

THE COMPLEXITY OF GENERAL-VALUED CONSTRAINT SATISFACTION PROBLEMS SEEN FROM THE OTHER SIDE*

CLÉMENT CARBONNEL[†], MIGUEL ROMERO[‡], AND STANISLAV ŽIVNÝ[§]

Abstract. The constraint satisfaction problem (CSP) is concerned with homomorphisms between two structures. For CSPs with restricted left-hand-side structures, the results of Dalmau, Kolaitis, and Vardi [*Proceedings of the 8th International Conference on Principles and Practice of Constraint Programming*, Springer, New York, 2002, pp. 310–326], Grohe [*J. ACM*, 54 (2007), 1], and Atserias, Bulatov, and Dalmau [*Proceedings of the 34th International Colloquium on Automata, Languages and Programming*, Springer, New York, 2007, pp. 279–290] establish the precise borderline of polynomial-time solvability (subject to complexity-theoretic assumptions) and of solvability by bounded-consistency algorithms (unconditionally) as bounded treewidth modulo homomorphic equivalence. The general-valued constraint satisfaction problem (VCSP) is a generalization of the CSP concerned with homomorphisms between two *valued* structures. For VCSPs with restricted left-hand-side valued structures, we establish the precise borderline of polynomial-time solvability (subject to complexity-theoretic assumptions) and of solvability by the k th level of the Sherali–Adams LP hierarchy (unconditionally). We also obtain results on related problems concerned with finding a solution and recognizing the tractable cases; the latter has an application in database theory.

Key words. valued constraint satisfaction, homomorphism problems, fractional homomorphism, treewidth, Sherali–Adams LP relaxation

AMS subject classifications. 68Q25, 68R01, 03B70, 90C05

DOI. 10.1137/19M1250121

1. Introduction.

1.1. Constraint satisfaction problems. The homomorphism problem for relational structures is a fundamental computer science problem: Given two relational structures \mathbf{A} and \mathbf{B} over the same signature, the goal is to determine the existence of a homomorphism from \mathbf{A} to \mathbf{B} (see, e.g., the book by Hell and Nešetřil on this topic [35]). The homomorphism problem is known to be equivalent to the evaluation problem and the containment problem for conjunctive database queries [13, 37] and also to the constraint satisfaction problem (CSP) [24], which originated in artificial intelligence [43] and provides a common framework for expressing a wide range of both theoretical and real-life combinatorial problems.

For a class \mathcal{C} of relational structures, we denote by $\text{CSP}(\mathcal{C}, -)$ the restriction of the homomorphism problem in which the input structure \mathbf{A} belongs to \mathcal{C} and the input structure \mathbf{B} is arbitrary (these types of restrictions are known as *structural* restrictions). Similarly, by $\text{CSP}(-, \mathcal{C})$ we denote the restriction of the homomorphism

*Received by the editors March 15, 2019; accepted for publication (in revised form) October 13, 2021; published electronically January 25, 2022. An extended abstract of this work appeared in the *Proceedings of the 59th Annual IEEE Symposium on Foundations of Computer Science*, 2018.

<https://doi.org/10.1137/19M1250121>

Funding: This research was supported by funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement 714532). The second author was supported by Fondecyt grant 11200956. The third author was supported by a Royal Society University Research Fellowship.

[†]CNRS, University of Montpellier, LIRMM UMR 5506, CC477, 161 rue Ada, 34095 Montpellier Cedex 5, France (clement.carbonnel@lirmm.fr).

[‡]Faculty of Engineering and Science, Universidad Adolfo Ibáñez, Diagonal Las Torres 2640, Peñalolén, Santiago, Chile (miguel.romero.o@uai.cl).

[§]Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, OX1 3QD Oxford, UK (standa.zivny@cs.ox.ac.uk).

problem in which the input structure \mathbf{A} is arbitrary and the input structure \mathbf{B} belongs to \mathcal{C} .

Feder and Vardi initiated the study of $\text{CSP}(-, \{\mathbf{B}\})$, also known as nonuniform CSPs, and famously conjectured that for every fixed finite structure \mathbf{B} , $\text{CSP}(-, \{\mathbf{B}\})$ is in PTIME or $\text{CSP}(-, \{\mathbf{B}\})$ is NP-complete. For example, if \mathbf{B} is a clique on k vertices then $\text{CSP}(-, \{\mathbf{B}\})$ is the well-known k -coloring problem, which is known to be in PTIME for $k \leq 2$ and NP-complete for $k \geq 3$. Most of the progress on the Feder–Vardi conjecture (e.g., [7, 3, 36, 11, 2]) is based on the algebraic approach [10], culminating in two (affirmative) solutions to the Feder–Vardi conjecture obtained independently by Bulatov [8] and Zhuk [52].

$\text{CSP}(\mathcal{C}, -)$ is interesting only if \mathcal{C} is an infinite class of structures as otherwise $\text{CSP}(\mathcal{C}, -)$ is always in PTIME. (This is, however, not the case for $\text{CSP}(-, \mathcal{C})$, as we have seen in the example of 3-coloring.) Freuder observed that $\text{CSP}(\mathcal{C}, -)$ is in PTIME if \mathcal{C} consists of trees [26] or, more generally, if it has bounded treewidth [27]. Later, Dalmau, Kolaitis, and Vardi showed that $\text{CSP}(\mathcal{C}, -)$ is solved by k -consistency, a fundamental local propagation algorithm [17], if \mathcal{C} is of bounded treewidth *modulo homomorphic equivalence*, i.e., if the treewidth of the cores of the structures from \mathcal{C} is at most k , for some fixed $k \geq 1$ [16]. Atserias, Bulatov, and Dalmau showed that this is precisely the class of structures solved by k -consistency [1]. Grohe proved [32] that the tractability result of Dalmau, Kolaitis, and Vardi [16] is optimal for classes \mathcal{C} of bounded arity: Under the assumption that $\text{FPT} \neq \text{W}[1]$, $\text{CSP}(\mathcal{C}, -)$ is tractable if and only if \mathcal{C} has bounded treewidth modulo homomorphic equivalence.

1.2. General-valued constraint satisfaction problems. General-valued constraint satisfaction problems (VCSPs) are generalizations of CSPs which allow for not only decision problems but also optimization problems (and the mix of the two) to be considered in one framework [15]. In the case of VCSPs we deal with *valued* structures. Regarding tractable restrictions, the situation of the nonuniform case is by now well-understood. Indeed, it holds that for any fixed valued structure \mathbb{B} , either $\text{VCSP}(-, \{\mathbb{B}\})$ is in PTIME or $\text{VCSP}(-, \{\mathbb{B}\})$ is NP-complete [41, 39].

For structural restrictions, it is a folklore result that $\text{VCSP}(\mathcal{C}, -)$ is tractable if \mathcal{C} is of bounded treewidth; see, e.g., [4]. So is the belief that the $(k+1)$ st level of the Sherali–Adams LP hierarchy [48] solves $\text{VCSP}(\mathcal{C}, -)$ to optimality if the treewidth of all structures in \mathcal{C} is at most k . (We are not aware of any reference for this. For certain special problems, it is discussed in [5]. For the extension complexity of such problems, see [38].) However, unlike the CSP case, the precise borderline of polynomial-time solvability and the power of fundamental algorithms (such as the Sherali–Adams LP hierarchy) for $\text{VCSP}(\mathcal{C}, -)$ is still unknown. Understanding these complexity and algorithmic frontiers for $\text{VCSP}(\mathcal{C}, -)$ is the main goal of this paper.

1.3. Contributions. We study the problem $\text{VCSP}(\mathcal{C}, -)$ for classes \mathcal{C} of valued structures over finite signatures and give three main results.

(1) *Complexity classification.* As our first result, we give (in Theorem 4.1) a complete complexity classification of $\text{VCSP}(\mathcal{C}, -)$ and identify the precise borderline of tractability for classes \mathcal{C} of bounded arity. A key ingredient in our result is a novel notion of *valued equivalence* and a characterization of this notion in terms of *valued cores*. More precisely, we show that $\text{VCSP}(\mathcal{C}, -)$ is tractable if and only if \mathcal{C} has bounded treewidth modulo valued equivalence. This latter notion *strictly* generalizes bounded treewidth and it is *strictly* weaker than bounded treewidth modulo homomorphic equivalence. Our proof builds on the characterization by Dalmau, Kolaitis, and Vardi [16] and Grohe [32] for CSPs. We show that the newly identified tractable classes are solvable by the Sherali–Adams LP hierarchy.

(2) *Power of Sherali–Adams.* Our second result (Theorem 5.4) gives a precise characterization of the power of Sherali–Adams for $\text{VCSP}(\mathcal{C}, -)$. In particular, we show that the $(k + 1)$ st level of the Sherali–Adams LP hierarchy solves $\text{VCSP}(\mathcal{C}, -)$ to optimality if and only if the valued cores of the structures from \mathcal{C} have treewidth *modulo scopes* at most k and the *overlaps* of scopes are of size at most $k + 1$. The proof builds on the work of Atserias, Bulatov, and Dalmau [1] and Thapper and Živný [51], as well as on an adaptation of the classical result connecting treewidth and brambles by Seymour and Thomas [47].

(3) *Search VCSP.* Our first two results are for the VCSP in which we ask for the cost of an optimal solution. It is also possible to define the VCSP as a *search* problem, in which one is additionally required to return a solution with the optimal cost. A complete characterization of tractable search cases in terms of structural properties of (a class of structures) \mathcal{C} is open even for CSPs and there is some evidence that the tractability frontier cannot be captured in simple terms.¹ Building on our first two results as well as on techniques from [50], we give in section 6 a characterization of the tractable cases for search $\text{VCSP}(\mathcal{C}, -)$ in terms of tractable core computation (Theorem 6.1).

(4) *Additional results.* In addition to our main results, we provide in section 7 tight complexity bounds for several problems related to our classification results, e.g., deciding whether the treewidth is at most k modulo valued equivalence, deciding solvability by the k th level of the Sherali–Adams LP hierarchy, and deciding valued equivalence for valued structures. These results have interesting consequences to database theory, specifically, to the evaluation and optimization of conjunctive queries (CQs) over *annotated* databases. In particular, we show that the containment problem of CQs over the tropical semiring is in NP, thus improving on the work of [40], which put it in Π_2^P .

1.4. Related work. In his Ph.D. thesis [22], Färnqvist studied the complexity of $\text{VCSP}(\mathcal{C}, -)$ and some other fragments of VCSPs (see also [23, 21]). He considered a very specific framework that only allows for particular types of classes \mathcal{C} 's to be classified. For these classes, he showed that only bounded treewidth gives rise to tractability (assuming bounded arity) and asked about more general classes. In particular, decision CSPs do *not* fit in his framework and Grohe's classification [32] is not implied by Färnqvist's work. In contrast, our characterization (of *all* classes \mathcal{C} 's of valued structures) gives rise to new tractable cases going beyond those identified by Färnqvist. Moreover, we can derive both Grohe's classification and Färnqvist's classification directly from our results, as explained in section 4.

It is known that Grohe's characterization applies only to classes \mathcal{C} of *bounded arity*, i.e., when the arities of the signatures are always bounded by a constant (for instance, CSPs over digraphs) and fails for classes of unbounded arity. In this direction, several hypergraph-based restrictions that lead to tractability have been proposed (for a survey see, e.g., [28]). Nevertheless, the precise tractability frontier for $\text{CSP}(\mathcal{C}, -)$ is not known. The situation is different for *fixed-parameter tractability*: Marx gave a complete classification of the fixed-parameter tractable restrictions $\text{CSP}(\mathcal{C}, -)$ for classes \mathcal{C} of structures of unbounded arity [42]. Gottlob, Greco, and Scarcello [29] and Färnqvist [21] applied well-known hypergraph-based tractable restrictions of CSPs to VCSPs.

¹In particular, [9, Lemma 1] shows that a description of tractable cases of $\text{SCSP}(\mathcal{C}, -)$, which is the search variant of $\text{CSP}(\mathcal{C}, -)$ defined in section 6, would imply a description of tractable cases of $\text{CSP}(-, \{\mathbf{B}\})$.

2. Preliminaries. We assume familiarity with relational structures and homomorphisms. Briefly, a *relational signature* is a finite set τ of relation symbols R , each with a specified arity $\text{ar}(R)$. A *relational structure* \mathbf{A} over a relational signature τ (or a relational τ -structure, for short) is a finite nonempty universe A together with one relation $R^{\mathbf{A}} \subseteq A^{\text{ar}(R)}$ for each symbol $R \in \tau$. A *homomorphism* from a relational τ -structure \mathbf{A} (with universe A) to a relational τ -structure \mathbf{B} (with universe B) is a mapping $h : A \rightarrow B$ such that for all $R \in \tau$ and all tuples $\mathbf{x} \in R^{\mathbf{A}}$ we have $h(\mathbf{x}) \in R^{\mathbf{B}}$. We refer the reader to [35] for more details.

We use $\overline{\mathbb{Q}}_{\geq 0}$ to denote the set of nonnegative rational numbers with positive infinity, i.e., $\overline{\mathbb{Q}}_{\geq 0} = \mathbb{Q}_{\geq 0} \cup \{\infty\}$. As usual, we assume that $\infty + c = c + \infty = \infty$ for all $c \in \overline{\mathbb{Q}}_{\geq 0}$, $\infty \times 0 = 0 \times \infty = 0$, $\infty \times c = c \times \infty = \infty$ for all $c > 0$, and $c/\infty = 0$ for all $c \in \overline{\mathbb{Q}}_{\geq 0}$.

Valued structures. A *signature* is a finite set σ of function symbols f , each with a specified arity $\text{ar}(f)$. A *valued structure* \mathbb{A} over a signature σ (or valued σ -structure, for short) is a finite nonempty universe A together with one function $f^{\mathbb{A}} : A^{\text{ar}(f)} \rightarrow \overline{\mathbb{Q}}_{\geq 0}$ for each symbol $f \in \sigma$. We define $\text{tup}(\mathbb{A})$ to be the set of all pairs (f, \mathbf{x}) such that $f \in \sigma$ and $\mathbf{x} \in A^{\text{ar}(f)}$. The set of *positive* tuples of \mathbb{A} is defined by $\text{tup}(\mathbb{A})_{>0} := \{(f, \mathbf{x}) \in \text{tup}(\mathbb{A}) \mid f^{\mathbb{A}}(\mathbf{x}) > 0\}$. If $\mathbb{A}, \mathbb{B}, \dots$ are valued structures, then A, B, \dots denote their respective universes.

VCSPs. We define VCSPs as in [49]. An instance of the VCSP is given by two valued structures \mathbb{A} and \mathbb{B} over the same signature σ . For a mapping $h : A \rightarrow B$, we define

$$\text{cost}(h) = \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{B}}(h(\mathbf{x})).$$

If needed, we will write $\text{cost}_{\mathbb{A} \rightarrow \mathbb{B}}(h)$ for $\text{cost}(h)$ to emphasize the source and target structures. Given a VCSP instance, the goal is to find the minimum cost over all possible mappings $h : A \rightarrow B$. We denote this cost by $\text{opt}(\mathbb{A}, \mathbb{B})$.

For a class \mathcal{C} of valued structures (not necessarily over the same signature), we denote by $\text{VCSP}(\mathcal{C}, -)$ the class of VCSP instances (\mathbb{A}, \mathbb{B}) such that $\mathbb{A} \in \mathcal{C}$. We say that $\text{VCSP}(\mathcal{C}, -)$ is in PTIME, the class of problems solvable in *polynomial time*, if there is a deterministic algorithm that solves any instance (\mathbb{A}, \mathbb{B}) of $\text{VCSP}(\mathcal{C}, -)$ in time $(|\mathbb{A}| + |\mathbb{B}|)^{O(1)}$. We also consider the parameterised version of $\text{VCSP}(\mathcal{C}, -)$, denoted by $p\text{-VCSP}(\mathcal{C}, -)$, where the parameter is $|\mathbb{A}|$. We say that $p\text{-VCSP}(\mathcal{C}, -)$ is in FPT, the class of problems that are *fixed-parameter tractable*, if there is a deterministic algorithm that solves any instance (\mathbb{A}, \mathbb{B}) of $p\text{-VCSP}(\mathcal{C}, -)$ in time $f(|\mathbb{A}|) \cdot |\mathbb{B}|^{O(1)}$, where $f : \mathbb{N} \rightarrow \mathbb{N}$ is an arbitrary computable function. The class $W[1]$, introduced in [19], can be seen as an analogue of NP in parameterized complexity theory. Proving $W[1]$ -hardness of $p\text{-VCSP}(\mathcal{C}, -)$ (under an fpt-reduction, formally defined in section 4.1) is a strong indication that $p\text{-VCSP}(\mathcal{C}, -)$ is not in FPT as it is believed that $\text{FPT} \neq W[1]$. We refer the reader to [25] for more details on parameterized complexity.

Treewidth of a valued structure. The notion of treewidth (originally introduced by Bertelé and Brioschi [4] and later rediscovered by Robertson and Seymour [44]) is a well-known measure of the tree-likeness of a graph [18]. Let $G = (V(G), E(G))$ be a graph. A *tree decomposition* of G is a pair (T, β) where $T = (V(T), E(T))$ is a tree and β is a function that maps each node $t \in V(T)$ to a subset of $V(G)$ such that

1. $V(G) = \bigcup_{t \in V(T)} \beta(t)$,
2. for every $u \in V(G)$, the set $\{t \in V(T) \mid u \in \beta(t)\}$ induces a connected subgraph of T , and

3. for every edge $\{u, v\} \in E(G)$, there is a node $t \in V(T)$ with $\{u, v\} \subseteq \beta(t)$. The *width* of the decomposition (T, β) is $\max\{|\beta(t)| \mid t \in V(T)\} - 1$. The *treewidth* $tw(G)$ of a graph G is the minimum width over all its tree decompositions.

