Rarity and Structure of Syntax Splittings for Ranking Functions and Total Preorders

Fabian Bräuer[1,*], Simon Bräuer[1]

[1]*Unaffiliated*

Abstract

Ranking functions, also known as ordinal conditional functions (OCFs), and total preorders on worlds (TPOs) are standard representations in belief change and nonmonotonic reasoning. In bounded unstructured spaces, we study how often these representations admit non-trivial syntax splittings, i.e., decompositions induced by partitions of the underlying propositional signature, and how the resulting splitting families are structured. For rank-bounded OCFs, coarsening closure makes the unique finest syntax splitting an exact counting handle. We derive exact counts, prevalence ratios, and finest-shape distributions. The resulting class is small relative to the ambient space and is strongly concentrated on coarse finest shapes. For depth-bounded TPOs, we report exact prevalence counts for the exhaustive bounded spaces considered here and analyze the subcorpus admitting non-trivial syntax splittings using Bell baselines, closure deficits, and relative splitting profiles. In the exhaustive $|\Sigma| = 3$ series, all such TPOs are Bell-complete up to depth 3, while closure-irregular families first occur at depth 4. A supplementary inspection interface connects these aggregate findings to concrete TPO instances by displaying splitting families, missing coarsenings, and bridge OCFs obtained from absolute marginal sums.

Keywords

Syntax Splitting, Total Preorders, Ranking Functions (OCFs), Empirical Analysis, Interactive Inspection

1. Introduction

Within belief change, one influential notion of relevance goes back to Parikh: if a propositional signature decomposes into syntactically independent blocks, then change driven by information formulated in one block should leave the others unaffected. Syntax splitting formalizes this idea by partitioning the signature into disjoint sublanguages and asking whether the represented structure factorizes accordingly [1, 2, 3, 4]. The notion was introduced for belief sets and later generalized to richer formalisms, including OCFs and TPOs [3, 5]. We study this generalized setting over finite propositional signatures.

Despite substantial theoretical work on syntax splitting, its empirical behaviour in bounded unstructured spaces is still poorly understood. In particular, little is known about the rarity of non-trivial splittings, their distribution across the partition lattice, and the structural contrast between OCF- and TPO-splitting families.

Beyond prevalence, we also report splitting-family structure, where the structural contrast between OCFs and TPOs leads to different counting problems. For OCFs, coarsening closure gives an exact structural handle, since the unique finest splitting already determines the full splitting family [6]. We use this fact below to reconstruct aggregate counts from the distribution of unique finest splittings. For TPOs, the corresponding finest-splitting theorem is weaker: every splitting lies above the finest splitting, but not every coarsening need be a splitting [6]. We therefore analyze which coarsenings above the finest splitting are themselves splittings. The bounded baseline spaces used below measure both prevalence and family structure. These spaces are not meant to model application-induced distributions directly. Instead, they serve as controlled reference classes with no additional modelling assumptions.

Joint proceedings of the Workshop on Theory and Methods for Abstraction (THEMA 2026) and the Workshop on Modularity and Splitting Techniques for Knowledge Representation and Reasoning (MoST 2026) at the 9th Federated Logic Conference (FLoC 2026), July 13–29, Lisbon, Portugal

*Corresponding author.

f.braeuer@mailbox.org (F. Bräuer); simon_braeuer@mailbox.org (S. Bräuer)

0009-0009-2507-7534 (F. Bräuer)

© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

published 2026-11-15

CEUR-WS.org/Vol-4331/paper5.pdf

Although the OCF and TPO analyses require different diagnostics, the two representations are not independent. Every OCF induces a TPO, and every OCF-splitting is inherited as a TPO-splitting of the induced TPO [5]. However, the induced TPO may admit additional splittings that are not present at the OCF level [7]. This asymmetry motivates the closure-sensitive diagnostics for TPOs and the parameterized TPO-to-OCF bridge inspected in the accompanying interface.

The OCF results combine exhaustive rank-bounded datasets, computed using the splitting procedure of [8], with the lattice-based reconstruction developed in Subsection 3.2. For TPOs, exhaustive generation within the reported bounds provides the total-space counts, which are checked against a closed-form enumeration. The subcorpus admitting non-trivial splittings is analyzed using the splitting procedure of [9]. The supplementary artifact repository provides the computational artifacts used for reproduction, including the Docker-packaged counter, generator/analyzer, local inspection interface, example files, and README files.[1] The interface is used to inspect representative witnesses underlying the aggregate TPO diagnostics, especially closure-irregular splitting families and the behaviour of bridge OCFs relative to selected partitions.

The main contributions are:

- exact rarity counts and finest-shape distributions for rank-bounded OCFs, reconstructed by a lattice-based method using admissible blockwise rank counts and Möbius inversion,
- exact exhaustive prevalence counts for depth-bounded TPO spaces in the reported finite ranges, together with closure-sensitive diagnostics and relative-profile counts on the subcorpus admitting non-trivial splittings, including the first closure-irregular depth in the exhaustive $|\Sigma| = 3$ series and the onset of all relative profiles over three finest blocks,
- an exploratory interface for inspecting concrete TPO-splitting families, missing coarsenings, and the extent to which a partition-dependent OCF bridge preserves ordering and splitting behaviour.

The paper is organized as follows. Section 2 recalls the formal background on OCFs, TPOs, and syntax splitting. Section 3 gives the main empirical and combinatorial results: the exact OCF counting results and the exhaustive depth-bounded TPO diagnostics. Section 4 presents the interface for inspecting selected instances and the bridge construction. Section 5 discusses limitations and further directions, and Section 6 concludes.

2. Preliminaries

Signatures, worlds, and formulas. A *(propositional) signature* is a finite set Σ of propositional variables. The propositional language over Σ is denoted by $\mathcal{L}_\Sigma$. Variables are denoted by lowercase letters $a, b, c, \ldots$, while formulas are denoted by uppercase letters $A, B, C, \ldots$. For brevity, conjunctions may be written multiplicatively, i.e. $A \wedge B$ may be abbreviated by AB, and negation $\neg A$ by $\bar{A}$.

An interpretation over Σ, also called *world*, is a truth assignment $\omega : \Sigma \to \{0, 1\}$. The set of all interpretations over Σ is denoted by Ω_Σ. Extending to formulas in the standard way, we write $\omega \models A$ if A is true under ω. The set of models of A is $\mathrm{Mod}_\Sigma(A) = \{\omega \in \Omega_\Sigma \mid \omega \models A\}$.

We represent worlds as complete conjunctions of literals. For example, over $\Sigma = \{a, b, c\}$, the world making a false and b, c true may be written as $\neg a \wedge b \wedge c$, or simply as $\bar{a}bc$. Up to permutation of literals, each world has a unique such representation.

For a subsignature $\Sigma' \subseteq \Sigma$, the projection of a world $\omega \in \Omega_\Sigma$ to Σ' is denoted by $\omega|_{\Sigma'}$. Projected worlds are identified with their corresponding complete conjunctions over the smaller signature when used as formulas. When the context is clear and $\Sigma_i \subseteq \Sigma$ is a designated block, we write ω_i for $\omega|_{\Sigma_i}$ and $\omega_{\neq i}$ for $\omega|_{\Sigma \setminus \Sigma_i}$.

OCFs and TPOs. An OCF over Σ is a map $\kappa : \Omega_\Sigma \to \mathbb{N}_0 \cup \{\infty\}$ such that $\kappa^{-1}(0) \neq \emptyset$, in the sense of Spohn's ordinal conditional functions [10]. Intuitively, smaller ranks correspond to more plausible worlds, so OCFs can be read as qualitative implausibility measures [10].

[1] https://codeberg.org/KRRR-Lab/MOST26_Supplementary_material

All OCFs considered in the remainder of the paper are finite-valued, i.e., $\kappa(\omega) \in \mathbb{N}_0$ for every $\omega \in \Omega_\Sigma$. For $K \in \mathbb{N}_0$, such an OCF κ is *rank-bounded by* K if $\kappa(\omega) \in \{0, \dots, K\}$ for all $\omega \in \Omega_\Sigma$. The empirical OCF spaces studied later are the corresponding rank-bounded spaces.

