the initial probability function P on some evidence E for all possible $e \in [0, 1]$ (Definition 3.5). Secondly, in Theorem 3.6 we actually prove that such set of preferential orders, qualitative defined from the initial order $\succ_P$, corresponds to the set of preferential orders generated by all the instances of Jeffrey conditioning P on E. Finally, we proceed to show what are the logical properties that are satisfied by all the belief change operators induced by all such order (Proposition 3.8). The following two Lemmas will come in handy when proving the theorem.

LEMMA 3.1. *Let P be a probability distribution over Ω and let $\succ_P$ be a semi-order defined over $\Omega_0 \subseteq \Omega$ generated by P for some real threshold $t > 1$. Then, for any $\alpha_i, \alpha_j \in \Omega_0$,*

$$P(\alpha_i) \geq P(\alpha_j) \;\Rightarrow\; r_{\succ_P}(\alpha_i) \geq r_{\succ_P}(\alpha_j).$$

PROOF. We prove the contrapositive: $r_{\succ_P}(\alpha_i) > r_{\succ_P}(\alpha_j) \;\Rightarrow\; P(\alpha_i) > P(\alpha_j)$.
Since $\succ_P$ is a semi-order, $r_P(\alpha_i) > r_P(\alpha_j), r_P(\alpha_i) \geq r_P(\alpha_j)$ implies that $Down_{\succ}(\alpha_i) \supseteq Down_{\succ}(\alpha_j)$ and $Up_{\succ}(\alpha_i) \subseteq Up_{\succ}(\alpha_j)$. At the same time, $r_P(\alpha_i) > r_P(\alpha_j)$ is equivalent to $|Down_{\succ_P}(\alpha_i)| - |Down_{\succ_P}(\alpha_j)| > |Up_{\succ_P}(\alpha_j)| - |Up_{\succ_P}(\alpha_i)|$. Therefore, from the two either (A) $Down_{\succ_P}(\alpha_i) \supset Down_{\succ_P}(\alpha_j)$ or (B) $Up_{\succ_P}(\alpha_i) \subset Up_{\succ_P}(\alpha_j)$. By cases.

(A) $Down_{\succ_P}(\alpha_i) \supset Down_{\succ_P}(\alpha_j)$
   $\Rightarrow$ there exists $\alpha_k \in \Omega$ s.t. $\alpha_i \succ_P \alpha_k$ and $\alpha_j \not\succ_P \alpha_k$
   $\Leftrightarrow$ $P(\alpha_i) \geq t \cdot P(\alpha_k) > P(\alpha_j) \;\Rightarrow\; P(\alpha_i) > P(\alpha_j)$.

(B) $Up_{\succ_P}(\alpha_i) \subset Up_{\succ_P}(\alpha_j)$
   $\Rightarrow$ there exists $\alpha_k \in \Omega$ s.t. $\alpha_k \succ_P \alpha_j$ and $\alpha_k \not\succ_P \alpha_i$
   $\Leftrightarrow$ $t \cdot P(\alpha_i) > P(\alpha_k) \geq t \cdot P(\alpha_j) \;\Rightarrow\; P(\alpha_i) > P(\alpha_j)$.

The two exhaustive cases (A) and (B) compatible with the assumption $r_P(\alpha_i) > r_P(\alpha_j)$ imply $P(\alpha_i) > P(\alpha_j)$. The contrapositive has thus been proven. □

LEMMA 3.2. *Let $\succ$ be a semi-order over $\Omega_0 \subseteq \Omega$ and let $\mathcal{P}(\succ)$ be the set of probability distributions generating $\succ$ for some contextually determined real threshold $t > 1$. Then, for any two atoms $\alpha_i, \alpha_j \in \Omega_0$, if for all $P \in \mathcal{P}(\succ)$ $P(\alpha_i) \geq P(\alpha_j)$, then $r_\succ(\alpha_i) > r_\succ(\alpha_j)$.*

PROOF. We prove the contrapositive. First, notice that all the inequalities induced by $\succ$ are of the form $P(\alpha_i) \gtreqless t \cdot P(\alpha_j)$ for any $\alpha_i, \alpha_j \in \Omega_0$ with $P(\alpha_j) > 0$. Because $\succ$ is a semi-order over $\Omega_0 \subseteq \Omega$, thus it can be generated by a probability function over Ω; equivalently, any system of inequalities induced by some $\succ$ has a solution. In addition, since all such primitive inequalities are linear and with a null constant term, also all the derivable inequalities are so. It follows that no inequality induced by $\succ$ reduces the dimension of the solution space, equal to $|\Omega_0| - 1$ as the probabilities sum to 1. Therefore, for any $\succ$ and any $\alpha_i, \alpha_j \in \Omega_0$, it is not the case that for all $P \in \mathcal{P}(\succ)$, $P(\alpha_i) = P(\alpha_j)$.

Now take any two distinct atoms $\alpha_1, \alpha_2 \in \Omega_0$ such that $r_\succ(\alpha_1) \leq r_\succ(\alpha_2)$. If $r_\succ(\alpha_1) < r_\succ(\alpha_2)$, then it follows by Lemma 3.1 that any probability function P generating $\succ$ is such that $P(\alpha_1) < P(\alpha_2)$, contradicting $P(\alpha_i) \geq P(\alpha_j)$ for any $\alpha_i, \alpha_j \in \Omega_0$, as desired. Then, we can assume that $r_\succ(\alpha_1) = r_\succ(\alpha_2)$. By the reasoning above we know that *not* for all $P \in \mathcal{P}(\succ)$, $P(\alpha_1) = P(\alpha_2)$; hence, there exists a probability function $P \in \mathcal{P}(\succ)$ such that, say, $P(\alpha_1) > P(\alpha_2)$. At the same time, by the properties of semi-orders, because the two atoms have same rank, we know that α_1 and α_2 have the same Up- and Down-sets, the same relations. Therefore, the difference between their probabilities is small enough that it has no impact in the preferential relations established with the other atoms. Hence, now consider the probability function P' identical to P but for the fact that α_1 and α_2 swap their probabilities. P' still generates the order $\succ$, but it holds that $P'(\alpha_1) < P'(\alpha_2)$. Thus, also for the case $r_\succ(\alpha_1) = r_\succ(\alpha_2)$, there exists a probability function $P' \in \mathcal{P}(\succ)$ s.t. $P'(\alpha_1) < P'(\alpha_2)$. The converse is so proven.

□

Before giving a qualitative, probability-independent characterization of the preferential orders that can be obtained by Jeffrey conditioning a probability function P on some evidence E for all $e \in [0,1]$, we need to introduce the notion of *relative distance interval.* As the name suggests, the relative distance interval between two elements α_i and α_j relative to the order $\succ$ captures in terms of integer numbers the distance between α_i and α_j. On the one hand, the longer is the $\succ$-chain between α_i and α_j, the bigger α_i needs to be in order to have that many elements of the order in between. For that reason $\underline{d}_{\succ}(\alpha_i, \alpha_j)$ is the lower bound of the relative distance interval between the two worlds. On the contrary, in general, if there is no preferential relation between two elements of the order, this means that the two quantities associated with the two elements are close. Extending the point to $\sim$-chains, the shorter is the $\sim$-chain between α_i and α_j of the order, the closer α_i needs to be to α_j. Hence, $\overline{d}_{\succ}(\alpha_i, \alpha_j)$ constitutes the upped bound of the relative distance interval between the two. The upper and lower bounds determine what we call the relative distance interval between two worlds. This notion plays an important role when defining the preference change of a preferential order for some evidence E. In particular, Definition 3.5 requires the initial order $\succ_P$ and any order obtained by preference change to have compatible relative distance interval for any couple of worlds both satisfying either E or $\neg E$ (Condition (3)).

Definition 3.3 (Relative lower and upper bounds). *Let $\succ$ be a semi-order over Ω and let α_i, $\alpha_j \in \Omega$ such that $r_{\succ}(\alpha_i) \geq r_{\succ}(\alpha_j)$. If $r_{\succ}(\alpha_i) > r_{\succ}(\alpha_j)$, the relative lower bound $\underline{d_{\succ}}(\alpha_i, \alpha_j)$ and the relative upper bound $\overline{d_{\succ}}(\alpha_i, \alpha_j)$ are defined as follows:*

$$\underline{d_{\succ}}(\alpha_i, \alpha_j) = \begin{cases} \textit{the length of the maximal } \succ\textit{-path from } \alpha_j \textit{ to } \alpha_i \\ 0 \textit{ if no such path exists} \end{cases}$$

$$\overline{d_{\succ}}(\alpha_i, \alpha_j) = \begin{cases} \textit{the length of the minimal rank-increasing} \sim\textit{-path from } \alpha_j \textit{ to } \alpha_i \\ +\infty \textit{ if no such path exists} \end{cases}$$

Otherwise, if $r_{\succ}(\alpha_i) = r_{\succ}(\alpha_j)$, then $\underline{d_{\succ}}(\alpha_i, \alpha_j) = -1$ and $\overline{d_{\succ}}(\alpha_i, \alpha_j) = +1$.

DEFINITION 3.4 (Relative distance interval). *Let $\succ$ be a semi-order over Ω and let $\alpha_i,\ \alpha_j \in \Omega$ are two atoms such that $r_\succ(\alpha_i) \geq r_\succ(\alpha_j)$. The Relative distance interval between α_i and α_j is defined as the left-closed right-open interval*

$$\Delta_\succ(\alpha_i, \alpha_j) = [\, \underline{d_\succ}(\alpha_i, \alpha_j), \overline{d_\succ}(\alpha_i, \alpha_j)\,).$$

Instead, the Relative distance interval between α_j and α_i is defined as follows:

$$\Delta_\succ(\alpha_j, \alpha_i) = -\Delta_\succ(\alpha_i, \alpha_j) = \ (\, -\overline{d_\succ}(\alpha_i, \alpha_j), -\underline{d_\succ}(\alpha_i, \alpha_j)\,].$$

DEFINITION 3.5 (Preference change). *Let $\succ$ be a semi-order over Ω_0 and E be a $\succ$-consistent formula. Then, an order $\succ^E$ is a result of an E-induced preference change on $\succ$ if and only if $\succ^E = (\succ |E)$ or $\succ^E = (\succ |\neg E)$ or $\succ^E$ satisfies properties $(0) - (3)$:*

- *(0) $\succ^E$ is a semi-order over Ω_0;*
- *(1) for all $\alpha_i, \alpha_j \Vdash E$ and for all $\alpha_i, \alpha_j \Vdash \neg E$, $\alpha_i \succ \alpha_j \iff \alpha_i \succ^E \alpha_j$;*
- *(2) for all $\alpha_i, \alpha_j \Vdash E$ and for all $\alpha_i, \alpha_j \Vdash \neg E$, $r_\succ(\alpha_i) > r_\succ(\alpha_j) \implies r_{\succ^E}(\alpha_i) \geq r_{\succ^E}(\alpha_j)$;*
- *(3) for all $\alpha_i, \alpha_j \Vdash E$ and for all $\alpha_i, \alpha_j \Vdash \neg E$, $\Delta_\succ(\alpha_i, \alpha_j) \cap \Delta_{\succ^E}(\alpha_i, \alpha_j) \neq \emptyset$ whenever defined.*

$\mathcal{C}(\succ, E)$ is the set of all the resulting orders $\succ^E$.

The set $\mathcal{C}(\succ, E)$ captures the impact that Jeffrey conditioning P on E has on the preferential order $\succ$. Theorem 3.6 states such correspondence. From its formulation however, it is clear that Theorem 3.6 proves something weaker: it proves that $\mathcal{C}(\succ_P, E)$ captures all the preferential orders that can be obtained by Jeffrey conditioning on E not just P, but all those probability functions P' that generate the same order $\succ_P$. Indeed, there is no way to qualitatively update $\succ_P$ so to obtain all *and only* those preferential orders $\mathcal{O}(\mathcal{J}(P, E))$ generated by any of the instances of JC of P on E, i.e. any $P^e_E \in \mathcal{J}(P, E)$. This is because any such belief revision method would need to rely on the information provided by the preferential order $\succ_P$ only. Such order $\succ_P$ gives a qualitative representation of the probabilities assigned to the worlds by P. At the same time however, shifting to this qualitative representation implies loosing some of the information held by the initial real valued function P. As a consequence, updating P will always give a more precise result than updating $\succ_P$. To better appreciate the point, consider Figure 4. In Figure 4a the black segment ranging from α_1 to $P|\alpha_2 \vee \alpha_3$ and passing thorough P denotes the set $J(P, \alpha_2 \vee \alpha_3)$, while the green area that includes the segment captures all the possible semi-orders that can be obtained by Jeffrey conditioning P on the evidence $\alpha_2 \vee \alpha_3$. It is clear that the two orders $\succ_1$ and $\succ_2$ in red cannot be generated by any $P_E^e \in \mathcal{J}(P, E)$; in other words, $\succ_1$ and $\succ_2$ are incompatible with any Jeffrey conditioning of P on E. On the contrary, there is a probability function P' such that both $\succ_1$ and $\succ_2$ belong to $\mathcal{O}(\mathcal{J}(P', E))$. Also, notice that we chose both P and P' so that they both generate the same preferential order $\succ_P$. Therefore, the example shows that any belief change method which relies on the initial preferential order $\succ_P$ would not be able to distinguish whether $\succ_P$ was generated by P or by P'. The loss of information that comes from dealing with $\succ_P$ instead of any of the two P and P' has a cost in terms of the ability of restricting which are the possible semi-orders resulting from preference change. As a result, any proper belief change method has to consider $\succ_1$ and $\succ_2$, and to add them to $\mathcal{C}(\succ_P, E)$ even though they are incompatible with $\mathcal{J}(\mathcal{P}, \mathcal{E})$. Indeed, it is worth noting that $\succ_P$ refers to P, but the index is used here only for notational clarity while for the belief change method it has no meaning.

It is clear that because the belief change method can only rely on the initial preferential order $\succ_P$ and not on the real-valued function P, $\mathcal{C}(\succ_P, E)$ will always be a supset of $\mathcal{O}(\mathcal{J}(P, E))$. If the strong correspondence $\mathcal{C}(\succ_P, E) = \mathcal{O}(\mathcal{J}(P, E))$ between belief change and Jeffrey conditioning cannot be established in general, we still hope to define a belief change method which is restrictive enough to capture all and only those semi-orders $\mathcal{O}(\mathcal{J}(P', E))$ for *any* probability function P' which generates the initial preferential order $\succ_P$, such as P' in Figure 4. Going back again to that example, we want a belief change method that gives no other semi-order that those in the blue area in Figure 4 (which includes $\succ_1$ and $\succ_2$). This is precisely what Theorem 3.6 proves about the belief change method defined by Definition 3.5. In order to get such result, Definition 3.5 imposes a set of conditions on the preferential orders that belong to $\mathcal{C}(\succ_P, E)$. Let us now go through them and explain their role.

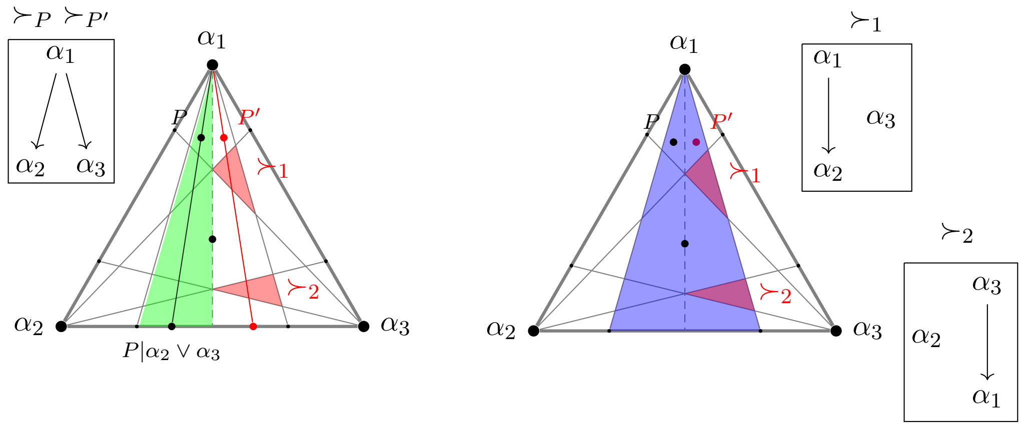

FIGURE 4. Updating a probability function and updating a semi-order.

Condition (0) is trivial as it requires $\succ_P^E$ to be a semi-order; were not that the case, $\succ_P^E$ would not be representable probabilistically and so it could not be obtained by JC in any manner. Condition (1) extends the approach that was taken by Lin and Kelly, which goes back to Yoav Shoham [Sho87]. In their case, Bayesian conditioning P on E has the same impact on $\succ_P$ as deleting all the worlds inconsistent with E, namely $\succ_P |E$. This is because all the worlds falsifying E get assigned probability 0 and $\succ_{P|E}$ is not defined on them, and at the same time the ratio between the probabilities of the worlds satisfying E remains constant. This last aspect still holds for Jeffrey conditioning, as we discussed at the beginning of the section, and here it also extends to all couple of worlds satisfying $\neg E$ as well. Accordingly, condition (1) requires to the updated order $\succ_P^E$ to preserve all the $\succ_P$-relations over E and $\neg E$ respectively.

Both conditions (2) and (3) have to do with the relative position between two atoms both of which either satisfy of falsify the evidence. As proven by Lemma 3.1, if the rank of an atom α_i is strictly greater than the rank of another α_j, then it must be the case that, for any P generating that order, α_i has a strictly greater probability than α_j. If in addition both α_i and α_j either satisfy or falsify E, then we know that their probabilities change with JC by maintaining the same proportion. As a result, by the same reason, the rank of α_j would never

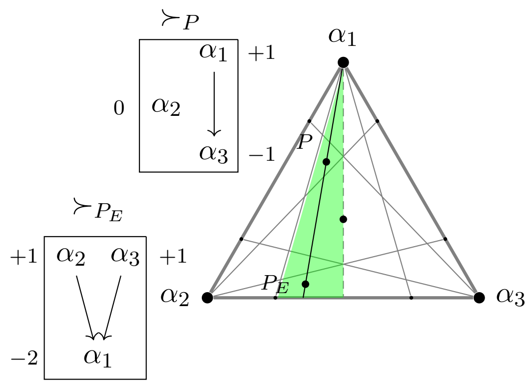

FIGURE 5. Condition (2).

be greater than the rank of α_i. It might be however the case that the two end up having the same rank, as α_2 and α_3 show in Figure 5. And it also compatible with (2) that if two worlds start with the same rank, in the resulting order one of the two has a higher rank that the other – consider again Figure 5 with P_E as prior probability and P as the probability function resulting from JC.