Let \mathbf{A} be a relational structure over relational signature τ . Its *Gaifman graph* (also known as its *primal graph*), denoted by $G(\mathbf{A})$, is the graph whose vertex set is the universe of \mathbf{A} and whose edges are the pairs $\{u, v\}$ for which there is a tuple \mathbf{x} and a relation symbol $R \in \tau$ such that u, v appear in \mathbf{x} and $\mathbf{x} \in R^{\mathbf{A}}$. We define the treewidth of \mathbf{A} to be $tw(\mathbf{A}) = tw(G(\mathbf{A}))$.

Let \mathbb{A} be a valued σ -structure. If \mathbb{A} is the left-hand side of an instance of the VCSP, the only tuples relevant to the problem are those in $\text{tup}(\mathbb{A})_{>0}$. Hence, in order to define structural restrictions and, in particular, the notion of treewidth, we focus on the structure induced by $\text{tup}(\mathbb{A})_{>0}$. Formally, we associate with the signature σ a relational signature $\text{rel}(\sigma)$ that contains, for every $f \in \sigma$, a relation symbol R_f of the same arity as f . We define $\text{Pos}(\mathbb{A})$ to be the relational structure over $\text{rel}(\sigma)$ with the same universe A of \mathbb{A} such that $\mathbf{x} \in R_f^{\text{Pos}(\mathbb{A})}$ if and only if $f^{\mathbb{A}}(\mathbf{x}) > 0$. We let the treewidth of \mathbb{A} be $tw(\mathbb{A}) = tw(\text{Pos}(\mathbb{A}))$.

Remark 2.1. Observe that, in the VCSP, we allow infinite costs not only in \mathbb{B} but also in the left-hand-side structure \mathbb{A} . This allows us to consider the VCSP as the minimum-cost mapping problem between two mathematical objects of the same nature. Intuitively, mapping the tuples of \mathbb{A} to infinity ensures that those are logically equivalent to *hard constraints*, as any minimum-cost solution of finite cost must map them to tuples of cost exactly 0 in \mathbb{B} . Thus, decision CSPs, which are $\{0, \infty\}$ -valued VCSPs, are a special case of our definition and all our results also apply to CSPs.

3. Equivalence for valued structures. We start by introducing the notion of valued equivalence that is crucial for our results.

DEFINITION 3.1. Let \mathbb{A}, \mathbb{B} be valued σ -structures. We say that \mathbb{A} improves \mathbb{B} , denoted by $\mathbb{A} \preceq \mathbb{B}$, if $\text{opt}(\mathbb{A}, \mathbb{C}) \leq \text{opt}(\mathbb{B}, \mathbb{C})$ for all valued σ -structures \mathbb{C} .

When two valued structures improve each other, we call them equivalent. In section 1, we used the term “valued equivalence.” In the rest of the paper, we drop the word “valued” unless needed for clarity.

DEFINITION 3.2. Let \mathbb{A}, \mathbb{B} be valued σ -structures. We say that \mathbb{A} and \mathbb{B} are equivalent, denoted by $\mathbb{A} \equiv \mathbb{B}$, if $\mathbb{A} \preceq \mathbb{B}$ and $\mathbb{B} \preceq \mathbb{A}$.

Hence, two valued σ -structures \mathbb{A} and \mathbb{B} are equivalent if they have the same optimal cost over all right-hand-side valued structures. Observe that equivalence implies *homomorphic* equivalence of $\text{Pos}(\mathbb{A})$ and $\text{Pos}(\mathbb{B})$. Indeed, whenever $\text{Pos}(\mathbb{A})$ is not homomorphic to $\text{Pos}(\mathbb{B})$, we can define a valued σ -structure \mathbb{C} as follows: \mathbb{C} and \mathbb{B} have the same universe B , and for every $f \in \sigma$ and $\mathbf{x} \in B^{\text{ar}(f)}$, $f^{\mathbb{C}}(\mathbf{x}) = 0$ if $(f, \mathbf{x}) \in \text{tup}(\mathbb{B})_{>0}$, and $f^{\mathbb{C}}(\mathbf{x}) = \infty$ otherwise. Since the identity mapping has cost 0, we have $\text{opt}(\mathbb{B}, \mathbb{C}) = 0$. On the other hand, $\text{opt}(\mathbb{A}, \mathbb{C}) = \infty$ as any finite-cost map from \mathbb{A} to \mathbb{C} is a homomorphism from $\text{Pos}(\mathbb{A})$ to $\text{Pos}(\mathbb{B})$. Consequently, $\mathbb{A} \not\equiv \mathbb{B}$. As the following example shows, the converse does not hold in general.

Example 3.3. Consider the valued σ -structures \mathbb{A} and \mathbb{B} from Figure 1, with $\sigma = \{f, \mu\}$, where f and μ are binary and unary function symbols, respectively. In Figure 1, μ is represented by the numbers labeling the nodes, and f is represented as follows: pairs receiving cost ∞ are depicted as edges, while all remaining pairs are mapped to 0. Observe that $\text{Pos}(\mathbb{A})$ and $\text{Pos}(\mathbb{B})$ are homomorphically equivalent. However, they are not (valued) equivalent. Indeed, consider the valued σ -structure \mathbb{C} with same universe




FIG. 1. The valued structures from Examples 3.3 and 3.14.

B as \mathbb{B} such that (i) for every $\mathbf{x} \in B^2$, $f^{\mathbb{C}}(\mathbf{x}) = 0$ if $f^{\mathbb{B}}(\mathbf{x}) = \infty$, otherwise $f^{\mathbb{C}}(\mathbf{x}) = \infty$, and (ii) $\mu^{\mathbb{C}} = \mu^{\mathbb{B}}$. It follows that $\text{opt}(\mathbb{A}, \mathbb{C}) = 9$ and $\text{opt}(\mathbb{B}, \mathbb{C}) = 5$, and thus $\mathbb{A} \not\equiv \mathbb{B}$.

In the rest of the section, we give characterizations of equivalence in terms of certain types of homomorphisms and (valued) cores. We conclude with a useful characterization of bounded treewidth modulo equivalence in terms of cores. In order to keep the flow uninterrupted, we defer (most of) the proofs from this section to Appendix A.

3.1. Inverse fractional homomorphisms. A homomorphism between two relational structures is a structure-preserving mapping; e.g., for graphs, this means preserving adjacency—edges are mapped to edges only. An inverse homomorphism between two relational structures is a mapping with the property that the preimage of any tuple in a relation in the target structure can only contain tuples in the corresponding relation in the source structure; e.g., for graphs, only edges (and not nonedges) from the source graph can be mapped to the edges of the target graph.

A *fractional* homomorphism between two valued structures has played an important role in the work of Thapper and Živný on the power of the basic LP relaxation for VCSPs with a fixed valued constraint language [49]. Intuitively, it is a probability distribution over mappings between the universes of the two structures with the property that the expected cost is not increased [49]. In this paper, we will use a different but related notion of *inverse* fractional homomorphism (IFH).

First, we need to define two elementary operations on valued structures (summation and scaling), as well as pointwise comparison. Let $\mathbb{A}_1, \dots, \mathbb{A}_k$ be valued σ -structures. The valued σ -structure $\mathbb{C} = \sum_i \mathbb{A}_i$ has universe $C = \cup_i A_i$, and for each $f \in \sigma$ and $\mathbf{x} \in C^{\text{ar}(f)}$, $f^{\mathbb{C}}(\mathbf{x}) = \sum_{i: \mathbf{x} \in A_i^{\text{ar}(f)}} f^{\mathbb{A}_i}(\mathbf{x})$. Similarly, for $c \in \mathbb{Q}_{\geq 0}$, the structure $c \cdot \mathbb{A}_1$ is obtained from \mathbb{A}_1 by scaling its functions by c , i.e., $f^{c \cdot \mathbb{A}_1}(\mathbf{x}) = c \cdot f^{\mathbb{A}_1}(\mathbf{x})$. Finally, if \mathbb{A} and \mathbb{B} are valued σ -structures with $A = B$, then we define the partial order $\mathbb{A} \leq \mathbb{B}$ pointwise, i.e., $\mathbb{A} \leq \mathbb{B}$ if and only if $f^{\mathbb{A}}(\mathbf{x}) \leq f^{\mathbb{B}}(\mathbf{x})$ for every symbol f and tuple \mathbf{x} .

For sets A and B , we denote by B^A the set of all mappings $g: A \rightarrow B$ from A to B . For a probability distribution ω on B^A , the *support* of ω is the set $\text{supp}(\omega) := \{g \in B^A \mid \omega(g) > 0\}$. The following notion of an “image valued structure” will be useful.

DEFINITION 3.4. Let \mathbb{A} be a valued σ -structure and $g: A \rightarrow B$ be a mapping. We define $g(\mathbb{A})$ to be the valued σ -structure over universe $g(A)$ such that, for all $f \in \sigma$ and $\mathbf{y} \in g(A)^{\text{ar}(f)}$,

$$f^{g(\mathbb{A})}(\mathbf{y}) = f^{\mathbb{A}}(g^{-1}(\mathbf{y})) := \sum_{\mathbf{x} \in A^{\text{ar}(f)}: g(\mathbf{x})=\mathbf{y}} f^{\mathbb{A}}(\mathbf{x}).$$

If ω is a probability distribution over B^A we put $\omega(\mathbb{A}) = \sum_{g \in B^A} \omega(g)g(\mathbb{A})$.

Observe that, in the definition above, the universe of $\omega(\mathbb{A})$ is always B . Consequently, if \mathbb{B} is a valued σ -structure over B , then pointwise comparison is possible between $\omega(\mathbb{A})$ and \mathbb{B} .

DEFINITION 3.5. *Let \mathbb{A}, \mathbb{B} be valued σ -structures. A probability distribution ω over B^A is an IFH from \mathbb{A} to \mathbb{B} if $\omega(\mathbb{A}) \leq \mathbb{B}$.*

Example 3.14 on page 26 describes an explicit example of an IFH.

The following result relates improvement from Definition 3.1 and IFHs from Definition 3.5. The proof is based on Farkas' lemma.

PROPOSITION 3.6. *Let \mathbb{A}, \mathbb{B} be valued σ -structures. Then, $\mathbb{A} \preceq \mathbb{B}$ if and only if there exists an IFH from \mathbb{A} to \mathbb{B} .*

Remark 3.7. Let us remark that an IFH ω from \mathbb{A} to \mathbb{B} is actually a distribution over the set of homomorphisms from $\text{Pos}(\mathbb{A})$ to $\text{Pos}(\mathbb{B})$, i.e., every $g \in \text{supp}(\omega)$ is a homomorphism from $\text{Pos}(\mathbb{A})$ to $\text{Pos}(\mathbb{B})$. Indeed, for every $\mathbf{x} \in R_f^{\text{Pos}(\mathbb{A})}$, where $f \in \sigma$ and $\mathbf{x} \in A^{\text{ar}(f)}$, it must be the case that $f^{\mathbb{B}}(g(\mathbf{x})) \geq \sum_{h \in B^A} \omega(h) f^{\mathbb{A}}(h^{-1}(g(\mathbf{x}))) \geq \omega(g) f^{\mathbb{A}}(\mathbf{x}) > 0$. Hence, $g(\mathbf{x}) \in R_f^{\text{Pos}(\mathbb{B})}$. In view of Proposition 3.6, this offers another explanation of the fact that equivalence implies homomorphic equivalence (of the positive parts).

Let ω and ω' be IFHs from \mathbb{A} to \mathbb{B} and from \mathbb{B} to \mathbb{C} , respectively. We define $\omega' \circ \omega : C^A \rightarrow \mathbb{Q}_{\geq 0}$ as

$$\omega' \circ \omega(h) = \sum_{\substack{h_1: A \rightarrow B, h_2: B \rightarrow C \\ h_2 \circ h_1 = h}} \omega(h_1) \omega'(h_2).$$

Observe that $\omega' \circ \omega$ is an IFH from \mathbb{A} to \mathbb{C} .

3.2. Cores. Appropriate notions of cores have played an important role in the complexity classifications of left-hand-side restricted CSPs [32], right-hand-side restricted CSPs [10, 8, 52], and right-hand-side restricted VCSPs [50, 39]. In this paper, we will define cores around IFHs.

For two valued σ -structures \mathbb{A} and \mathbb{B} , we say that an IFH ω from \mathbb{A} to \mathbb{B} is *surjective* if every $g \in \text{supp}(\omega)$ is surjective.

DEFINITION 3.8. *A valued σ -structure \mathbb{A} is a core if every IFH from \mathbb{A} to \mathbb{A} is surjective.*

Next we show that equivalent valued structures that are cores are in fact isomorphic.

DEFINITION 3.9. *Let \mathbb{A}, \mathbb{B} be valued σ -structures. An isomorphism from \mathbb{A} to \mathbb{B} is a bijective mapping $h : A \rightarrow B$ such that $f^{\mathbb{A}}(\mathbf{x}) = f^{\mathbb{B}}(h(\mathbf{x}))$ for all $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$. If such an h exists, we say that \mathbb{A} and \mathbb{B} are isomorphic.*

PROPOSITION 3.10. *If \mathbb{A}, \mathbb{B} are core valued σ -structures such that $\mathbb{A} \equiv \mathbb{B}$, then \mathbb{A} and \mathbb{B} are isomorphic.*

We now introduce the central notion of a core of a valued structure and show that every valued structure has a unique core (up to isomorphism).

DEFINITION 3.11. *Let \mathbb{A}, \mathbb{B} be valued σ -structures. We say that \mathbb{B} is a core of \mathbb{A} if \mathbb{B} is a core and $\mathbb{A} \equiv \mathbb{B}$.*

PROPOSITION 3.12. *Every valued structure \mathbb{A} has a core and all cores of \mathbb{A} are isomorphic. Moreover, for a given valued structure \mathbb{A} , it is possible to effectively compute a core of \mathbb{A} and all cores of \mathbb{A} are over a universe of size at most $|A|$.*

Proposition 3.12 allows us to speak about *the* core of a valued structure. It follows then that equivalence can be characterized in terms of cores: \mathbb{A} and \mathbb{B} are equivalent if and only if their cores are isomorphic.

We conclude with a technical characterization of the property of being a core that will be important in the paper. Intuitively, it states that every nonsurjective mapping from a core \mathbb{A} to itself is suboptimal with respect to a fixed weighting of the tuples of \mathbb{A} .

PROPOSITION 3.13. *Let \mathbb{A} be a valued σ -structure. Then, \mathbb{A} is a core if and only if there exists a mapping $c : \text{tup}(\mathbb{A}) \rightarrow \mathbb{Q}_{\geq 0}$ such that for every nonsurjective mapping $g : A \rightarrow A$,*

$$\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) c(f, \mathbf{x}) < \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) c(f, g(\mathbf{x})).$$

Moreover, such a mapping $c : \text{tup}(\mathbb{A}) \rightarrow \mathbb{Q}_{\geq 0}$ is computable, whenever \mathbb{A} is a core.

Example 3.14. Let \mathbb{A} and \mathbb{A}' be the valued σ -structures depicted in Figure 1. Recall that $\sigma = \{f, \mu\}$, where f is a $\{0, \infty\}$ -valued binary function (represented by edges) and μ is unary (represented by node labels). Also, the elements of \mathbb{A} are denoted by a_{ij} , where i and j indicate the corresponding row and column of the grid, respectively (for readability, only a_{11} is depicted in Figure 1). We claim that \mathbb{A}' is the core of \mathbb{A} . Indeed, since $\text{Pos}(\mathbb{A}')$ is a relational core, it follows that \mathbb{A}' is a core. To see that $\mathbb{A} \preceq \mathbb{A}'$, let $g : A \rightarrow A'$ be the mapping that maps all elements in the i th diagonal of \mathbb{A} (first diagonal is $\{a_{11}\}$, second diagonal is $\{a_{21}, a_{12}\}$, and so on) to a'_i . Assigning $\omega(g) = 1$ gives an IFH from \mathbb{A} to \mathbb{A}' , and thus $\mathbb{A} \preceq \mathbb{A}'$ by Proposition 3.6.

Conversely, we will consider the six mappings $g_1, g_2, g_3, g_4, g_5, g_6$ from A' to A that map $(a'_1, a'_2, a'_3, a'_4, a'_5)$ to $(a_{11}, a_{21}, a_{31}, a_{32}, a_{33})$, $(a_{11}, a_{21}, a_{22}, a_{32}, a_{33})$, $(a_{11}, a_{12}, a_{22}, a_{23}, a_{33})$, $(a_{11}, a_{12}, a_{13}, a_{23}, a_{33})$, $(a_{11}, a_{12}, a_{22}, a_{32}, a_{33})$, and $(a_{11}, a_{21}, a_{22}, a_{23}, a_{33})$, respectively. We define the distribution $\omega'(g_1) = 1/3$, $\omega'(g_2) = 1/12$, $\omega'(g_3) = 1/12$, $\omega'(g_4) = 1/3$, $\omega'(g_5) = 1/12$, and $\omega'(g_6) = 1/12$. Then, we have

$$\sum_k \omega'(g_k) \mu^{\mathbb{A}'}(g_k^{-1}(a_{ij})) = \begin{cases} (2 \times 1/3 + 4 \times 1/12) \times 1 & \text{if } a_{ij} \in \{a_{11}, a_{33}\}, \\ (1/3 + 2 \times 1/12) \times 2 & \text{if } a_{ij} \in \{a_{12}, a_{21}, a_{23}, a_{32}\}, \\ (1/3) \times 3 & \text{if } a_{ij} \in \{a_{13}, a_{31}\}, \\ (4 \times 1/12) \times 3 & \text{if } a_{ij} = a_{22} \end{cases}$$

and hence for all i, j , $\sum_k \omega'(g_k) \mu^{\mathbb{A}'}(g_k^{-1}(a_{ij})) = 1 \leq \mu^{\mathbb{A}}(a_{ij})$. It follows that ω' is an IFH from \mathbb{A}' to \mathbb{A} . Therefore, $\mathbb{A}' \preceq \mathbb{A}$ and \mathbb{A}' is the core of \mathbb{A} . In particular, \mathbb{A} is not a core. As we explain later in Example 4.5, it is possible to modify \mathbb{A} (more precisely, $\mu^{\mathbb{A}}$) so that it becomes a core.