For a set $X \subseteq \Omega_\Sigma$, we write $\kappa(X) = \min_{\omega \in X} \kappa(\omega)$, with the usual convention $\kappa(\emptyset) = \infty$. OCFs are extended to formulas by $\kappa(A) = \kappa(\mathrm{Mod}_\Sigma(A))$.

A TPO over Σ is a binary relation $\preceq \subseteq \Omega_\Sigma \times \Omega_\Sigma$ that is reflexive, transitive, and total. Its intended meaning is that $\omega_1 \preceq \omega_2$ says that ω_1 is at least as plausible as ω_2. The associated strict relation $\prec$ is defined as usual by $\omega_1 \prec \omega_2$ iff $\omega_1 \preceq \omega_2$ and $\omega_2 \not\preceq \omega_1$. For formulas $A, B \in \mathcal{L}_\Sigma$, we write $A \preceq B$ if there exists a world $\omega_1 \in \mathrm{Mod}_\Sigma(A)$ such that $\omega_1 \preceq \omega_2$ for every world $\omega_2 \in \mathrm{Mod}_\Sigma(B)$. A TPO is equivalently described by its ordered plausibility layers, i.e., by the unique ordered partition of Ω_Σ into nonempty $\preceq$-equivalence classes arranged from most to least plausible. We use this layer representation throughout: layers are indexed contiguously from 0 (most plausible) upward, and the depth of a TPO is later measured by the largest layer index used.

Every OCF κ induces a TPO $\preceq_\kappa$ over Σ, defined on Ω_Σ by $\omega_1 \preceq_\kappa \omega_2$ iff $\kappa(\omega_1) \leq \kappa(\omega_2)$. Thus OCFs provide a ranked refinement of the qualitative ordering captured by TPOs [7].

Marginalisation and syntax splitting. We recall marginalisation and syntax splitting for OCFs and TPOs, following [5, 6].

Definition 1 (Marginalisation of OCFs (adapted from [5])). *Let κ be an OCF over Σ, and let $\Sigma' \subseteq \Sigma$. The* marginal $\kappa|_{\Sigma'} : \Omega_{\Sigma'} \to \mathbb{N}_0$ *is defined by*

$$\kappa|_{\Sigma'}(\omega') = \min\{\kappa(\omega) \mid \omega \in \Omega_\Sigma,\ \omega|_{\Sigma'} = \omega'\}$$

for all $\omega' \in \Omega_{\Sigma'}$.

Definition 2 (Marginalisation of TPOs (adapted from [5, 6])). *Let $\preceq$ be a TPO over Σ, and let $\Sigma' \subseteq \Sigma$. The* marginal $\preceq|_{\Sigma'}$ *is the TPO on $\Omega_{\Sigma'}$ defined, for all $\omega'_1, \omega'_2 \in \Omega_{\Sigma'}$, by*

$$\omega'_1 \preceq|_{\Sigma'} \omega'_2 \quad \textit{iff} \quad \omega'_1 \preceq \omega'_2,$$

where the projected worlds on the right-hand side are read as formulas and compared by the formula comparison induced by $\preceq$.

Marginals capture how an OCF or TPO behaves over a restricted subsignature. Syntax splitting asks whether, relative to a partition of the signature into blocks, the global behaviour is blockwise independent. In the OCF case, independence is expressed by marginal additivity: the global OCF is exactly the sum of its blockwise marginals. In the TPO case, the condition is comparison-based rather than reconstructive: once the values outside a block are fixed, comparisons inside that block must agree with the corresponding marginal TPO. Thus TPO-splitting does not in general reconstruct the full TPO from its marginals, but it does identify partitions along which plausibility comparisons decouple blockwise [5].

Definition 3 (Syntax splitting of OCFs (adapted from [5])). *Let κ be an OCF over Σ, and let $P = \{\Sigma_1, \dots, \Sigma_k\}$ be a partition of Σ. We say that κ* admits a syntax splitting along P *if*

$$\forall \omega \in \Omega_\Sigma : \kappa(\omega) = \sum_{i=1}^{k} \kappa|_{\Sigma_i}(\omega|_{\Sigma_i}).$$

We also call such a splitting an OCF-splitting. *Equivalently, there are OCFs κ_i over Σ_i such that $\kappa = \kappa_1 \oplus \cdots \oplus \kappa_k$ in the usual marginal-sum sense [5, 6].*

Definition 4 (Syntax splitting of TPOs (adapted from [5])). *Let $\preceq$ be a TPO over Σ, and let $P = \{\Sigma_1, \dots, \Sigma_k\}$ be a partition of Σ. We say that $\preceq$* admits a syntax splitting along P *if, for every block Σ_i and all $\omega_1, \omega_2 \in \Omega_\Sigma$,*

$$\omega_{1,\neq i} = \omega_{2,\neq i} \quad \Longrightarrow \quad \big(\omega_1 \preceq \omega_2 \iff \omega_{1,i} \preceq|_{\Sigma_i} \omega_{2,i}\big).$$

We also call such a splitting a TPO-splitting.

Partition structure and coarsening closure. Syntax splittings are naturally organized in the partition lattice of the underlying signature.

Let $\Pi(\Sigma)$ denote the lattice of all set partitions of Σ, ordered by refinement. For partitions $P, Q \in \Pi(\Sigma)$, we write

$$P \sqsubseteq Q$$

iff every block of P is contained in some block of Q, i.e. P is a refinement of Q and Q is a coarsening of P. The finest partition is the singleton partition $\hat{0} = \big\{\{v\} \mid v \in \Sigma\big\}$, and the coarsest partition is the one-block partition $\hat{1} = \{\Sigma\}$.

For $E \in \{\kappa, \preceq\}$, let

$$\mathsf{Spl}(E) = \{P \in \Pi(\Sigma) \mid E \text{ admits a syntax splitting along } P\}.$$

Thus $\mathsf{Spl}(E)$ denotes the full splitting family of E.

Remark 1 (Trivial splittings and full splitting families)**.** *The one-block partition $\hat{1}$ is always a splitting and is called the trivial splitting: for OCFs, marginalisation to the full signature returns the original OCF, and for TPOs, marginalisation to the full signature returns the original TPO. Throughout the paper, $\mathsf{Spl}(E)$ denotes the full splitting family of E, including $\hat{1}$. By contrast, E is said to* admit a non-trivial splitting *if $\mathsf{Spl}(E) \setminus \{\hat{1}\} \neq \emptyset$. Prevalence counts use this non-trivial notion, while the structural diagnostics below are computed on full splitting families.*

We will also refer to the *shape* of a partition, namely the multiset of its block sizes written in decreasing order. Thus a partition of a four-variable signature into one block of size 3 and one singleton has shape $3 + 1$, while the singleton partition has shape $1 + 1 + 1 + 1$. Shapes allow us later to aggregate finest splittings by block-size pattern rather than by individual set partition.

Definition 5 (Coarsening-closed families)**.** *Let $\mathcal{S} \subseteq \Pi(\Sigma)$. We call $\mathcal{S}$* coarsening-closed *if*

$$P \in \mathcal{S} \text{ and } P \sqsubseteq Q \quad \Longrightarrow \quad Q \in \mathcal{S}.$$

Equivalently, $\mathcal{S}$ is an upper set in the refinement order. Thus, once P is a splitting, every partition obtained by merging blocks of P must also be a splitting. For $E \in \{\kappa, \preceq\}$, we say that E has a coarsening-closed splitting family if $\mathsf{Spl}(E)$ is coarsening-closed.