THEOREM 3.6. *Let P be a probability distribution over Ω, $\succ_P$ the semi-order generated by P and $\mathcal{P}(\succ_P)$ the set of probability distributions generating $\succ_P$. Let $\mathcal{J}(P, E)$ be the set of possible probability distributions P^e_E resulting from Jeffrey conditioning P on E. $\mathcal{O}(\mathcal{J}(P, E))$ is the set of odds-threshold orders generated by the probability distributions in $\mathcal{J}(P, E)$. Finally, let $\mathcal{C}(\succ_P, E)$ be the set of orders resulting from $\succ_P$ by preference change on E, as defined in Definition 3.5. Then, the following holds:*

$$\mathcal{C}(\succ_P, E) = \mathcal{O}\left(\bigcup \{\mathcal{J}(P', E) : \ P' \in \mathcal{P}(\succ_P)\}\right).$$

PROOF.

To prove the equality, we proceed by showing that either side is included in the other.

$\mathcal{C}(\succ_P, E) \supseteq \mathcal{O}\left(\bigcup \{\mathcal{J}(P', E) : \ P' \in \mathcal{P}(\succ_P)\}\right)$

We show that for each $P' \in \mathcal{P}(\succ_P)$, the order $\succ_{P'_E}$ generated by the probability distribution P'_E obtained by an instance of Jeffrey conditioning on E, belongs to $\mathcal{C}(\succ_P, E)$, as required.

First, we address the two extreme cases of Jeffrey conditioning where E is assigned respectively $P'_E(E) = 1$ and $P'_E(E) = 0$. The two cases are specular; we only address the former.

$P'_E(E) = 1$. Hence $\succ_{P'_E} = (\succ_P |E)$. Since $\succ_{P'} = \succ_P$, the two orders are defined over the same set of atoms $A \subseteq \Omega$. By the definition of odds-threshold order, this is the case if and only if all and only the atoms in A get strictly positive probability. Hence, for all and only $\alpha_i \Vdash \neg A$, both $P(\alpha_i) = 0$ and $P'(\alpha_i) = 0$ hold. Coming to $\succ_{P'_E}$, $P'_E(E) = 1$ implies that, $\alpha_i \Vdash \neg(A \vee E)$ if and only if $P'_E(\alpha_i) = P'(\alpha_i|E) = \frac{P'(\alpha_i \cap E)}{P'(E)} = 0$. On the other hand, for any $\alpha_i \Vdash A \wedge E$, $P'(\alpha_i) > 0$ and $P'(\alpha_i|E) > 0$. Therefore, $\succ_{P'_E}$ is defined over $\widehat{A \wedge E}$. Likewise, it is immediate to see that $\succ_{P'} |E$ is defined over $\widehat{A \wedge E}$. Finally, the two orders have the same relations: for any $\alpha_i, \alpha_j \Vdash A \wedge E$,

$$\alpha_i(\succ_{P'} |E)\alpha_j \iff \alpha_i \succ_{P'} \alpha_j \iff P'(\alpha_i) \geq t \cdot P'(\alpha_j) \iff P'(\alpha_i) \geq t \cdot P'(\alpha_j)$$
$$\iff P'_E(\alpha_i) = P'(\alpha_i|E) \geq t \cdot P'(\alpha_j|E) = t \cdot P'_E(\alpha_j) \iff \alpha_i \succ_{P'_E} \alpha_j$$

So, we proved that $\succ_{P'_E} = (\succ_P |E)$ for $P'_E(E) = 1$. Specularly, it can be proven that $\succ_{P'_E} = (\succ_P |\neg E)$ for $P'_E(E) = 0$. Since both $\succ_P |E$ and $\succ_P |\neg E$ belong to $\mathcal{C}(\succ_P, E)$, it follows that for the two extreme cases of Jeffrey conditioning, $\succ_{P'_E}$ belongs to $\mathcal{C}(\succ_P, E)$, as desired.

We can now address the non-extreme case where P'_E assigns a probability value $0 < e < 1$ to the evidence E. We show that $\succ_{P'_E}$ satisfies all conditions (0)-(3).

(0). Like in the previous case, since $\succ_P = \succ_{P'}$ defined over $A \subseteq \Omega$, both P and P' assign non-zero probability to all and only $\alpha_i \in A$. Since in this case $0 < e < 1$, for any $\alpha_i \in A$ the posterior probability $P'_E(\alpha_i)$ remains strictly positive:

$$P'(\alpha_i) > 0 \iff P'_E(\alpha_i) = e \cdot P'(\alpha_i|E) + (1 - e) \cdot P'(\alpha_i|\neg E) > 0$$

Therefore, $\succ_{P'_E}$ is defined for all and only $\alpha_i \in A$. Also, since $\succ_{P'_E}$ is generated by a probability function over Ω, we know by Corollary 3.5 that $\succ_{P'_E}$ is a semi-order. Therefore, $\succ_{P'_E}$ satisfies property (0) of Def. 3.5 of being a semi-order over A.

(1). Let $\alpha_i, \alpha_j \in E$. It has been established in point (0) that both the orders $\succ'_P$ and $\succ_{P'_E}$ are defined for all and only $\alpha_k \in A$. Therefore, if any of the two atoms α_i or α_j satisfies $\neg A$, then the orders have no relation concerning them. Hence, condition (1) is trivially satisfied.

Then, we can assume that both $\alpha_i, \alpha_j \Vdash A \wedge E$. Since $\succ_{P'} = \succ_P$ by assumption, condition (1) is satisfied by $\succ_{P'_E}$:

$$\alpha_i \succ_{P'} \alpha_j \iff P'(\alpha_i) \geq t \cdot P'(\alpha_j) \iff e \cdot P'(\alpha_i) \geq t \cdot e \cdot P'(\alpha_j)$$
$$\iff P'_E(\alpha_i) \geq t \cdot P'_E(\alpha_j) \iff \alpha_i \succ_{P'_E} \alpha_j$$

The proof for $\alpha_i, \alpha_j \Vdash \neg E$ is analogous.

(2). Let $\alpha_i, \alpha_j \Vdash E\,(\alpha_i, \alpha_j \Vdash \neg E)$. As in the previous case, if any of the two atoms α_i or α_j belongs to $\neg A$, the two ranks $\succ_{P'}$ and $\succ_{P'_E}$ are undefined for at least one of the two, and so are the rank functions induced by the orders. Therefore, under such assumption, condition (2) is trivially satisfied. We can assume that both $\alpha_i, \alpha_j \Vdash A \wedge E\,(\alpha_i, \alpha_j \Vdash A \wedge \neg E)$. Since Jeffrey conditioning increases or decreases their probability proportionally, their relative position within the order captured by their rank cannot swap. This follow by the properties of Jeffrey conditioning and two applications of Lemma 3.1.

$$r_P(\alpha_i) > r_P(\alpha_j) \;\Rightarrow\; P(\alpha_i) > P(\alpha_j) \;\Leftrightarrow\; P_E(\alpha_i) > P_E(\alpha_j) \;\Rightarrow\; r_{P_E}(\alpha_i) \geq r_{P_E}(\alpha_j).$$

(3). Take any two $\alpha_i, \alpha_j \Vdash E$ and assume $P'(\alpha_i) \geq P'(\alpha_j)$, both strictly positive otherwise condition (3) is trivially satisfied because $\Delta_{\succ_{P'}}(\alpha_i, \alpha_j)$ and $\Delta_{\succ_{P'_E}}(\alpha_i, \alpha_j)$ are undefined. By Lemma 3.1, we know that $r_{\succ_{P'}}(\alpha_i) \geq r_{\succ_{P'}}(\alpha_j)$, and since their probabilities change proportionally in P'_E, we also know that $r_{\succ_{P'_E}}(\alpha_i) \geq r_{\succ_{P'_E}}(\alpha_j)$. Then, $\Delta_{\succ_{P'}}(\alpha_i, \alpha_j)$ and $\Delta_{\succ_{P'_E}}(\alpha_i, \alpha_j)$ are both defined. We shall see whether condition (3) is satisfied.

If $P'(\alpha_i) = P'(\alpha_j)$, then $P'_E(\alpha_i) = P'_E(\alpha_j)$. Hence, $r_{\succ_{P'}}(\alpha_i) = r_{\succ_{P'}}(\alpha_j)$ and $r_{\succ_{P'_E}}(\alpha_i) = r_{\succ_{P'_E}}(\alpha_j)$. From their definition, we get $\Delta_{\succ P'}(\alpha_i, \alpha_j) = \Delta_{\succ P'_E}(\alpha_i, \alpha_j) = [0, 1)$, making the two intervals compatible as their intersection is non-empty.

Otherwise, $P'(\alpha_i) > P'(\alpha_j)$. Then there is a maximal $m \in \mathbb{N}$ for which the following conditions hold:

$$P'(\alpha_i) > t^m P'(\alpha_j) \;\text{ and }\; P'(\alpha_i) < t^{m+1} P'(\alpha_j)$$

Since $\frac{P'(\alpha_i)}{P'(\alpha_j)} = \frac{P'_E(\alpha_i)}{P'_E(\alpha_j)}$, the same conditions hold for P'_E for the same m:

$$P'_E(\alpha_i) > t^m P'_E(\alpha_j) \;\text{ and }\; P'_E(\alpha_i) < t^{m+1} P'_E(\alpha_j)$$

Since m is maximal, we know that $m \geq \underline{\delta_{\succ P'}}(\alpha_i, \alpha_j)$ and $m+1 \leq \overline{\delta_{\succ P'}}(\alpha_i, \alpha_j)$ and $m \geq \underline{\delta_{\succ P'_E}}(\alpha_i, \alpha_j)$ and $m+1 \leq \overline{\delta_{\succ P'_E}}(\alpha_i, \alpha_j)$. Consequently, $[m, m+1) \subseteq \Delta_{\succ P'}(\alpha_i, \alpha_j)$ and $[m, m+1) \subseteq \Delta_{\succ P'_E}(\alpha_i, \alpha_j)$, from which $\Delta_{\succ P'}(\alpha_i, \alpha_j) \cap \Delta_{\succ P'_E}(\alpha_i, \alpha_j) \neq \emptyset$ as desired.

$\mathcal{C}(\succ_P, E) \subseteq \mathcal{O}\left(\bigcup\{\mathcal{J}(P', E) : P' \in \mathcal{P}(\succ_P)\}\right)$

We prove that, if $\succ_P^E \notin \mathcal{O}\left(\bigcup\{\mathcal{J}(P', E) : P' \in \mathcal{P}(\succ_P)\}\right)$, then $\succ_P^E \notin \mathcal{C}(\succ_p, E)$. We can assume $\succ_P^E$ to be a semi-order, otherwise the condition trivially holds.

$\succ_P^E \notin \mathcal{O}\left(\bigcup\{\mathcal{J}(P', E) : P' \in \mathcal{P}(\succ_P)\}\right)$ iff $\mathcal{P}(\succ_P^E)$ has empty intersection with the rigid extension of the set $\mathcal{P}(\succ_P)$ giving by Jeffrey conditioning every probability function in $\mathcal{P}(\succ_P)$ for every parametric posterior probability of E, $e \in [0,1]$ – "rigid" meaning that for every $P' \in \mathcal{P}(\succ_P)$, the extensions of P' preserve the conditional probability assignments P'_E and $P'_{\neg E}$. A point belongs to $\bigcup\{\mathcal{J}(P', E) : P' \in \mathcal{P}(\succ_P)\}$ if and only if its conditional probabilities for E and $\neg E$ are the same assignments as P'_E and $P'_{\neg E}$, for some point $P' \in \mathcal{P}(\succ_P)$. Also, $\mathcal{P}(\succ_P^E)$ and $\bigcup\{\mathcal{J}(P', E) : P' \in \mathcal{P}(\succ_P)\}$ have empty intersection if and only if the constraints given by $\succ_P^E$ are incompatible with $\bigcup\{\mathcal{J}(P', E) : P' \in \mathcal{P}(\succ_P)\}$. But by the properties of $\bigcup\{\mathcal{J}(P', E) : P' \in \mathcal{P}(\succ_P)\}$, this can be the case if and only if the constraints given by $\succ_P^E$ on the conditional probabilities for E and $\neg E$ are incompatible with the constraint on those same conditional probabilities given by $\succ_P$. Remember that the only kind of constraints induced by $\succ_P$ and $\succ_P^E$ are as follows: for any two $\alpha_i, \alpha_j \in \Omega$, if $\alpha_i \succ_P \alpha_j$ then $P(\alpha_i) \geq t \cdot P(\alpha_j)$, otherwise $P(\alpha_i) < t \cdot P(\alpha_j)$; likewise, if $\alpha_i \succ_P^E \alpha_j$ then $R(\alpha_i) \geq t \cdot R(\alpha_i)$, otherwise $R(\alpha_i) < t \cdot R(\alpha_i)$. Thus, all the derived inequalities are linear and with a null constant term. We said that the constraints given by $\succ_P$ and $\succ_P^E$ on at least one between the posterior probabilities on E or $\neg E$ are incompatible (possibly both). Since Bayesian conditioning updates proportionally the probabilities of the atoms respectively satisfying and falsifying the evidence, and since all the constraints on the posterior probabilities are linear and with a null constant term, then the constraints given by $\succ_P$ and $\succ_P^E$ on the posterior probabilities on E or $\neg E$ are incompatible if and only if the constraints given by $\succ_P$ and $\succ_P^E$ on the atoms satisfying E or the atoms satisfying $\neg E$ are incompatible. Assume that the two systems of inequalities respectively induced by $\succ_P$ and $\succ_P^E$ are incompatible over the atoms satisfying E – the case for $\neg E$ is specular. Again, all such constraints are linear inequalities with null constant term. Then, since the two systems are incompatible, there are $\alpha_i, \alpha_j \in \widehat{E}$ such that for some $m \in \mathbb{N}$

- either $P'(\alpha_i) \geq t^m P'(\alpha_j)$ and $R(\alpha_i) < t^m R(\alpha_j)$,
- or $P'(\alpha_i) < t^m P'(\alpha_j)$ and $R(\alpha_i) \geq t^m R(\alpha_j)$.

The two cases are analogous; we consider the former. For $m = 0$ we get $P'(\alpha_i) \geq P'(\alpha_j)$ and $R(\alpha_i) < R(\alpha_j)$. Since the inequalities hold in general, for two generic probability functions P' and R, we can apply Lemma 3.2 and obtain that $r_{\succ_P}(\alpha_i) > r_{\succ_P}(\alpha_j)$ and $r_{\succ_P^E}(\alpha_i) < r_{\succ_P^E}(\alpha_j)$. This contradicts condition (2) of Definition 3.5. Hence, $\succ_P^E \notin \mathcal{C}(\succ_P, E)$ as required.

For $m > 0$, the inequality $P'(\alpha_i) \geq t^m P'(\alpha_j)$ follows from the system of inequalities induced by $\succ_P$. As a consequence, it must be the case that the longest $\succ_P$-chain between α_i and α_j is not shorter than m, i.e. $m \leq \underline{\delta}_{\succ_P}(\alpha_i, \alpha_j)$. Analogously, $Q(\alpha_i) < t^m Q(\alpha_j)$ can be derived from the system of inequalities induced by $\succ_P^E$. It follows that the shortest rank-increasing $\sim_P^E$-chain has length at most m, i.e. $\overline{\delta}_{\succ_P^E}(\alpha_i, \alpha_j) \leq m$. Since by definition $\overline{\delta}_{\succ_P^E}(\alpha_i, \alpha_j)$ is not included in $\Delta_{\succ_P^E}(\alpha_i, \alpha_j)$, we get that $\Delta_{\succ_P^E}(\alpha_i, \alpha_j) \cap \Delta_{\succ_P^E}(\alpha_i, \alpha_j) = \emptyset$. Condition (3) is so falsified and $\succ_P^E \notin \mathcal{C}(\succ_P, E)$, as desired. □

Theorem 3.6 allows us to characterize in terms of purely order-theoretic properties what is the impact of Jeffrey conditioning a probability function P on the preferential order $\succ_P$. Then, we are now able to say what are the properties that our preferences will certainly satisfy when we apply JC to our credences, regardless of what degree of uncertainty we assign to the evidence. Because in the present framework our beliefs are determined by our preferences over the set of alternatives, we can now address how our beliefs change when the belief change method is applied to our preferences. Specifically, $\mathsf{C}_{\succ}(\cdot)$ is the belief change operator over the language $\mathcal{SL}$ which takes a piece of evidence E, applying it to the initial strongest belief $\mathsf{C}_{\succ}(\top)$, and giving as a result the updated strongest belief $\mathsf{C}_{\succ}(E)$. Proposition 3.8 addresses what are the logical properties of the nonmonotonic system **D** that are satisfied by the class of belief change operators.

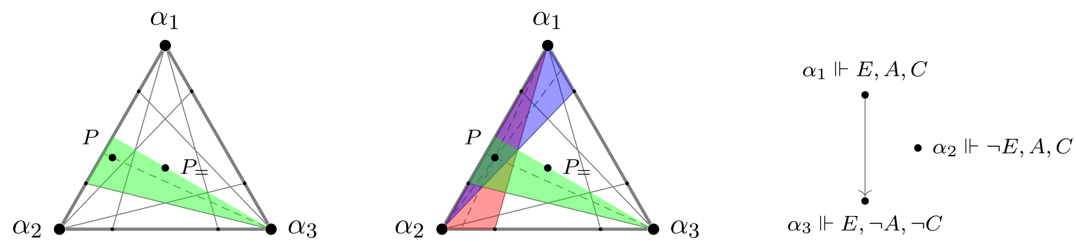

(A) Counterexample to Ref. (B) Counterexample to Or. (C) Counterexample to CMon.

FIGURE 6. C_P does not satisfy three properties of system **P**.

DEFINITION 3.7 (Belief change operator). *Let $\succ$ be a semi-order over Ω. Let E be a $\succ$-consistent formula. Then, the belief change operator $\mathsf{C}_\succ$ is defined as:*

$$\mathsf{C}_\succ(E) \equiv \bigvee \left\{\alpha_i \in max(\succ^E) : \succ^E \in \mathcal{C}(\succ, E)\right\}$$

PROPOSITION 3.8. *Let $\mathsf{C}_\succ(E)$ be a belief change operator defined as in Definition 3.7 and E some $\succ$-consistent formula. Then, C_P satisfies the rules of inference LLE, RW and And, while it does not satisfy the rules Ref, Or, CMon, DRat nor PlaMon of the nonmonotonic system* **D**.