3.3. Treewidth modulo equivalence. In this section we show an elementary property of cores that is crucial for our purposes: the treewidth of the core of a valued structure \mathbb{A} is the lowest possible among all structures equivalent to \mathbb{A} .

PROPOSITION 3.15. *Let \mathbb{A} be a valued σ -structure and \mathbb{A}' be its core. Then, $\text{tw}(\mathbb{A}') \leq \text{tw}(\mathbb{A})$.*

Proof. Since treewidth is preserved under relational substructures, it suffices to show that $\text{Pos}(\mathbb{A}')$ is isomorphic to a substructure of $\text{Pos}(\mathbb{A})$, i.e., there is an injective homomorphism from $\text{Pos}(\mathbb{A}')$ to $\text{Pos}(\mathbb{A})$. Let ω and ω' be IFHs from \mathbb{A}' to \mathbb{A} and from \mathbb{A} to \mathbb{A}' , respectively. Pick any mapping $g \in \text{supp}(\omega)$. As observed at the end of section 3.1, g has to be a homomorphism from $\text{Pos}(\mathbb{A}')$ to $\text{Pos}(\mathbb{A})$. It suffices to show

that g is injective. But this follows immediately from the fact that $\omega' \circ \omega$ is an IFH from \mathbb{A}' to itself and the assumption that \mathbb{A}' is a core. \square

We conclude section 3 with the following useful characterization of “being equivalent to a bounded treewidth structure” in terms of cores.

PROPOSITION 3.16. *Let \mathbb{A} be a valued σ -structure and $k \geq 1$. Then, the following are equivalent:*

1. *There is a valued σ -structure \mathbb{A}' such that $\mathbb{A}' \equiv \mathbb{A}$ and \mathbb{A}' has treewidth at most k .*
2. *The treewidth of the core of \mathbb{A} is at most k .*

Proof. $2 \Rightarrow 1$ is immediate. For $1 \Rightarrow 2$, let \mathbb{A}' be of treewidth at most k such that $\mathbb{A}' \equiv \mathbb{A}$. Let \mathbb{B} and \mathbb{B}' be the cores of \mathbb{A} and \mathbb{A}' , respectively. Since $\mathbb{B} \equiv \mathbb{B}'$, \mathbb{B} and \mathbb{B}' are isomorphic, by Proposition 3.10. By Proposition 3.15, the treewidth of \mathbb{B}' is at most k , and so is the treewidth of \mathbb{B} . \square

4. Complexity of $\text{VCSP}(\mathcal{C}, -)$. Let \mathcal{C} be a class of valued structures. We say that \mathcal{C} has *bounded arity* if there is a constant $r \geq 1$ such that for every valued σ -structure $\mathbb{A} \in \mathcal{C}$ and $f \in \sigma$, we have that $\text{ar}(f) \leq r$. Similarly, we say that \mathcal{C} has *bounded treewidth modulo equivalence* if there is a constant $k \geq 1$ such that every $\mathbb{A} \in \mathcal{C}$ is equivalent to a valued structure \mathbb{A}' with $\text{tw}(\mathbb{A}') \leq k$. The following is our first main result.

THEOREM 4.1 (complexity classification). *Assume $\text{FPT} \neq \text{W}[1]$. Let \mathcal{C} be a recursively enumerable class of valued structures of bounded arity. Then, the following are equivalent:*

1. *$\text{VCSP}(\mathcal{C}, -)$ is in PTIME .*
2. *$p\text{-VCSP}(\mathcal{C}, -)$ is in FPT .*
3. *\mathcal{C} has bounded treewidth modulo equivalence.*

Remark 4.2. Although Grohe’s result [32] for CSPs looks almost identical to Theorem 4.1, we emphasize that his result involves a different type of equivalence. In Grohe’s case, the equivalence in question is *homomorphic equivalence*, whereas in our case the equivalence in question involves *improvement* (cf. Definition 3.2). As we will explain later in this section, Grohe’s classification follows as a special case of Theorem 4.1.

Remark 4.3. As in [32], we can remove the condition in Theorem 4.1 of \mathcal{C} being recursively enumerable by assuming a stronger hypothesis than $\text{FPT} \neq \text{W}[1]$ regarding nonuniform complexity classes.

Bounded treewidth implies bounded treewidth modulo equivalence but we will show in Example 4.4 that the converse is not true in general. Thus Theorem 4.1 gives new tractable cases compared to classes of VCSPs of (previously known to be tractable) bounded treewidth.

Bounded treewidth modulo equivalence implies bounded treewidth modulo homomorphic equivalence (of the positive parts) but we will show in Example 4.5 below that the converse is not true in general either. Therefore, Theorem 4.1 tells us that the tractability frontier for $\text{VCSP}(\mathcal{C}, -)$ lies *strictly* between bounded treewidth and bounded treewidth modulo homomorphic equivalence.

Example 4.4. We construct a class of structures of unbounded treewidth whose cores are of bounded treewidth (in fact treewidth 1).

By Proposition 3.16, a class \mathcal{C} has bounded treewidth modulo equivalence if and only if the class given by the *cores* of the valued structures in \mathcal{C} has bounded treewidth.

Thus we obtain a separation of bounded treewidth and bounded treewidth modulo equivalence.

Consider the signature $\sigma = \{f, \mu\}$, where f and μ are binary and unary function symbols, respectively. For $n \geq 1$, let \mathbb{A}_n be the valued σ -structure with universe $A_n = \{1, \dots, n\} \times \{1, \dots, n\}$ such that (i) $f^{\mathbb{A}_n}((i, j), (i', j')) = \infty$ if $i \leq i', j \leq j'$, and $(i' - i) + (j' - j) = 1$, otherwise $f^{\mathbb{A}_n}((i, j), (i', j')) = 0$, and (ii) $\mu^{\mathbb{A}_n}((i, j)) = 1$ for all $(i, j) \in A_n$. Also, for $n \geq 1$, let \mathbb{A}'_n be the valued σ -structure with universe $A'_n = \{1, \dots, 2n - 1\}$ such that (i) $f^{\mathbb{A}'_n}(i, j) = \infty$ if $j = i + 1$, otherwise $f^{\mathbb{A}'_n}(i, j) = 0$, and (ii) $\mu^{\mathbb{A}'_n}(i) = i$ for $1 \leq i \leq n$, and $\mu^{\mathbb{A}'_n}(i) = 2n - i$ for $n + 1 \leq i \leq 2n - 1$. The structures \mathbb{A} and \mathbb{A}' from Figure 1 correspond to \mathbb{A}_3 and \mathbb{A}'_3 , respectively; informally \mathbb{A}_n is a crisp directed grid of size $n \times n$ with a unary function μ with weight 1 applied to each element. Generalizing the reasoning behind Example 3.14, we argue in Appendix B that for each $n \geq 1$ the valued structure \mathbb{A}'_n is the core of \mathbb{A}_n . Since $tw(\mathbb{A}'_n) = 1$, the class $\mathcal{C} := \{\mathbb{A}_n \mid n \geq 1\}$ has bounded treewidth modulo equivalence. However, \mathcal{C} has unbounded treewidth as the Gaifman graphs in $\{G(\text{Pos}(\mathbb{A}_n)) \mid n \geq 1\}$ correspond to the class of (undirected) *grids*, which is a well-known family of graphs with unbounded treewidth (see, e.g., [18]). We also describe in Appendix B how to alter the definition of \mathcal{C} to obtain a class of *finite-valued* structures (taking on finite values in $\mathbb{Q}_{\geq 0}$) that has bounded treewidth modulo equivalence but Gaifman graphs of unbounded treewidth.

Example 4.5. We construct a class of structures of unbounded treewidth modulo equivalence but bounded treewidth modulo homomorphic equivalence, thus separating bounded treewidth modulo equivalence from bounded treewidth modulo homomorphic equivalence.

For $n \geq 3$, let \mathbb{A}_n be the valued σ -structure from Example 4.4. Let \mathbb{C}_n be the valued σ -structure with the same universe as \mathbb{A}_n , i.e., $C_n = \{1, \dots, n\} \times \{1, \dots, n\}$, such that $f^{\mathbb{C}_n} = f^{\mathbb{A}_n}$ and $\mu^{\mathbb{C}_n}$ is defined as follows. Let D_1, \dots, D_n be the n first diagonals of C_n starting from the bottom left corner $(1, 1)$ (see Figure 2 for an illustration of \mathbb{C}_4). For $1 \leq i \leq n$, let E_i be the top-left to bottom-right enumeration of D_i . In particular, $E_1 = ((1, 1))$, $E_2 = ((2, 1), (1, 2))$, $E_3 = ((3, 1), (2, 2), (1, 3))$, and $E_n = ((n, 1), (n - 1, 2), \dots, (1, n))$. Fix an integer $M = M(n)$ such that $M > n^2$. The values assigned by $\mu^{\mathbb{C}_n}$ to E_1 , E_2 , and E_3 are (1) , $(M, 1)$, and (M^3, M^2, M^4) , respectively, and for E_i , with $4 \leq i \leq n$, is $(M, 1, \dots, 1, M)$. All remaining elements in $C_n \setminus \bigcup_{1 \leq i \leq n} D_i$ receive cost 1. Figure 2 depicts the case of \mathbb{C}_4 .




FIG. 2. The valued σ -structure \mathbb{C}_4 from Example 4.5 ($M > 16$).

Let $\mathcal{C} := \{\mathbb{C}_n \mid n \geq 3\}$. Note first that $\text{Pos}(\mathbb{C}_n)$ is homomorphically equivalent to the relational structure \mathbf{P}_{2n-1} over relational signature $\text{rel}(\sigma) = \{R_f, R_\mu\}$ (recall the definition of $\text{rel}(\sigma)$ from section 2), whose universe is $P_{2n-1} = \{1, \dots, 2n-1\}$, $R_\mu^{\mathbf{P}_{2n-1}} = P_{2n-1}$, and $R_f^{\mathbf{P}_{2n-1}} = \{(i, i+1) \mid 1 \leq i \leq 2n-2\}$. Since $\text{tw}(\mathbf{P}_{2n-1}) = 1$, for all $n \geq 3$, it follows that $\{\text{Pos}(\mathbb{C}_n) \mid n \geq 3\}$ has bounded treewidth modulo homomorphic equivalence. We claim that \mathcal{C} has unbounded treewidth modulo (valued) equivalence. It suffices to show that \mathbb{C}_n is a core for all $n \geq 3$. In order to prove this, we apply Proposition 3.13. Fix $n \geq 3$ and define $c : \text{tup}(\mathbb{C}_n) \rightarrow \mathbb{Q}_{\geq 0}$ such that (i) $c(f, \mathbf{x}) = 0$ if $f^{\mathbb{C}_n}(\mathbf{x}) = \infty$, otherwise $c(f, \mathbf{x}) = 1$, and (ii) $c(\mu, x) = 1/\mu^{\mathbb{C}_n}(x)$ for all $x \in C_n$. We have $\sum_{(p, \mathbf{y}) \in \text{tup}(\mathbb{C}_n)} p^{\mathbb{C}_n}(\mathbf{y}) c(p, \mathbf{y}) = |\mathbb{C}_n| = n^2$. Next we show that if $g : C_n \rightarrow C_n$ satisfies that $v(g) = \sum_{(p, \mathbf{y}) \in \text{tup}(\mathbb{C}_n)} p^{\mathbb{C}_n}(\mathbf{y}) c(p, g(\mathbf{y})) \leq n^2$, then g is the identity mapping. Using Proposition 3.13, this implies that \mathbb{C}_n is a core.

Let $g : C_n \rightarrow C_n$ such that $v(g) \leq n^2$. The mapping g must satisfy the following two conditions: (a) g is a homomorphism from $\text{Pos}(\mathbb{C}_n)$ to $\text{Pos}(\mathbb{C}_n)$ (otherwise $v(g) = \infty$), and (b) for every $x \in C_n$, $\mu^{\mathbb{C}_n}(x) \leq \mu^{\mathbb{C}_n}(g(x))$, otherwise $v(g) \geq \mu^{\mathbb{C}_n}(x) c(\mu, g(x)) = \mu^{\mathbb{C}_n}(x) / \mu^{\mathbb{C}_n}(g(x)) \geq M > n^2$. We can argue inductively, and show that g is the identity over D_i , for all $1 \leq i \leq n$. Condition (a) implies that g is the identity over the remaining elements in $C_n \setminus \bigcup_{1 \leq i \leq n} D_i$, as required. For D_1 , we have that $g((1, 1)) = (1, 1)$ by condition (a). For D_2 , note that (a) implies that $\{g((2, 1)), g((1, 2))\} \subseteq \{(2, 1), (1, 2)\}$. By condition (b), $g((2, 1)) = (2, 1)$. To see that $g((1, 2)) = (1, 2)$, suppose by contradiction that $g((1, 2)) = (2, 1)$; then condition (a) implies that $g((1, 3)) \in \{(3, 1), (2, 2)\}$, which violates (b). For the case $3 \leq i \leq n$, recall that $E_i = x_1, x_2, \dots, x_{|D_i|}$ is the above-defined enumeration of D_i . Since g is the identity over D_{i-1} and by condition (a), we have that g is the identity over $\{x_2, \dots, x_{|D_i|-1}\}$. As $\mu^{\mathbb{C}_n}(x_1) > \mu^{\mathbb{C}_n}(x_2)$ and $\mu^{\mathbb{C}_n}(x_{|D_i|}) > \mu^{\mathbb{C}_n}(x_{|D_i|-1})$, conditions (a) and (b) imply that $g(x_1) = x_1$ and $g(x_{|D_i|}) = x_{|D_i|}$, as required.

To conclude this example we note that the class of valued σ -structures $\{\mathbb{B}_n \mid n \geq 3\}$ where each \mathbb{B}_n is derived from \mathbb{C}_n by setting $B_n = C_n$, $\mu^{\mathbb{B}_n} = \mu^{\mathbb{C}_n}$ and $f^{\mathbb{B}_n}(\mathbf{x}) = \min(1, f^{\mathbb{C}_n}(\mathbf{x}))$ for all $\mathbf{x} \in (B_n)^2$ is an example of a *finite-valued* class of structures that has bounded treewidth modulo homomorphic equivalence but unbounded treewidth modulo valued equivalence, since each \mathbb{B}_n is a core (this follows from exactly the same argument used for \mathbb{C}_n).

Corollaries of the complexity classification. We can obtain the classification for CSPs of Dalmau, Kolaitis, and Vardi [16] and Grohe [32] as a special case of Theorem 4.1. Indeed, we can associate with a relational τ -structure \mathbf{A} a valued σ_τ -structure $\mathbb{A}_{0,\infty}$ such that (i) $\sigma_\tau = \{f_R \mid R \in \tau, \text{ar}(f_R) = \text{ar}(R)\}$, (ii) \mathbf{A} and $\mathbb{A}_{0,\infty}$ have the same universe A , and (iii) if $\mathbf{x} \in R^{\mathbf{A}}$, then $f_R^{\mathbb{A}_{0,\infty}}(\mathbf{x}) = \infty$, otherwise $f_R^{\mathbb{A}_{0,\infty}}(\mathbf{x}) = 0$, for every $R \in \tau$ and $\mathbf{x} \in A^{\text{ar}(R)}$. For a class \mathcal{C} of relational structures, we define the class of valued structures $\mathcal{C}_{0,\infty} := \{\mathbb{A}_{0,\infty} \mid \mathbf{A} \in \mathcal{C}\}$. It is not hard to check that, when \mathcal{C} is of bounded arity, $\text{CSP}(\mathcal{C}, -)$ reduces in polynomial time to $\text{VCSP}(\mathcal{C}_{0,\infty}, -)$ and vice versa. Hence, a classification of $\text{CSP}(\mathcal{C}, -)$, for \mathcal{C} 's of bounded arity, is equivalent to a classification of $\text{VCSP}(\mathcal{C}_{0,\infty}, -)$. Finally, note that \mathcal{C} has bounded treewidth modulo homomorphic equivalence if and only if $\mathcal{C}_{0,\infty}$ has bounded treewidth modulo (valued) equivalence. This implies the known CSP classification from [16] and [32].

In his Ph.D. thesis [22], Färnqvist also considered the complexity of $\text{VCSP}(\mathcal{C}, -)$. However, he considered a different definition of the problem, that we denote by $\text{VCSP}_F(\mathcal{C}, -)$. Formally, for a relational τ -structure \mathbf{A} , let \mathbb{A}_F be the valued $\sigma_{\mathbf{A}}$ -structure such that (i) $\sigma_{\mathbf{A}} = \{f_{R,\mathbf{x}} \mid R \in \tau, \mathbf{x} \in R^{\mathbf{A}}, \text{ar}(f_{R,\mathbf{x}}) = \text{ar}(R)\}$, (ii) \mathbf{A} and \mathbb{A}_F have the same universe A , and (iii) for every $f_{R,\mathbf{x}} \in \sigma_{\mathbf{A}}$ and $\mathbf{x} \in A^{\text{ar}(f_{R,\mathbf{x}})}$, we have

that $f_{R,\mathbf{x}}^{\mathbb{A}_F}(\mathbf{x}) = 1$ and $f_{R,\mathbf{x}}^{\mathbb{A}_F}(\mathbf{y}) = 0$, for all $\mathbf{y} \neq \mathbf{x}$. For a class of relational structures \mathcal{C} , $\text{VCSP}_F(\mathcal{C}, -)$ is precisely the problem $\text{VCSP}(\mathcal{C}_F, -)$, where $\mathcal{C}_F := \{\mathbb{A}_F \mid \mathbb{A} \in \mathcal{C}\}$. It was shown in [22] that for a class \mathcal{C} of relational structures of bounded arity, $\text{VCSP}_F(\mathcal{C}, -)$ is tractable if and only if \mathcal{C} has bounded treewidth. This result follows directly from Theorem 4.1 as every valued structure in a class of the form \mathcal{C}_F is a (valued) core, and hence, bounded treewidth modulo equivalence boils down to bounded treewidth.