Proposition 1 (Finest splittings and coarsening closure (adapted from [6]))**.** *Let Σ be a propositional signature and let $P \in \Pi(\Sigma)$ be some partition of Σ.*

1. *Every OCF κ over Σ has a unique finest syntax splitting $\mathcal{F}_\kappa$. Moreover,*

   $$P \in \mathsf{Spl}(\kappa) \quad \textit{iff} \quad \mathcal{F}_\kappa \sqsubseteq P.$$

   Equivalently, every coarsening of $\mathcal{F}_\kappa$ is again a splitting.
2. *Every TPO $\preceq$ over Σ has a unique finest syntax splitting $\mathcal{F}_\preceq$. Moreover,*

   $$P \in \mathsf{Spl}(\preceq) \quad \Longrightarrow \quad \mathcal{F}_\preceq \sqsubseteq P.$$

   The converse need not hold.

The contrast in Proposition 1 can already be seen on a three-variable signature. For OCFs, coarsening closure makes the full splitting family an upper set above the finest splitting. For TPOs, the finest splitting still constrains all possible splittings, but some coarsenings may be missing. The following example illustrates this failure of coarsening closure and, at the same time, shows that an OCF and its induced TPO may have different splitting families.

Example 1. *Let $\Sigma = \{a, b, c\}$, and abbreviate the three two-block partitions by*

$$P_{ab|c} = \{\{a, b\}, \{c\}\}, \qquad P_{ac|b} = \{\{a, c\}, \{b\}\}, \qquad P_{a|bc} = \{\{a\}, \{b, c\}\}.$$

Consider the TPOs

$$\preceq\colon \bar{a}\bar{b}\bar{c} \prec \bar{a}b\bar{c},\ a\bar{b}\bar{c},\ \bar{a}\bar{b}c \prec a\bar{b}c,\ \bar{a}bc,\ ab\bar{c} \prec abc,$$

$$\preceq'\colon \bar{a}\bar{b}\bar{c} \prec' a\bar{b}\bar{c} \prec' \bar{a}b\bar{c} \prec' \bar{a}\bar{b}c \prec' a\bar{b}c,\ ab\bar{c} \prec' \bar{a}bc \prec' abc.$$

Commas separate worlds in the same plausibility layer. Both TPOs are induced by the following OCFs:

ω	$\bar{a}\bar{b}\bar{c}$	$a\bar{b}\bar{c}$	$\bar{a}b\bar{c}$	$\bar{a}\bar{b}c$	$a\bar{b}c$	$ab\bar{c}$	$\bar{a}bc$	abc
$\kappa(\omega)$	0	1	1	1	2	2	2	3
$\kappa'(\omega)$	0	1	2	3	4	4	5	7

Thus $\preceq = \preceq_\kappa$ and $\preceq' = \preceq_{\kappa'}$.

For κ, the singleton partition $\hat{0}$ is the unique finest OCF-splitting, since κ is the sum of the three singleton OCFs assigning rank 1 to the positive literal and rank 0 to the negative literal. Thus

$$\mathsf{Spl}(\kappa) = \mathsf{Spl}(\preceq) = \{\hat{0}, P_{ab|c}, P_{ac|b}, P_{a|bc}, \hat{1}\}.$$

For κ', the partition $P_{ab|c}$ is an OCF-splitting, since

$$\kappa' = \kappa'|_{\{a,b\}} \oplus \kappa'|_{\{c\}}, \quad \kappa'|_{\{a,b\}}(\bar{a}\bar{b}, a\bar{b}, \bar{a}b, ab) = (0, 1, 2, 4), \quad \kappa'|_{\{c\}}(\bar{c}, c) = (0, 3).$$

The singleton partition is not an OCF-splitting, for example

$$\kappa'(ab\bar{c}) = 4 \neq 1 + 2 + 0 = \kappa'|_{\{a\}}(a) + \kappa'|_{\{b\}}(b) + \kappa'|_{\{c\}}(\bar{c}).$$

Hence

$$\mathsf{Spl}(\kappa') = \{P_{ab|c}, \hat{1}\}.$$

For the induced TPO, however,

$$\mathsf{Spl}(\preceq') = \{\hat{0}, P_{ab|c}, P_{ac|b}, \hat{1}\}.$$

Thus $P_{a|bc}$ is missing above the finest TPO-splitting $\hat{0}$, so $\mathsf{Spl}(\preceq')$ is not coarsening-closed.

This is the structural reason for using different diagnostics in the two cases. For OCFs, coarsening closure makes the distribution of unique finest splittings sufficient for reconstructing the aggregate splitting-family structure. For TPOs, the occupied part of the coarsening lattice must be inspected directly.

3. Bounded-space analysis

This section studies non-trivial splittings in bounded unstructured spaces. Coarsening closure reduces the rank-bounded OCF analysis to exact counting by finest splittings. For depth-bounded TPOs, we first report prevalence in the full bounded space and then apply closure-sensitive diagnostics to the cases admitting non-trivial splittings.

3.1. Rank-bounded OCFs: rarity and finest shapes

On the OCF side, coarsening closure allows the class admitting non-trivial splittings to be analyzed exactly via finest splittings. For fixed signature size n and rank bound K, the underlying space consists of all normalized rank vectors in $\{0, \dots, K\}^{2^n}$ with at least one world of rank 0.

Tables 1 and 2 summarize the resulting counts: the first reports prevalence in the bounded OCF space, and the second reports finest-shape distributions among the cases admitting non-trivial splittings.

The prevalence values in Table 1 decrease steeply as K grows and, more sharply, from $n = 3$ to $n = 4$. Among the OCFs admitting non-trivial splittings, Table 2 shows a strong concentration on coarse finest shapes. For example, at $n = 4, K = 2$, shape $3 + 1$ accounts for more than 96% of these OCFs, while the singleton shape $1 + 1 + 1 + 1$ accounts for only 0.15%. Thus, even within the already sparse class, fine-grained decompositions are exceptionally rare.

Table 1
Exact numbers of rank-bounded OCFs and rank-bounded OCFs admitting non-trivial syntax splittings.

n	K	total OCFs	admitting non-trivial splittings	ratio
3	1	255	37	14.5098%
3	2	6,305	241	3.8224%
3	3	58,975	885	1.5006%
4	1	65,535	941	1.4359%
4	2	42,981,185	26,893	0.0626%
4	3	4,251,920,575	288,193	0.0068%

Table 2
Finest-shape distribution among rank-bounded OCFs admitting non-trivial syntax splittings.

n	K	shape	count	share
3	1	$2+1$	30	81.0811%
3	1	$1+1+1$	7	18.9189%
3	2	$2+1$	216	89.6266%
3	2	$1+1+1$	25	10.3734%
3	3	$2+1$	822	92.8814%
3	3	$1+1+1$	63	7.1186%
4	1	$3+1$	872	92.6674%
4	1	$2+1+1$	60	6.3762%
4	1	$1+1+1+1$	9	0.9564%
4	2	$3+1$	26,000	96.6794%
4	2	$2+1+1$	552	2.0526%
4	2	$2+2$	300	1.1155%
4	2	$1+1+1+1$	41	0.1525%
4	3	$3+1$	282,616	98.0648%
4	3	$2+1+1$	2,628	0.9119%
4	3	$2+2$	2,820	0.9785%
4	3	$1+1+1+1$	129	0.0448%

3.2. Counting by partition-lattice inversion

We now describe a direct lattice-based counting method for rank-bounded OCFs. It combines admissible blockwise rank counts with Möbius inversion on the partition lattice [11, 12] and gives an independent combinatorial reconstruction of the exact values reported in Tables 1 and 2.

We work over a fixed propositional signature Σ with $|\Sigma| = n$ and a fixed rank bound K. The aim is to count rank-bounded OCFs according to their unique finest splitting. The counting argument proceeds in two steps. First, for each partition $P \in \Pi(\Sigma)$, we count all OCFs that admit P as a splitting. Second, using the coarsening-closure property from Proposition 1, we recover the classes with prescribed unique finest splitting by inversion over the partition lattice.

Let $T_n^{\mathrm{OCF}}(K)$ denote the total number of rank-bounded OCFs over n variables. For each partition $P \in \Pi(\Sigma)$, let $A_P(K)$ denote the number of rank-bounded OCFs admitting P as a syntax splitting, and let $G_P(K)$ denote the number of rank-bounded OCFs whose unique finest splitting is exactly P.