That LLE, RW and And are satisfied directly follows from the definition of $\mathsf{C}_\succ(E)$. At the same time, Theorem 3.6 allows us to abstract away from the probabilities generating $\succ$ and still to consider the logical properties induced by preference change. Nevertheless, we can exploit the correspondence between preference change and Jeffrey conditioning in order to define the countermodels to the logical properties Ref, Or and Cmon. Specifically, in Figure 6 we consider again the convex set of probability distribution over three worlds. That Reflexivity is invalid is clear because belief chance is compatible with the updated preferential order ($\succ |\neg E$), whose maximal elements do not satisfy E. Analogously, in Figure 6a consider the initial preferential order $\succ_P$ and the evidence $E \equiv \alpha_1 \vee \alpha_2$. In green we highlight the preferential orders that can be obtained by Jeffrey conditioning P on E for every possible value of the parameter $e = P_E(E)$. The green area includes the point $(0, 0, 1)$, corresponding to α_3. For such point, both α_1 and α_2 get probability zero; hence, $\mathsf{C}_P(E) \vDash \alpha_3$.

As for Or, consider again the initial order $\succ_P$ in Figure 6b. The red area captures the orders that we could obtain by Jeffrey conditioning P on the evidence α_1, while the blue area the orders generated by Jeffrey conditioning P on the evidence α_2. In all such orders, the maximal elements are always α_1 or α_2 (or both). This is because, both α_1 and α_2 have an high enough probability with respect to the probability of α_3 that even when their probability decreases, they remain preferred to α_3. As a result, $\mathsf{C}_P(\alpha_1) \vDash \alpha_1 \vee \alpha_2$ and $\mathsf{C}_P(\alpha_2) \vDash \alpha_1 \vee \alpha_2$. However, consider finally the area in green again obtained by Jeffrey conditioning P on the evidence $\alpha_1 \vee \alpha_2$. As we discussed before, $\mathsf{C}_P(\alpha_1 \vee \alpha_2) \nvDash \alpha_1 \vee \alpha_2$ because Jeffrey conditioning on $\alpha_1 \vee \alpha_2$ might even eliminate the two atoms from the preferential order so giving α_3 as the only remaining belief.

Concerning the counterexample for CMon, Figure 6c still considers a case with three worlds, each corresponding to a distinct maximal conjunction of literals. It is easy to check that for such model $\mathsf{C}_\succ(E) \vDash A$ and $\mathsf{C}_\succ(E) \vDash C$ since the world $\alpha_1 \Vdash A \wedge C \wedge E$ always remains preferred to $\alpha_3 \Vdash \neg A \wedge \neg C \wedge E$, by Theorem 3.6 and by property (1) of belief change, as they both satisfy E and α_1 is initially preferred to α_3. However, $\mathsf{C}_\succ(A \wedge E) \nvDash C$ because, once the evidence switches from E to $E \wedge A$, α_1 is the only world satisfying the evidence, while both α_2 and α_3 satisfy its negation. At the same time, α_2 and α_3 are initially unrelated and are both maximal in the order $\succ |\neg(E \wedge A)$. Since α_3 falsifies C, it follows that there is at least one order in $\mathcal{C}_\succ(\succ, E \wedge A)$, namely $\succ |\neg(E \wedge A)$, whose maximal elements do not all validate C. CMon is invalidated.

Turning to DRat, Figure 7a gives a counterexample to the validity of the rule. Consider the initial probability distribution P. The red area captures the orders that we could obtain by Jeffrey conditioning P on the evidence $\alpha_1 \vee \alpha_2$. The maximal elements of all those orders are α_2 or α_3, possibly both, thus it holds that $\mathsf{C}_P(\alpha_1 \vee \alpha_2) \vDash \alpha_2 \vee \alpha_3$, first premise of DRat. At the same time, the blue area captures the orders that we could obtain by Jeffrey conditioning P on α_1, among which we have $\succ_P |\alpha_1$. Then, it clearly holds that $\mathsf{C}_P(\alpha_1) \nvDash \alpha_2 \vee \alpha_3$ since α_1 is among the possible outcomes. Thus, the second condition of DRat is satisfied. Finally, the green area captures the orders obtained by Jeffrey conditioning P on α_2. Among those,

(A) Counterexample to DRat. (B) Counterexample to PlaMon and RMon.

FIGURE 7. C_P does not satisfy DRat, PlaMon nor RMon.

we have $\succ_P |\alpha_1 \vee \alpha_3$, whose maximal elements are precisely α_1 and α_3. As a consequence, it holds that $\mathsf{C}_P(\alpha_2) \not\vDash \alpha_2 \vee \alpha_3$, falsifying DRat.

Finally, Figure 7b offers a counterexample to PlaMon and to RMon. Consider again the same initial point P. For the first condition of RMon ($\mathsf{C}_P(A) \vDash C$), consider the evidence $\alpha_1 \vee \alpha_2$ and all the orders in the red area obtained by Jeffrey conditioning P on $\alpha_1 \vee \alpha_2$. As we previously saw, $\mathsf{C}_P(\alpha_1 \vee \alpha_2) \vDash \alpha_2 \vee \alpha_3$ holds. As a trivial consequence, $\mathsf{C}_P(\alpha_1 \vee \alpha_2) \not\vDash \neg(\alpha_1 \vee \alpha_2 \vee \alpha_3)$ holds as well, PlaMon second premise ($\mathsf{C}_P(A) \not\vDash \neg B$). Then, assume $C \equiv D$, and from the first two conditions $\mathsf{C}_P(\alpha_1 \vee \alpha_2) \vDash \alpha_2 \vee \alpha_3$ holds as PlaMon third condition. Finally, consider the green area and all the orders obtained by Jeffrey conditioning P on $A \wedge B \wedge C$, here $(\alpha_1 \vee \alpha_2) \wedge (\alpha_2 \vee \alpha_3) \equiv \alpha_2$. As before, among those, we have $\succ_P |\alpha_1 \vee \alpha_3$, whose maximal elements are precisely α_1 and α_3. As a consequence, it holds that $\mathsf{C}_P(\alpha_2) \not\vDash \alpha_2 \vee \alpha_3$, falsifying PlaMon. In conclusion, to see that also RMon is invalid, consider again that $\mathsf{C}_P(\alpha_1 \vee \alpha_2) \vDash \alpha_2 \vee \alpha_3$ holds, from which $\mathsf{C}_P(\alpha_1 \vee \alpha_2) \not\vDash \neg(\alpha_2 \vee \alpha_3)$. Thus, $\alpha_2 \vee \alpha_3$ can be conjuncted to the initial premises. As we have seen however, $\mathsf{C}_P(\alpha_2) \not\vDash \alpha_2 \vee \alpha_3$ holds in the counterexample, making RMon invalid. In this case, it is clear that Belief change is so weak that even adding the same conclusion as an additional premise results in its own demise.

4. Belief revision and successful Jeffrey conditioning

In the previous section we saw how our preferences over the set of alternatives change when we update the underlying probability distribution by Jeffrey conditioning. Imagine that you

Systems			Rules	Conditions on C	C
			LLE	$\mathsf{C}(A) \vDash C, \vDash A \leftrightarrow B \Rightarrow \mathsf{C}(B) \vDash C$	✓
			RW	$\mathsf{C}(A) \vDash B, \vDash B \to C \Rightarrow \mathsf{C}(A) \vDash C$	✓
			Ref	$\mathsf{C}(A) \vDash A$	×
		P	And	$\mathsf{C}(A) \vDash B,\ \mathsf{C}(A) \vDash C \Rightarrow \mathsf{C}(A) \vDash B \wedge C$	✓
			CMon	$\mathsf{C}(A) \vDash B,\ \mathsf{C}(A) \vDash C \Rightarrow \mathsf{C}(A \wedge B) \vDash C$	×
			Or	$\mathsf{C}(A) \vDash C,\ \mathsf{C}(B) \vDash C \Rightarrow \mathsf{C}(A \vee B) \vDash C$	×
	D		DRat	$\mathsf{C}(A \vee B) \vDash C,\ \mathsf{C}(A) \nvDash C \Rightarrow \mathsf{C}(B) \vDash C$	×
			PlaMon	$\mathsf{C}(A) \vDash C,\ \mathsf{C}(A) \nvDash \neg B,\ \mathsf{C}(A \wedge B) \vDash D \Rightarrow \mathsf{C}(A \wedge B \wedge C) \vDash D$	×
R			RMon	$\mathsf{C}(A) \vDash C,\ \mathsf{C}(A) \nvDash \neg B \Rightarrow \mathsf{C}(A \wedge B) \vDash C$	×

TABLE 2. Belief revision induced by Bayesian conditioning.

start from a semi-order $\succ$ describing you preferences over the alternatives and we stick to the preferences without bothering to understand which precise probability distribution generated your preferences – such an additional task might be too demanding in terms of say time or cognitive energy. Theorem 3.6 establishes that even when that is the case, because the set $\mathcal{C}(\succ, E)$ does not depend on any prior probability distribution but on you preferences $\succ$ only, you can be certain that your updated preferences will be given by one of the orders in $\mathcal{C}(\succ, E)$ when you acquire some evidence concerning E. In other terms, Theorem 3.6 allows us to abstract form the single probability functions, to focus only on our preferences and how they change when some evidence about E is learned.

We use the expression that some evidence *about* E is learned because up to this point the updates that we are taking into account, Jeffrey conditioning and belief change, might as well confirm or disconfirm E. Specifically, for Jeffrey conditioning, the probability of E might increase or decrease, while for preference change, the position of the worlds consistent with E might move up or down along the preferential order. When we move to consider the beliefs that are generated by the preferential orders, the variability that both the quantitative and the qualitative update methods allow results in beliefs which are very weak to new evidence. Indeed, Proposition 3.8 shows that the belief change operator $\mathcal{C}(E)$ does not validate many of the logical properties which have proven to be crucial in the field of nonmonotonic logics and belief revision. As an example, Reflexivity is not satisfied as changing your credence about E by Jeffrey conditioning does not guarantee that you end up believing in E to be true.

As a consequence, if on the one hand we want to generalize the approch given by Lin and Kelly assuming Jeffrey conditioning as a probabilistic update, it is clear that the general definition of Jeffrey conditioning with $e \in [0, 1]$ is too broad, it captures a notion of evidence which is too weak and it does not capture a robust enough notion of belief update. In the present section, we give a first attempt to restrict the general notion of Jeffrey conditioning by imposing E is among the beliefs generated by P_E^e. Any instance of JC P_E^e which satisfies such constraint is referred to in the literature as *Successful Jeffrey conditioning*. The success comes from the fact that the new evidence concerning E is strong enough that the agent ends up believing E to be true. More precisely, the uncertainty $1 - e$ that the agents assigns to E is low enough that the agent believes E to be true. Analogously, one can interpret such restricted version of JC as the fact that the agent did not just learn some information about E; rather they learned that E *is true*, but they don't fully trust the information or the source of information. As a result, the agent does indeed believe E to be true, without however assigning to E maximal probability, keeping some non-null degree of uncertainty $1 - e$ about E.

Definition 4.1 (Successful Jeffrey Conditioning). *Let P a probability distribution over Ω and E a proposition such that $P(E) > 0$. Then, P_E^e is an instance of successful Jeffrey conditioning if $\mathsf{B}_{P_E^e}(\top) \vDash E$. $\mathcal{J}^s(P, E)$ is the set of all probability distributions resulting from successful Jeffrey conditioning P on E.*

As we will see later, this is a quite strong restriction on Jeffrey conditioning. In accordance with the present approach, we assume that the beliefs generated by the probability function P^e_E are every proposition that is implied by all the maximal elements of the semi-order $\succ_{P^e_E}$ induced by P^e_E, i.e. $\mathsf{B}_{P^e_E}(\top) \equiv \bigvee max(\succ_{P^e_E})$ and $\mathsf{B}_{P^e_E}(\top) \vDash E$. Notice that such definition depends on the semi-order induced by P^e_E, which, as we know, is generated here by the odds-threshold method. As a result, we could avoid mentioning the belief operator $\mathsf{B}_{P^e_E}$ at all and use the following equivalent property to define successful JC: an instance of JC is successful if and only if

$$max_{\alpha_i \in \widehat{E}} P^e_E(\alpha_i) \geq t \cdot max_{\alpha_j \in \widehat{\neg E}} P^e_E(\alpha_j).$$

Because every world satisfying $\neg E$ is defeated by some world satisfying E (specifically the one which maximizes probability), the beliefs generated by the posterior probability P^e_E imply E.

Because we already dispose of a method to generate beliefs from a probability distribution, we can consider some examples which illustrate the impact of imposing the restriction $\mathsf{B}_{P^e_E}(\top) \vDash E$ on Jeffrey conditioning. In Figure 8 we consider again the convex set of probability distributions over a set of three worlds. In particular, in Figure 8a, $P(E)$ is low enough that the agent does not initially believe E to be true. On the contrary, after learning E, because we are now restricted to cases of successful JC, we hit the blue area, all the points P' of which generate a strongest belief which implies E, i.e. $\mathsf{B}_{P'}(\top) \vDash E$. Those points on the green segment which is included in the blue area are all and only the probability distributions in $\mathcal{J}^s(P, E)$. The same holds in Figure 8b with the only difference that there the prior probability P assigns to E a probability high enough that $\mathsf{B}_P(\top) \vDash E$. We chose the two points in the two figures so that one is an instance of JC of the other. As a consequence, notice that, despite the differences between the two initial points P, the two green segments representing all the instance of successful JC are identical. This matches the fact that even under the new restriction for successful JC, if two probability distributions are one an instance of JC of the other, then the probability distributions resulting from successful JC for the two are exactly the same.

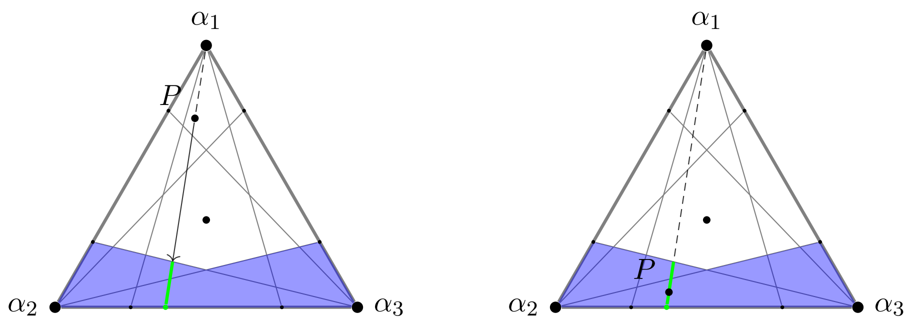

(A) $\mathsf{B}_P(\top) \not\vDash E$. (B) $\mathsf{B}_P(\top) \vDash E$.

FIGURE 8. Successful Jeffrey conditioning, $\mathcal{J}^s(P, E)$.

The whole point of restricting Jeffrey conditioning to its successful instances is to obtain a stronger probabilistic update which, when analyzed in terms of the impact it has on our beliefs, induces an update of our beliefs which satisfies all the desired properties of belief revision, namely the rules of system **P**. As we did for the general version of JC and preference change in the previous section, here as well we qualitatively characterize in Definition 4.2 the set of semi-orders that can be obtained by all the instances of successful Jeffrey conditioning. Such set is directly obtained from the set of preference change that was defined in the previous section by imposing the condition (Success). Theorem 4.3 proves the desired correspondence between preference revision and successful Jeffrey conditioning,

DEFINITION 4.2 (Preference revision). *Let $\succ$ be a semi-order over Ω_0 and E be a $\succ$-consistent formula. Then, an order $\succ^E$ is a result of an E-induced preference revision on $\succ$ if and only if $\succ^E \in \mathcal{C}(\succ, E)$ and $\succ^E$ satisfies (Success):*

(4) for all $\alpha_j \Vdash \neg E$ there exists $\alpha_i \Vdash E$ s.t. $\alpha_i \succ^E \alpha_j$. *(Success)*

$\mathcal{R}(\succ, E)$ is the set of all the resulting orders $\succ^E$.

THEOREM 4.3. *Let P be a probability distribution over Ω. Let $\mathcal{J}^s(P, E)$ the set of probability distributions P^e_E resulting from successful Jeffrey conditioning P on E. $\mathcal{O}(\mathcal{J}^s(P, E))$ is the set of odds-threshold orders generated by the probability distributions in $\mathcal{J}^s(P, E)$. Finally, let $\mathcal{R}(\succ_P, E)$ be the set of orders resulting from $\succ_P$ by belief revision on E, as defined* *in Definition 4.2. Then, the following holds:*

$$\mathcal{R}(\succ_P, E) = \mathcal{O}\left(\bigcup \{\mathcal{J}^s(P', E) : \ P' \in \mathcal{P}(\succ_P)\}\right).$$

Proof. The result follows from Theorem 3.6, from which the following inequality holds: $\mathcal{R}(\succ_P, E) = \mathcal{O}\left(\bigcup \{\mathcal{J}^s(P', E) : \ P' \in \mathcal{P}(\succ_P)\}\right)$. $\mathcal{R}(\succ_P, E)$ is a subset of $\mathcal{C}(\succ_P, E)$, specifically it is the set of the orders in $\mathcal{C}(\succ_P, E)$ that additionally satisfy condition (4); at the same time, $\bigcup \{\mathcal{O}(\mathcal{J}^s(P', E)) : \succ_{P'} = \succ_P\}$ is analogously obtained by imposing (Success) on the possible instances of JC. Therefore, it is sufficient to show that the two conditions are equivalent. This is immediate as also (Success) is a condition on the order generated by the posterior distribution obtained by JC. Specifically, (Success) imposes that $max(\succ_{P_E}) \subseteq \widehat{E}$. By definition of B, this is the case if and only if all the maximal elements of $\succ_{P_E}$ belong to E, if and only if for any atom $\alpha_j \Vdash \neg E$ there exists $\alpha_i \Vdash E$ s.t. $\alpha_i \succ_{P_E} \alpha_j$. The restrictions on $\mathcal{C}(\succ_P, E)$ and $\mathcal{O}(\mathcal{J}(P', E))$ are equivalent; given the initial equality between the two sets, the desired result follows. □

Just as in the case of Jeffrey conditioning and preference change, the correspondence established by Theorem 4.3 is the strongest one can hope for. It states that, starting from a semi-order $\succ$, the set of semi-orders satisfying the properties (0)-(4) are precisely those semi-orders that can be obtained from any of the probability functions that generate the same $\succ$ by successful JC. Since the preference revision only relies on the initial order $\succ$, it cannot determine which probability distribution actually generated $\succ$. As a result, it has to consider as possible revisions of the preferences all the orders in $\mathcal{J}^s(P, E)$ for all $P \in \mathcal{P}(\succ)$. This is true because any order in $\mathcal{O}\left(\bigcup \{\mathcal{J}^s(P', E) : \ P' \in \mathcal{P}(\succ_P)\}\right)$ is a member of $\mathcal{R}(\succ, E)$. At the same time, it is desirable that the conditions imposed by Definition 4.2 are strong enough not to allows any other order in $\mathcal{R}(\succ, E)$. This is also true by the converse direction of set inclusion. Therefore, Theorem 4.3 tells us that Definition 4.2 captures precisely the properties corresponding to the impact of successful Jeffrey conditioning on our preference over alternatives, and any additional proper restriction would be unfounded.