Intuitively, $\text{VCSP}_F(\mathcal{C}, -)$ restricts VCSP only based on the (multiset of) tuples appearing in the structures from \mathcal{C} . In contrast, our definition of $\text{VCSP}(\mathcal{C}, -)$ considers directly the structures in \mathcal{C} . This allows us for a more fine-grained analysis of structural restrictions and, in particular, provides us with new tractable classes beyond bounded treewidth. Indeed, as Example 4.4 illustrates, we can find simple tractable classes of valued structures with unbounded treewidth.

Finally, let us note that since Theorem 4.1 applies to *all* valued structures, it in particular covers the *finite-valued* VCSP, where all functions are restricted to take finite values in $\mathbb{Q}_{\geq 0}$, and hence the tractability part of Theorem 4.1 directly applies to the finite-valued case. The hardness part also applies to the finite-valued case. Indeed, the right-hand-side structure \mathbb{B} constructed in the reduction of Proposition 4.6 can be made to be finite-valued as explained at the end of the proof. Therefore, Theorem 4.1 also gives a classification for finite-valued VCSPs. Moreover, Examples 4.4 and 4.5 demonstrate that already for finite-valued structures the tractability frontier is strictly between bounded treewidth and bounded treewidth modulo homomorphic equivalence.

The rest of this section is devoted to proving the hardness part of Theorem 4.1, i.e., the implication (2) \Rightarrow (3). The tractability part of Theorem 4.1 (implication (3) \Rightarrow (1)) is established in section 5. In particular, it will follow from Theorem 5.4 that, if there is a constant $k \geq 1$ such that every valued structure in the class \mathcal{C} is equivalent to a valued structure of treewidth at most k , then $\text{VCSP}(\mathcal{C}, -)$ can be solved in polynomial time using the $(k+1)$ st level of the Sherali–Adams LP hierarchy. The remaining implication (1) \Rightarrow (2) is immediate.

4.1. Hardness. We start with the notion of fpt-reductions [25] tailored to our setting. Formally, a decision problem \mathcal{P} with parameter κ over Σ is a subset of Σ^* , the set of all strings over the alphabet Σ , describing the “yes” instances of \mathcal{P} , and $\kappa : \Sigma^* \rightarrow \mathbb{N}$. It is known that each optimization problem has an equivalent decision problem. We denote by $p\text{-VCSP}_d(\mathcal{C}, -)$ the decision version of $p\text{-VCSP}(\mathcal{C}, -)$. Formally, $p\text{-VCSP}_d(\mathcal{C}, -)$ with parameter κ' over Σ' is a subset of $(\Sigma')^*$ such that $x = ((\mathbb{A}, \mathbb{B}), c) \in p\text{-VCSP}_d(\mathcal{C}, -)$ if and only if (\mathbb{A}, \mathbb{B}) is a $\text{VCSP}(\mathcal{C}, -)$ instance such that $\text{opt}(\mathbb{A}, \mathbb{B}) \leq c$, and $\kappa' : (\Sigma')^* \rightarrow \mathbb{N}$ is defined by $\kappa'(x) = |\mathbb{A}|$.

An fpt-reduction from (\mathcal{P}, κ) to $p\text{-VCSP}_d(\mathcal{C}, -)$ is a mapping $\text{red} : \Sigma^* \rightarrow (\Sigma')^*$ such that (i) for all $x \in \Sigma^*$ we have $x \in \mathcal{P}$ if and only if $\text{red}(x) \in p\text{-VCSP}_d(\mathcal{C}, -)$; (ii) there is a computable function $f : \mathbb{N} \rightarrow \mathbb{N}$ and an algorithm that, given $x \in \Sigma^*$, computes $\text{red}(x)$ in time $f(\kappa(x)) \cdot |x|^{O(1)}$; and (iii) there is a computable function $g : \mathbb{N} \rightarrow \mathbb{N}$ such that for all instances $x \in \Sigma^*$, we have $\kappa'(\text{red}(x)) \leq g(\kappa(x))$.

Let us mention that our hardness result does not follow directly from Grohe’s result for CSPs [32]. The natural approach is to define, for a class of valued structures \mathcal{C} , the class of relational structures $\text{Pos}(\mathcal{C}) = \{\text{Pos}(\mathbb{A}) \mid \mathbb{A} \in \mathcal{C}\}$. Then one can observe that $p\text{-CSP}(\text{Pos}(\mathcal{C}), -)$ fpt-reduces to $p\text{-VCSP}_d(\mathcal{C}, -)$, and hence W[1]-hardness of the former problem implies hardness for the latter. However, if \mathcal{C} has unbounded treewidth modulo equivalence, the class $\text{Pos}(\mathcal{C})$ does not necessarily have unbounded treewidth modulo homomorphic equivalence (see Example 4.5), and hence $\text{Pos}(\mathcal{C})$ is not necessarily hard according to Grohe’s classification.

We instead adapt Grohe's proof to the case of VCSPs. We need some notation. For $k \geq 1$, a k -clique of a graph is a clique on k vertices. A graph H is a *minor* of a graph G if H is isomorphic to a graph that can be obtained from a subgraph of G by contracting edges (for more details see, e.g., [18]). For $k, \ell \geq 1$, the $(k \times \ell)$ -grid is the graph with vertex set $\{1, \dots, k\} \times \{1, \dots, \ell\}$ and an edge between (i, j) and (i', j') if $|i - i'| + |j - j'| = 1$. The parameterized problem p -CLIQUE asks, given instance (G, k) , whether there is a k -clique in G , and has parameter κ such that $\kappa(G, k) = k$. It is a well-known result that p -CLIQUE is complete for $W[1]$ under fpt-reductions [20]. The implication (2) \Rightarrow (3) in Theorem 4.1 follows from the following proposition.

PROPOSITION 4.6. *Let \mathcal{C} be a recursively enumerable class of valued structures of bounded arity. Suppose \mathcal{C} is of unbounded treewidth modulo equivalence. If p -VCSP($\mathcal{C}, -$) is fixed-parameter tractable then $FPT = W[1]$.*

Proof. We present an fpt-reduction from p -CLIQUE to p -VCSP $_d(\mathcal{C}, -)$. More precisely, given an instance (G, k) of p -CLIQUE, we shall construct valued structures $\mathbb{A}' \in \mathcal{C}$ and \mathbb{B} , together with a threshold $M^* \geq 0$, such that G contains a k -clique if and only if $\text{opt}(\mathbb{A}', \mathbb{B}) \leq M^*$. As in [32], we rely on the excluded grid theorem [45], which states that there is a function $w : \mathbb{N} \rightarrow \mathbb{N}$ such that every graph H of treewidth at least $w(k)$ contains the $(k \times k)$ -grid as a minor. Given an instance (G, k) of p -CLIQUE, with $k \geq 2$, we start by enumerating the class \mathcal{C} until we obtain a valued structure $\mathbb{A}' \in \mathcal{C}$ with core \mathbb{A} such that $\text{tw}(\mathbb{A}) \geq w(K)$, where $K = \binom{k}{2}$. By Proposition 3.16, such $\mathbb{A}' \in \mathcal{C}$ always exists. Let $\mathbf{A} := \text{Pos}(\mathbb{A})$. By the excluded grid theorem, the Gaifman graph $G(\mathbf{A})$ of \mathbf{A} (see the definition of the Gaifman graph in section 2) contains the $(K \times K)$ -grid and, hence, the $(k \times K)$ -grid as a minor. Since \mathcal{C} is recursively enumerable and cores are computable (by Proposition 3.12), the valued structure \mathbb{A}' and its core \mathbb{A} can be effectively computed in time $\alpha(k)$, where α is a computable function.

In order to define \mathbb{B} , we exploit the main construction in [32], which defines a relational structure \mathbf{B} from G, k , and \mathbf{A} . The key property of \mathbf{B} is that, assuming \mathbf{A} is a relational core, then G contains a k -clique if and only if there is a homomorphism from \mathbf{A} to \mathbf{B} . Since in our case \mathbf{A} is not necessarily a relational core, we restate in the following lemma the properties of \mathbf{B} simply in terms of surjective homomorphisms. Together with our characterization of (valued) cores in Proposition 3.13, this will allow us to define our required valued structure \mathbb{B} and threshold M^* .

LEMMA 4.7. *Given $k \geq 2$, $K = \binom{k}{2}$, graph G , and relational τ -structure \mathbf{A} such that the $(k \times K)$ -grid is a minor of its Gaifman graph $G(\mathbf{A})$, there is a relational τ -structure \mathbf{B} such that*

1. *there exists a fixed homomorphism π from \mathbf{B} to \mathbf{A} such that the following are equivalent:*
 - (a) *G contains a k -clique,*
 - (b) *there is a homomorphism h from \mathbf{A} to \mathbf{B} such that $\pi \circ h$ is the identity mapping from A to itself, where A is the universe of \mathbf{A} ;*
2. *\mathbf{B} can be computed in time $\beta(|\mathbf{A}|, k) \cdot |G|^{O(r(\mathbf{A}))}$, where β is a computable function and $r(\mathbf{A}) \geq 1$ is the arity of the relational signature τ .*

Proof. We can separate \mathbf{A} into two relational τ -structures \mathbf{A}' and $\mathbf{A} \setminus \mathbf{A}'$ such that the disjoint union of their Gaifman graphs $G(\mathbf{A}')$ and $G(\mathbf{A} \setminus \mathbf{A}')$ is precisely $G(\mathbf{A})$, and $G(\mathbf{A}')$ is a connected graph containing the $(k \times K)$ -grid as a minor. The main hardness reduction in [32] produces for the connected structure \mathbf{A}' a relational τ -structure \mathbf{B}' and a homomorphism π' from \mathbf{B}' to \mathbf{A}' such that conditions (1) and (2) hold with respect to \mathbf{A}' , \mathbf{B}' , and π' . (This is shown in section 4 from [32] as a

preparation for the proof of Theorem 4.1.) We define our required structure \mathbf{B} to be the disjoint union of \mathbf{B}' and $\mathbf{A} \setminus \mathbf{A}'$. We define π to be the homomorphism from \mathbf{B} to \mathbf{A} such that $\pi(b) = \pi'(b)$, if b belongs to \mathbf{B}' and $\pi(b) = b$, otherwise. It suffices to check condition (1).

Suppose that G contains a k -clique. Then there is a homomorphism h' from \mathbf{A}' to \mathbf{B}' such that $\pi' \circ h' = \text{id}$. We can define the mapping h from \mathbf{A} to \mathbf{B} such that $h(a) = h'(a)$, if a belongs to \mathbf{A}' , and $h(a) = a$, otherwise. It follows that h is a homomorphism and $\pi \circ h = \text{id}$. Assume now that there is a homomorphism h from \mathbf{A} to \mathbf{B} such that $\pi \circ h = \text{id}$. Then there exists a homomorphism h' from \mathbf{A}' to \mathbf{B}' such that $\pi' \circ h' = \text{id}$ (h' is simply the restriction of h to \mathbf{A}'). Since \mathbf{A}' and \mathbf{B}' satisfy condition (1), we conclude that G contains a k -clique. \square

Let \mathbf{B} be the relational structure from Lemma 4.7 applied to G, k , and $\mathbf{A} = \text{Pos}(\mathbb{A})$. Recall that \mathbf{A} , and hence \mathbf{B} , is defined over the relational signature $\text{rel}(\sigma)$, where σ is the signature of \mathbb{A} . Since \mathbb{A} is a core, by Proposition 3.13 we can compute a function $c^* : \text{tup}(\mathbb{A}) \rightarrow \mathbb{Q}_{\geq 0}$ such that for every nonsurjective mapping $g : A \rightarrow A$, it is the case that

$$\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) c^*(f, \mathbf{x}) < \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) c^*(f, g(\mathbf{x})).$$

From \mathbf{B} and c^* we construct a valued structure \mathbb{B} over the same signature σ of \mathbb{A} as follows. Let π be the homomorphism from \mathbf{B} to \mathbf{A} given by Lemma 4.7. Let $M^* := \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) c^*(f, \mathbf{x}) < \infty$. The universe of \mathbb{B} is the universe B of \mathbf{B} . For each $f \in \sigma$ and $\mathbf{x} \in B^{\text{ar}(f)}$, we define

$$f^{\mathbb{B}}(\mathbf{x}) = \begin{cases} c^*(f, \pi(\mathbf{x})) & \text{if } \mathbf{x} \in R_f^{\mathbf{B}}, \\ \infty & \text{otherwise.} \end{cases}$$

We show that G contains a k -clique if and only if $\text{opt}(\mathbb{A}', \mathbb{B}) \leq M^*$. Before doing so, note that the total running time of the reduction is $\alpha(k) + \beta(\alpha(k), k) |G|^{O(r(\mathbf{A}))}$, where β is from Lemma 4.7 and $r(\mathbf{A})$ is the arity of $\text{rel}(\sigma)$ and hence the arity of σ . Since the class \mathcal{C} has bounded arity, there is a constant $r \geq 1$ such that $r(\mathbf{A}) \leq r$. Thus the running time of the reduction is $\beta'(k) |G|^{O(1)}$ for a computable function β' . Also, since \mathbb{A}' is computed in time $\alpha(k)$, we have that $|\mathbb{A}'| \leq \alpha(k)$. It follows that our reduction is actually an fpt-reduction, and hence, $p\text{-VCSP}_d(\mathcal{C}, -)$ is W[1]-hard. If $p\text{-VCSP}(\mathcal{C}, -)$ is in FPT, then $p\text{-VCSP}_d(\mathcal{C}, -)$ is in FPT, and consequently $\text{FPT} = \text{W}[1]$, as required.

Assume that G contains a k -clique. By Lemma 4.7, there is a homomorphism h from \mathbf{A} to \mathbf{B} such that $\pi \circ h$ is the identity. We have that

$$\begin{aligned} \text{opt}(\mathbb{A}', \mathbb{B}) &= \text{opt}(\mathbb{A}, \mathbb{B}) \leq \text{cost}_{\mathbb{A} \rightarrow \mathbb{B}}(h) \\ &= \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{>0}} f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{B}}(h(\mathbf{x})) = \sum_{\mathbf{x} \in R_f^{\mathbf{A}}} f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{B}}(h(\mathbf{x})) \\ &= \sum_{\mathbf{x} \in R_f^{\mathbf{A}}} f^{\mathbb{A}}(\mathbf{x}) c^*(f, \pi(h(\mathbf{x}))) \\ &= \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{>0}} f^{\mathbb{A}}(\mathbf{x}) c^*(f, \mathbf{x}) = M^*. \end{aligned}$$

Suppose now that $\text{opt}(\mathbb{A}', \mathbb{B}) \leq M^*$. In particular, $\text{opt}(\mathbb{A}, \mathbb{B}) \leq M^*$. Let $h^* : A \rightarrow B$ be a mapping with cost $\text{opt}(\mathbb{A}, \mathbb{B})$. Note that h^* is a homomorphism from \mathbf{A} to \mathbf{B}

(otherwise $\text{cost}_{\mathbb{A} \rightarrow \mathbb{B}}(h^*) = \infty$). We show that $\pi \circ h^*$ is surjective. Assume to the contrary. By the definition of c^* , it follows that $M^* < \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) c^*(f, \pi(h^*(\mathbf{x})))$. On the other hand, using the fact that h^* is a homomorphism, we have that

$$\begin{aligned} \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) c^*(f, \pi(h^*(\mathbf{x}))) &= \sum_{\mathbf{x} \in R_f^{\mathbb{A}}} f^{\mathbb{A}}(\mathbf{x}) c^*(f, \pi(h^*(\mathbf{x}))) \\ &= \sum_{\mathbf{x} \in R_f^{\mathbb{A}}} f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{B}}(h^*(\mathbf{x})) \\ &= \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{>0}} f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{B}}(h^*(\mathbf{x})) = \text{cost}(h^*). \end{aligned}$$

Hence, $M^* < \text{cost}(h^*)$, a contradiction. It follows that $\pi \circ h^*$ is an isomorphism of \mathbf{A} . We can define $g = h^* \circ (\pi \circ h^*)^{-1}$ and obtain a homomorphism g from \mathbf{A} to \mathbf{B} such that $\pi \circ g = \text{id}$. By Lemma 4.7 we obtain that G contains a k -clique.

Finally, note that we can make \mathbb{B} in the reduction to be finite-valued. We only need to replace ∞ by a sufficiently large N in its definition. More precisely, we need $\text{cost}_{\mathbb{A} \rightarrow \mathbb{B}}(h) > M^*$, whenever $h : A \rightarrow B$ is not a homomorphism from \mathbf{A} to \mathbf{B} , i.e., we need $f^{\mathbb{A}}(\mathbf{x}) \cdot N > M^*$ for any $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{>0}$. (We can take $N = 1 + (M^* / \min\{f^{\mathbb{A}}(\mathbf{x}) : (f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{>0}\})$.) \square

5. Power of Sherali–Adams for $\text{VCSP}(\mathcal{C}, -)$. In this section we will present and prove our second main result, Theorem 5.4. First, in section 5.1, we will define the Sherali–Adams LP relaxation for VCSPs and state Theorem 5.4. Second, in section 5.2, we will define the key concept of *treewidth modulo scopes*, which essentially captures the applicability of the Sherali–Adams for $\text{VCSP}(\mathcal{C}, -)$. Section 5.3 proves the sufficiency part of Theorem 5.4, whereas sections 5.4 and 5.5 prove the necessity part of Theorem 5.4.