Lemma 1. *For every $n \geq 1$ and $K \geq 0$, the number of OCFs over n variables with rank bound K is*

$$T_n^{\mathrm{OCF}}(K) = (K+1)^{2^n} - K^{2^n}.$$

Proof. There are $(K+1)^{2^n}$ functions $\Omega_\Sigma \to \{0, \ldots, K\}$, and K^{2^n} of them avoid 0. □

Lemma 2. *For every $m \geq 1$ and $r \geq 0$, let $E_m(r)$ denote the number of rank-bounded OCFs over m variables whose maximum rank is exactly r. Then*

$$E_m(0) = 1, \qquad E_m(r) = T_m^{\mathrm{OCF}}(r) - T_m^{\mathrm{OCF}}(r-1) \quad (r \geq 1).$$

Proof. The case $r = 0$ is immediate. For $r \geq 1$, an OCF has maximum exactly r iff it is bounded by r but not by $r - 1$. □

Proposition 2. *Let $P = \{\Sigma_1, \ldots, \Sigma_s\} \in \Pi(\Sigma)$ with block sizes $n_i = |\Sigma_i|$. Then*

$$A_P(K) = \sum_{\substack{r_1+\cdots+r_s \leq K \\ r_i \geq 0}} \prod_{i=1}^{s} E_{n_i}(r_i).$$

Proof. An OCF κ admits P as a syntax splitting if and only if it can be written as $\kappa = \kappa_1 \oplus \cdots \oplus \kappa_s$, where each κ_i is an OCF over the block Σ_i. For each i, let

$$r_i = \max \kappa_i.$$

If the tuple $(r_1, \ldots, r_s)$ is fixed, then the number of possible choices for the blockwise components is

$$\prod_{i=1}^{s} E_{n_i}(r_i),$$

because the blocks are chosen independently. The representation is injective: because each component has minimum 0, the marginal of $\kappa_1 \oplus \cdots \oplus \kappa_s$ over Σ_i is exactly κ_i.

Moreover,

$$\max \kappa = r_1 + \cdots + r_s.$$

The right-hand side is an upper bound by construction, and equality is attained by choosing, on each block, a world realizing the corresponding local maximum. Hence the global rank bound $\max \kappa \leq K$ is equivalent to

$$r_1 + \cdots + r_s \leq K.$$

Summing over all admissible tuples yields the claim. □

Proposition 3. *For every partition $P \in \Pi(\Sigma)$,*

$$A_P(K) = \sum_{Q \sqsubseteq P} G_Q(K).$$

Proof. A rank-bounded OCF is counted by $A_P(K)$ iff it admits P as a syntax splitting. By Proposition 1, this is equivalent to saying that its unique finest syntax splitting is some partition Q refining P. Summing over all such Q yields the identity. □

Corollary 1. *For every partition $P \in \Pi(\Sigma)$,*

$$G_P(K) = A_P(K) - \sum_{Q \sqsubset P} G_Q(K).$$

Proof. Separate the term $Q = P$ from the sum in Proposition 3 and rearrange. □

Corollary 1 is the recursive form of Möbius inversion on the partition lattice $\Pi(\Sigma)$ [11, 12].

Theorem 1. *The number of rank-bounded OCFs over Σ admitting a non-trivial syntax splitting is*

$$S_n(K) = \sum_{P \neq \hat{1}} G_P(K).$$

Proof. By Proposition 1, every rank-bounded OCF has a unique finest splitting. It admits a non-trivial syntax splitting iff this finest splitting is not the one-block partition $\hat{1}$. Summing the recovered finest-splitting classes $G_P(K)$ over all $P \neq \hat{1}$ gives the claim. □

Corollary 2. *The proportion of rank-bounded OCFs over n variables with rank bound K that admit a non-trivial syntax splitting is*

$$\rho_n(K) = \frac{S_n(K)}{T_n^{\mathrm{OCF}}(K)}.$$

The computation evaluates $T_n^{\mathrm{OCF}}(K)$, $E_m(r)$, and $A_P(K)$, recovers $G_P(K)$ by inversion, and then sums the non-trivial classes. Since finest shapes are unions of finest-partition classes, the same computation also yields exact class sizes for finest shapes by summing the corresponding values $G_P(K)$ over all partitions P of the same shape.

3.3. Depth-bounded TPOs: closure deficits and profiles

For TPOs, the OCF reconstruction strategy does not directly apply, because the finest splitting does not determine which coarsenings are splittings. This is the main reason why we do not obtain a TPO analogue of the lattice reconstruction theorem for OCFs. In the OCF case, once the finest splitting is fixed, coarsening closure determines the whole splitting family, and the additive form gives blockwise counting handles. In the TPO case, syntax splitting is comparison-based rather than reconstructive: two TPOs with the same finest splitting may occupy different subsets of the coarsenings above it. Consequently, the splittable TPO subcorpus cannot be reconstructed from finest-splitting counts alone, and the present paper therefore uses exhaustive enumeration within feasible bounds. The depth bound d restricts the maximum layer index used. Throughout this subsection, $n = |\Sigma|$ denotes the signature size. Because $|\Omega_\Sigma| = 2^n$, a TPO can have at most 2^n nonempty plausibility layers, and therefore maximal depth $2^n - 1$.

The ambient depth-bounded TPO spaces also have a closed-form count via their ordered-layer representation. We use this count as an independent check on the total-space sizes obtained by exhaustive generation. The formula is the standard enumeration by Stirling numbers of the second kind and, in the unbounded case, by Fubini numbers [12, 13].

Lemma 3. *Let Σ be a propositional signature with $|\Sigma| = n$, and write $N = |\Omega_\Sigma| = 2^n$. For $d \in \mathbb{N}_0$, let $T_n^{\mathrm{TPO}}(d)$ denote the number of TPOs over Σ whose maximum layer index is at most d. Then*

$$T_n^{\mathrm{TPO}}(d) = \sum_{j=1}^{\min(d+1,N)} j! \left\{ {N \atop j} \right\},$$

where $\left\{ {N \atop j} \right\}$ is the Stirling number of the second kind, counting set partitions of the N worlds into j nonempty layers.

Proof. By the ordered-layer characterization of TPOs from Section 2, TPOs on Ω_Σ are in bijection with ordered partitions of the $N = 2^n$ worlds into nonempty plausibility layers. If exactly j layers are used, the worlds are first partitioned into j nonempty blocks, counted by $\left\{ {N \atop j} \right\}$, and these blocks are then ordered, contributing a factor of $j!$. Thus the number of TPOs with exactly j layers is $j! \left\{ {N \atop j} \right\}$.

Under the depth bound d, at most $d + 1$ nonempty layers can occur, and in any case no more than N. Summing over $1 \leq j \leq \min(d + 1, N)$ yields the claim. □

Table 3 reports the ambient depth-bounded spaces and the number of TPOs in them that admit non-trivial syntax splittings. The closed-form count from Lemma 3 is used as an independent check on the exhaustive generation.

Since the finest splitting does not determine the full TPO-splitting family, prevalence counts are only the first step. For the cases admitting at least one non-trivial splitting, we inspect the internal

Table 3
Exhaustive depth-bounded TPO spaces and TPOs admitting non-trivial syntax splittings. Depth d is the maximum layer index used. For $n = 2$ and $n = 3$, the largest reported d equals $2^n - 1$ and thus yields the full unstructured TPO space.

n	d	total TPOs	admitting non-trivial splittings	ratio
2	1	15	5	33.3333%
2	2	51	9	17.6471%
2	3	75	17	22.6667%
3	1	255	37	14.5098%
3	2	6,051	205	3.3879%
3	3	46,875	597	1.2736%
3	4	172,875	1,773	1.0256%
3	5	364,395	4,101	1.1254%
3	6	505,515	6,885	1.3620%
3	7	545,835	8,613	1.5770%
4	1	65,535	941	1.4359%
4	2	42,915,651	25,953	0.0605%

family structure using closure diagnostics and relative profiles. The most detailed structural conclusions come from the exhaustive $n = 3$ depth series. The available $n = 4$ results are limited to shallow depth bounds. The $n = 2$ rows give the smallest complete reference case. They are structurally forced to be coarsening-closed, but the ratio of cases admitting non-trivial splittings is already nonmonotone.