As an example consider Figure 9. On the left the green segment captures the instances obtained from P by successful JC on $\alpha_2 \vee \alpha_3$. The order $\succ_2$ is clearly incompatible with

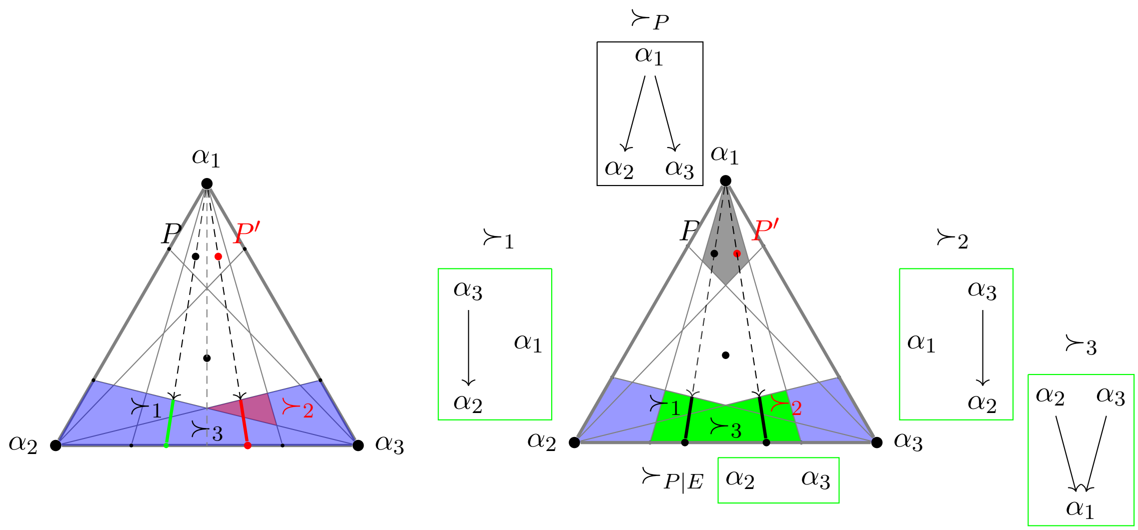

(A) $\mathcal{J}^s(P, E)$.

(B) $\mathcal{R}(\succ_P, E)$.

FIGURE 9. Successful Jeffrey conditioning, $\mathcal{J}^s(P, E)$.

$\mathcal{J}(P, E)$ since $\succ_2$ assigns to α_3 a higher rank than α_2. By Lemma 3.1, it follows that any probability distribution generating $\succ_2$ has to assign to α_3 a strictly higher probability than α_2. Since on the contrary P makes α_2 strictly more likely than α_3, it follows that $\alpha_3 \notin \mathcal{J}(P, E)$. A fortiori, it holds that $\alpha_3 \notin \mathcal{J}^s(P, E)$. If we consider Figure 9b instead, the green area captures all the semi-orders in $\mathcal{R}(\succ_P, E)$. Indeed, $\succ_1$, $\succ_2$, $\succ_3$ and $\succ_4$ are the only four semi-orders satisfying all the four conditions imposed by Definition 4.2.

The correspondence proven by Theorem 4.3 allows to consider the impact of successful JC on our beliefs. This turns out to be a quite easy task. Indeed, Proposition 4.4 proves that, fixed a prior probability P, all the instances of successful Jeffrey conditioning generate distinct semi-orders which share the same maximal elements. In particular then, all such instances generate semi-orders with the same maximal elements as the one generated by Bayesian conditioning, i.e. $e = 1$, which at its own turn has the same maximal elements of $\succ_P |E$ by [LK12b] Theorem 2 (see Figure 2). As a consequence, as soon as Jeffrey conditioning "hits *Success*", that is, as soon as it assigns high enough probability to E that E is believed to be true, then the agent's beliefs remain the same even if the uncertainty about E, $1-e$, reduces and eventually goes to 0.

PROPOSITION 4.4. *Let P be a probability distribution over Ω. For every $P_E^e \in \mathcal{J}^s(P,E)$, it holds that $max(\succ_{P_E^e}) = max(\succ_P |E)$. It follows that $\mathsf{B}_{P_E^e}(\top) \equiv \mathsf{B}_{P_E^1}(\top)$.*

PROOF. Since $P_E^e \in \mathcal{J}^s(P,E)$, we know by Definition 4.1 that $max(\succ_{P_E^e}) \subseteq \widehat{E}$. Let $\alpha_i, \alpha_k \in \widehat{E}$. Then, $\alpha_i \succ_{P_E^e} \alpha_k$ iff $P_E^e(\alpha_i) \geq t \cdot P_E^e(\alpha_k)$ iff $P(\alpha_i) \geq t \cdot P(\alpha_k)$ iff $\alpha_i \succ_P \alpha_k$ iff $\alpha_i(\succ_P |E)\alpha_k$. The second part directly follows as each strongest belief B is just the disjunction of the maximal worlds of the respective semi-order. □

Theorem 4.3 proves a correspondence between successful JC and preference revision. We discussed of how such correspondence shows that there is no possible way to qualitatively update our preferences so that the resulting semi-order always matches the change of our preference when the underlying probability distribution is updated by successful JC. For that reason, Definition 4.2 revises the initial preferential order $\succ_P$ so to allow all semi-orders that can be generated by the instances of successful JC for all probability functions $P' \in \mathcal{P}(\succ_P)$. Thus, the correspondence between preference revision and successful Jeffrey conditioning is weaker than the one established by Lin and Kelly between $\succ_{P|E}$ and $\succ_P |E$. Nevertheless, Proposition 4.4 shows that a stronger consonance between the qualitative and the quantitative update methods is preserved at the level of beliefs since, as soon as e is high enough that $\mathsf{B}_{P_E^e}(\top) \vDash E$, the maximal elements of $\succ_{P_E^e}$ are exactly the same of $\succ_P |E$. Therefore, we can now give the following definition of belief revision operator, and show that any such operator satisfy all and only the inference rules of system **D**, i.e. system **P** together with Disjunctive Rationality and Plausible Monotonicity.

DEFINITION 4.5. *Let $\succ$ be a semi-order over a set Ω_0 and E be a $\succ$-consistent formula. Then, the belief revision operator $\mathsf{B}_\succ$ is defined as:*

$$\mathsf{B}_\succ(E) \equiv \bigvee \left\{\alpha_i \in max(\succ^E) : \succ^E \in \mathcal{R}(\succ, E)\right\}.$$

THEOREM 4.6. *$\mathsf{B}_\succ$ is a belief revision operator defined as in Definition 4.5 if and only if $\mathsf{B}_\succ$ satisfies system **D**.*

Proof. The proof is immediate. By Proposition 4.4, we know that $max(\succ_{P^e_E}) = max(\succ_P |E)$. Given by Theorem 4.3 that $\mathcal{R}(\succ_P, E) = \mathcal{O}\left(\bigcup\{\mathcal{J}^s(P', E) : P' \in \mathcal{P}(\succ_P)\}\right)$, we can apply Proposition 4.4 to all such $P' \in \mathcal{P}(\succ_P)$. Because they all share the same initial order $\succ_P = \succ_{P'}$, it immediately follows that for all $\succ^E \in \mathcal{R}(\succ, E)$, $max(\succ^E) = max(\succ |E)$. Therefore, $\mathsf{B}_\succ(E) = max(\succ |E)$. Theorem 4.1 proves that a belief revision operator can be defined by $\succ |E$, with $\succ$ a semi-order and E a $\succ$-consistent formula, if and only if it satisfies the rules of system **D**. Since both $\mathcal{R}(\succ, E)$ and $\succ |E$ are defined under the same conditions of $\succ$ being a semi-order and E a $\succ$-consistent formula, we can conclude that $\mathsf{B}_\succ$ is a belief revision operator defined as in Definition 4.5 if and only if it satisfies system **D**. □

5. Iterated belief revision

In the previous section we proved what are the rules of simple, non-iterated belief revision that are satisfied when the underlying probability function is updated by successful Jeffrey conditioning. Such rules correspond to the nonmonotonic system **D**, as proven by Theorem 4.6. We should not stop at simple belief revision. Indeed, one crucial aspect of Jeffrey conditioning is that, by not assuming the evidence to be certain, it allows to consider sequences of evidences, the bits of which can even disconfirm or contradict each other. Thus, it is natural to consider the extension of the belief revision operators defined in Definition 4.5 to allow sequences of evidences.

Definition 5.1 (Iterated preference revision). *Let $\succ$ be a semi-order over Ω_0. Let $\mathcal{E} = (E_1, \dots, E_n)$ be a list of formulas and let $\mathcal{E}_i = (E_1, \dots, E_i)$ be the first i elements of $\mathcal{E}$, with $i, n \in \mathbb{N}$ and $i \leq n$. The preference revision $\mathcal{R}$ defined in Definition 4.2 is extended to extended to sequences $\mathcal{E}$ as follows. First, define $\mathcal{R}(\succ, ()) := \{\succ\}$. Then, $\mathcal{R}(\succ, \mathcal{E})$ is defined by the following:*

$$\mathcal{R}(\succ, \mathcal{E}_i) := \bigcup\{\mathcal{R}(\succ^{\mathcal{E}_{i-1}}, E_i) \ : \succ^{\mathcal{E}_{i-1}} \in \mathcal{R}(\succ, \mathcal{E}_{i-1})\}.$$

Definition 5.2 (Iterated belief revision operators). *Let $\succ$ be a semi-order over Ω_0. Let $\mathcal{E}$ be a sequence of n formulas in $\mathcal{SL}$. The iterated belief revision operator $\mathsf{B}_\succ$ is defined as* *follows:*

$$\mathsf{B}_{\succ}(\mathcal{E}) \equiv \bigvee\{\alpha_i \in max(\succ^{\mathcal{E}}) : \succ^{\mathcal{E}} \in \mathcal{R}(\succ, \mathcal{E})\}.$$

The iterated belief revision operator $\mathsf{B}_{\succ}(\mathcal{E})$ is the disjunction of the maximal elements of all the semi-orders that can be obtained by applying the preference revision defined in Definition 4.2 on the formula E_i for each of the semi-orders $\succ^{\mathcal{E}_{i-1}} \in \mathcal{R}(\succ, \mathcal{E}_{i-1})$. Because in general $\mathcal{R}(\succ, E)$ contains more than one semi-order, it is clear that at each iteration of preference revision the number of semi-orders in $\mathcal{R}(\succ^{\mathcal{E}_{i-1}}, E_i)$ increase. Nevertheless, as the following proposition shows, the iterated belief revision operator $\mathsf{B}_{\succ}(\mathcal{E})$ satisfies two of the desired logical properties for iterated belief revision introduced in [DP97] by Adnan Darwiche and Judea Pearl. The main reason why the growth of the number of semi-order does not entirely invalidate the logical properties of the iterated belief revision operator lies in the fact that, fixed some $\succ^{\mathcal{E}_{i-1}}$, all the semi-orders in $\mathcal{R}(\succ^{\mathcal{E}_{i-1}}, E_i)$ still share the same maximal elements. As a result, differently from the preferential orders, the beliefs are less impacted by such variability.

Following the original paper [DP97], the explanations behind the rules (C1)-(C4) are the following. (C1) states that when two pieces of information arrive and the second one is more specific than the first, then the first is redundant, in the sense that the second evidence alone would yield the same beliefs. (C2) states instead that when two pieces of evidence are acquired in sequence, the last one prevails. (C4) instead states that no evidence can contribute to its own demise; if C is not contradicted after learning A, then it should remain uncontradicted when A is preceded by C itself. We kept (C3) for last because it is the only one of the four main rules proposed by Darwiche and Pearl which is not validated by $\mathsf{B}_{\succ}(\mathcal{E})$. In the words of [DP97], (C3) states that C should be retained after accommodating a more recent evidence A that implies C given the current beliefs.

One can also interpret (C3) also as a variation of (C4), stating that no evidence can contribute to its own withdraw. In confirmation of this interpretation, let us consider again the counterexample to (C3) given in Figure 10. Taking P as initial probability function, successful Jeffrey conditioning does require E to be implied by the beliefs generated by the posterior probability P_E^e. However, it does not require the probability of E to increase. In other words, successful JC allows the probability of the evidence itself E to decrease. Indeed, in Figure 10, the instance of successful JC that generates the counterexample, P_C assigns to the evidence C a lower probability – this can be verified by the fact that P_C is closer than P to the vertex α_3, the only one satisfying $\neg C$, and analogously P_E is more distant that P from the segment $\alpha_1 - \alpha_2$, the probability distributions over which assign to C value 1. Thus, even though P is revised after learning that E is true, the uncertainty $1-e$ concerning E increase, i.e. $e < P(E)$. And the uncertainty of E increases to the point that the two alternative $\neg A \wedge C$ and $A \wedge \neg C$ swap, the latter one becoming preferred to the former despite the fact that only the former is logically compatible with the evidence C – the right of Figure 10 highlights in orange such swap.

PROPOSITION 5.3. *Let* $\mathsf{B}_\succ$ *be a iterated belief revision operator defined in Definition 5.2.* $\mathsf{B}_\succ$ *satisfies system* **D** *together with the rules (C1), (C2) and (C4), while it falsifies (C3).*

$$A \vDash C \ \Rightarrow\ \mathsf{B}(C, A) \equiv \mathsf{B}(A) \tag{C1}$$

$$A \vDash \neg C \ \Rightarrow\ \mathsf{B}(C, A) \equiv \mathsf{B}(A) \tag{C2}$$

$$\mathsf{B}(A) \vDash C \ \Rightarrow\ \mathsf{B}(C, A) \vDash C \tag{C3}$$

$$\mathsf{B}(A) \nvDash \neg C \ \Rightarrow\ \mathsf{B}(C, A) \nvDash \neg C \tag{C4}$$

PROOF.

That $\mathsf{B}_\succ$ satisfies system **D** is a direct consequence of Definition 6.7 and Theorem 4.6. We address now the rules of iterated belief revision.

(C1). For every $\succ^{(A)} \in \mathcal{R}(\succ, (A))$, it holds by (Success) that $max(\succ^{(A)}) \subseteq \widehat{A}$. Additionally, by property (1) of Def. 3.5, we know that all the $\succ$-relations between atoms satisfying A are preserved in $\succ^{(A)}$. Therefore, $max(\succ^{(A)}) = max(\succ |A)$. In the same way, we can obtain that $max(\succ^{(C)}) = max(\succ |C)$. Additionally, since $A \vDash C$, we have that $max(\succ^{(C)}) = max(\succ |C) \supseteq max(\succ |A)$. Therefore, for every $\succ^{(C,A)} \in \mathcal{R}(\succ, (C, A))$, by (Success) and (1) applied again, we obtain that $max(\succ^{(C,A)}) = max(\succ|A)$. Hence, for every $\succ^{(A)} \in \mathcal{R}(\succ,(A))$ and for every $\succ^{(C,A)} \in \mathcal{R}(\succ,(C,A))$, $max(\succ^{(A)}) = max(\succ|A) = max(\succ^{(C,A)})$. The desired result holds: $\mathsf{B}_{\succ}(C,A) \equiv \bigvee max(\succ|A) \equiv \mathsf{B}_{\succ}(A)$.

(C2). It was established in (C1) that, for every $\succ^{(A)} \in \mathcal{R}(\succ,(A))$, $max(\succ^{(A)}) = max(\succ|A)$. Take any $\succ^{(C)} \in \mathcal{R}(\succ,(C))$. By property (1) of Def. 3.5, all the $\succ$-relations between atoms satisfying $\neg C$ are preserved in $\succ^{(C)}$; since by assumption $A \vDash \neg C$, all the $\succ$-relations between atoms satisfying A are preserved in $\succ^{(C)}$. As a consequence, it holds that $max(\succ|A) = max(\succ^{(C)}|A)$. Now take any $\succ^{(C,A)} \in \mathcal{R}(\succ,(C,A))$; by definition of $\mathcal{R}(\succ,(C,A))$, it holds that $\succ^{(C,A)} \in \mathcal{R}(\succ^{(C)},A)$ for some $\succ^{(C)} \in \mathcal{R}(\succ,(C))$. Then, we already know that $max(\succ^{(C,A)}) = max(\succ^{(C)}|A)$. As $max(\succ|A) = max(\succ^{(C)}|A)$ holds for any $\succ^{(C)} \in \mathcal{R}(\succ,(C))$, it is a general condition that $max(\succ^{(C,A)}) = max(\succ|A)$ for any $\succ^{(C,A)} \in \mathcal{R}(\succ,(C,A))$. Hence, for every $\succ^{(A)} \in \mathcal{R}(\succ,(A))$ and for every $\succ^{(C,A)} \in \mathcal{R}(\succ,(C,A))$, $max(\succ^{(A)}) = max(\succ|A) = max(\succ^{(C,A)})$. The desired result holds: $\mathsf{B}_{\succ}(C,A) \equiv \bigvee max(\succ|A) \equiv \mathsf{B}_P(A)$.

(C3). A counterexample is given in Figure 10.