5.1. Sherali–Adams LP relaxations. Before we define the Sherali–Adams LP relaxation formally, we give an informal explanation that might be helpful to the reader who has not seen this before. Any VCSP instance (\mathbb{A}, \mathbb{B}) over signature σ has a natural integer linear programming formulation (ILP) with 0/1-variables. The variables of the ILP are pairs $(f^{\mathbb{A}}, \mathbf{x})$, where $f^{\mathbb{A}}$ is a “constraint” and $\mathbf{x} \in A^{\text{ar}(f)}$ is a “local assignment to the variables on which the constraint $f^{\mathbb{A}}$ depends,” denoted by s in Figure 3. The ILP has several constraints. First, local assignments with infinite costs are forbidden by setting the corresponding ILP variables to 0. Second, the integrality constraints (i.e., having 0/1 variables) ensure that for every constraint exactly one assignment is selected. Finally, “marginalization constraints” ensure that the local assignments are consistent in the sense that if $x \in A$ is mapped by a solution to $b \in B$ then it is the same b independently of the constraint that depends on x . The objective function of the ILP is the cost of an assignment given by the ILP variables.

By relaxing the integrality constraints, one obtains the so-called basic LP relaxation. One can think of the LP variables, for every constraint, as a probability distribution over the set of local assignments. Looking at Figure 3, conditions (SA2) and (SA4) ensure that the variables are indeed probability distributions. Conditions (SA3) correspond to the feasibility part (the underlying CSP) and rule out locally infeasible assignments. Conditions (SA1) ensure consistency between pairs of constraints, called “marginalization constraints” above. Finally, the objective function of the LP is the expected cost of an assignment given by the LP variables with respect to the local probability distributions.

$$\begin{aligned}
& \min \sum_{\substack{(f, \mathbf{x}) \in \text{tup}(\mathbb{A}_k)_{>0}, s: \text{Set}(\mathbf{x}) \rightarrow B_k \\ f^{\mathbb{A}_k}(\mathbf{x}) \times f^{\mathbb{B}_k}(s(\mathbf{x})) < \infty}} \lambda(f, \mathbf{x}, s) f^{\mathbb{A}_k}(\mathbf{x}) f^{\mathbb{B}_k}(s(\mathbf{x})), \\
& \lambda(f, \mathbf{x}, s) = \sum_{r: \text{Set}(\mathbf{y}) \rightarrow B_k, r|_{\text{Set}(\mathbf{x})} = s} \lambda(p, \mathbf{y}, r) \quad \text{(SA1)} \quad \begin{array}{l} \forall (f, \mathbf{x}), (p, \mathbf{y}) \in \text{tup}(\mathbb{A}_k)_{>0} : \\ \text{Set}(\mathbf{x}) \subseteq \text{Set}(\mathbf{y}) \text{ and } |\text{Set}(\mathbf{x})| \leq k; \\ \forall s: \text{Set}(\mathbf{x}) \rightarrow B_k \end{array} \\
& \sum_{s: \text{Set}(\mathbf{x}) \rightarrow B_k} \lambda(f, \mathbf{x}, s) = 1 \quad \text{(SA2)} \quad \forall (f, \mathbf{x}) \in \text{tup}(\mathbb{A}_k)_{>0} \\
& \lambda(f, \mathbf{x}, s) = 0 \quad \text{(SA3)} \quad \forall (f, \mathbf{x}) \in \text{tup}(\mathbb{A}_k)_{>0}, s: \text{Set}(\mathbf{x}) \rightarrow B_k : f^{\mathbb{A}_k}(\mathbf{x}) \times f^{\mathbb{B}_k}(s(\mathbf{x})) = \infty \\
& \lambda(f, \mathbf{x}, s) \geq 0 \quad \text{(SA4)} \quad \forall (f, \mathbf{x}) \in \text{tup}(\mathbb{A}_k)_{>0}, s: \text{Set}(\mathbf{x}) \rightarrow B_k
\end{aligned}$$

FIG. 3. The Sherali–Adams relaxation $SA_k(\mathbb{A}, \mathbb{B})$ of level k of (\mathbb{A}, \mathbb{B}) .

The idea of the k th level of the Sherali–Adams relaxation is to strengthen the LP by introducing constraints on all sets of variables of size up to k and imposing (local) consistency among these constraints. A further (notational) complication in the definition is the possible repetition of a variable in a constraint scope, which is dealt with using the $\text{Set}()$ operator, which we now define.

Given a tuple \mathbf{x} , we write $\text{Set}(\mathbf{x})$ to denote the set of elements appearing in \mathbf{x} . Let (\mathbb{A}, \mathbb{B}) be an instance of the VCSP over a signature σ and $k \geq 1$. We define a new signature $\sigma_k = \sigma \cup \{\rho_k\}$, where ρ_k is a new function symbol of arity k . Then, we create from (\mathbb{A}, \mathbb{B}) an instance $(\mathbb{A}_k, \mathbb{B}_k)$ over σ_k such that $A_k = A$, $B_k = B$, $\rho_k^{\mathbb{A}_k}(\mathbf{x}) = 1$ for any $\mathbf{x} \in A_k^k$, $\rho_k^{\mathbb{B}_k}(\mathbf{x}) = 0$ for any $\mathbf{x} \in B_k^k$, and for every $f \in \sigma$ we have $f^{\mathbb{A}_k} = f^{\mathbb{A}}$ and $f^{\mathbb{B}_k} = f^{\mathbb{B}}$. Because the new function ρ_k is identically zero in \mathbb{B}_k , we have that for any mapping $h: A \rightarrow B$, $\text{cost}(h)$ is the same in both instances (\mathbb{A}, \mathbb{B}) and $(\mathbb{A}_k, \mathbb{B}_k)$. The Sherali–Adams relaxation of level k [48] of (\mathbb{A}, \mathbb{B}) , denoted by $SA_k(\mathbb{A}, \mathbb{B})$, is the linear program given in Figure 3, which has one variable $\lambda(f, \mathbf{x}, s)$ for each $(f, \mathbf{x}) \in \text{tup}(\mathbb{A}_k)_{>0}$ and $s: \text{Set}(\mathbf{x}) \rightarrow B_k$.

The variables are indexed not only by \mathbf{x} and s but also by f . This would not be necessary if $k \geq \text{ar}(f)$ but we are also interested in the case of $k < \text{ar}(f)$.

DEFINITION 5.1. Let \mathbb{A}, \mathbb{B} be valued σ -structures and $k \geq 1$.

- We denote by $\text{opt}_k^{SA}(\mathbb{A}, \mathbb{B})$ the minimum cost of a solution to $SA_k(\mathbb{A}, \mathbb{B})$.
- We write $\mathbb{A} \preceq_k^{SA} \mathbb{B}$ if $\text{opt}_k^{SA}(\mathbb{A}, \mathbb{C}) \leq \text{opt}_k^{SA}(\mathbb{B}, \mathbb{C})$ for all valued σ -structures \mathbb{C} .

The proof of the following can be found in Appendix C.

PROPOSITION 5.2. Let \mathbb{A}, \mathbb{B} be valued σ -structures and $k \geq 1$. If there exists an IFH from \mathbb{A} to \mathbb{B} , then $\mathbb{A} \preceq_k^{SA} \mathbb{B}$.

We also have the following.

PROPOSITION 5.3. Let \mathbb{A} be a valued σ -structure, \mathbb{A}' be the core of \mathbb{A} , and $k \geq 1$. Then, $\text{opt}_k^{SA}(\mathbb{A}, \mathbb{C}) = \text{opt}_k^{SA}(\mathbb{A}', \mathbb{C})$ for all valued σ -structures \mathbb{C} .

Proof. Since $\mathbb{A} \equiv \mathbb{A}'$, by Proposition 3.6, there exist IFHs from \mathbb{A} to \mathbb{A}' and from \mathbb{A}' to \mathbb{A} . Therefore, by Proposition 5.2, $\text{opt}_k^{SA}(\mathbb{A}, \mathbb{C}) = \text{opt}_k^{SA}(\mathbb{A}', \mathbb{C})$ for all valued σ -structures \mathbb{C} . \square

Given a valued σ -structure \mathbb{A} , the *overlap* of \mathbb{A} is the largest integer m such that there exist $(f, \mathbf{x}), (p, \mathbf{y}) \in \text{tup}(\mathbb{A})_{>0}$ with $(f, \mathbf{x}) \neq (p, \mathbf{y})$ and $|\text{Set}(\mathbf{x}) \cap \text{Set}(\mathbf{y})| = m$.

The following is our second main result; $tw_{ms}(\mathbb{A})$ is defined in section 5.2 and Theorem 5.4 is implied by Theorems 5.11, 5.19, and 5.24 proved in sections 5.3, 5.4, and 5.5, respectively.

THEOREM 5.4 (power of Sherali–Adams). *Let \mathbb{A} be a valued σ -structure and let $k \geq 1$. Let \mathbb{A}' be the core of \mathbb{A} . The Sherali–Adams relaxation of level k is always tight for \mathbb{A} , i.e., for every valued σ -structure \mathbb{B} , we have that $opt_k^{SA}(\mathbb{A}, \mathbb{B}) = opt(\mathbb{A}, \mathbb{B})$, if and only if (i) $tw_{ms}(\mathbb{A}') \leq k - 1$ and (ii) the overlap of \mathbb{A}' is at most k .*

We remark that although Theorem 5.4 deals with valued structures, it also shows the same result for finite-valued structures, where all functions are restricted to finite values in $\mathbb{Q}_{\geq 0}$. In particular, the sufficiency part of Theorem 5.4, i.e., if the core satisfies conditions (i) and (ii) then the k th level of Sherali–Adams is tight, applies directly to the finite-valued case. It will follow from the proofs of Theorems 5.12 and 5.20 that whenever the core violates condition (i) or (ii) then the k th level of Sherali–Adams is not tight even for finite-valued structures. Hence, Theorem 5.4 also characterizes the tightness of Sherali–Adams for finite-valued VCSPs.

Let us note that the characterization given by Theorem 5.4 for levels $k \geq r$, where r is the arity of the signature of \mathbb{A} , boils down to the notion of treewidth. That is, if $k \geq r$, the k th level of Sherali–Adams is tight if and only if the treewidth of the core of \mathbb{A} is at most $k - 1$. Interestingly, Theorem 5.4 tells us *precisely* under which conditions the k th level works even for $k < r$.

Finally, we remark that the tractability part of Theorem 4.1, which we obtain as a corollary of Theorem 5.4, does not immediately follow from a naive algorithm that would compute, for $\mathbb{A} \in \mathcal{C}$, $opt(\mathbb{A}, \mathbb{B})$ using dynamic programming along a tree decomposition of the core \mathbb{A}' of \mathbb{A} . Such an algorithm would first need to compute \mathbb{A}' , and it is not clear that it can be done in polynomial time, even with the promise that \mathbb{A}' has bounded treewidth. (The situation is different for relational structures, where this promise problem is known to be solvable in polynomial time [14, Lemma 25].) Theorem 5.4 gives a way to circumvent this issue since the linear program $SA_k(\mathbb{A}, \mathbb{B})$ does not depend on \mathbb{A}' in any way.

5.2. Treewidth modulo scopes. Let \mathbf{A} be a relational structure with universe A over a relational signature τ . Recall from section 2 that $G(\mathbf{A})$ denotes the Gaifman graph of \mathbf{A} . A *scope* of $G(\mathbf{A})$ is a set X for which there is relation symbol $R \in \tau$ and a tuple $\mathbf{x} \in R^A$ such that $X = \text{Set}(\mathbf{x})$. In other words, the scopes of $G(\mathbf{A})$ are the sets that appear precisely in the tuples of \mathbf{A} .² Observe that every scope X of $G(\mathbf{A})$ induces a clique in $G(\mathbf{A})$.

DEFINITION 5.5. *Let \mathbf{A} be a relational structure and $G(\mathbf{A})$ its Gaifman graph. Let (T, β) be a tree decomposition of $G(\mathbf{A})$, where $T = (V(T), E(T))$. The width modulo scopes of (T, β) is defined by*

$$\max\{|\beta(t)| - 1 \mid t \in V(T) \text{ and } \beta(t) \text{ is not a scope of } G(\mathbf{A})\}.$$

If $\beta(t)$ is a scope for all nodes $t \in V(T)$ then we set the width modulo scopes of (T, β) to be 0. The treewidth modulo scopes of $G(\mathbf{A})$, denoted by $tw_{ms}(G(\mathbf{A}))$, is the minimum width modulo scopes over all its tree decompositions. The treewidth modulo scopes of \mathbf{A}

²In a (V)CSP instance, the term *scope* usually refers to the list of variables a (valued) constraint depends on.

is $tw_{ms}(\mathbf{A}) = tw_{ms}(G(\mathbf{A}))$. For a valued structure \mathbb{A} , we define the treewidth modulo scopes of \mathbb{A} as $tw_{ms}(\mathbb{A}) = tw_{ms}(Pos(\mathbb{A}))$.

Unlike treewidth, the notion of treewidth modulo scopes is not *monotone*, i.e., it can increase after taking substructures. To see this, take, for instance, the relational structure \mathbf{A} that corresponds to the undirected $k \times k$ grid. We have $tw_{ms}(\mathbf{A}) = k$. However, adding a new relation with only one tuple containing all elements of \mathbf{A} lowers the treewidth modulo scopes to 0. Let us also remark that the relational structures with treewidth modulo scopes 0 are precisely the relational structures whose underlying hypergraphs are α -acyclic (see, e.g., [28]).

Let $G = (V, E)$ be a graph. A *bramble* \mathcal{B} of G is a collection of subsets of V such that (i) each $B \in \mathcal{B}$ is a connected set, and (ii) every pair of sets $B, B' \in \mathcal{B}$ touches, i.e., they have a vertex in common or G contains an edge between them. A subset of V is a *cover* of \mathcal{B} if it intersects every set in \mathcal{B} . There is a well-known connection between treewidth and brambles.

THEOREM 5.6 (see [47]). *Let G be a graph and $k \geq 1$. Then the treewidth of G is at most k if and only if any bramble in G can be covered by a set of size at most $k + 1$.*

We show an analogous characterization for treewidth modulo scopes, which will be important later in the proof of Theorem 5.12.

THEOREM 5.7. *Let \mathbf{A} be a relational structure and $k \geq 0$. Then $tw_{ms}(\mathbf{A}) \leq k$ if and only if any bramble in $G(\mathbf{A})$ can be covered by a set of size at most $k + 1$ or by a scope in $G(\mathbf{A})$.*

Proof. The proof is an adaptation of the proof of [18, Theorem 12.3.9]. Suppose first that $tw_{ms}(\mathbf{A}) \leq k$ and let (T, β) be a tree decomposition of $G(\mathbf{A})$ of width modulo scopes at most k , where $T = (V(T), E(T))$. Let \mathcal{B} be any bramble of $G(\mathbf{A})$. We show that there is $t \in V(T)$ such that $\beta(t)$ covers \mathcal{B} . If there is an edge $\{t_1, t_2\}$ of T such that $\beta(t_1) \cap \beta(t_2)$ covers \mathcal{B} , then we are done. Otherwise, we can define an orientation for each edge $\{t_1, t_2\}$ of T as follows. Let $X := \beta(t_1) \cap \beta(t_2)$ and $\mathcal{B}_X = \{B \in \mathcal{B} \mid X \cap B = \emptyset\}$. By assumption, X does not cover \mathcal{B} and then \mathcal{B}_X is not empty. If we remove the edge $\{t_1, t_2\}$ from T , we obtain exactly two connected components $T_1 = (V(T_1), E(T_1))$ and $T_2 = (V(T_2), E(T_2))$ containing t_1 and t_2 , respectively. Let $U_1 := \bigcup_{t \in V(T_1)} \beta(t)$ and $U_2 := \bigcup_{t \in V(T_2)} \beta(t)$. It is a well-known property of tree decompositions (e.g., [18, Lemma 12.3.1]) that X separates U_1 and U_2 in $G(\mathbf{A})$. Take $B \in \mathcal{B}_X$. Since B is connected, it follows that $B \subseteq U_i \setminus X$ for some $i \in \{1, 2\}$. Since all sets in \mathcal{B}_X mutually touch, we can choose the same $i \in \{1, 2\}$ for all $B \in \mathcal{B}_X$. We then orient $\{t_1, t_2\}$ toward t_i .

Let t be a sink node in this orientation of T , i.e., t has only incoming arcs (note that there must be at least one sink). We claim that $\beta(t)$ covers \mathcal{B} . Toward a contradiction, suppose there exists $B \in \mathcal{B}$ with $B \cap \beta(t) = \emptyset$. Take an arbitrary $b \in B$ and a node $t_b \in V(T)$ such that $b \in \beta(t_b)$. Let $\{t', t\}$ be the last edge in the unique simple path from t_b to t . Observe that $(\beta(t') \cap \beta(t)) \cap B = \emptyset$, and hence $B \in \mathcal{B}_{\beta(t') \cap \beta(t)}$. Since the edge $\{t', t\}$ is oriented toward t , we have that $B \subseteq U_t \setminus \beta(t') \cap \beta(t)$, where $U_t = \bigcup_{t \in V(T_t)} \beta(t)$, and $T_t = (V(T_t), E(T_t))$ is the connected component of $T - \{t', t\}$ containing t . In particular, there is $t^* \in V(T_t)$ such that $b \in \beta(t^*)$. By the connectedness property of tree decompositions, we have that $b \in \beta(t)$, a contradiction.