Within the reported bounded spaces, TPOs admitting non-trivial splittings are already rare from $n = 3$ onward. At $n = 4, d = 2$, the ratio is already only 0.0605%. For $K = 1$ and $d = 1$, the rank-bounded OCF space and the depth-1 TPO space coincide, which explains the identical totals and splittable counts in the corresponding OCF and TPO rows. In the exhaustive $n = 3$ series, the ratio decreases sharply from $d = 1$ to $d = 4$, then rises mildly for larger depths. As the closure diagnostics below show, this rise does not indicate a return to the low-depth regime.

Closure-sensitive diagnostics. For TPOs, Proposition 1 gives only an upper-set constraint: every syntax splitting lies above the unique finest splitting, but some coarsenings may be missing. We measure these missing coarsenings against the Bell baseline. If the finest splitting has k blocks, its coarsenings are in bijection with the partitions of a k-element index set, giving $\mathcal{B}_k$ possible coarsenings. The following quantities record whether this upper set is fully occupied and, if not, how large the missing part is.

Definition 6 (Bell baseline and closure deficit)**.** *Let $\preceq$ be a TPO over Σ, and let $\mathcal{F}_{\preceq}$ be its unique finest splitting. The* Bell baseline *of $\preceq$ is the number $\mathcal{B}_{|\mathcal{F}_{\preceq}|}$, and its* closure deficit *is*

$$\Delta(\preceq) = \mathcal{B}_{|\mathcal{F}_{\preceq}|} - |\mathsf{Spl}(\preceq)|.$$

We call $\preceq$ *Bell-complete* if $\Delta(\preceq) = 0$. Since every TPO-splitting is a coarsening of the finest splitting, this is equivalent to coarsening closure of $\mathsf{Spl}(\preceq)$.

Definition 7 (Corpus-level closure measures)**.** *Let $\mathcal{D}$ be a finite corpus of TPOs admitting at least one non-trivial syntax splitting. For each $\preceq \in \mathcal{D}$, coarsening closure is evaluated on the full splitting family $\mathsf{Spl}(\preceq)$, including the trivial one-block partition.*

The coarsening-closure rate *(*CCR*) of $\mathcal{D}$ is*

$$\mathrm{CCR}(\mathcal{D}) = \frac{|\{\preceq \in \mathcal{D} \mid \Delta(\preceq) = 0\}|}{|\mathcal{D}|}.$$

We also report the restricted coarsening-closure rate

$$\mathrm{CCR}_{|\mathcal{F}|>2}(\mathcal{D}) = \frac{|\{\preceq \in \mathcal{D} \mid |\mathcal{F}_{\preceq}| > 2 \text{ and } \Delta(\preceq) = 0\}|}{|\{\preceq \in \mathcal{D} \mid |\mathcal{F}_{\preceq}| > 2\}|},$$

Table 4
Closure-sensitive diagnostics for depth-bounded TPOs over $n = 3$ variables that admit non-trivial syntax splittings. The column *cases* gives the size of this restricted class. CCR and NCD are computed on it at the given depth. The column $|\mathcal{F}| > 2$ gives the number of cases whose finest splitting has more than two blocks, and $\mathrm{CCR}_{|\mathcal{F}|>2}$ is computed on that stratum. NCD is the normalized closure deficit from Definition 7.

d	cases	CCR	NCD	$\lvert\mathcal{F}\rvert > 2$	$\mathrm{CCR}_{\lvert\mathcal{F}\rvert>2}$
1	37	1.0000	0.0000	7	1.0000
2	205	1.0000	0.0000	19	1.0000
3	597	1.0000	0.0000	51	1.0000
4	1,773	0.9459	0.0473	171	0.4386
5	4,101	0.8713	0.1126	675	0.2178
6	6,885	0.8187	0.1353	1,443	0.1351
7	8,613	0.8217	0.1289	1,827	0.1593

whenever $\mathcal{D}$ contains TPOs with $|\mathcal{F}_{\preceq}| > 2$. The normalized closure deficit *(*NCD*) of $\mathcal{D}$ is*

$$\mathrm{NCD}(\mathcal{D}) = \frac{\sum_{\preceq \in \mathcal{D}} \Delta(\preceq)}{\sum_{\preceq \in \mathcal{D}} \mathcal{B}_{|\mathcal{F}_{\preceq}|}}.$$

The restricted rate $\mathrm{CCR}_{|\mathcal{F}|>2}$ isolates the stratum in which nontrivial closure failure can occur beyond the forced extremal coarsenings. Together, CCR, $\mathrm{CCR}_{|\mathcal{F}|>2}$, and NCD distinguish Bell-completeness in the whole class, Bell-completeness among finest splittings with more than two blocks, and the aggregate size of the missing part of the Bell baseline.

The exhaustive $n = 3$ depth series also reveals a qualitative transition in family structure. As Table 4 shows, all observed splitting families are Bell-complete up to $d = 3$, whereas closure failure first appears at $d = 4$. For larger depths, the closure rates and normalized deficits show that these failures remain a persistent part of the splittable subcorpus.

The widening gap between CCR and $\mathrm{CCR}_{|\mathcal{F}|>2}$ shows that closure failure is concentrated among cases whose finest splitting has more than two blocks. This motivates a finer profile analysis of the occupied coarsenings above the finest splitting.

Definition 8 (Relative splitting profile)**.** *Let $\preceq$ be a TPO with unique finest splitting $\mathcal{F}_{\preceq} = \{\Sigma_1, \ldots, \Sigma_k\}$. For each splitting $P \in \mathsf{Spl}(\preceq)$, let $\pi_P \in \Pi(\{1, \ldots, k\})$ be the partition of the index set induced by merging exactly those indices i, j for which Σ_i and Σ_j lie in the same block of P. The* relative splitting profile *of $\preceq$ is*

$$\mathsf{Prof}(\preceq) = \{\pi_P \mid P \in \mathsf{Spl}(\preceq)\} \subseteq \Pi(\{1, \ldots, k\}).$$

Combinatorially, $\Pi(\{1, \ldots, k\})$ has $\mathcal{B}_k$ elements, and $\mathsf{Prof}(\preceq)$ contains both extremal elements of the lattice: the finest k-block partition and the trivial one-block partition. For $k = 3$, a relative profile is therefore determined by which of the three intermediate two-block coarsenings are splittings, so at most eight profiles are possible.

For $n = 3$, profile variation occurs only in the $|\mathcal{F}_{\preceq}| = 3$ stratum, which is the restricted stratum reported in Table 4. Table 5 decomposes this stratum by relative profile size and by the number of distinct relative profiles realized. At full depth $d = 7$, all closure-irregular profile types are still present, and NCD remains well above zero. The late-depth variation in the aggregate diagnostics therefore reflects a change in the relative frequencies of splitting-family types, rather than a restoration of coarsening closure.

The transition at $d = 4$ is therefore structural, not merely numerical: the first irregular families are precisely the three size-3 profiles, each retaining one intermediate two-block coarsening. For every $5 \leq d \leq 7$, the full profile spectrum appears: one profile of size 2, three of size 3, three of size 4, and the Bell-complete profile of size 5.

Table 5
Exact relative profile structure for the $|\mathcal{F}_{\preceq}| > 2$ stratum from Table 4. Each entry is of the form *count/profiles*, where *count* is the number of TPOs with the given profile size and *profiles* is the number of distinct relative profiles over the indexed finest blocks. Bell-complete profiles have size $\mathcal{B}_3 = 5$.

d	$\lvert\mathcal{F}\rvert > 2$ cases	size 2	size 3	size 4	size 5
1	7	0/0	0/0	0/0	7/1
2	19	0/0	0/0	0/0	19/1
3	51	0/0	0/0	0/0	51/1
4	171	0/0	96/3	0/0	75/1
5	675	192/1	240/3	96/3	147/1
6	1,443	384/1	432/3	432/3	195/1
7	1,827	432/1	528/3	576/3	291/1

3.4. Scope and OCF–TPO contrast

The exhaustive results are limited by rapid combinatorial growth. The reported rank-bounded OCF cases are reconstructed exactly by the lattice-based counting method above. For TPOs, full exhaustive analysis is feasible for $n = 2$ and $n = 3$, whereas the $n = 4$ results are limited to $d \leq 2$. Thus the observed first closure-irregular depth is an exact result for the exhaustive $n = 3$ series, not a claim about all larger signatures. Within the shallow $n = 4$ slice, all observed splitting families are Bell-complete, so the data do not determine at which later depth closure-irregular splitting families first appear for $n = 4$.