(C4). The premise is equivalent to $max(\succ|A) \cap \widehat{C} \neq \emptyset$. Thus we know that there is $\alpha_i \in \Omega_0$ that satisfies both A and C. Among $\mathcal{R}(\succ,C)$ there is $\succ|C$. Since α_i is an element of $\succ|C$, A is a $\succ|C$-consistent formula, and $\succ|C$ can be revised on A. Hence, $(\succ|C)|A = \succ|A \wedge C$ belongs to $\mathcal{R}(\succ,(C,A))$. Then, $max(\succ^{(C,A)}) \cap \widehat{C} \neq \emptyset$ for some $\succ^{(C,A)} \in \mathcal{R}(\succ,(C,A))$, namely $\succ|A \wedge C$. As a result, $\mathsf{B}_{\succ}(C,A) \nvDash \neg C$. □

6. Positive evidence and confirmatory iterated belief revision

We showed that successful Jeffrey conditioning is much stronger than the general version of Jeffrey conditioning. As a result, it induces on the beliefs a revision which satisfies the rules of system **D** for the simple revision, and rules (C1), (C2) and (C4) for the iterated belief revision. Nevertheless, it does not satisfy (C3), a rule over which there is a widespread consensus in the literature o iterated belief revision. And the reason because (C3) is not satisfied seems to lie in the fact that, despite updating on the evidence E, successful JC allows the probability of E to decrease. Therefore, in the following, we will impose the additional constraint on Jeffrey

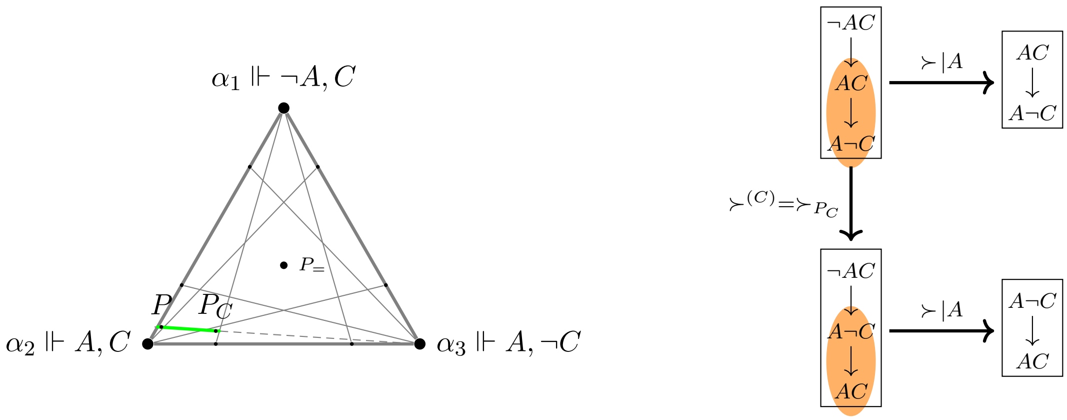

FIGURE 10. Counterexample to C3.

conditioning which requires the posterior probability of the evidence to be greater than its prior probability, i.e. $e \geq P(E)$. Precisely because the probability of E can only increase, we dubbed this version of Jeffrey conditioning *confirmatory*. The label is derived from Bayesian confirmation theory – see [CTG07], [Fit99] and [Fit06] for an introduction on the topic and for many aspects which make confirmation a crucial notion in philosophy, in logic and in psychology. In a gist, generally it is said that evidence E confirms or supports a hypothesis H when E and H are positively probabilistically correlated. In the present case however, the proposition that is positively confirmed is E itself. then, a reasonable interpretation is the following: the agent receives from some source the information that E is true; differently from the case of (successful) JC, here the agent always believes the source to be reliable, reliable at least to the degree that their uncertainty about E does not increase.

DEFINITION 6.1 (Confirmatory successful Jeffrey Conditioning). *Let Ω be a non-empty set and P a probability distribution over Ω. Then, P^e_E is an instance of confirmatory successful Jeffrey conditioning if respectively $e \geq P(E)$ and $\mathsf{B}_{P^e_E}(\top) \vDash E$.*

Once the notion of confirmatory Jeffrey conditioning is defined, we proceed as before by giving an analogous condition which further restricts preference revision. This condition is named (Success) and once it is stated out, we first prove in Theorem 6.3 the correspondence at the level of preference of confirmatory preference revision and confirmatory successful Jeffrey conditioning. Once this desired correspondence is prove, we will extend again confirmatory preference revision to sequences of evidences so to make an iterated version of it. Finally, we will fined the iterated belief revision operators and show that they satisfy all the properties of system **D** for the simple revision, and (C1)-(C4) for the iterated part, as desired.

Definition 6.2 (Confirmatory preference revision). *Let $\succ$ be a semi-order defined over Ω_0 and E a $\succ$-consistent formula. Then, an order $\succ^E$ is a result of an E-induced confirmatory preference revision on $\succ$ if and only if $\succ^E = (\succ |E)$ or $\succ^E$ belongs to $\mathcal{R}(\succ, E)$ and satisfies (5):*

(5) for all $\alpha_i \Vdash E$, $Up_{\succ^E}(\alpha_i) \subseteq Up_{\succ}(\alpha_i)$ and $Down_{\succ^E}(\alpha_i) \supseteq Down_{\succ}(\alpha_i)$. (Order Confirmation)

$\mathcal{R}^c(\succ, E)$ is the set of all the resulting orders $\succ^E$.

Theorem 6.3. *Let P be a probability distribution over Ω. Let $\mathcal{J}^c_s(P, E)$ the set of possible probability distributions P^e_E resulting from confirmatory and successful Jeffrey conditioning P on E. $\mathcal{O}(\mathcal{J}^c_s(P, E))$ is the set of odds-threshold orders generated by the probability distributions in $\mathcal{J}^c_s(P, E)$. Finally, let $\mathcal{R}^c(\succ_P, E)$ be the set of orders resulting from $\succ_P$ by confirmatory preference revision on E, as defined in Definition 6.2. Then, the following holds:*

$$\mathcal{R}^c(\succ_P, E) = \mathcal{O}\left(\bigcup\left\{\mathcal{J}^c_s(P', E) : \ P' \in \mathcal{P}(\succ_P)\right\}\right).$$

Proof.

Theorem 4.3 proves that $\mathcal{R}(\succ_P, E) = \mathcal{O}\left(\bigcup\{\mathcal{J}^s(P', E) : \ P' \in \mathcal{P}(\succ_P)\}\right)$. Here we want to prove an analogous correspondence between two sets obtained by respectively imposing (Order Confirmation) on $\mathcal{R}(\succ_P, E)$, and $e \geq P'(E)$ on $\mathcal{J}(P', E)$. Hence, to obtain the result, it is enough to prove that the two conditions are equivalent given the other assumptions. By [LK12b] Theorem 2, we already know that, for every $P' \in \mathcal{P}(\succ_P)$, when $e = 1$ we get that $\succ_{P'^1_E} = (\succ_P |E)$. Hence, we can focus on the non extreme cases where, on the one hand $\succ^E_P \neq (\succ_P |E)$, and on the other $e \in [P(E), 1)$ is taken as E's posterior probability given by Jeffrey conditioning.

First, we show that by imposing $e \in [P'(E), 1)$ on $\mathcal{J}(P', E)$, every preferential order so obtained satisfies (Order Confirmation), i.e. $Up_{\succ_P^E}(\alpha_i) \subseteq Up_{\succ_P}(\alpha_i)$ and $Down_{\succ_P^E}(\alpha_i) \supseteq Down_{\succ_P}(\alpha_i)$. Take any $\alpha_i \Vdash E$. Jeffrey conditioning preserves all the $\succ_{P'}$ relations among atoms satisfying E. Additionally, for any $\alpha_j \Vdash \neg E$ and any $e \geq P'(E)$, we get

$$\alpha_i \succ_{P'} \alpha_j \;\Leftrightarrow\; P'(\alpha_i) \geq t \cdot P'(\alpha_j) \;\Rightarrow\; P'_E(\alpha_i) \geq t \cdot P'_E(\alpha_j) \;\Leftrightarrow\; \alpha_i \succ_{P'_E} \alpha_j.$$

Therefore, if $\alpha_i \succ_{P'} \alpha_j$, then $\alpha_i \succ_{P'^E} \alpha_j$. Likewise, we get that if $\alpha_j \not\succ_{P'} \alpha_i$, then $\alpha_j \not\succ_{P'_E} \alpha_i$. From the two, (Order Confirmation) follows.

For the left to right direction of the theorem, we need to prove that if $\succ_P^E \notin \mathcal{O}(\bigcup\{\mathcal{J}_s^c(P', E) : P' \in \mathcal{P}(\succ_P)\})$, then $\succ_P^E \notin \mathcal{R}^c(\succ_P, E)$. If the premise is satisfied because $\succ_P^E \notin \mathcal{O}(\bigcup\{\mathcal{J}^s(P', E) : P' \in \mathcal{P}(\succ_P)\})$, then by Theorem 4.3 we obtain that $\succ_P^E \notin \mathcal{R}(\succ_P, E) \supseteq \mathcal{R}^c(\succ_P, E)$ as desired. Hence, assume that $\succ_P^E \in \mathcal{O}(\bigcup\{\mathcal{J}^s(P', E) : P' \in \mathcal{P}(\succ_P)\}) = \mathcal{R}(\succ_P, E)$.
Since $\succ_P^E \in \mathcal{O}(\bigcup\{\mathcal{J}^s(P', E) : P' \in \mathcal{P}(\succ_P)\}) \setminus \mathcal{O}(\bigcup\{\mathcal{J}_s^c(P', E) : P' \in \mathcal{P}(\succ_P)\})$, we know that there is no $Q \in \mathcal{P}(\succ_P)$ whose posterior probabilities Q_E^e obtained by Jeffrey Conditioning for any $e \geq Q(E)$ generate $\succ_P^E$, i.e. $\succ_P^E \neq \succ_{Q_E^e}$ for any $e \geq Q(E)$. Since $\succ_P^E \in \mathcal{R}(\succ_P, E)$, we know by Theorem 4.3 that there is $Q \in \mathcal{P}(\succ_P)$ such that Q_E^e for some $e < Q(E)$ generates $\succ_P^E$. Let Q be any such probability function. Remember that $\succ_Q = \succ_P$. Because $e < Q(E)$, we know that for every $\alpha_i \Vdash E$, $Up_{\succ Q}(\alpha_i) \supseteq Up_{\succ Q_E^e}(\alpha_i)$ and $Down_{\succ Q}(\alpha_i) \subseteq Down_{\succ Q_E^e}(\alpha_i)$. At the same time, $\succ_{Q_E^e} \neq \succ_Q$ as $\succ_Q \in \mathcal{O}(\mathcal{J}^c(P, E))$ while $\succ_{Q_E^e} \in \mathcal{O}(\bigcup\{\mathcal{J}^c(P', E) : P' \in \mathcal{P}(\succ_P)\})$. Therefore, $\succ_Q$ and $\succ_{Q_E^e}$ cannot have the same relations. Because Jeffrey conditioning preserves the preferential relations over $\widehat{E}$ and $\widehat{\neg E}$, there must be $\alpha_i \Vdash E$ and $\alpha_j \Vdash \neg E$ such that $\succ_Q$ and $\succ_{Q_E^e}$ relate the two differently. Hence, by combing the two facts, one of the following cases must hold: $Up_{\succ Q}(\alpha_i) \supset Up_{\succ Q_E^e}(\alpha_i)$ and $Down_{\succ Q}(\alpha_i) \subseteq Down_{\succ Q_E^e}(\alpha_i)$, or $Up_{\succ Q}(\alpha_i) \supseteq Up_{\succ Q_E^e}(\alpha_i)$ and $Down_{\succ Q}(\alpha_i) \subset Down_{\succ Q_E^e}(\alpha_i)$. Both cases contradict (Order Confirmation). The contrapositive has thus been obtained, and the proof is concluded. □

Theorem 6.3 shows that the condition (Order Confirmation) added in Definition 6.2 captures the impact that confirmatory successful JC has on our preferences. In particular, the theorem states that a semi-order $\succ^E$ results from confirmatory successful JC from a probability function generating $\succ$ if and only if it is the result of preference revision and for every worlds α_i satisfying E, the Up-set of α_i decreases while its Down-set increases. Keeping in mind condition (1) of Definition 4.2 which keeps fixed the relations between atoms respectively satisfying E and $\neg E$, (Order Confirmation) implies that for every α_i satisfying E there is no α_j satisfying $\neg E$ for which initially $\alpha_j \not\succ \alpha_i$ but then $\alpha_j \succ^E \alpha_i$, or for which initially $\alpha_i \succ \alpha_j$ but after the revision $\alpha_j \not\succ^E \alpha_j$. Put in short, (Order confirmation) implies that the position of every world satisfying E improves in the preferential order, while the position of every worlds falsifying E worsens relatively. After considering the correspondence established by Theorem 6.3, Corollary 6.5 states what are the logical properties of single belief revision satisfied by the belief revision operators defined by confirmatory preference revision, namely again system **D**. Such corollary will conclude the result of the present work concerning single, non-iterated belief revision. Right after, we will extend the preference confirmatory revisions $\mathcal{R}^c$ to sequences of evidences and conclude the section.

DEFINITION 6.4. *Let $\succ$ be a semi-order over a set Ω_0 and E be a $\succ$-consistent formula. Then, the belief revision operator $\mathsf{B}^c_\succ$ is defined as:*

$$\mathsf{B}^c_\succ(E) \equiv \bigvee \left\{\alpha_i \in max(\succ^E) : \succ^E \in \mathcal{R}^c(\succ, E)\right\}.$$

COROLLARY 6.5. *$\mathsf{B}^c_\succ$ is an iterated belief revision operator defined as in Definition 6.4 if and only if $\mathsf{B}^c_\succ$ satisfies system* **D**.

PROOF. The result directly follows from the proof for Theorem 4.6 since $\mathcal{R}^c(\succ, E) \subseteq \mathcal{R}(\succ, E)$ for every $\succ$ and E. □

DEFINITION 6.6 (Iterated confirmatory preference revision). *Let $\succ$ be a semi-order over Ω_0. Let $\mathcal{E} = (E_1, \ldots, E_n)$ be a list of formulas and let $\mathcal{E}_i = (E_1, \ldots, E_i)$ be the first i elements of $\mathcal{E}$, with $i, n \in \mathbb{N}$ and $i \leq n$. The preference revision $\mathcal{R}^c$ defined in Definition 4.2 is extended to sequences $\mathcal{E}$ as follows. First, define $\mathcal{R}(\succ, ()) := \{\succ\}$ and $\mathcal{R}^c(\succ, ()) := \{\succ\}$. Then, $\mathcal{R}(\succ, \mathcal{E})$ and $\mathcal{R}^c(\succ, \mathcal{E})$ are defined by the following:*

$$\mathcal{R}^c(\succ, \mathcal{E}_i) = \bigcup\{\mathcal{R}^c(\succ^{\mathcal{E}_{i-1}}, E_i) \ : \succ^{\mathcal{E}_{i-1}} \in \mathcal{R}^c(\succ, \mathcal{E}_{i-1}))\}.$$

Systems			Rules	Conditions on B	B	B^c
			LLE	$\mathsf{B}(A) \vDash C, \vDash A \leftrightarrow B \ \Rightarrow \ \mathsf{B}(B) \vDash C$		
			RW	$\mathsf{B}(A) \vDash B, \vDash B \rightarrow C \ \Rightarrow \ \mathsf{B}(A) \vDash C$		
			Ref	$\mathsf{B}(A) \vDash A$		
		P	And	$\mathsf{B}(A) \vDash B,\ \mathsf{B}(A) \vDash C \ \Rightarrow \ \mathsf{B}(A) \vDash B \wedge C$	✓	✓
			CMon	$\mathsf{B}(A) \vDash B,\ \mathsf{B}(A) \vDash C \ \Rightarrow \ \mathsf{B}(A \wedge B) \vDash C$		
			Or	$\mathsf{B}(A) \vDash C,\ \mathsf{B}(B) \vDash C \ \Rightarrow \ \mathsf{B}(A \vee B) \vDash C$		
	D		DRat	$\mathsf{B}(A \vee B) \vDash C,\ \mathsf{B}(A) \nvDash C \ \Rightarrow \ \mathsf{B}(B) \vDash C$		
			PlaMon	$\mathsf{B}(A) \vDash C,\ \mathsf{B}(A) \nvDash \neg B,\ \mathsf{B}(A \wedge B) \vDash D \ \Rightarrow \ \mathsf{B}(A \wedge B \wedge C) \vDash D$		
R			RMon	$\mathsf{B}(A) \vDash C,\ \mathsf{B}(A) \nvDash \neg B \ \Rightarrow \ \mathsf{B}(A \wedge B) \vDash C$	×	×

TABLE 3. Rules of belief revision.

DEFINITION 6.7 (Iterated confirmatory preference revision operators). *Let $\succ$ be a semi-order over Ω_0. Let $\mathcal{E}$ be a sequence of n formulas in $\mathcal{SL}$. The confirmatory iterated belief revision operator $\mathsf{B}^c_{\succ}$ is defined as follows:*

$$\mathsf{B}^c_{\succ}(\mathcal{E}) \equiv \bigvee\{\alpha_i \in max(\succ^{\mathcal{E}}) : \succ^{\mathcal{E}} \in \mathcal{R}^c(\succ, \mathcal{E})\}.$$

PROPOSITION 6.8. *Let $\succ$ be a order and $\mathcal{E}$ a sequence of non-empty subsets of Ω. Let $\mathsf{B}^c_{\succ}$ be the iterated belief revision operators defined in Definition 6.7. $\mathsf{B}^c_{\succ}$ satisfies system **D** together with the rules (C1), (C2), (C3) and (C4).*

PROOF. That $\mathsf{B}^c_{\succ}$ satisfies system **D** directly follows by Definition 6.7 and Corollary 6.5.

Properties (C1), (C2) and (C4) hold for $\mathsf{B}_{\succ}$. Since they are respectively defined by the maximal elements of the orders in $\mathcal{R}(\succ, \mathcal{E})$ and $\mathcal{R}^c(\succ, \mathcal{E})$ and $\mathcal{R}^c(\succ, \mathcal{E}) \subseteq \mathcal{R}(\succ, \mathcal{E})$, a fortiori those same properties are satisfies by $\mathsf{B}^c_{\succ}$. It is left to prove (C3).

Rules	Conditions on B	B	B^c
C1	$A \vDash C \;\Rightarrow\; \mathsf{B}(C, A) \equiv \mathsf{B}(A)$	✓	✓
C2	$A \vDash \neg C \;\Rightarrow\; \mathsf{B}(C, A) \equiv \mathsf{B}(A)$		
C3	$\mathsf{B}(A) \vDash C \;\Rightarrow\; \mathsf{B}(C, A) \vDash C$	×	
C4	$\mathsf{B}(A) \nvDash \neg C \;\Rightarrow\; \mathsf{B}(C, A) \nvDash \neg C$	✓	
C1$^+$	$A \vDash C \;\Rightarrow\; \mathsf{B}(E_1, \ldots, C, A, \ldots, E_n) \equiv \mathsf{B}(E_1, \ldots, A, \ldots, E_n)$	×	×
CB	$\mathsf{B}(A) \vDash \neg C \;\Rightarrow\; \mathsf{B}(C, A) \equiv \mathsf{B}(A)$		
Rec	$A \nvDash \neg C \;\Rightarrow\; \mathsf{B}(C, A) \vDash C$		
MinInf	$\mathsf{B}_1(\top) \vDash \neg A$, $\mathsf{B}_2(\top) \vDash \neg A \;\Rightarrow\; \mathsf{B}_1(A) \equiv \mathsf{B}_2(A)$		

FIGURE 11. Rules of iterated belief revision.