We now prove the other direction. Assume that any bramble in $G(\mathbf{A})$ can be covered by a set of size at most $k + 1$ or by a scope. We say that a set of vertices satisfying this property is *good*; otherwise it is *bad*. Let G' be a subgraph of $G(\mathbf{A})$

and \mathcal{B} be a bramble of $G(\mathbf{A})$. A \mathcal{B} -admissible tree decomposition for G' is a tree decomposition (T, β) of G' with $T = (V(T), E(T))$ such that every $\beta(t)$, with $t \in V(T)$, that is a bad set of $G(\mathbf{A})$ does not cover \mathcal{B} . We will show that for every bramble \mathcal{B} of $G(\mathbf{A})$, there is a \mathcal{B} -admissible tree decomposition for $G(\mathbf{A})$. Since every set covers the empty bramble $\mathcal{B} = \emptyset$, the result follows as any \mathcal{B} -admissible tree decomposition for $\mathcal{B} = \emptyset$ has only good bags, and then its width modulo scopes is at most k .

We now fix a bramble \mathcal{B} of $G(\mathbf{A})$ and assume inductively that $G(\mathbf{A})$ has a \mathcal{B}' -admissible tree decomposition for every bramble \mathcal{B}' with more sets than \mathcal{B} . (The induction starts, since every bramble of $G(\mathbf{A})$ has at most $2^{|V(G(\mathbf{A}))|}$ sets.) By hypothesis, we can take a good set X covering \mathcal{B} . We show the following: For every connected component $C = (V(C), E(C))$ of $G(\mathbf{A}) - X$, there is a \mathcal{B} -admissible tree decomposition of $G(\mathbf{A})[X \cup V(C)]$ (i.e., the subgraph of $G(\mathbf{A})$ induced by $X \cup V(C)$) which has X as one of its bags. These decompositions can be glued to define a \mathcal{B} -admissible tree decomposition of $G(\mathbf{A})$ as required.

Let $C = (V(C), E(C))$ be a fixed connected component of $G(\mathbf{A}) - X$. Since $X \cap V(C) = \emptyset$ and X covers \mathcal{B} , it follows that $V(C) \notin \mathcal{B}$. Then if we define $\mathcal{B}' := \mathcal{B} \cup \{V(C)\}$, we have that $|\mathcal{B}'| > |\mathcal{B}|$. Suppose first that \mathcal{B}' is not a bramble. (Note that in the base case of the induction we always have this case.) This means that $V(C)$ fails to touch a set $B^* \in \mathcal{B}$. Let $N(C)$ be the union of $V(C)$ and all vertices in $G(\mathbf{A})$ adjacent to some vertex in $V(C)$. We have $N(C) \cap B^* = \emptyset$, and hence $N(C)$ does not cover \mathcal{B} . We also have $N(C) \subseteq X \cup V(C)$. Then our desired \mathcal{B} -admissible tree decomposition for $G(\mathbf{A})[X \cup V(C)]$ consists of two adjacent bags X and $N(C)$.

We now assume that \mathcal{B}' is a bramble. By induction, there is a \mathcal{B}' -admissible tree decomposition (T, β) for $G(\mathbf{A})$ with $T = (V(T), E(T))$. We start by considering the case when $|X| \leq k + 1$. Without loss of generality, we can assume that X is a cover of \mathcal{B} of minimum size. In this case we can argue as in the proof of [18, Theorem 12.3.9]. Let $\ell := |X| \leq k + 1$. If (T, β) is \mathcal{B} -admissible, then we are done, so we assume that there exists $s \in V(T)$ such that $\beta(s)$ is bad and covers \mathcal{B} . Every separator of two covers of a bramble is also a cover of that bramble [18, Lemma 12.3.8]. It follows that the minimum size of a separator of X and $\beta(s)$ is ℓ . By Menger's theorem [18, Theorem 3.3.1], there exist ℓ disjoint paths P_1, \dots, P_ℓ linking X and $\beta(s)$ such that each P_i intersects X and $\beta(s)$ exactly in its initial node and final node, respectively. We denote by x_i the initial node of P_i . We have $X = \{x_1, \dots, x_\ell\}$. Observe that since (T, β) is \mathcal{B}' -admissible, the bag $\beta(s)$ fails to cover \mathcal{B}' , that is, $\beta(s) \cap V(C) = \emptyset$. It follows that the path P_i intersects $X \cup V(C)$ exactly in its initial node x_i .

Let (T, β') be the restriction of (T, β) to the set of nodes $X \cup V(C)$, i.e., $\beta'(t) = \beta(t) \cap (X \cup V(C))$, for all $t \in V(T)$. The pair (T, β') is a tree decomposition of $G(\mathbf{A})[X \cup V(C)]$. The desired tree decomposition (T, β'') for $G(\mathbf{A})[X \cup V(C)]$ is the result of adding some of the nodes x_i 's to some particular bags of (T, β') . Let us fix for each x_i a node $t_i \in V(T)$ such that $x_i \in \beta(t_i)$. Formally, (T, β'') is given by

$$\beta''(t) = \beta'(t) \cup \{x_i \mid t \text{ is in the unique simple path from } t_i \text{ to } s \text{ in } T\}.$$

Observe that (T, β'') still satisfies the connectedness property of tree decompositions. Also, observe that $|\beta(t)| \geq |\beta''(t)|$ for all $t \in V(T)$. Indeed, for each $x_i \in \beta''(t) \setminus \beta(t)$, with $i \in \{1, \dots, \ell\}$, the node t is in the unique simple path from t_i to s in T . It follows that $\beta(t)$ contains a node w_i from the path P_i that is different from x_i . Since P_i meets $X \cup V(C)$ only in x_i , we have that $w_i \notin \beta''(t)$. The claim follows since all the w_i 's are distinct.

Since $\beta''(s) = X$, it remains to show that (T, β'') is \mathcal{B} -admissible. Pick $t \in V(T)$ such that $\beta''(t)$ is bad, that is, $|\beta''(t)| > k + 1$ and $\beta''(t)$ is not a scope. We need

to show that $\beta''(t)$ does not cover \mathcal{B} . We claim first that $\beta(t)$ is also a bad set. By contradiction, suppose that $\beta(t)$ is good. As $|\beta(t)| \geq |\beta''(t)| > k + 1$, we have that $\beta(t)$ is a scope. Since $|\beta(t)| \geq |\beta''(t)| > k + 1$ and $|X| = \ell \leq k + 1$, it is the case that $\beta(t) \cap V(C) \neq \emptyset$. It follows that $\beta(t) \subseteq X \cup V(C)$ as $\beta(t)$ is a clique in $G(\mathbf{A})$. Consequently, $\beta(t) \subseteq \beta''(t)$, and since $|\beta(t)| \geq |\beta''(t)|$, we conclude that $\beta(t) = \beta''(t)$. But $\beta''(t)$ is not a scope, which is a contradiction.

Since (T, β) is \mathcal{B}' -admissible and $\beta(t)$ is bad, we have that there is $B^* \in \mathcal{B}'$ such that $\beta(t) \cap B^* = \emptyset$. As discussed in the previous paragraph, we know that $\beta(t) \cap V(C) \neq \emptyset$. It follows that $B^* \in \mathcal{B}$. We claim that $\beta''(t) \cap B^* = \emptyset$, which implies that $\beta''(t)$ does not cover \mathcal{B} , as required. Toward a contradiction, suppose that $x_i \in B^*$, for $x_i \in \beta''(t) \setminus \beta(t)$, with $i \in \{1, \dots, \ell\}$. By definition, t is in the unique simple path from t_i to s in T . Also, by definition, $\beta(s)$ covers \mathcal{B} , and then there is $b \in B^* \cap \beta(s)$. Since B^* is connected, there is a path P from x_i to b in $G(\mathbf{A})$ whose nodes belong to B^* . We have that $\beta(t)$ must contain a node from P , and then a node from B^* , which is a contradiction.

It remains to consider the case when $|X| > k + 1$, and hence X is a scope. Without loss of generality, we assume that X is a maximal scope covering \mathcal{B} . Again, let (T, β') be the restriction of (T, β) to the set of nodes $X \cup V(C)$. For $t \in V(T)$, we say that $\beta'(t)$ is *maximal* if there is no $t' \in V(T)$ such that $\beta'(t) \subsetneq \beta'(t')$. We define $(\tilde{T}, \tilde{\beta})$, where $\tilde{T} = (V(\tilde{T}), E(\tilde{T}))$, to be a decomposition of $G(\mathbf{A})[X \cup V(C)]$ whose bags are precisely the maximal bags of (T, β') , or more formally,

$$\{\tilde{\beta}(t) \mid t \in V(\tilde{T})\} = \{\beta'(t) \mid t \in V(T) \text{ and } \beta'(t) \text{ is maximal}\}.$$

In order to obtain $(\tilde{T}, \tilde{\beta})$ we can iteratively remove nonmaximal bags from (T, β') as follows: if $\beta'(t)$ is not maximal and it is strictly contained in $\beta'(t')$, for $t, t' \in V(T)$, and $\{t, t'\}$ is the first edge in the unique simple path from t to t' in T , then remove $\beta'(t)$ by contracting the edge $\{t, t'\}$ into a new node s defining $\beta'(s) = \beta'(t')$.

We claim that $(\tilde{T}, \tilde{\beta})$ satisfies the required conditions. Since X is a scope and then a clique in $G(\mathbf{A})$, there is a node $t_X \in V(T)$ such that $X \subseteq \beta'(t_X)$. We show first that there is no $t \in V(T)$ such that $X \subsetneq \beta'(t)$. Toward a contradiction, suppose such a t exists. Note that $X \subsetneq \beta'(t) \subseteq \beta(t)$. This implies that $\beta(t)$ covers \mathcal{B}' . Moreover, since $|X| > k + 1$ and X is a maximal scope, we deduce that $\beta(t)$ is a bad set. This contradicts the \mathcal{B}' -admissibility of (T, β) and the claim follows. Using this, we obtain that $\beta'(t_X) = X$ and moreover, $\beta'(t_X)$ is a maximal bag for (T, β') . By construction of $(\tilde{T}, \tilde{\beta})$, there exists $\tilde{t} \in V(\tilde{T})$ such that $\tilde{\beta}(\tilde{t}) = X$. It only remains to show that $(\tilde{T}, \tilde{\beta})$ is \mathcal{B} -admissible.

Suppose $\tilde{\beta}(t)$ is a bad set for a node $t \in V(\tilde{T})$. We prove first that $\tilde{\beta}(t) \cap V(C) \neq \emptyset$. By contradiction, assume that $\tilde{\beta}(t) \subseteq X$. It follows that $\tilde{\beta}(t) = X$; otherwise $\tilde{\beta}(t)$ would not be a maximal bag of (T, β') . But X is a scope and then $\tilde{\beta}(t)$ cannot be bad. This is a contradiction. Let $t^* \in V(T)$ be a node such that $\beta'(t^*) = \tilde{\beta}(t)$. We claim that $\beta(t^*)$ is a bad set. Since $\beta'(t^*) \subseteq \beta(t^*)$, we know that $|\beta(t^*)| \geq |\beta'(t^*)| > k + 1$. For the sake of contradiction, suppose that $\beta(t^*)$ is a scope. In particular, $\beta(t^*)$ is a clique in $G(\mathbf{A})$. Since $\beta'(t^*) \cap V(C) \neq \emptyset$, and hence $\beta(t^*) \cap V(C) \neq \emptyset$, we have that $\beta(t^*) \subseteq X \cup V(C)$. Therefore, $\beta'(t^*) = \beta(t^*)$. This is a contradiction since $\beta'(t^*)$ is bad and then not a scope. Thus $\beta(t^*)$ is a bad set. As (T, β) is \mathcal{B}' -admissible, there is $B^* \in \mathcal{B}'$ with $\beta(t^*) \cap B^* = \emptyset$. Since $\beta(t^*) \cap V(C) \neq \emptyset$, it follows that $B^* \in \mathcal{B}$. Finally, since $\beta'(t^*) \subseteq \beta(t^*)$, we have that $\beta'(t^*) \cap B^* = \emptyset$. Hence, $\beta'(t^*) = \tilde{\beta}(t)$ does not cover \mathcal{B} . We conclude that $(\tilde{T}, \tilde{\beta})$ is \mathcal{B} -admissible. \square

5.3. Sufficiency. We show the following.

THEOREM 5.8. *Let \mathbb{A} be a valued σ -structure and let $k \geq 1$. Suppose that (i) $tw_{ms}(\mathbb{A}) \leq k - 1$ and (ii) the overlap of \mathbb{A} is at most k . Then the Sherali–Adams relaxation of level k is always tight for \mathbb{A} , i.e., for every valued σ -structure \mathbb{B} , we have that $opt_k^{SA}(\mathbb{A}, \mathbb{B}) = opt(\mathbb{A}, \mathbb{B})$.*

Proof. Let \mathbb{B} be an arbitrary valued σ -structure with universe B . Let $\mathbf{A} := \text{Pos}(\mathbb{A})$ and let (T, β) be a tree decomposition of the Gaifman graph $G(\mathbf{A})$ of width modulo scopes $\leq k - 1$, where $T = (V(T), E(T))$. As usual, we denote by A the universe of \mathbf{A} , \mathbf{A} and $G(\mathbf{A})$. Recall that the solutions for $SA_k(\mathbb{A}, \mathbb{B})$ are indexed by the set

$$\begin{aligned} \mathcal{I} := & \{(f, \mathbf{x}, s) : (f, \mathbf{x}) \in \text{tup}(\mathbb{A}_k)_{>0}, s : \text{Set}(\mathbf{x}) \rightarrow B_k\} \\ = & \{(f, \mathbf{x}, s) : (f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{>0}, s : \text{Set}(\mathbf{x}) \rightarrow B\} \\ & \cup \{(\rho_k, \mathbf{x}, s) : \mathbf{x} \in A^k, s : \text{Set}(\mathbf{x}) \rightarrow B\}. \end{aligned}$$

Let $\text{BTup} \subseteq \text{tup}(\mathbb{A}_k)_{>0}$ be the set

$$\text{BTup} := \{(f, \mathbf{x}) \in \text{tup}(\mathbb{A}_k)_{>0} \mid \text{Set}(\mathbf{x}) \subseteq \beta(t) \text{ for some } t \in V(T)\}.$$

We have $\text{tup}(\mathbb{A})_{>0} \subseteq \text{BTup}$. We define $\mathcal{T} := \{(f, \mathbf{x}, s) \in \mathcal{I} : (f, \mathbf{x}) \in \text{BTup}\}$. Let $P(\mathbb{A}, \mathbb{B})$ be the system of linear inequalities given by the constraints of $SA_k(\mathbb{A}, \mathbb{B})$ restricted to the variables indexed by \mathcal{T} . More precisely, $P(\mathbb{A}, \mathbb{B})$ is the following system over variables $\{\lambda(f, \mathbf{x}, s) : (f, \mathbf{x}, s) \in \mathcal{T}\}$:

$$(5.1) \quad \begin{aligned} \lambda(f, \mathbf{x}, s) = & \sum_{r: \text{Set}(\mathbf{y}) \rightarrow B, r|_{\text{Set}(\mathbf{x})} = s} \lambda(p, \mathbf{y}, r) & \forall (f, \mathbf{x}), (p, \mathbf{y}) \in \text{BTup} : \\ & \text{Set}(\mathbf{x}) \subseteq \text{Set}(\mathbf{y}) \text{ and } |\text{Set}(\mathbf{x})| \leq k; & \forall s : \text{Set}(\mathbf{x}) \rightarrow B, \end{aligned}$$

$$(5.2) \quad \sum_{s: \text{Set}(\mathbf{x}) \rightarrow B} \lambda(f, \mathbf{x}, s) = 1 \quad \forall (f, \mathbf{x}) \in \text{BTup},$$

$$(5.3) \quad \lambda(f, \mathbf{x}, s) = 0 \quad \forall (f, \mathbf{x}, s) \in \mathcal{T}, f^{\mathbb{A}_k}(\mathbf{x}) \times f^{\mathbb{B}_k}(s(\mathbf{x})) = \emptyset,$$

$$(5.4) \quad \lambda(f, \mathbf{x}, s) \geq 0 \quad \forall (f, \mathbf{x}, s) \in \mathcal{T}.$$

From the definition of the Sherali–Adams hierarchy, we have $opt_k^{SA}(\mathbb{A}, \mathbb{B}) \leq opt(\mathbb{A}, \mathbb{B})$. We need to prove that $opt(\mathbb{A}, \mathbb{B}) \leq opt_k^{SA}(\mathbb{A}, \mathbb{B})$. Let $\lambda = \{\lambda(f, \mathbf{x}, s) : (f, \mathbf{x}, s) \in \mathcal{I}\}$ be an optimal solution to $SA_k(\mathbb{A}, \mathbb{B})$. Let $c = \{c(f, \mathbf{x}, s) : (f, \mathbf{x}, s) \in \mathcal{I}\}$ be the vector defining the objective function of $SA_k(\mathbb{A}, \mathbb{B})$. Consider the projection $\lambda|_{\mathcal{T}} = \{\lambda(f, \mathbf{x}, s) : (f, \mathbf{x}, s) \in \mathcal{T}\}$ of λ to \mathcal{T} . Similarly, consider the projection $c|_{\mathcal{T}}$ of c to \mathcal{T} . The restriction of c to $\mathcal{I} \setminus \mathcal{T}$ is the vector 0, and hence, $c|_{\mathcal{T}} \cdot \lambda|_{\mathcal{T}} = c \cdot \lambda$. Also, $\lambda|_{\mathcal{T}}$ is a solution to $P(\mathbb{A}, \mathbb{B})$. By Lemma 5.9 proved below, the polytope $P(\mathbb{A}, \mathbb{B})$ is integral and thus $\lambda|_{\mathcal{T}}$ is a convex combination of integral solutions I^{g_1}, \dots, I^{g_n} of $P(\mathbb{A}, \mathbb{B})$, for assignments $g_i : A \rightarrow B$. It follows that there exists $i \in \{1, \dots, n\}$ such that $c|_{\mathcal{T}} \cdot I^{g_i} \leq c|_{\mathcal{T}} \cdot \lambda|_{\mathcal{T}} = c \cdot \lambda$. In particular, the cost of the assignment g_i is $\leq c \cdot \lambda = opt_k^{SA}(\mathbb{A}, \mathbb{B})$. We conclude that $opt(\mathbb{A}, \mathbb{B}) \leq opt_k^{SA}(\mathbb{A}, \mathbb{B})$. \square

This is the last missing piece in the proof of Theorem 5.8.