The OCF results primarily show scarcity: non-trivial decomposability is rare, and when it occurs the finest splitting is usually coarse. The TPO results show that scarcity is only part of the picture. From depth 4 onward in the three-variable series, the finest splitting no longer determines which coarsenings are splittings, so the relevant structure lies in the internal occupancy of the coarsening lattice.

4. Interactive splitting inspection

The empirical results above identify structural patterns at aggregate level. To connect these patterns with concrete instances, we provide an interface for constructing TPOs, inspecting their splitting families, and exploring the parameterized TPO-to-OCF bridge. The interface is used here as an inspection tool: it displays concrete TPOs, their partition-lattice structure, missing coarsenings, and bridge OCFs.

A related software effort in the broader area of ranking-based reasoning is InfOCF and its web front end InfOCF-Web [14, 15], which support reasoning with conditional belief bases and OCFs. The focus of the present interface is different: it is tailored to syntax-splitting inspection for TPOs, partition-lattice occupancy, and the TPO-to-OCF bridge studied here.

Figure 1 illustrates the interface on the closure-irregular TPO $\preceq'$ from Example 1. For the selected partition $P_{ab|c}$, the bridge construction recovers the inducing OCF κ' displayed in that example.

4.1. Constructing and inspecting TPOs

The interface operates on a user-specified finite propositional signature Σ. It displays the worlds in Ω_Σ as draggable objects, allowing to construct a TPO by arranging worlds into ordered plausibility layers. The bridge OCF is displayed with the same layer-based layout, where layers encode integer ranks rather than TPO equivalence classes. Elementary operations such as moving worlds between layers, merging and splitting layers, reversing the order, and importing or exporting instances support reproducible local perturbations.

For each constructed TPO, the interface computes its syntax splittings and displays them in the full partition lattice $\Pi(\Sigma)$. Splitting nodes, closure-violating splittings, and non-splittings are visually distinguished, so that missing coarsenings above the finest splitting can be inspected directly.

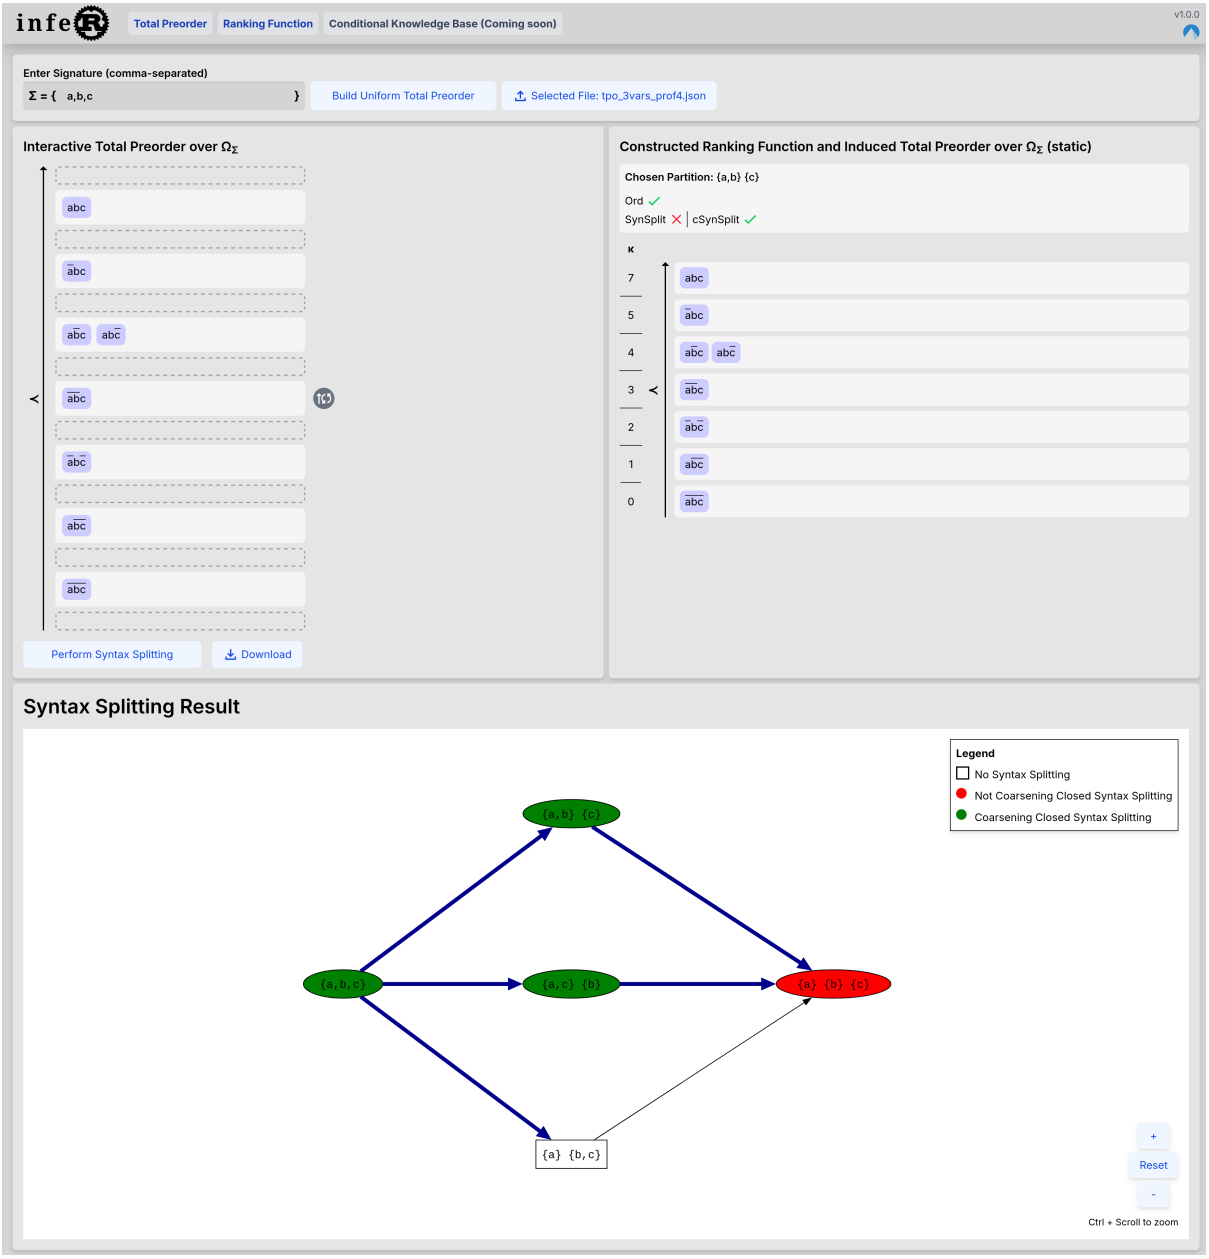


Figure 1: Interface screenshot. Left: TPO construction by plausibility layers. Bottom: partition-lattice view of $\mathrm{Spl}(\preceq')$ from Example 1, with locally coarsening-closed candidates and missing coarsenings highlighted. Right: bridge view for $P_{ab|c} = \{\{a, b\}, \{c\}\}$, displaying the bridge OCF $\kappa_{\preceq', P_{ab|c}}$, its induced TPO, and the faithfulness indicators from Subsection 4.2. The bridge OCF is shown in hybrid form: layers encode integer ranks, while the upward arrow abbreviates the induced TPO. Here $\mathrm{Spl}(\preceq')$ has two locally coarsening-closed two-block splittings, $P_{ab|c}$ and $P_{ac|b}$, but only the selected one yields the faithful bridge shown. For this selected partition, the bridge OCF coincides with κ' from Example 1: it induces $\preceq'$ and has $P_{ab|c}$ as its unique finest splitting. Its occupied ranks are $0, 1, 2, 3, 4, 5, 7$, producing an empty rank layer at 6.