(C3). For every $\succ_P^{(A)} \in \mathcal{R}^c(\succ, (A))$, $max(\succ^{(A)}) = max(\succ |A)$. Then, the premise implies $max(\succ |A) \subseteq \widehat{C}$; this is equivalent to the fact that for any $\alpha_j \Vdash A \wedge \neg C$, there exists $\alpha_i \Vdash A \wedge C$ such that $\alpha_i \succ \alpha_j$. So, let α_k be any atom in $max(\succ |A) \cap \widehat{C}$ with *highest rank*. When $\succ$ is updated on C, by (Order Confirmation) the relations between any such two α_i, α_j are preserved. Therefore, for any $\alpha_j \in A \wedge \neg C$, there exists $\alpha_i \Vdash A \wedge C$ s.t. $\alpha_i \succ^C \alpha_j$. Finally, when any such $\succ^C$ is updated on A, the relation between α_i and α_j remains fixed by property (1) of Def. 3.5. Hence, $\max(\succ^{(C,A)}) \subseteq \widehat{A \wedge C}$. (C3) has thus been obtained. □

7. Conclusion

The present chapter relates probabilistic updates with belief revision operators. It does so by identifying the logical properties that are satisfied by the belief revision operators induced by Bayesian and Jeffrey conditioning on the probability functions that generate the beliefs.

By relying on the results from the previous chapters, Theorem 2.1 proves that a belief revision operator B can be defined, or tracks, Bayesian conditioning if and only if B satisfies the rules of system **D**. Such results gives the best possible answer to the tracking problem initially proposed by Lin and Kelly by providing the necessary and sufficient conditions for a belief revision operator to track Bayesian conditioning.

We then extend such result to Jeffrey conditioning. Jeffrey conditioning assigns to the evidence a posterior probability e which can range over probability values $[0, 1]$. Then, because the probability of the evidence can even decrease, i.e. $e < P(E)$, the belief update induced by Jeffrey conditioning modulo the odds-threshold method, referred to as *belief change*, does not correspond to any belief revision operator. Specifically, Proposition 3.8 shows that a belief change operator C does not satisfy the rules Ref, Or and CMon of the nonmonotonic system **P**.

In response, we restrict Jeffrey conditioning by requiring that e must be high enough so that the evidence E is implied by the beliefs generated by the posterior probability, formally $\mathsf{B}_{P_E^e}(\top) \vDash E$. Such stronger notion of probabilistic update is called *successful* Jeffrey conditioning. Proposition 4.4 shows that for every instance of successful Jeffrey conditioning $P_E^e(\cdot)$ generates an odds-threshold order which has the same maximal element as Bayesian conditioning P on E, i.e. $max(\succ_{P_E^e}) = max(\succ_{P|E})$. An immediate consequence is that successful Jeffrey conditioning and Bayesian conditioning generate the same beliefs. This means that the odds-threshold method is indifferent to the amount of uncertainty $1 - e$ that is assigned to the evidence: as long as they are successful, all instances of Jeffrey conditioning generate the same beliefs.

Another consequence of Proposition 4.4 is that successful Jeffrey conditioning has the same impact on our qualitative beliefs as Bayesian conditioning. More formally, Theorem 4.3 states that a belief revision operator B can be defined by successful Jeffrey conditioning for some P if and only if B satisfies system D. Thus, Theorem 4.3 gives a sound and complete characterization of the belief revision induced by successful Jeffrey conditioning.

One great advantage of Jeffrey conditioning is that it does not assume the evidence to be certain. In terms of semi-orders, this translates into the worlds inconsistent with E not being necessarily eliminated from the order once this is revised on E. As a result, we considered what are the properties of iterated belief revision that are satisfied by the belief revision operators induced by successful Jeffrey conditioning. Proposition 5.3 shows that in general a belief revision operator induced by successful Jeffrey conditioning satisfies the rules $C1$, $C2$ and $C4$, while it falsifies $C3$, all introduced in [DP97]. In order for such an iterated belief revision to satisfy $C3$ as well, another condition has to be imposed on Jeffrey conditioning. Such conditioning, called *Confirmation*, requires the posterior probability of the evidence e to be higher than the prior probability of the evidence, formally $e > P(E)$. Proposition 6.8 proves that the iterated belief revision operators defined by confirmatory successful Jeffrey conditioning satisfy all properties $C1 - C4$.

Conclusion

The aim of this work was to study what it is like to reason by comparing alternative options. Whether you are at the supermarket picking the next shampoo, or you are a doctor trying to figure what's the best diagnosis that explains the patient's symptoms, many are the situations in which we are given a set of alternative options and we need to understand which of them is the best or the correct one.

In order to study such phenomenon, firstly we assumed that the agent has already determined a finite set of alternatives Ω to consider, and that the agent's degrees of belief about such alternatives can be captured by a probability function P. It is a robust result in psychometrics that we humans are not always able to distinguish real values when they are too close to each other. As a consequence, we introduced the notion of odds-threshold order $\succ_P$, a preferential order generated by P over Ω as follows: for any two $\alpha_i, \alpha_j \in \Omega$, α_i is preferred to α_j, $\alpha_i \succ_P \alpha_j$, if and only if $P(\alpha_i) \geq t \cdot P(\alpha_j)$, where t is some contextually determined real threshold $t > 1$. The notion of odds-threshold order captures those preferences that we would be certainly willing to accept since the values associated to the relative options are sufficiently distant. We argued in the introduction that the odds-threshold order is the adequate definition of the *definitely greater* kind of preference since its definition is in accordance with a well established law in psychology of perception, namely the Weber-Fechner's Law. Once our degrees of belief have generated the odds-threshold order, we believe that the best option is any among the mostly preferred one with respect to the odds-threshold order. However, since we cannot determine among such maximal elements which is actually the best choice, we cannot but believe that either of them is the best choice. In other words, our strongest belief, generated by our preferences, is going to be equivalent to the disjunction of such maximal elements. The relation between our degrees of belief P and our qualitative beliefs that was just established is referred to as odds-threshold method.

The odds-threshold method determines the relation between probabilities and belief at a given moment in time. In order to understand however how we reason by comparing hypotheses, we need to understand how our beliefs change when new evidence is acquired. Because our beliefs are determined by our credences, such inquiry translates into understanding our updating the underlying probability function P in light of new evidence impacts our beliefs.

In order to give an answer to such question, the first step that we took in Chapter 1 was to understand what kind of preferential orders odds-threshold orders actually are, what properties they satisfy. What we wanted was a probability-independent, qualitative characterization of odds-threshold orders. To such end, the notion of *semi-order* was introduced: a strict partial order with the *2v2* and *3v1* as forbidden configurations. As the first main result of this work, Theorem 3.4 proved that the classes of semi-orders and odds-threshold orders coincide.

Chapter 2 bridges from the qualitative characterization of semi-orders to the properties that reasoning with such order would entail. Specifically, the chapter studies what are the logical properties satisfied by the nonmonotonic consequence relations defined on the class of semi-orders. Interestingly, each of the two forbidden configurations which characterize semi-orders corresponds to an inference rule satisfied by such consequence relations. In particular, the strict partial orders with no *2v2* suborder, referred to as *filtered orders*, were already studied in [Fre93]. There, the authors proves (here Theorem 2.4) that a consequence relation can be defined by a filtered order if and only if it satisfies the nonmonotonic system **P** from [KLM90] together with the rules Disjunctive Rationality. We then addressed the class of strict partial order with no *3v1* suborder, dubbed here *mezzanine orders*. As a new result, we proved in Theorem 3.11 that a nonmonotonic consequence relation can be defined by a mezzanine order if and only if it satisfies **P** together with the new rule Plausible Monotonicity. By combining the two results, Theorem 4.1 shows that the nonmonotonic system **D**, composed of system **P** plus Disjunctive Rationality and Plausible Monotonicity, is sound and complete with respect to the class of consequence relations defined by semi-orders.

In order to prove such results we assumed as it is standard practice in the field of nonmonotonic logics that, given a preferential order $\succ$, when a new piece of evidence is acquired, the preferential order is updated by deleting all those worlds which are incompatible with the evidence. Crucially, in [LK12b] the authors prove that, if the preferential orders are odds-threshold orders, then updating the underlying probability function by Bayesian conditioning corresponds exactly to eliminating all the worlds which are incompatible with the evidence, the same kind of updating which is applied for the nonmonotonic consequence relations. Therefore, since odds-threshold order and semi-order are the equivalent notions, in Chapter 3 we proved in Theorem 2.1 that Bayesian conditioning P impacts our beliefs in such a way that in general our beliefs satisfy all and only the rules of the system **D**.

One of the objectives of the present work is to study reasoning with beliefs which are not believed to be certain. Indeed, the agent's beliefs are only their best attempt to determine the best choice. While Bayesian conditioning assigns to the evidence maximal posterior probability, $P(E|E) = 1$, we extended our results to Jeffrey conditioning. Following the example of Lin and Kelly, we first addressed how our preferences change when the underlying probability function is updated by Jeffrey conditioning. While Lin and Kelly prove that Bayesian conditioning has the same impact on preferences as eliminating all the worlds incompatible with the evidence, formally, $\succ_{P|E} = \succ_P |E$, the same nice correspondence cannot be proven for Jeffrey conditioning as well – see Figure 4 for an example. Nevertheless, in Theorem 3.6 we proved a strong enough correspondence, which captures the order-theoretic conditions satisfied by all the odds-threshold orders generated by the instances of Jeffrey conditioning, for all the probability functions that generate the initial order $\succ_P$.

Proposition 3.8 proved that, because of its generality, Jeffrey conditioning can change to such an extent our beliefs that most of the rules of system **P** are invalid. In response, we strengthened its definition and focused on *successful* Jeffrey conditioning, which requires the evidence E to be among the agent's posterior beliefs. Theorem 4.6 shows that once such restriction is done, the impact of Jeffrey conditioning on our beliefs satisfies all and only the rules of system **D**.

Another advantage of Jeffrey conditioning is that it does not necessarily assign to the evidence maximal probability. As a result, we could take into account sequences of evidences and which rules of iterated belief revision are satisfied by our beliefs when generated by the odds-threshold method. Proposition 5.3 showed that the instances of successful Jeffrey conditioning induce a change in our belief which satisfy the rules (C1), (C2) and (C4), proposed in [DP97]. In order for (C3) to be satisfied as well, we had to impose another conditioning and further restrict the definition of Jeffrey conditioning. Such condition consists of Confirmation, which requires the posterior probability of the evidence to be higher than its prior, formally $P_E(E) \geq P(E)$. Proposition 6.8 shows that, under such condition, (C3) is satisfied as well.

There are still many aspects which need to be taken into account in the future. First of all, we still miss a sound and complete characterization of the iterated belief revision operators characterized by (confirmatory) successful Jeffrey conditioning. Secondly, in the literature, several approaches relied on ranking models to study iterated belief revisions. One main reference is [Spo88], where the author studies the use of ordinal conditional functions to capture epistemic updates, also in relation to probabilistic updates. The present work took a different approach and, when it came to defining a method to update the preferential orders in light of new (uncertain) information, we did not specify how the orders were to change, but only gave general conditions which the updated preferential orders had to satisfy. The reason for our choice is that, while Spohn's work often lacks a sufficiently robust theoretical justification for the ordinal conditional functions chosen for the updates, the generality of our approach did not require us to provide additional motivations. Nevertheless, especially with respect to iterated belief revision, it would be interesting to provide further theoretical background and to redefine the present work in terms of ordinal conditional functions.

More broadly, this work is meant to be just the starting point. Here, Theorem 3.4 proved the correspondence between probability distributions and semiorders under the relation given by the odds-threshold method. At the same time, Theorem 4.1 proved the correspondence between semiorders and the belief revision operators (and nonmonotonic consequence relations) satisfying system **D**. These are all static correspondences between formal objects belonging to certain classes, which further allow to establish a correspondence between functions over the same objects. As an example, in [LK12b] the authors prove the correspondence between Bayesian conditioning a probability function P on the evidence E, and restricting the semiorder $\succ_P$ on the worlds satisfying E. In the present work, Theorem 2.1 extends Lin and Kelly's result by showing that the impact of Bayesian conditioning P on our beliefs can be exactly captured by a belief revision method satisfying the rules of the nonmonotonic system **D**. Additionally, Theorems 4.3 and 6.3 extend even further the previous result, so to address the correspondence between progressively stronger versions of Jeffrey conditioning with progressively stronger operators of iterated belief revision.

The approach of relating objects and functions over objects is precisely the main interest of the area of mathematics referred to as category theory. The present work must be considered as the study of just one possible correspondence between probabilities and beliefs, despite the fact that the correspondence, given by the odds-threshold method, was well motivated both in [LK12a], [LK12b] and [KL21]. In the future, by relying on the powerful tools of category theory, one could investigate different ways to relate credence functions and beliefs, shedding light on the relations between logical properties satisfied by our beliefs with different kinds of probabilistic update. This line of research might provide further motivation to study probabilistic updates which have not received enough attention so far.

In relation with probability, another point which must be addressed in a future work is the relation between the present approach and the work done by Branden Fitelson, Kevin Dorst and Ted Shear, among others, on the Lockean Thesis and belief revision. The Lockean Thesis states that the agent must believe A if and only if they are sufficiently confident of A. In [Dor19], Dorst shows that the Lockean Thesis follows from epistemic utility theory. Specifically, the Lockean Thesis is shown to be the only method to relate qualitative beliefs and credence functions that maximizes expected accuracy. Since the present work does not assume the Lockean Thesis but rather the odds-threshold method to relate beliefs and credence functions, it follows from Dorst's results that in the present case the agent will be less accurate in expectation than if they were to satisfy the Lockean Thesis.

However, at the same time, adopting the odds-threshold method instead of the Lockean Thesis to relate beliefs and credences gives a series of advantages from a logical perspective

– see again [LK12a], [LK12b], [KL21]. It is a well known fact that assuming the Lockean Thesis implies giving up some logical properties satisfied by our beliefs: most notably, the lottery paradox shows a case of tension between the Lockean Thesis and the assumptions of consistency and logical closure of the agent's beliefs. In [SF19], Fitelson and Shear give an exhaustive study of the logical properties satisfied by the Lockean Thesis, putting those properties in relation with the threshold chosen for the Lockean Thesis. In the present work instead, we proved that assuming the odds-threshold method implies that the agent's beliefs satisfy the system **D**, an extension of the system **P**, considered to be the standard for defeasible reasoning. Thus, it seems that a tradeoff between the logical strength of the beliefs and their expected accuracy has to be made. In that regard, in the future we must study what is the loss in expected accuracy when we assume the odds-threshold method with respect to the expected accuracy given by the Lockean Thesis. That would quantify the price in accuracy resulting from having a set of beliefs satisfying stronger logical properties.

In conclusion, the present work addressed the question of understanding how we reason by comparing alternative options. By employing a series of techniques and results, we managed to characterize how our confidence, our degrees of belief, impacts what we believe to be the best or the true alternatives. The odds-threshold method resulted in very strong beliefs, stable to the change of the underlying degrees of belief.

Appendix

PROOF OF PROPOSITION 2.9, CASE (B).

(B) $r(a) \geq r(b)$ and $Up(a) \nsubseteq Up(b)$.

$Up(a) \nsubseteq Up(b) \;\Leftrightarrow\; Up(a) \setminus Up(b) \neq \emptyset$

$\Rightarrow$ there exists $a' \in \Omega$ s.t. $a' \in Up(a)$ and $a' \notin Up(b)$.

If $a' \in Down(b)$, then $b \succ a$. This is in contradiction with $r(a) \geq r(b)$ by Proposition 2.8. Therefore, $a' \notin Down(b)$. Since $a' \notin Up(b)$, $a' \sim b$.

Given $Up(a) \nsubseteq Up(b)$, it either holds that (B.1) $Up(b) \setminus Up(a) \neq \emptyset$, or (B.2) $Up(a) \supset Up(b)$.

(B.1) $Up(b) \setminus Up(a) \neq \emptyset$.

Therefore, there exists $b' \in \Omega$ s.t. $b' \in Up(b)$ and $b' \notin Up(a)$.

In the following we prove that $a' \succ a$ and $b' \succ b$ are the only relations of the order, making it an instance on a *2v2* suborder. $a' \sim a$ has already been proven. In order, we show that $b' \sim a$, $a \sim b$ and $a' \sim b'$.

If $b' \in Down(a)$, then $a \succ b$ follows from $b' \in Up(b)$. Since $\succ$ is transitive, $Up(a) \subseteq Up(b)$. This is in plain contradiction with the assumption $Up(a) \nsubseteq Up(b)$. Therefore, $b' \notin Down(a)$. As $b' \notin Up(a)$ also holds, $b' \sim a$ follows.

If $a \succ b$, given $a' \succ a$, then by transitivity $a' \not\succ b$. This contradicts $a' \sim b$; therefore, $a \not\succ b$.

Likewise, $b \not\succ a$ follows from $b' \succ b$ and $b' \sim a$. Then, in short, $a \sim b$ holds.

Likewise, $b' \not\succ a'$ follows from $a' \succ a$ and $b' \sim a$.

Likewise, $a' \not\succ b'$ follows from $b' \succ b$ and $a' \sim b$. Then, in short, $a' \sim b'$ holds.

Therefore, $a' \succ a$ and $b' \succ b$ are the only relations of the order. It is thus obtained an instance of a *2v2* order. This contradicts (1); *not* (1) holds.

(B.2) $Up(a) \supset Up(b)$.

$r(a) \geq r(b) \Rightarrow |Down(a)| - |Down(b)| \geq |Up(a)| - |Up(b)|$.

$Up(a) \supset Up(b) \Rightarrow |Up(a)| - |Up(b)| > 0$.

$\Rightarrow |Down(a)| - |Down(b)| > 0$

$\Rightarrow$ there exists $a'' \in \Omega$ s.t. $a'' \in Down(a)$ and $b' \notin Down(b)$.

If $a'' \in Up(b)$, then the following chain is obtained: $a' \succ a \succ a'' \succ b$. This contradicts the initial assumption $a' \notin Up(b)$. Therefore $a'' \notin Up(b)$. Since $b' \notin Down(b)$, $b' \sim b$.
From $a' \sim b$, $a' \succ a \succ a''$ and $a'' \sim b$, it follows that $a \sim b$. It is thus obtained an instance of a *3v1* order. This contradicts (1); *not* (1) holds.

□

Proof of Lemma 3.2 (⇒). If $\mathcal{R} = \{R\}$, both properties are vacuously satisfied: when R_i and R_j coincide, all the four terms of the implications are false. Then, we can assume that $|\mathcal{R}| \geq 2$, from which $|\Omega| \geq 2$.

⇒. It has to be proven that, for any two rank sets $R_i, R_j \in \mathcal{R}$, if $r_i > r_j$, then $Q(R_i) > Q(R_j)$.

To obtain such result, we show that for any rank set $R_i \in \mathcal{R} \setminus \{R_{min}\}$,

$$Q(R_i) - Q(R_{i-1}) \geq \epsilon_s \tag{10}$$

where ϵ_s is the ϵ value at the time of definition of $Q(R_i)$; in other words, ϵ_s is either the ϵ value used by **if** of Def. 3.1 to define $Q(R_i)$ as $Q(R_i) = Q(R_{i-1}) + \epsilon_s$, or the ϵ value used by the last iteration of **if** prior to the definition $Q(R_i)$.