LEMMA 5.9. *The polytope described by $P(\mathbb{A}, \mathbb{B})$ is integral.*

Proof. We start with some notation. Let ω be a probability distribution over W^X . For $X' \subseteq X$, we define the *marginal* of ω , denoted by $\omega_{X'}$, to be the probability distribution over $W^{X'}$ defined by $\omega_{X'}(h) = \sum_{g: X \rightarrow W, g|_{X'} = h} \omega(g)$. We have the following.

CLAIM 5.10. *Let ω and δ be probability distributions over W^X and W^Y , respectively. Suppose $\omega_{X \cap Y} = \delta_{X \cap Y}$. Then there is a probability distribution ϵ over $W^{X \cup Y}$ with $\epsilon_X = \omega$ and $\epsilon_Y = \delta$.*

Proof. We set $\epsilon(h) = \frac{\omega(h|_X)\delta(h|_Y)}{\omega(h|_{X \cap Y})}$ for every $h : X \cup Y \rightarrow W$. ($\epsilon(h) = 0$ if $\omega(h|_{X \cap Y}) = 0$.) \square

Let us fix a solution λ of $P(\mathbb{A}, \mathbb{B})$. We need to show that λ is a convex combination of integral solutions. Since the width modulo scopes of the decomposition (T, β) is $\leq k - 1$, we know that either $|\beta(t)| \leq k$ or $\beta(t)$ is a scope of $G(\mathbf{A})$ for every $t \in V(T)$. It follows that for each $t \in V(T)$, we can find $(f_t, \mathbf{x}_t) \in \text{BTup}$ such that $\text{Set}(\mathbf{x}_t) = \beta(t)$.

For $t \in V(T)$, let ω^t be the probability distribution over $B^{\beta(t)}$ given by $\omega^t(h) = \lambda(f_t, \mathbf{x}_t, h)$. Note that (\star) for every edge $\{t, t'\} \in E(T)$, we have that $\omega_{\beta(t) \cap \beta(t')}^t = \omega_{\beta(t) \cap \beta(t')}^{t'}$. Indeed, since the overlap of \mathbb{A} is $\leq k$, we can pick a tuple $(\rho_k, \mathbf{z}) \in \text{BTup}$ such that $\text{Set}(\mathbf{z}) = \beta(t) \cap \beta(t')$. By restriction (5.1) in $P(\mathbb{A}, \mathbb{B})$, we obtain that $\omega_{\beta(t) \cap \beta(t')}^t = \lambda(\rho_k, \mathbf{z}, \cdot) = \omega_{\beta(t) \cap \beta(t')}^{t'}$. We can then define a distribution ϵ over B^A such that $\epsilon_{\beta(t)} = \omega^t$ for every $t \in V(T)$. To do this, we start at a particular node $r \in V(T)$. We apply iteratively Claim 5.10 to extend the current distribution (ω^r at the beginning) one bag at a time, until all bags are covered. Note that the hypothesis of Claim 5.10 (same marginals) is always satisfied due to (\star) .

For $g : A \rightarrow B$, we define $I^g = \{I^g(f, \mathbf{x}, s) : (f, \mathbf{x}, s) \in \mathcal{T}\}$ to be the 0/1-vector such that $I^g(f, \mathbf{x}, s) = 1$ if and only if $g|_{\text{Set}(\mathbf{x})} = s$. We show that $\lambda = \sum_{g \in \text{supp}(\epsilon)} \epsilon(g) I^g$. Take $(f, \mathbf{x}, s) \in \mathcal{T}$. By definition of \mathcal{T} , there is $t \in V(T)$ such that $\text{Set}(\mathbf{x}) \subseteq \beta(t) = \text{Set}(\mathbf{x}_t)$. If $(f, \mathbf{x}) = (f_t, \mathbf{x}_t)$, then

$$\sum_{g \in \text{supp}(\epsilon)} \epsilon(g) I^g(f_t, \mathbf{x}_t, s) = \sum_{g \in \text{supp}(\epsilon), g|_{\beta(t)} = s} \epsilon(g) = \omega^t(s) = \lambda(f_t, \mathbf{x}_t, s)$$

and we are done. In case $(f, \mathbf{x}) \neq (f_t, \mathbf{x}_t)$, since the overlap is $\leq k$, we have that $|\text{Set}(\mathbf{x})| \leq k$, and by condition (5.1) in $P(\mathbb{A}, \mathbb{B})$, we obtain

$$\lambda(f, \mathbf{x}, s) = \sum_{r: \beta(t) \rightarrow B, r|_{\text{Set}(\mathbf{x})} = s} \omega^t(r) = \epsilon_{\text{Set}(\mathbf{x})}(s) = \sum_{g \in \text{supp}(\epsilon)} \epsilon(g) I^g(f, \mathbf{x}, s).$$

It remains to show that I^g is a solution to $P(\mathbb{A}, \mathbb{B})$ for all $g \in \text{supp}(\epsilon)$. Conditions (5.1), (5.2), and (5.4) hold by construction. For condition (5.3), take $(f, \mathbf{x}, s) \in \mathcal{T}$ such that $f^{\mathbb{A}_k}(\mathbf{x}) \times f^{\mathbb{B}_k}(s(\mathbf{x})) = \infty$. Since $0 = \lambda(f, \mathbf{x}, s) = \sum_{g \in \text{supp}(\epsilon)} \epsilon(g) I^g(f, \mathbf{x}, s)$, it follows that $I^g(f, \mathbf{x}, s) = 0$ for all $g \in \text{supp}(\epsilon)$, and hence condition (5.3) is satisfied. \square

THEOREM 5.11. *Let \mathbb{A} be a valued σ -structure and \mathbb{A}' be its core. Suppose that (i) $\text{tw}_{ms}(\mathbb{A}') \leq k - 1$ and (ii) the overlap of \mathbb{A}' is at most k . Then the Sherali–Adams relaxation of level k is always tight for \mathbb{A} , i.e., for every valued σ -structure \mathbb{B} , we have that $\text{opt}_k^{\text{SA}}(\mathbb{A}, \mathbb{B}) = \text{opt}(\mathbb{A}, \mathbb{B})$.*

Proof. Let \mathbb{B} be a valued σ -structure. We can apply Theorem 5.8 to \mathbb{A}' and obtain that $\text{opt}_k^{\text{SA}}(\mathbb{A}', \mathbb{B}) = \text{opt}(\mathbb{A}', \mathbb{B})$. Since $\mathbb{A}' \equiv \mathbb{A}$, we know that $\text{opt}(\mathbb{A}, \mathbb{B}) = \text{opt}(\mathbb{A}', \mathbb{B})$ (by the definition of equivalence and cores) and $\text{opt}_k^{\text{SA}}(\mathbb{A}, \mathbb{B}) = \text{opt}_k^{\text{SA}}(\mathbb{A}', \mathbb{B})$ (by Proposition 5.3). Hence, $\text{opt}_k^{\text{SA}}(\mathbb{A}, \mathbb{B}) = \text{opt}(\mathbb{A}, \mathbb{B})$. \square

5.4. Necessity of treewidth modulo scopes. In this and the following section we will denote by \oplus the addition modulo 2. We show the following.

THEOREM 5.12. *Let \mathbb{A} be a valued σ -structure and let $k \geq 1$. Suppose that \mathbb{A} is a core and $tw_{ms}(\mathbb{A}) \geq k$. Then there exists a valued σ -structure \mathbb{B} such that $opt_k^{SA}(\mathbb{A}, \mathbb{B}) < opt(\mathbb{A}, \mathbb{B})$.*

Proof. Let $\mathbf{A} := \text{Pos}(\mathbb{A})$. As usual, we denote by A the universe of \mathbb{A} , \mathbf{A} and the vertex set of $G(\mathbf{A})$. Since $tw_{ms}(\mathbb{A}) = tw_{ms}(G(\mathbf{A})) \geq k$, Theorem 5.7 implies that there exists a bramble \mathcal{B} of $G(\mathbf{A})$ that cannot be covered by any scope nor subset of size at most k in $G(\mathbf{A})$. Every $B \in \mathcal{B}$ must belong to the same connected component of $G(\mathbf{A})$, which we denote by $G_0 = (A_0, E_0)$. Fix any $a_0 \in A_0$. For each $a \in A$, let d_a denote the degree of a in $G(\mathbf{A})$, and let $e_1^a, \dots, e_{d_a}^a$ be a fixed enumeration of all the edges incident to a in $G(\mathbf{A})$.

Recall that \mathbf{A} is defined over the relational signature $\text{rel}(\sigma) = \{R_f \mid f \in \sigma, \text{ar}(R_f) = \text{ar}(f)\}$. We define a relational structure \mathbf{B} over $\text{rel}(\sigma)$ as in [1]. The universe of \mathbf{B} , denoted by B , contains precisely all the tuples $(a, (b_1, \dots, b_{d_a}))$ such that

1. $a \in A$ and $b_1, \dots, b_{d_a} \in \{0, 1\}$;
2. $b_1 \oplus \dots \oplus b_{d_a} = 0$ if $a \neq a_0$;
3. $b_1 \oplus \dots \oplus b_{d_a} = 1$ if $a = a_0$.

Let $R_f \in \text{rel}(\sigma)$ be of arity n . A tuple $((a^1, (b_1^1, \dots, b_{d_{a^1}}^1)), \dots, (a^n, (b_1^n, \dots, b_{d_{a^n}}^n)))$ belongs to $R_f^{\mathbf{B}}$ if and only if

1. (a^1, \dots, a^n) belongs to $R_f^{\mathbf{A}}$,
2. if $\{a^\ell, a^m\} = e_i^{\ell} = e_j^m$, for some $\ell, m \in \{1, \dots, n\}$, $i \in \{1, \dots, d_{a^\ell}\}$ and $j \in \{1, \dots, d_{a^m}\}$, then $b_i^\ell = b_j^m$.

Let $\pi : B \rightarrow A$ be the first projection, i.e., $\pi((a, (b_1, \dots, b_{d_a}))) = a$, for all $(a, (b_1, \dots, b_{d_a})) \in B$. By definition of \mathbf{B} , π is a homomorphism from \mathbf{B} to \mathbf{A} .

The following lemma follows directly from the proof of [1, Lemma 1].

LEMMA 5.13. *There is no homomorphism h from \mathbf{A} to \mathbf{B} such that $\pi \circ h(A) = A$.*

By Proposition 3.13, since \mathbb{A} is a core, there is a function $c^* : \text{tup}(\mathbb{A}) \rightarrow \mathbb{Q}_{\geq 0}$ such that every nonsurjective mapping $g : A \rightarrow A$ satisfies

$$\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) c^*(f, \mathbf{x}) < \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) c^*(f, g(\mathbf{x})).$$

We denote $M^* := \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) c^*(f, \mathbf{x}) < \infty$.

Now we are ready to define \mathbb{B} . The universe of \mathbb{B} is B , i.e., the same as \mathbf{B} . For each $f \in \sigma$ and tuple $\mathbf{x} \in B^{\text{ar}(f)}$ we define

$$f^{\mathbb{B}}(\mathbf{x}) = \begin{cases} c^*(f, \pi(\mathbf{x})) & \text{if } \mathbf{x} \in R_f^{\mathbf{B}}, \\ \infty & \text{otherwise.} \end{cases}$$

CLAIM 5.14. $opt(\mathbb{A}, \mathbb{B}) > M^*$.

Proof. Let h be a homomorphism from \mathbf{A} to \mathbf{B} (if it is not, then $\text{cost}_{\mathbb{A} \rightarrow \mathbb{B}}(h) = \infty$). Then

$$\text{cost}_{\mathbb{A} \rightarrow \mathbb{B}}(h) = \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{>0}} f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{B}}(h(\mathbf{x})) = \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{>0}} f^{\mathbb{A}}(\mathbf{x}) c^*(f, \pi(h(\mathbf{x}))).$$

By Lemma 5.13, $\pi \circ h$ is a nonsurjective mapping from A to A . By definition of c^* , we have that $\text{cost}_{\mathbb{A} \rightarrow \mathbb{B}}(h) > M^*$. \square

In the rest of the proof, we will show that $\text{opt}_k^{\text{SA}}(\mathbb{A}, \mathbb{B}) \leq M^*$. Together with Claim 5.14, this establishes Theorem 5.12.

Let s be a partial homomorphism from \mathbf{A} to \mathbf{B} . We denote by $\text{dom}(s)$ the domain of s . We say that s is an *identity partial homomorphism* from \mathbf{A} to \mathbf{B} if $\pi(s(a)) = a$ for all $a \in \text{dom}(s)$. We denote by $\text{IPHom}(\mathbf{A}, \mathbf{B})$ the set of all identity partial homomorphisms from \mathbf{A} to \mathbf{B} . Let λ be a feasible solution of $\text{SA}_k(\mathbb{A}, \mathbb{B})$ satisfying the following property (\dagger): if $\lambda(f, \mathbf{x}, s) > 0$ for $(f, \mathbf{x}) \in \text{tup}(\mathbb{A}_k)_{>0}, s : \text{Set}(\mathbf{x}) \rightarrow B$, then s is an identity partial homomorphism from \mathbf{A} to \mathbf{B} . We claim that the cost of any such λ is precisely M^* . Indeed, we have that

$$\begin{aligned}
& \sum_{\substack{(f, \mathbf{x}) \in \text{tup}(\mathbb{A}_k)_{>0}, s : \text{Set}(\mathbf{x}) \rightarrow B \\ f^{\mathbb{A}_k}(\mathbf{x}) \times f^{\mathbb{B}_k}(s(\mathbf{x})) < \infty}} \lambda(f, \mathbf{x}, s) f^{\mathbb{A}_k}(\mathbf{x}) f^{\mathbb{B}_k}(s(\mathbf{x})) \\
&= \sum_{\substack{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{>0}, s : \text{Set}(\mathbf{x}) \rightarrow B \\ f^{\mathbb{A}}(\mathbf{x}) \times f^{\mathbb{B}}(s(\mathbf{x})) < \infty}} \lambda(f, \mathbf{x}, s) f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{B}}(s(\mathbf{x})) \\
&= \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{>0}} f^{\mathbb{A}}(\mathbf{x}) \sum_{s : \text{Set}(\mathbf{x}) \rightarrow B, f^{\mathbb{A}}(\mathbf{x}) \times f^{\mathbb{B}}(s(\mathbf{x})) < \infty} \lambda(f, \mathbf{x}, s) f^{\mathbb{B}}(s(\mathbf{x})) \\
&= \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{>0}} f^{\mathbb{A}}(\mathbf{x}) \sum_{\substack{s : \text{Set}(\mathbf{x}) \rightarrow B, f^{\mathbb{A}}(\mathbf{x}) \times f^{\mathbb{B}}(s(\mathbf{x})) < \infty \\ s \in \text{IPHom}(\mathbf{A}, \mathbf{B})}} \lambda(f, \mathbf{x}, s) f^{\mathbb{B}}(s(\mathbf{x})) \\
&= \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{>0}} f^{\mathbb{A}}(\mathbf{x}) \sum_{\substack{s : \text{Set}(\mathbf{x}) \rightarrow B, f^{\mathbb{A}}(\mathbf{x}) \times f^{\mathbb{B}}(s(\mathbf{x})) < \infty \\ s \in \text{IPHom}(\mathbf{A}, \mathbf{B})}} \lambda(f, \mathbf{x}, s) c^*(f, \mathbf{x}) \\
&= \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{>0}} f^{\mathbb{A}}(\mathbf{x}) c^*(f, \mathbf{x}) = M^*.
\end{aligned}$$

Thus it suffices to show the existence of a feasible solution to $\text{SA}_k(\mathbb{A}, \mathbb{B})$ satisfying (\dagger). In [1, Lemma 2] it is shown that if the treewidth of \mathbf{A} is at least k , then the $(k-1)$ -consistency test succeeds over \mathbf{A}, \mathbf{B} . To do this, the authors of [1] exhibit a winning strategy \mathcal{H}_k for the duplicator in the existential k -pebble game over \mathbf{A}, \mathbf{B} . We built on their construction to define a set \mathcal{H} of identity partial homomorphisms from \mathbf{A} to \mathbf{B} , which will allow us to define our required λ .

Recall that we fixed a bramble \mathcal{B} of $G(\mathbf{A})$ that cannot be covered by any scope nor subset of size at most k in $G(\mathbf{A})$, whose existence is guaranteed by Theorem 5.7. Recall also that every set in \mathcal{B} belongs to the connected component $G_0 = (A_0, E_0)$ of $G(\mathbf{A})$ and that $a_0 \in A_0$. Let P be a path in $G(\mathbf{A})$ starting at a_0 . For every edge e of $G(\mathbf{A})$, we define

$$x_e^P = \begin{cases} 1 & \text{if } e \text{ appears an odd number of times in } P, \\ 0 & \text{otherwise.} \end{cases}$$

For $a \in A$, we define $h^P(a) = (a, (x_{e_1^a}^P, \dots, x_{e_{d_a}^a}^P))$. We have $h^P(a) = (a, (0, \dots, 0))$, for all $a \in A \setminus A_0$. The following claim is shown in [1] (the second claim in the proof of Lemma 2).