4.2. A TPO-to-OCF bridge by absolute marginal sums

Beyond displaying the splitting family of a TPO $\preceq$, the interface lets the user select a partition-lattice node P and construct an associated OCF $\kappa_{\preceq, P}$. This construction is not intended as a canonical representation theorem from TPOs to OCFs. Its purpose is local and diagnostic: it turns the selected partition into an additive OCF decomposition and thereby makes visible which order relations and splitting properties survive when the TPO is compared with the induced OCF representation. In

particular, it helps connect the closure deficits observed in the aggregate TPO counts with concrete witnesses in the partition lattice.

Bridge construction for an arbitrary partition. Let $P = \{\Sigma_1, \ldots, \Sigma_k\}$ be *any* partition of Σ (not necessarily a splitting of $\preceq$). Write $\ell_{\preceq}(\omega) \in \mathbb{N}_0$ for the layer index of a world $\omega \in \Omega_\Sigma$ in $\preceq$ (with 0 the most plausible layer). For each block Σ_i, define the *absolute marginal rank* $\kappa_{\preceq,\Sigma_i}$ on Ω_{Σ_i} by

$$\kappa_{\preceq,\Sigma_i}(\omega_i) \;=\; \min\{\, \ell_{\preceq}(\nu) \mid \nu \in \Omega_\Sigma,\ \nu|_{\Sigma_i} = \omega_i \,\}.$$

The associated *bridge OCF along* P is then the marginal sum

$$\kappa_{\preceq,P}(\omega) \;=\; \sum_{i=1}^{k} \kappa_{\preceq,\Sigma_i}(\omega|_{\Sigma_i}) \qquad (\omega \in \Omega_\Sigma).$$

By construction, $\kappa_{\preceq,P}$ admits P as an OCF-splitting in the sense of Definition 3, since it is defined as a sum of blockwise components. This is true *for every partition* P. The interface deliberately does not restrict P to be a splitting of $\preceq$, so that users can inspect how the additive surrogate changes across arbitrary lattice nodes and observe where order preservation or splitting preservation is lost. When P is itself a TPO-splitting, the comparison becomes preservation-oriented: one can ask whether the bridge preserves the behaviour of $\preceq$ on the coarsenings above P.

Local closure above a pivot. For the bridge construction, the relevant question is local: whether all coarsenings of a selected splitting remain splittings. If $P \in \mathsf{Spl}(\preceq)$, then the bridge OCF $\kappa_{\preceq,P}$ has P as an OCF-splitting, and hence, by OCF coarsening closure, all coarsenings of P become splittings of the bridge. Thus, for splitting behaviour to be preserved above P, those coarsenings must already be present on the TPO side. A splitting $P \in \mathsf{Spl}(\preceq)$ is called a *coarsening-closed pivot* if its upper coarsening set is already contained in the TPO-splitting family, i.e.,

$$\{Q \in \Pi(\Sigma) \mid P \sqsubseteq Q\} \subseteq \mathsf{Spl}(\preceq).$$

Proposition 4 (Missing coarsenings become bridge splittings)**.** *Let $\preceq$ be a TPO over Σ, let $P \in \mathsf{Spl}(\preceq)$, and let $\kappa_{\preceq,P}$ be the bridge OCF along P. If there exists $Q \in \Pi(\Sigma)$ with $P \sqsubseteq Q$ and $Q \notin \mathsf{Spl}(\preceq)$, then Q is a splitting of the TPO induced by $\kappa_{\preceq,P}$. Consequently, a bridge along P cannot preserve splitting membership on the whole upper coarsening set of P unless P is a coarsening-closed pivot.*

Proof. By construction, $\kappa_{\preceq,P}$ admits P as a splitting. Since OCF-splitting families are coarsening-closed, every $Q \in \Pi(\Sigma)$ with $P \sqsubseteq Q$ is also a splitting of $\kappa_{\preceq,P}$. Every OCF-splitting is inherited as a TPO-splitting by the TPO induced by that OCF. Hence any missing coarsening $Q \notin \mathsf{Spl}(\preceq)$ above P becomes a splitting of the induced TPO. Thus the bridge creates a splitting false positive above P, so upper-set preservation can hold only for coarsening-closed pivots. □

The preservation-oriented bridge comparisons should use locally coarsening-closed pivots. Such pivots need not be unique: a closure-irregular splitting family may contain several incomparable splittings whose individual upper sets are already fully occupied. The interface exposes this local choice explicitly, so that bridge behaviour can be compared across different admissible pivots rather than only at the finest splitting.

Faithfulness indicators from epistemic-state mappings. For exploratory feedback, the interface reports indicators derived from the postulate framework for epistemic-state mappings introduced by Haldimann et al. [7]. On the ordering side, it compares $\preceq$ with the TPO induced by $\kappa_{\preceq,P}$. If exact order preservation holds, the display reports *Ord* only. Otherwise it reports which, if any, of the one-sided weakenings $wOrd^{\Rightarrow}$ and $wOrd^{\Leftarrow}$ holds, distinguishing pure coarsening from pure refinement of the induced order. On the splitting side, it compares the TPO-side and OCF-side splitting behaviour relative

to the selected partition P. If full syntax-splitting preservation holds, the display reports *SynSplit* only. Otherwise it reports whether the weaker pivot-relative indicator *cSynSplit* holds. This records whether splitting preservation still holds on the upper coarsening set above the selected coarsening-closed pivot.

These indicators are exploratory diagnostics rather than a complete postulate analysis. In particular, selecting a coarsening-closed pivot does not remove the known global incompatibility results for TPO-to-OCF mappings: there is no such mapping satisfying both (Ord) and (SynSplit), and incompatibilities persist for weaker variants as well [7]. What the pivot condition guarantees is only a local property relative to the chosen partition, namely that all coarsenings above P already belong to the TPO-splitting family.

For the witness $\preceq'$ from Example 1, Figure 1 illustrates both aspects. The TPO has two locally coarsening-closed two-block pivots, $P_{ab|c}$ and $P_{ac|b}$, but only the selected pivot $P_{ab|c}$ yields the faithful bridge shown. For this pivot, the bridge OCF coincides with the OCF κ' displayed in Example 1: it induces $\preceq'$ and has $P_{ab|c}$ as its unique finest splitting. Its occupied ranks are $0, 1, 2, 3, 4, 5, 7$, so rank 6 is empty. Thus even when *Ord* and *cSynSplit* hold, the bridge may select an OCF representative with quantitative rank gaps that are abstracted away by the TPO representation. This connects the bridge behaviour to the distinction between TPO and OCF representations studied in [16].

5. Discussion and further directions

The bounded baseline spaces make the contrast in Proposition 1 measurable. OCF-splitting families are determined by their finest splittings, while TPO-splitting families require inspection of the occupied coarsenings above the finest splitting.

From a KRR perspective, these results provide baseline information about when syntax-based decomposition should be expected without additional structural assumptions. The rarity results indicate that non-trivial splittings should not be treated as accidental background structure in unstructured bounded spaces. The contrast between OCF- and TPO-splitting families also shows which modularity phenomena are representation-sensitive: for OCFs, the finest splitting determines the whole coarsening family, whereas for TPOs, additional closure diagnostics are needed. This helps separate decomposition properties caused by the representation formalism from those caused by structured epistemic input.

These findings are limited to controlled baseline spaces: rank-bounded OCFs and depth-bounded TPOs over fixed signatures. Epistemic states induced by structured knowledge bases may form highly nonuniform subfamilies, and their modularity behaviour may therefore differ from the baseline frequencies reported here. The interface links the aggregate patterns to concrete witnesses, missing coarsenings, pivot choices, and bridge instances.