Only **if** changes the value of ϵ, and at each iteration of **if** ϵ is divided by $|\Omega|$. Since the first iteration of **if** assigns $\epsilon = \frac{(t-1)\cdot Q(R_{min})}{|\Omega|}$, we can state that the sth iteration of **if** assigns to ϵ the value

$$\epsilon_s = \frac{(t-1)\cdot Q(R_{min})}{(|\Omega|)^s}.$$

Since s is a counter, $s \in \mathbb{N}$; additionally, t and $Q(R_{min})$ are assumed to be $t,\ Q(R_{min}) \in \mathbb{R}$ such that $t > 1$ and $Q(R_{min}) > 0$; finally, $|\Omega| \geq 2$. It follows that for any $s \in \mathbb{N}$, $\epsilon_s > 0$.

The proof is going to be by induction on the size of the *Out* sets, starting from $R^{\downarrow}_{min} = \emptyset$ and progressively moving to those rank sets with larger *Down* sets.

(Base Case). As the proof goes by induction on the size of $Down(R_i)$, the base proves 10 for any rank set $R_i \in \mathcal{R}$ with rank $r(R_{min}) < r(R_i) \leq r(R_k)$, where R_k is the rank set with lowest rank and such that $Down(R_k) \neq \emptyset$. The existence of R_k is guaranteed by the assumption $|\mathcal{R}| \geq 2$:

$|\mathcal{R}| \geq 2$.

$r(R_{min}) \geq r(R_i) \;\Rightarrow\; Up(R_{min}) \supseteq Up(R_i)$, for any $R_i \in \mathcal{R}$ (by (2) of Prop. 2.9).

$R^{\uparrow}_{min} = \emptyset$ (by absurdum);

$\Rightarrow R^{\uparrow}_{i} = \emptyset$, for any $R_i \in \mathcal{R}$;

$\Leftrightarrow Down(R_i) = \emptyset$, for any $R_i \in \mathcal{R}$.

$\Rightarrow r(R_i) = |Down(R_i)| - |Up(R_i)| = 0 - 0 = 0$, for any $R_i \in \mathcal{R}$;

$\Rightarrow R_i = R_j$, since $r_i \neq r_j$ for any two distinct $R_i, R_j \in \mathcal{R}$;

$\Rightarrow |\mathcal{R}| = 1$.

$\Rightarrow\ Up(R_{min}) \neq \emptyset$;

$\Leftrightarrow\ Down(R_k) \neq \emptyset$, for some $R_k \in \mathcal{R}$, as desired.

Inequality 10 holds for all R_i such that $r(R_{min}) < r(R_i) < r(R_k)$:
R_k is the rank set with lowest rank and such that $Down(R_k) \neq \emptyset$;
$\Rightarrow Down(R_i) = Down(R_{min}) = \emptyset$, for all R_i s.t. $r(R_{min}) < r(R_i) < r(R_k)$;
$\Rightarrow Down(R_i) = Down(R_{i-1})$, for all R_i s.t. $r(R_{min}) < r(R_i) < r(R_k)$;
$\Rightarrow Q(R_i)$ is defined by **if** of Def. 3.1 as:

$$Q(R_i) = Q(R_{i-1}) + \epsilon_s \implies Q(R_i) - Q(R_{i-1}) \geq \epsilon_s,$$

where ϵ_s is the ϵ value at the time of definition of R_i, as desired.

Inequality 10 holds for $R_i = R_k$:
R_k is the rank set with the lowest rank and such that $Down(R_k) \neq \emptyset$;
$R_k \neq R_{min}$;
$\Rightarrow Down(R_k) \neq Down(R_{k-1}) = \emptyset$;
$\Rightarrow Q(R_k)$ is defined by **else** of Def. 3.1 as:

$$Q(R_k) = t \cdot Q(R'_k),$$

where R'_k is the rank set with the lowest rank in $Down(R_k)$.
$R_k \succ R'_k \ \Rightarrow\ r(R_k) > r(R'_k)$ (by Prop. 2.8) $\ \Leftrightarrow\ r(R_{k-1}) \geq r(R'_k)$.
We need to distinguish two cases, :

- $R_{k-1} = R_{min}$. Then **if** of Def. 3.1 is *not* applied before the definition of $Q(R_k)$. Therefore, at the time of definition of $Q(R_k)$, $s = 0$ and $\epsilon_0 = (t-1) \cdot Q(R_{min})$ as initially assigned. Additionally, $r(R'_k) \leq r(R_{k-1}) = r(R_{min})$; since R_{min} is the only such rank set, we get that $R'_k = R_{min}$. The following holds:

$$Q(R_k) - Q(R_{k-1}) = t \cdot Q(R_{min}) - Q(R_{min}) = (t-1) \cdot Q(R_{min})$$

$$\Rightarrow \quad Q(R_k) - Q(R_{k-1}) \geq (t-1) \cdot Q(R_{min}) = \epsilon_0;$$

- Otherwise, $R_{k-1} \neq R_{min}$. Then $|\mathcal{R}| \geq 3$, from which $|\Omega| \geq 3$. It is known that $Q(R_{k-1})$ is defined by **if** of Def. 3.1, say the sth iteration of **if** with $r \geq 1$. At the same time, the only constraint on R'_k is that $r(R'_k) \leq r(R_{k-1})$; then, either $R'_k = R_{min}$, or $r(R_{min}) < r(R'_k) \leq r(R_{k-1})$.
  - $R'_k = R_{min}$.

$$Q(R_k) - Q(R_{k-1}) = t \cdot Q(R_{min}) - \left(Q(R_{min}) + \sum_{n=1}^{s} \epsilon_n \right)$$

$$= (t-1) \cdot Q(R_{min}) - (t-1) \cdot Q(R_{min}) \cdot \sum_{n=1}^{s} \left(\frac{1}{|\Omega|} \right)^n = (t-1) \cdot Q(R_{min}) \cdot \left(1 - \sum_{n=1}^{s} \left(\frac{1}{|\Omega|} \right)^n \right)$$

Notice that the last sum is a finite geometric series:

$$\sum_{n=1}^{s}\left(\frac{1}{|\Omega|}\right)^n = \sum_{n=0}^{s}\left(\frac{1}{|\Omega|}\right)^n - 1 = \frac{1-\left(\frac{1}{|\Omega|}\right)^{s+1}}{1-\left(\frac{1}{|\Omega|}\right)} - 1 = \left(\frac{1}{|\Omega|}\right)\cdot\frac{1-\left(\frac{1}{|\Omega|}\right)^{s}}{1-\left(\frac{1}{|\Omega|}\right)} = \frac{1-\left(\frac{1}{|\Omega|}\right)^{s}}{|\Omega|-1}.$$

It was assumed that $|\mathcal{R}| \geq 3$, from which $|\Omega| \geq 3$, and $s \geq 1$. This implies that the finite geometrical series always has a value in $]0,1[$. Since $t > 1$ and $Q(R_{min}) > 0$ by assumption, the desired result follows:

$$Q(R_k) - Q(R_{k-1}) = t \cdot Q(R_{min}) - \left(Q(R_{min}) + \sum_{n=1}^{s}\epsilon_n\right) = (t-1)\cdot Q(R_{min}) \cdot \left(1 - \frac{1-\left(\frac{1}{|\Omega|}\right)^{s}}{|\Omega|-1}\right)$$
$$> (t-1)\cdot Q(R_{min}) > t \cdot Q(R_{min}) \cdot \left(\frac{1}{|\Omega|}\right)^s = \epsilon_s.$$

- $r(R_{min}) < r(R'_k) \leq r(R_{k-1})$.

  In this case, $Q(R_{k-1})$ is defined by the sth iteration of **if**, and $Q(R'_k)$ by the qth, with $1 \leq q \leq s$. Since, each ϵ value is strictly positive, the desired result follows from the last inequality:

$$Q(R_k) - Q(R_{k-1}) = t \cdot \left(Q(R_{min}) + \sum_{n=1}^{q}\epsilon_n\right) - \left(Q(R_{min}) + \sum_{n=1}^{s}\epsilon_n\right)$$
$$= t \cdot Q(R_{min}) - \left(Q(R_{min}) + \sum_{n=1}^{s}\epsilon_n\right) + t \cdot \sum_{n=1}^{q}\epsilon_n$$
$$> t \cdot Q(R_{min}) - \left(Q(R_{min}) + \sum_{n=1}^{s}\epsilon_n\right) > \epsilon_s.$$

We have proven the base case, that 10 holds for any rank set $R_i \in \mathcal{R}$ with rank $r(R_{min}) < r(R_i) \leq r(R_k)$, where R_k is the rank set with lowest rank and such that $Down(R_k) \neq \emptyset$.

(Inductive Step). The *inductive hypothesis* (IH): Take a rank set R_j such that $Down(R_j) \supset Down(R_{j-1})$. For any the rank set R_i such that $r(R_{min}) < r(R_i) \leq r(R_j)$, 10 holds:

$$Q(R_i) - Q(R_{i-1}) \geq \epsilon_s \tag{10}$$

where ϵ_s is the ϵ value at the time of definition of $Q(R_i)$.
The (Inductive Step) shows that 10 holds as well for all rank sets R_i such that $r(R_{min}) < r(R_i) \leq r(R_k)$, when there exists the rank set R_k with the lowest rank r_k such that $Down(R_j) \subset Down(R_k)$, or for all rank sets R_i such that $r(R_{min}) < r(R_i) \leq r(R_{max})$, i.e. $R_i \in \mathcal{R} \setminus \{R_{min}\}$, when there is no such R_k.

Notice that, differently from the (Base Case), the existence of such R_k is not guaranteed.

- There is no R_k.

  $\Rightarrow$ For all rank sets R_i, $Down(R_i) \subseteq Down(R_j)$.

- $R_j = R_{max}$;

  $\Rightarrow$ (by IH) 10 holds for all R_i such that $r(R_{min}) < r(R_i) \leq r(R_{max})$, as desired.

- $R_j \neq R_{max}$;

  $\Rightarrow$ there exists some R_i such that $r(R_j) < r(R_i) \leq r(R_{max})$.

  For all such R_i:

  $r(R_i) > r(R_j)$;

  $\Rightarrow$ $Down(R_i) \supseteq Down(R_j)$ (by (2) of Prop. 2.9);

  $Down(R_i) \subseteq Down(R_j)$;

  $\Rightarrow Down(R_i) = Down(R_j)$;

  $r(R_j) < r(R_i)$;

  $\Rightarrow$ $Down(R_i) = Down(R_{i-1})$ (by (2) of Prop. 2.9);

  $\Leftrightarrow$ $Q(R_i)$ is defined by **if** of Def. 3.1 as

$$Q(R_i) = Q(R_{i-1}) + \epsilon_s,$$

  where ϵ_s is the ϵ value at the time of definition of $Q(R_i)$;

  $\Rightarrow$ $Q(R_i) - Q(R_{i-1}) \geq \epsilon_s$.

  $Q(R_i) - Q(R_{i-1}) \geq \epsilon_s$, for all R_i such that $r(R_j) < r(R_i) \leq r(R_{max})$;

  $Q(R_i) - Q(R_{i-1}) \geq \epsilon_s$, for all R_i such that $r(R_{min}) < r(R_i) \leq r(R_j)$ (by IH);

  $\Rightarrow Q(R_i) - Q(R_{i-1}) \geq \epsilon_s$, for all R_i such that $r(R_{min}) < r(R_i) \leq r(R_{max})$;

  $\Leftrightarrow$ 10 holds for all $R_i \in \mathcal{R} \setminus \{R_{min}\}$, as desired.

- There exists some rank set R_k with the lowest rank $r(R_k)$ such that $Down(R_j) \subset Down(R_k)$.

  $\Rightarrow$ $Down(R_{k-1}) = Down(R_j)$;

  $\Rightarrow$ $Q(R_i) - Q(R_{i-1}) \geq \epsilon_s$, for all rank sets R_i such that $r(R_{min}) < r(R_i) \leq r(R_{k-1})$;

  this is obtained by an analogous proof to the one we have just given with R_{k-1} instead of R_{max}.

It is left to prove that 10 holds for $R_i = R_k$.

$Down(R_j) \supset Down(R_{j-1})$ (by IH);

$\Rightarrow$ $Q(R_j)$ is defined by **else** of Def. 3.1 as

$$Q(R_j) = t \cdot Q(R'_j),$$

where R'_j is the rank set with the highest rank in $Down(R_j)$.

R_k is the rank set with the lowest rank $r(R_k)$ such that $Down(R_j) \subset Down(R_k)$;

$\Rightarrow$ $Down(R_k) \supset Down(R_{k-1}) = Down(R_j)$;

$\Rightarrow$ $Q(R_k)$ is defined by **else** of Def. 3.1 as

$$Q(R_k) = t \cdot Q(R'_k),$$

where R'_k is the rank set with the highest rank in $Down(R_k)$.

We prove by absurdum that $R'_k \notin Down(R_j)$.

$R'_k \in Down(R_j)$.

$Down(R_k) \supset Down(R_j) \;\Rightarrow\;$ there exists R_l such that $R_l \in Down(R_k)$ and $R_l \notin Down(R_j)$.

$R'_k \in Down(R_j)$ and $R_l \notin Down(R_j) \;\Rightarrow\; R'_k \neq R_l$;

$\Rightarrow$ either $r(R'_k) < r(R_l)$ or $r(R'_k) > r(R_l)$;

the former case is false because both R_l and R'_k belong to $Down(R_k)$, and R'_k is assumed to be the rank set with the highest rank in $Down(R_k)$;

$\Rightarrow\; r(R'_k) > r(R_l) \;\Rightarrow\; R^{\uparrow}_{k'} \subseteq Up(R_i)$ (by (2) of Prop. 2.9).

$R'_k \in Down(R_j)$ and $R_l \notin Down(R_j) \;\Leftrightarrow\; R_j \in Up(R_{k'})$ and $R_j \notin Up(R_l) \;\Rightarrow\; Up(R_{k'}) \nsubseteq Up(R_l)$.

$\Rightarrow R'_k \notin Down(R_j)$;

$R'_j \in Down(R_j)$;

$\Rightarrow r(R'_k) > r(R'_j)$ (by (2) of Prop. 2.9);

$R'_k \in Down(R_k) \;\Rightarrow\;$ (by Prop. 2.8) $r(R'_k) < r(R_k) \;\Rightarrow\; r(R'_k) \leq r(R_{k-1})$;

$\Rightarrow\; r(R'_j) < r(R'_k) \leq r(R_{k-1})$.

Nevertheless, it is not possible to establish any inequality comparing $r(R_j)$ and $r(R'_k)$; we need to distinguish three cases.

(1) $r(R_j) < r(R'_k) \leq r(R_{k-1})$.

$Down(R_j) \supset Down(R_{j-1}) \;\Rightarrow\; R_j \neq R_{min}$;

$Down(R_j) \subset Down(R_k) \;\Rightarrow\; R_k \neq R_j$ and $R_k \neq R_{min}$;

$\Rightarrow\; |\mathcal{R}| \geq 3 \;\Rightarrow\; |\Omega| \geq 3$.

$Down(R_j) = Down(R_{j+1}) = Down(R'_k) = Down(R_j)$;

$\Rightarrow\; Q(R_{j+1}), Q(R'_k)$ and $Q(R_{k-1})$ are defined by **if** of Def. 3.1, respectively by the pth, qth and rth iteration, with $p, q, r \in \mathbb{N}_{>0}$ such that $r \geq q \geq p$:

$$Q(R'_k) = Q(R_j) + \sum_{n=p}^{q} \epsilon_n = Q(R_j) + (t-1) \cdot Q(R_{min}) \cdot \sum_{n=p}^{q} \left(\frac{1}{|\Omega|}\right)^n,\ q \geq p$$

$$Q(R_{k-1}) = Q(R_j) + \sum_{n=p}^{r} \epsilon_n = Q(R_j) + (t-1) \cdot Q(R_{min}) \cdot \sum_{n=p}^{r} \left(\frac{1}{|\Omega|}\right)^n,\ r \geq q.$$

We now show that $Q(R_k) - Q(R_{k-1}) \geq \epsilon_r$, which gives the desired result as ϵ_r is the ϵ value at the time of definition of $Q(R_k)$:

$$Q(R_k) - Q(R_{k-1}) \geq \epsilon_r \quad \Leftrightarrow \quad t \cdot \left(Q(R_j) + \sum_{n=p}^{q} \epsilon_n\right) - \left(Q(R_j) + \sum_{n=p}^{r} \epsilon_n\right) \geq \epsilon_r$$

$$\Leftrightarrow \quad (t-1)\cdot Q(R_j) + t\cdot\left(\sum_{n=p}^{q}\epsilon_n\right) \geq \sum_{n=p}^{r}\epsilon_n + \epsilon_r$$

$$\Leftrightarrow \quad Q(R_j) + t\cdot Q(R_{min})\cdot\sum_{n=p}^{q}\left(\frac{1}{|\Omega|}\right)^n \geq Q(R_{min})\cdot\left(\sum_{n=p}^{r}\left(\frac{1}{|\Omega|}\right)^n + \left(\frac{1}{|\Omega|}\right)^r\right)$$

$$\Leftrightarrow \quad \frac{Q(R_j)}{Q(R_{min})} + t\cdot\sum_{n=p}^{q}\left(\frac{1}{|\Omega|}\right)^n \geq \sum_{n=p}^{r}\left(\frac{1}{|\Omega|}\right)^n + \left(\frac{1}{|\Omega|}\right)^r$$

We first consider the right term of the last inequality:

$$\sum_{n=p}^{r}\left(\frac{1}{|\Omega|}\right)^n + \left(\frac{1}{|\Omega|}\right)^r < \sum_{n=p}^{\infty}\left(\frac{1}{|\Omega|}\right)^n + \left(\frac{1}{|\Omega|}\right)^r$$

$$= \left(\frac{1}{|\Omega|}\right)^p\cdot\sum_{n=0}^{\infty}\left(\frac{1}{|\Omega|}\right)^n + \left(\frac{1}{|\Omega|}\right)^r = \left(\frac{1}{|\Omega|}\right)^p\cdot\frac{1}{1-\frac{1}{|\Omega|}} + \left(\frac{1}{|\Omega|}\right)^r = \left(\frac{1}{|\Omega|}\right)^{p-1}\cdot\frac{1}{|\Omega|-1} + \left(\frac{1}{|\Omega|}\right)^r$$

The restrictions on $|\Omega|$, p, q, r are $|\Omega|$ and $r \geq q \geq p \geq 1$. For $|\Omega| = 3$ and $p = q = r = 1$, their lowest values, the last term assumes value 5/6. Additionally, as the values of $|\Omega|$, p, q, r increase, the value of the whole term decreases. Therefore, the term always has value lower than 1.

$$1 > \left(\frac{1}{|\Omega|}\right)^p\cdot\sum_{n=0}^{\infty}\left(\frac{1}{|\Omega|}\right)^n + \left(\frac{1}{|\Omega|}\right)^r > \sum_{n=p}^{r}\left(\frac{1}{|\Omega|}\right)^n + \left(\frac{1}{|\Omega|}\right)^r$$

Hence, given $t > 1$, $|\Omega| \geq 3$ and $r \geq q \geq p \geq 1$, the desired result can be obtained.

$R_j \neq R_{min} \;\Rightarrow\; r(R_j) > r(R_{min}) \;\Rightarrow\;$ (by IH and Def. 3.1) $Q(R_j) > Q(R_{min}) > 0$;

$\Rightarrow\; \dfrac{Q(R_j)}{Q(R_{min})} > 1 \;\Rightarrow\; \dfrac{Q(R_j)}{Q(R_{min})} + t\cdot\displaystyle\sum_{n=p}^{q}\left(\frac{1}{|\Omega|}\right)^n > 1.$

$\Rightarrow \dfrac{Q(R_j)}{Q(R_{min})} + t\cdot\displaystyle\sum_{n=p}^{q}\left(\frac{1}{|\Omega|}\right)^n > 1 > \displaystyle\sum_{n=p}^{r}\left(\frac{1}{|\Omega|}\right)^n + \left(\frac{1}{|\Omega|}\right)^r \;\Rightarrow\; Q(R_k) - Q(R_{k-1}) \geq \epsilon_r.$

(2) $r(R'_k) = r(R_j)$.

$\Leftrightarrow$ $R'_k = R_j$ (by Def. 2.10).

- $R_{k-1} = R_j$;

  $Down(R_j) \supset Down(R_{j-1}) \;\Rightarrow\; Q(R_j)$ is defined by **else** of Def. 3.1;

  $\Rightarrow \epsilon_r$ with $r \geq 0$ is the ϵ value at the time of definition of $Q(R_j)$;

  $|\Omega| \geq 3$;

  $\Rightarrow\; 1 \geq \left(\frac{1}{|\Omega|}\right)^r$;

  $R_j \neq R_{min} \;\Rightarrow\; \dfrac{Q(R_j)}{Q(R_{min})} > 1$ (see case (1));

$$\Rightarrow \quad \frac{Q(R_j)}{Q(R_{min})} > 1 \geq \left(\frac{1}{|\Omega|}\right)^r \quad \Rightarrow \quad \frac{Q(R_j)}{Q(R_{min})} \geq \left(\frac{1}{|\Omega|}\right)^r$$

$$\Leftrightarrow \quad (t-1) \cdot Q(R_j) \geq (t-1) \cdot Q(R_{min}) \cdot \left(\frac{1}{|\Omega|}\right)^r \quad \Leftrightarrow \quad Q(R_k) - Q(R_{k-1}) \geq \epsilon_r.$$

- $R_{k-1} \neq R_j$;

$\Rightarrow$ $r(R_j) < r(R_{j+1}) \leq r(R_{k-1}) < r(R_k)$;

$\Rightarrow$ $|\mathcal{R}| \geq 3 \;\Rightarrow\; |\Omega| \geq 3$.

$\Rightarrow$ $Down(R_j) = Down(R_{j+1}) = Down(R_{k-1})$;

$\Rightarrow$ $Q(R_{j+1})$ and $Q(R_{k-1})$ are respectively defined by the pth and rth iteration of **if** of Def. 3.1, with $1 \leq p \leq r$.

Analogously to case (1), for $|\Omega| \geq 3$ and $1 \leq p \leq r$ it holds that:

$$1 > \left(\frac{1}{|\Omega|}\right)^p \cdot \sum_{n=0}^{\infty}\left(\frac{1}{|\Omega|}\right)^n + \left(\frac{1}{|\Omega|}\right)^r > \sum_{n=p}^{r}\left(\frac{1}{|\Omega|}\right)^n + \left(\frac{1}{|\Omega|}\right)^r.$$

$R_j \neq R_{min} \;\Rightarrow\; \dfrac{Q(R_j)}{Q(R_{min})} > 1$ (see case (1)).

Hence, given $t > 1$, $|\Omega| \geq 3$ and $r \geq p \geq 1$, the desired result can be derived.

$$\Rightarrow \quad \frac{Q(R_j)}{Q(R_{min})} \geq \sum_{n=p}^{r}\left(\frac{1}{|\Omega|}\right)^n + \left(\frac{1}{|\Omega|}\right)^r$$

$$\Leftrightarrow \quad (t-1) \cdot Q(R_j) \geq (t-1) \cdot Q(R_{min}) \cdot \left(\sum_{n=p}^{r}\left(\frac{1}{|\Omega|}\right)^n + \left(\frac{1}{|\Omega|}\right)^r\right)$$

$$\Leftrightarrow \quad t \cdot Q(R_j) - \left(Q(R_j) + \sum_{n=p}^{r}\epsilon_n\right) \geq \epsilon_r \quad \Leftrightarrow \quad Q(R_k) - Q(R_{k-1}) \geq \epsilon_r.$$

(3) $r(R'_k) < r(R_j)$.

$r(R'_j) < r(R'_k) \leq r(R_{k-1})$ (see beginning of (Inductive Step)).

$\Rightarrow$ $r(R'_j) < r(R'_k) < r(R_j)$;

$\Rightarrow$ (by IH) $Q(R'_k) - Q(R'_j) \geq \epsilon_m$, where ϵ_m with $m \geq 0$ is the ϵ value at the time of definition of $Q(R'_k)$.

- $R_{k-1} = R_j$;

ϵ_r is the ϵ value at the time of definition of $Q(R_{k-1})$;

$\Rightarrow$ ϵ_r is the ϵ value at the time of definition of $Q(R_j)$.

$r(R'_k) < r(R_j)$;

$\Rightarrow$ $r \geq m \geq 0$;

ϵ strictly decreases at each iteration of **if** by a factor of $(|\Omega|)^{-1}$ with $|\Omega| > 0$;

$\Rightarrow$ $\epsilon_m \geq \epsilon_r$.

Since $t > 1$, we are able to derive the desired result:

$$Q(R_k) - Q(R_{k-1}) = Q(R_k) - Q(R_j) = t \cdot Q(R'_k) - t \cdot Q(R'_j) \geq t \cdot \epsilon_m \geq \epsilon_r.$$

∘ $R_{k-1} \neq R_j$;

$\Rightarrow$ $(R_j) < r(R_{j+1}) \leq r(R_{k-1}) < r(R_k)$;

$\Rightarrow$ $\{R_j, R_{k-1}, R_k\} \subseteq \mathcal{R}$ $\Rightarrow$ $|\Omega| \geq 3$.

ϵ_p is the ϵ value at the time of definition of $Q(R_{j+1})$;

ϵ_r is the ϵ value at the time of definition of $Q(R_{k-1})$;

$\Rightarrow$ $1 \leq p \leq r$.

$Down(R_j) = Down(R_{j+1}) = Down(R_{k-1})$ (by def. of R_k);

$\Rightarrow$ $Q(R_{j+1})$ is defined by **if** of Def. 3.1;

$\Rightarrow$ ϵ_{p-1} is the ϵ value at the time of definition of $Q(R_j)$.

$r(R'_j) < r(R'_k) < r(R_j)$;

$\Rightarrow$ (by IH) $Q(R'_k) - Q(R'_j) \geq \epsilon_m$, where ϵ_m with $m \geq 0$ is the ϵ value at the time of definition of $Q(R'_k)$;

$\Rightarrow$ $0 \leq m \leq p - 1 < p$.

With $|\Omega| \geq 3$ and $t > 1$, we are now able to derive the desired result:

$$1 > \frac{1}{|\Omega|} \cdot \frac{2|\Omega| - 1}{|\Omega| - 1} \quad \Rightarrow \quad 1 > \left(\frac{1}{|\Omega|}\right)^{p-m} \cdot \frac{2|\Omega| - 1}{|\Omega| - 1} \quad \Leftrightarrow \quad \left(\frac{1}{|\Omega|}\right)^{m} > \left(\frac{1}{|\Omega|}\right)^{p} \cdot \frac{2|\Omega| - 1}{|\Omega| - 1}$$

$$\Leftrightarrow \quad \left(\frac{1}{|\Omega|}\right)^{m} > \left(\frac{1}{|\Omega|}\right)^{p} \cdot \frac{|\Omega|}{|\Omega| - 1} + \left(\frac{1}{|\Omega|}\right)^{p} \quad \Rightarrow \quad \left(\frac{1}{|\Omega|}\right)^{m} > \left(\frac{1}{|\Omega|}\right)^{p} \cdot \frac{1}{1 - \frac{1}{|\Omega|}} + \left(\frac{1}{|\Omega|}\right)^{r}$$

$$\Leftrightarrow \quad \epsilon_m > \sum_{n=p}^{\infty} \epsilon_n + \epsilon_r \quad \Rightarrow \quad \epsilon_m > \sum_{n=p}^{r} \epsilon_n + \epsilon_r \quad \Leftrightarrow \quad t \cdot \epsilon_m \geq \sum_{n=p}^{r} \epsilon_n + \epsilon_r$$

$$\Rightarrow \quad t \cdot \left(Q(R'_k) - Q(R'_j)\right) \geq \sum_{n=p}^{r} \epsilon_n + \epsilon_r \quad \Leftrightarrow \quad t \cdot Q(R'_k) - \left(t \cdot Q(R'_j) + \sum_{n=p}^{r} \epsilon_n\right) \geq \epsilon_r$$

$$\Leftrightarrow \quad Q(R_k) - \left(Q(R_j) + \sum_{n=p}^{r} \epsilon_n\right) \geq \epsilon_r \quad \Leftrightarrow \quad Q(R_k) - Q(R_{k-1}) \geq \epsilon_r.$$

We have established that for all three cases it holds that $Q(R_k) - Q(R_{k-1}) \geq \epsilon_r$;

$Q(R_i) - Q(R_{i-1}) \geq \epsilon_s$, for all rank sets R_i such that $r(R_{min}) < r(R_i) \leq r(R_{k-1})$;

$\Rightarrow Q(R_i) - Q(R_{i-1}) \geq \epsilon_s$, for all rank sets R_i such that $r(R_{min}) < r(R_i) \leq r(R_k)$.

The (Base Case) and the (Inductive Step) together establish that for each rank set $R_i \in \mathcal{R} \setminus \{R_{min}\}$, it holds that

$$Q(R_i) - Q(R_{i-1}) \geq \epsilon_s,$$

where ϵ_s is the ϵ value at the time of definition of $Q(R_i)$. Given $t > 1$, $Q(R_{min}) > 0$ by Def. 3.1 and $|\Omega| \geq 2$ (see beginning of the proof), it is immediate to see that for each $s \geq 0$,

$$\epsilon_s = t \cdot Q(R_{min}) \cdot \left(\frac{1}{|\Omega|}\right)^{s} > 0.$$

Therefore, the following obtains: for each $R_i \in \mathcal{R} \setminus \{R_{min}\}$,

$$Q(R_i) - Q(R_{i-1}) > 0 \quad \Leftrightarrow \quad Q(R_i) > Q(R_{i-1}).$$

Since each rank set is uniquely assigned an integer rank value, the rank sets are completely ordered by the rank. If $r_i > r_j$, there are only a finite number of rank sets between R_i and R_j. Therefore, we have established that $Q(R_i) > Q(R_{i-1}) > \cdots > Q(R_{j+1}) > Q(R_j)$, from which $Q(R_i) > Q(R_j)$, as desired. $\square$

Bibliography

[Arm50] Wallace Edwin Armstrong. "A note on tue theory of consumer's behaviour". In: *Oxford Economic Papers* 2.1 (1950), pp. 119–122.

[BMP97] Hassan Bezzazi, David Makinson, and Ramón Pino Pérez. "Beyond rational monotony: Some strong non-horn rules for nonmonotonic inference relations". In: *Journal of Logic and Computation* 7.5 (1997), pp. 605–631.

[Bou96] Craig Boutilier. "Iterated revision and minimal change of conditional beliefs". In: *Journal of Philosophical Logic* 25 (1996), pp. 263–305.

[CR19] John Cantwell and Hans Rott. "Probability, coherent belief and coherent belief changes". In: *Annals of Mathematics and Artificial Intelligence* 87 (2019), pp. 259–291.

[CTG07] Vincenzo Crupi, Katya Tentori, and Michel Gonzalez. "On Bayesian measures of evidential support: Theoretical and empirical issues". In: *Philosophy of Science* 74.2 (2007), pp. 229–252.

[Dan17] Kahneman Daniel. *Thinking, fast and slow*. 2017.

[Doh+79] Michael E Doherty et al. "Pseudodiagnosticity". In: *Acta psychologica* 43.2 (1979), pp. 111–121.

[Dor19] Kevin Dorst. "Lockeans maximize expected accuracy". In: *Mind* 128.509 (2019), pp. 175–211.

[DP02] Brian A Davey and Hilary A Priestley. *Introduction to lattices and order*. Cambridge university press, 2002.

[DP97] Adnan Darwiche and Judea Pearl. "On the logic of iterated belief revision". In: *Artificial intelligence* 89.1-2 (1997), pp. 1–29.

[Fis70] Peter C Fishburn. "Intransitive indifference with unequal indifference intervals". In: *Journal of Mathematical Psychology* 7.1 (1970), pp. 144–149.

[Fis73] Peter C Fishburn. "Interval representations for interval orders and semiorders". In: *Journal of Mathematical Psychology* 10.1 (1973), pp. 91–105.

[Fit06] Branden Fitelson. "Logical foundations of evidential support". In: *Philosophy of Science* 73.5 (2006), pp. 500–512.

[Fit99] Branden Fitelson. "The plurality of Bayesian measures of confirmation and the problem of measure sensitivity". In: *Philosophy of science* 66.S3 (1999), S362–S378.

[FLM91] Michael Freund, Daniel Lehmann, and Paul Morris. "Rationality, transitivity, and contraposition". In: *Artificial Intelligence* 52.2 (1991), pp. 191–203.

[Fol09] Richard Foley. "Beliefs, degrees of belief, and the Lockean thesis". In: *Degrees of belief*. Springer, 2009, pp. 37–47.

[Fre93] Michael Freund. "Injective models and disjunctive relations". In: *Journal of Logic and Computation* 3.3 (1993), pp. 231–247.

[GC82] Richard A Griggs and James R Cox. "The elusive thematic-materials effect in Wason's selection task". In: *British journal of psychology* 73.3 (1982), pp. 407–420.

[GG11] Gerd Gigerenzer and Wolfgang Gaissmaier. "Heuristic decision making". In: *Annual review of psychology* 62.2011 (2011), pp. 451–482.

[Haw09] James Hawthorne. "The Lockean thesis and the logic of belief". In: *Degrees of belief*. Springer, 2009, pp. 49–74.

[Joy21] James Joyce. "Bayes' Theorem". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Fall 2021. Metaphysics Research Lab, Stanford University, 2021.

[KL21] Kevin T Kelly and Hanti Lin. "Beliefs, Probabilities, and Their Coherent". In: *Lotteries, Knowledge, and Rational Belief: Essays on the Lottery Paradox* (2021), p. 185.

[KLM90] Sarit Kraus, Daniel Lehmann, and Menachem Magidor. "Nonmonotonic reasoning, preferential models and cumulative logics". In: *Artificial intelligence* 44.1-2 (1990), pp. 167–207.

[Lei17] Hannes Leitgeb. *The stability of belief: How rational belief coheres with probability*. Oxford University Press, 2017.

[Lev96] Isaac Levi. *For the sake of the argument: Ramsey test conditionals, inductive inference and nonmonotonic reasoning*. Cambridge University Press, 1996.

[LK12a] Hanti Lin and Kevin T Kelly. "A geo-logical solution to the lottery paradox, with applications to conditional logic". In: *Synthese* 186 (2012), pp. 531–575.

[LK12b] Hanti Lin and Kevin T Kelly. "Propositional reasoning that tracks probabilistic reasoning". In: *Journal of philosophical logic* 41.6 (2012), pp. 957–981.

[LM92] Daniel Lehmann and Menachem Magidor. "What does a conditional knowledge base entail?" In: *Artificial intelligence* 55.1 (1992), pp. 1–60.

[Loc47] John Locke. *An essay concerning human understanding*. Kay & Troutman, 1847.

[Luc56] R Duncan Luce. "Semiorders and a theory of utility discrimination". In: *Econometrica, Journal of the Econometric Society* (1956), pp. 178–191.

[McK78] Stuart J McKelvie. "Graphic rating scales—How many categories?" In: *British Journal of psychology* 69.2 (1978), pp. 185–202.

[Ram73] James O Ramsay. "The effect of number of categories in rating scales on precision of estimation of scale values". In: *Psychometrika* 38.4 (1973), pp. 513–532.

[Rob70] Fred S Roberts. "On nontransitive indifference". In: *Journal of Mathematical Psychology* 7.2 (1970), pp. 243–258.

[Rob71] Fred S Roberts. "Homogeneous families of semiorders and the theory of probabilistic consistency". In: *Journal of Mathematical Psychology* 8.2 (1971), pp. 248–263.

[Rot17] Hans Rott. "Preservation and postulation: Lessons from the new debate on the Ramsey test". In: *Mind* 126.502 (2017), pp. 609–626.

[Sco64] Dana Scott. "Measurement structures and linear inequalities". In: *Journal of mathematical psychology* 1.2 (1964), pp. 233–247.

[SF19] Ted Shear and Branden Fitelson. "Two approaches to belief revision". In: *Erkenntnis* 84 (2019), pp. 487–518.

[Sho87] Yoav Shoham. "A semantical approach to nonmonotonic logics". In: *Readings in nonmonotonic reasoning*. 1987, pp. 227–250.

[Spo88] Wolfgang Spohn. "Ordinal conditional functions: A dynamic theory of epistemic states". In: *Causation in decision, belief change, and statistics: Proceedings of the Irvine Conference on Probability and Causation*. Springer. 1988, pp. 105–134.

[SS58] Dana Scott and Patrick Suppes. "Foundational aspects of theories of measurement1". In: *The journal of symbolic logic* 23.2 (1958), pp. 113–128.

[Sta09] Robert Stalnaker. "Iterated belief revision". In: *Erkenntnis* 70.2 (2009), pp. 189–209.

[SZ62] Patrick Suppes and Joseph L Zinnes. *Basic measurement theory*. Vol. Techincal report n.45. Stanford University, 1962.

[TCR13] Katya Tentori, Vincenzo Crupi, and Selena Russo. "On the determinants of the conjunction fallacy: probability versus inductive confirmation." In: *Journal of Experimental Psychology: General* 142.1 (2013), p. 235.