Several directions remain for future work. The finite OCF counts suggest studying how quickly the fraction admitting non-trivial splittings decreases as the signature size and rank bound grow. Such an analysis should ideally include shape-sensitive bounds, not only aggregate rates. For TPOs, constructive generators for prescribed finest splittings or closure profiles would help move beyond shallow exhaustive depth bounds.

For the bridge construction, a pivot-sensitive postulate analysis remains open. Since strong preservation requirements such as (Ord) and (SynSplit) are incompatible in the TPO-to-OCF direction [7], the relevant next question is which weakened properties remain achievable relative to a chosen pivot. The exploratory interface could also be extended from TPO inspection to direct construction of OCFs and conditional knowledge bases. This would make it possible to compare syntax-splitting behaviour across structured input and induced OCFs, thereby connecting the present baseline analysis with ranking-based resources such as CLKR and InfOCF(-Web) [14, 15, 17]. Finally, the approach could be adapted to related formalisms, such as conditional syntax splitting [18, 19].

6. Conclusion

We investigated syntax splitting for rank-bounded OCFs and depth-bounded TPOs over finite propositional signatures. For OCFs, coarsening closure makes the unique finest splitting an exact counting handle, yielding rarity counts and finest-shape distributions. For TPOs, the finest splitting does not by itself determine the full splitting family, so we use closure deficits and relative profiles to record the occupied coarsenings above it. In the exhaustive three-variable depth series, closure-irregular profiles first occur at depth 4, and the full spectrum of relative profiles over three finest blocks occurs for every $5 \leq d \leq 7$.

The bounded baseline spaces show that non-trivial syntax splittings are rare, but rarity is only part of the picture. The OCF counting results give exact prevalence values and finest-shape distributions. For TPOs, the closure deficits and relative profiles expose structure that is invisible from prevalence alone.

The accompanying interface connects the aggregate enumeration with individual witnesses by displaying concrete splitting families, missing coarsenings, and bridge OCFs.

Declaration on Generative AI

During the preparation of this work, the authors used Mistral and DeepL for grammar and spelling checks, text translation, peer review simulation based on related reviews, assistance in writing parts of the code for dataset generation and analysis, and assistance in developing the inspection interface. All generated or suggested text, code, and interface components were reviewed, edited, tested, and validated by the authors as needed. The authors take full responsibility for the publication's content, the reported results, and the accompanying computational artifacts.

References

[1] R. Parikh, Beliefs, belief revision, and splitting languages, in: L. S. Moss, J. Ginzburg, M. de Rijke (Eds.), Logic, Language and Computation, Volume 2, number 96 in CSLI Lecture Notes, CSLI Publications, Stanford, CA, 1999, pp. 266–278. Earlier version appeared in the preliminary proceedings of the 1996 conference on information-theoretic approaches to logic, language, and computation.

[2] D. C. Makinson, G. Kourousias, Respecting relevance in belief change, Análisis Filosófico 26 (2006) 53–61. URL: https://analisisfilosofico.org/index.php/af/article/view/194.

[3] P. Peppas, M. Williams, S. Chopra, N. Y. Foo, Relevance in belief revision, Artif. Intell. 229 (2015) 126–138. doi:10.1016/J.ARTINT.2015.08.007.

[4] G. Kern-Isberner, C. Beierle, G. Brewka, Syntax splitting = relevance + independence: New postulates for nonmonotonic reasoning from conditional belief bases, in: D. Calvanese, E. Erdem, M. Thielscher (Eds.), Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR, Rhodes, Greece, September 12-18, 2020, pp. 560–571. doi:10.24963/KR.2020/56.

[5] G. Kern-Isberner, G. Brewka, Strong syntax splitting for iterated belief revision, in: C. Sierra (Ed.), Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017, ijcai.org, 2017, pp. 1131–1137. doi:10.24963/IJCAI.2017/157.

[6] J. P. Haldimann, C. Beierle, Finest syntax splittings of ranking functions and total preorders on worlds, in: P. Marquis, T. C. Son, G. Kern-Isberner (Eds.), Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning, KR 2023, Rhodes, Greece, September 2-8, 2023, 2023, pp. 747–751. doi:10.24963/KR.2023/75.

[7] J. P. Haldimann, C. Beierle, G. Kern-Isberner, Epistemic state mappings among ranking functions and total preorders, FLAP 10 (2023) 155–191. URL: https://collegepublications.co.uk/downloads/ifcolog00058.pdf.

[8] F. Bräuer, J. P. Haldimann, Computing the finest syntax splitting of ranking functions from bipartitions, 2026. Unpublished Manuscript.

[9] F. Bräuer, Implementing a System for the Computation of Syntax Splittings for Total Preorders and for Conditional Knowledge Bases, Master's thesis, FernUniversität in Hagen, 2022. doi:10.5281/zenodo.7274824.

[10] W. Spohn, Ordinal conditional functions. a dynamic theory of epistemic states, in: W. L. Harper, B. Skyrms (Eds.), Causation in Decision, Belief Change, and Statistics, vol. II, Kluwer Academic Publishers, 1988, pp. 105–134.

[11] G.-C. Rota, On the foundations of combinatorial theory i. theory of möbius functions, Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete 2 (1964) 340–368. URL: https://api.semanticscholar.org/CorpusID:121334025.

[12] R. P. Stanley, Enumerative Combinatorics, Volume 1, volume 49 of *Cambridge Studies in Advanced Mathematics*, 2 ed., Cambridge University Press, Cambridge, 2012. URL: https://doi.org/10.1017/CBO9781139058520.

[13] L. Comtet, Advanced Combinatorics, Reidel, Dordrecht, 1974.

[14] S. Kutsch, C. Beierle, Infocf-web: An online tool for nonmonotonic reasoning with conditionals and ranking functions, in: Z.-H. Zhou (Ed.), Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21, International Joint Conferences on Artificial Intelligence Organization, 2021, pp. 4996–4999. URL: https://doi.org/10.24963/ijcai.2021/711. doi:10.24963/ijcai.2021/711, demo Track.

[15] C. Beierle, J. Haldimann, A. Sanin, A. Spang, L.-P. Spiegel, M. von Berg, The infocf library for reasoning with conditional belief bases, in: Logics in Artificial Intelligence: 19th European Conference, JELIA 2025, Kutaisi, Georgia, September 1–4, 2025, Proceedings, Part II, Springer-Verlag, Berlin, Heidelberg, 2025, p. 19–27. URL: https://doi.org/10.1007/978-3-032-04590-4_2. doi:10.1007/978-3-032-04590-4_2.

[16] G. Kern-Isberner, A. Hahn, J. Haldimann, C. Beierle, Total preorders vs. ranking functions under belief revision - the dynamics of empty layers, in: Proceedings of the 21st International Conference on Principles of Knowledge Representation and Reasoning, 2024, pp. 498–508. doi:10.24963/kr.2024/47.

[17] C. Beierle, J. Haldimann, L. Schwarzer, CLKR: Conditional logic and knowledge representation, Künstliche Intelligenz 38 (2024) 61–67. doi:10.1007/s13218-024-00842-z.

[18] J. Heyninck, G. Kern-Isberner, T. Meyer, J. P. Haldimann, C. Beierle, Conditional syntax splitting for non-monotonic inference operators, Proceedings of the AAAI Conference on Artificial Intelligence 37 (2023) 6416–6424. URL: https://ojs.aaai.org/index.php/AAAI/article/view/25789. doi:10.1609/aaai.v37i5.25789.

[19] L.-P. Spiegel, J. Haldimann, J. Heyninck, G. Kern-Isberner, C. Beierle, Generalized safe conditional syntax splitting of belief bases, in: J. Kwok (Ed.), Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25, International Joint Conferences on Artificial Intelligence Organization, 2025, pp. 4678–4686. URL: https://doi.org/10.24963/ijcai.2025/521. doi:10.24963/ijcai.2025/521, main Track.