CLAIM 5.15. *Let $X \subseteq A$, and let P be a path in $G(\mathbf{A})$ from a_0 to b , where $b \notin X$. Then, the restriction $h^P|_X$ of h^P to X is a partial homomorphism from \mathbf{A} to \mathbf{B} .*

Notice that in the previous claim, $h^P|_X$ is actually an identity partial homomorphism. Let $\mathcal{X} := \{X \subseteq A : X \text{ does not cover } \mathcal{B}\}$. We define a set \mathcal{H} of identity partial

homomorphisms from \mathbf{A} to \mathbf{B} as follows. For every $X \in \mathcal{X}$ and every path P in $G(\mathbf{A})$ from a_0 to b , where b belongs to a set $B \in \mathcal{B}$ disjoint from X , we include $h^P|_X$ to \mathcal{H} . By Claim 5.15, \mathcal{H} contains only identity partial homomorphisms from \mathbf{A} to \mathbf{B} . For $X \in \mathcal{X}$, let $\mathcal{H}(X) := \{h \in \mathcal{H} \mid \text{dom}(h) = X\}$. We have $|\mathcal{H}(X)| > 0$ for all $X \in \mathcal{X}$, as X does not cover \mathcal{B} and then there is a set $B \in \mathcal{B}$ disjoint from X , and by connectivity, there is at least one path P from a_0 to some node $b \in B$.

It follows from the proof of Lemma 2 in [1] that \mathcal{H} has the following closure properties.

CLAIM 5.16. *Let $X, X' \in \mathcal{X}$ such that $X \subseteq X'$, then*

1. *if $h \in \mathcal{H}(X)$, then there exists $h' \in \mathcal{H}(X')$ such that $h'|_X = h$,*
2. *if $h' \in \mathcal{H}(X')$, then $h'|_X \in \mathcal{H}(X)$.*

For $X, X' \in \mathcal{X}$ such that $X \subseteq X'$ and $h \in B^X$, let $\mathcal{H}(X')^{X,h} := \{h' \in \mathcal{H}(X') \mid h'|_X = h\}$. Claim 5.16, item 1 states that $|\mathcal{H}(X')^{X,h}| > 0$ for all $h \in \mathcal{H}(X)$.

The arguments below are an adaptation of the argument from [51] for proving the existence of gap instances for Sherali–Adams relaxations of $\text{VCSP}(E_{\mathcal{G},3})$, and a refinement of an argument from [1, Lemma 2] (specifically Claim 5.18 below). In particular, we shall define a well-behaved subset $\mathcal{S} \subseteq \mathcal{X}$. A key property of \mathcal{S} is that the family of distributions $\{U(\mathcal{H}(S)) \mid S \in \mathcal{S}\}$, where $U(\mathcal{H}(S))$ is the uniform distribution over $\mathcal{H}(S)$, is *consistent* in the following sense: for every $S \subseteq S'$, the marginal distribution of $U(\mathcal{H}(S'))$ over S coincides with $U(\mathcal{H}(S))$. As it turns out, we will be able to extend $\{U(\mathcal{H}(S)) \mid S \in \mathcal{S}\}$ to a consistent family $\{\mu(\mathcal{H}(X)) \mid X \in \mathcal{X}\}$ over the whole set \mathcal{X} . We shall define our required λ from this latter family.

We say that $X \subseteq A$ *separates* G_0 if there exist nodes $u, u' \in A_0 \setminus X$ such that X separates u and u' , i.e., every path connecting u and u' in G_0 intersects X . Let $\mathcal{S} := \{S \in \mathcal{X} \mid S \text{ does not separate } G_0\}$. Below we state two important properties of \mathcal{S} .

CLAIM 5.17. *\mathcal{S} is closed under intersection.*

Proof. Let $S, S' \in \mathcal{S}$. Clearly, $S \cap S'$ does not cover \mathcal{B} . To see that $S \cap S'$ does not separate G_0 , take $a, a' \in A_0 \setminus (S \cap S')$. We need to find a path in G_0 between a and a' avoiding $S \cap S'$. If $a, a' \in A_0 \setminus S$ or $a, a' \in A_0 \setminus S'$, we are done as S and S' do not separate G_0 . Thus, without loss of generality we can assume that $a \in (A_0 \cap S) \setminus S'$ and $a' \in (A_0 \cap S') \setminus S$. Let $B, B' \in \mathcal{B}$ such that $B \cap S = \emptyset$ and $B' \cap S' = \emptyset$, and pick $b \in B$ and $b' \in B'$. Since S' does not separate a and b' , there exists a path P from a to b' avoiding S' , and hence $S \cap S'$. Since B, B' are connected and they touch, there is also a path P' from b' to b avoiding $S \cap S'$. Finally, since S does not separate a' and b , there is a path P'' from b to a' avoiding S , and then $S \cap S'$. The claim follows by taking the concatenation of P, P' , and P'' . \square

CLAIM 5.18. *Let $S, S' \in \mathcal{S}$ such that $S \subseteq S'$ and let $h \in \mathcal{H}(S)$. Then $|\mathcal{H}(S')^{S,h}| = |\mathcal{H}(S')|/|\mathcal{H}(S)|$.*

Proof. The proof is by induction on $\ell := |S' \setminus S|$. If $\ell = 0$, we are done. Suppose that $\ell = 1$ and $S' \setminus S = \{x\}$. Assume first that $x \notin A_0$. Let $g \in \mathcal{H}(S)$ and P be any path starting at a_0 and ending at some node in $B \in \mathcal{B}$, with $B \cap S = \emptyset$, such that $h^P|_S = g$. Since $B \cap S' = \emptyset$ as $B \subseteq A_0$ and thus $h^P|_{S'} \in \mathcal{H}(S')^{S,g}$. Notice also that for every $g' \in \mathcal{H}(S')^{S,g}$, it must be the case that $g'(x) = (x, (0, \dots, 0))$, as a path P' starting at a_0 cannot visit an edge e incident to x , and hence $x_e^{P'} = 0$. It follows that $\mathcal{H}(S')^{S,g} = \{h^P|_{S'}\}$. In particular, $|\mathcal{H}(S')^{S,h}| = 1$ and $|\mathcal{H}(S')| = \sum_{g \in \mathcal{H}(S)} |\mathcal{H}(S')^{S,g}|$. By Claim 5.16, item 2, $\sum_{g \in \mathcal{H}(S)} |\mathcal{H}(S')^{S,g}| = \sum_{g \in \mathcal{H}(S)} |\mathcal{H}(S')^{S,g}|$. This implies that $|\mathcal{H}(S')| = |\mathcal{H}(S)|$, and then the claim follows.

We now assume that $x \in A_0$. We show that there is an edge $\{x, u\}$ in $G(\mathbf{A})$ such that $u \notin S'$. Assume to the contrary, and take $B' \in \mathcal{B}$ such that $S' \cap B' = \emptyset$ and pick $b' \in B'$. We have $x, b' \in A_0 \setminus S$. Since every neighbor of x lies in S , we have that S separates x and b' , which contradicts the fact that $S \in \mathcal{S}$. Let e_1, \dots, e_m , with $m \geq 1$, be an enumeration of all the edges from x to a node outside S' . Recall that $e_1^x, \dots, e_{d_x}^x$ is our fixed enumeration of all the edges incident to x . Without loss of generality, suppose that $e_1^x, \dots, e_{d_x}^x = e_1, \dots, e_m, e'_1, \dots, e'_n$, where e'_1, \dots, e'_n is an enumeration ($n \geq 0$) of all the edges from x to a node inside S . We claim that $|\mathcal{H}(S')^{S,g}| = 2^{m-1}$ for all $g \in \mathcal{H}(S)$.

We start by proving $|\mathcal{H}(S')^{S,g}| \leq 2^{m-1}$.

We define an injective mapping η from $\mathcal{H}(S')^{S,g}$ to $\{0, 1\}^{m-1}$. Let $g' \in \mathcal{H}(S')^{S,g}$ and suppose that we have $g'(x) = (x, (y_1, \dots, y_m, y'_1, \dots, y'_n))$. We define $\eta(g') = (y_1, \dots, y_{m-1})$. Suppose that $\eta(g'_1) = \eta(g'_2)$ for $g'_1, g'_2 \in \mathcal{H}(S')^{S,g}$. Thus, $g'_1(x)$ is of the form $g'_1(x) = (x, (\eta(g'_1), y_{1m}, y'_{11}, \dots, y'_{1n}))$ and $g'_2(x)$ is of the form $g'_2(x) = (x, (\eta(g'_2), y_{2m}, y'_{21}, \dots, y'_{2n}))$. We have $(y'_{11}, \dots, y'_{1n}) = (y'_{21}, \dots, y'_{2n})$ as these values are already determined by g . It follows that $y_{1m} = y_{2m}$ as the parity of the number of ones in $(\eta(g'_1), y_{1m}, y'_{11}, \dots, y'_{1n})$ and $(\eta(g'_2), y_{2m}, y'_{21}, \dots, y'_{2n})$ must be the same. In particular, $g'_1 = g'_2$, and thus η is injective.

Now we show that $|\mathcal{H}(S')^{S,g}| \geq 2^{m-1}$. As we mentioned above, if $g' \in \mathcal{H}(S')^{S,g}$ and $g'(x) = (x, (y_1, \dots, y_m, y'_1, \dots, y'_n))$, then (y'_1, \dots, y'_n) is already determined by g . Let $\bar{z} = (z_1, \dots, z_m) \in \{0, 1\}^m$ be an arbitrary 0/1-vector such that

- $\bigoplus_{1 \leq i \leq m} z_i \oplus \bigoplus_{1 \leq j \leq n} y'_j = 1$ if $x = a_0$;
- $\bigoplus_{1 \leq i \leq m} z_i \oplus \bigoplus_{1 \leq j \leq n} y'_j = 0$ if $x \neq a_0$.

The number of such \bar{z} 's is precisely 2^{m-1} , and hence, it suffices to show that, for every such \bar{z} , there exists $g'_{\bar{z}} \in \mathcal{H}(S')^{S,g}$ such that $g'_{\bar{z}}(x) = (x, (\bar{z}, y'_1, \dots, y'_n))$. Let P be any path in $G(\mathbf{A})$ from a_0 to a node $b \in B$, where $B \in \mathcal{B}$ and $B \cap S = \emptyset$, such that $h^P|_S = g$. Let $B' \in \mathcal{B}$ such that $B' \cap S' = \emptyset$ and pick any $b' \in B'$. Since B, B' are connected and they touch, there exists a path P' from b to b' completely contained in $B \cup B'$ and, hence, avoiding S . Let W be the concatenation of P and P' . By construction, $g' := h^W|_{S'} \in \mathcal{H}(S')$. Since P' avoids S , $x_e^W = x_e^P$, for every edge e incident to a node in S . Therefore, $h^W(a) = h^P(a)$ for all $a \in S$. It follows that $g' \in \mathcal{H}(S')^{S,g}$.

Let $g'(x) = (x, (y_1, \dots, y_m, y'_1, \dots, y'_n))$. Suppose without loss of generality that $\bar{z} = (z_1, \dots, z_r, z_{r+1}, \dots, z_m)$, where $r \in \{0, \dots, m\}$ and $z_i \neq y_i$ if and only if $i \in \{1, \dots, r\}$. If $r = 0$, we are done as we can set $g'_{\bar{z}} = g'$. Suppose then that $r \geq 1$. Observe that $\bigoplus_{1 \leq i \leq r} y_i = \bigoplus_{1 \leq i \leq r} z_i$. This implies that r is even. Recall that e_1, \dots, e_m is our fixed enumeration of all the edges from x to a node outside S' . Let u_1, \dots, u_r be the nodes outside S' such that $e_i = \{x, u_i\}$ for all $i \in \{1, \dots, r\}$. We have $u_i, b' \in A_0 \setminus S'$. Since $S' \in \mathcal{S}$, S' cannot separate u_i and b' , there is a path P_i from b' to u_i avoiding S' , for all $i \in \{1, \dots, r\}$. Let W_i be the extension of P_i with the edge $\{u_i, x\}$. We denote by W_i^{-1} the reverse path of W_i , in particular, W_i^{-1} is a path from x to b' . Let W' be the extension of the path W with the concatenation Z of the paths $W_1, W_2^{-1}, \dots, W_{r-1}, W_r^{-1}$. Since W' ends at b' we have $h^{W'}|_{S'} \in \mathcal{H}(S')$. Since the subpath Z of W' avoids S , $h^{W'}(a) = h^W(a)$ for all $a \in S$, and then, $h^{W'}|_{S'} \in \mathcal{H}(S')^{S,g}$. Notice also that the subpath Z visits only the e_i 's with $i \in \{1, \dots, r\}$, and it visits each such e_i exactly once. This implies that $h^{W'}(x) = (x, (y_1 \oplus 1, \dots, y_r \oplus 1, y_{r+1}, \dots, y_m, y'_1, \dots, y'_n))$. In other words, $h^{W'}(x) = (x, (z_1, \dots, z_m, y'_1, \dots, y'_n))$. The claim follows by taking $g'_{\bar{z}} = h^{W'}|_{S'}$. Therefore, $|\mathcal{H}(S')^{S,g}| = 2^{m-1}$ for all $g \in \mathcal{H}(S)$. In particular, $|\mathcal{H}(S')^{S,h}| = 2^{m-1}$. On the

other hand, $|\mathcal{H}(S')| = \sum_{g \in B^S} |\mathcal{H}(S')^{S,g}|$. By Claim 5.16, item 2, $\sum_{g \in B^S} |\mathcal{H}(S')^{S,g}| = \sum_{g \in \mathcal{H}(S)} |\mathcal{H}(S')^{S,g}|$. This implies that $|\mathcal{H}(S')| = 2^{m-1} |\mathcal{H}(S)|$, and then our claim holds for $\ell = 1$.

Assume now that $\ell \geq 2$. We claim that there is a node $x^* \in S' \setminus S$ such that $S' \setminus \{x^*\} \in \mathcal{S}$. By contradiction, suppose that this is not the case. If $x \notin A_0$ for $x \in S' \setminus S$, then $S' \setminus \{x\} \in \mathcal{S}$ as $S' \setminus \{x\}$ cannot separate G_0 . Similarly, if there is an edge $\{x, u\}$ in $G(\mathbf{A})$ with $x \in S' \setminus S$ and $u \notin S'$, then $S' \setminus \{x\} \in \mathcal{S}$ for the same reason as before. It follows that $x \in A_0$ and every neighbor of x lies inside S' for all $x \in S' \setminus S$. Let $B' \in \mathcal{B}$ be disjoint from S' . Then any path from a node in $S' \setminus S$ to a node in B' must intersect S . This contradicts the fact that $S \in \mathcal{S}$. Let $S^* := S' \setminus \{x^*\} \in \mathcal{S}$.

We have $|\mathcal{H}(S')^{S,h}| = \sum_{g \in \mathcal{H}(S^*)^{S,h}} |\mathcal{H}(S')^{S^*,g}|$. By the inductive hypothesis, we know that $|\mathcal{H}(S')^{S^*,g}| = |\mathcal{H}(S')|/|\mathcal{H}(S^*)|$, for every $g \in \mathcal{H}(S^*)$, and $|\mathcal{H}(S^*)^{S,h}| = |\mathcal{H}(S^*)|/|\mathcal{H}(S)|$. It follows that

$$|\mathcal{H}(S')^{S,h}| = (|\mathcal{H}(S')|/|\mathcal{H}(S^*)|)|\mathcal{H}(S^*)^{S,h}| = |\mathcal{H}(S')|/|\mathcal{H}(S)|. \quad \square$$

Now we are ready to define our vector λ satisfying (\dagger) . Fix $(f, \mathbf{x}) \in \text{tup}(\mathbb{A}_k)_{>0}$. $\text{Set}(\mathbf{x})$ is either a scope or a subset of size at most k of $G(\mathbf{A})$. In particular, $\text{Set}(\mathbf{x})$ cannot cover \mathcal{B} and thus $\text{Set}(\mathbf{x}) \in \mathcal{X}$. Let $\overline{\text{Set}(\mathbf{x})} := \bigcap_{S \in \mathcal{S} | \text{Set}(\mathbf{x}) \subseteq S} S$. Let us note that $\overline{\text{Set}(\mathbf{x})}$ is well-defined, i.e., there is $S^* \in \mathcal{S}$ such that $\text{Set}(\mathbf{x}) \subseteq S^*$. Indeed, we can take $S^* = A \setminus B \in \mathcal{S}$, where B is any set in \mathcal{B} disjoint from $\text{Set}(\mathbf{x})$. By Claim 5.17, $\overline{\text{Set}(\mathbf{x})} \in \mathcal{S}$. Observe then that $\overline{\text{Set}(\mathbf{x})}$ is the inclusionwise minimal set of \mathcal{S} containing $\text{Set}(\mathbf{x})$. For every mapping $s : \text{Set}(\mathbf{x}) \rightarrow B$, we define

$$\lambda(f, \mathbf{x}, s) = \frac{\Pr}{h \sim U(\mathcal{H}(\overline{\text{Set}(\mathbf{x})}))} [h|_{\text{Set}(\mathbf{x})} = s],$$

where $U(\mathcal{H}(\overline{\text{Set}(\mathbf{x})}))$ denotes the uniform distribution over $\mathcal{H}(\overline{\text{Set}(\mathbf{x})})$.

By Claim 5.16, we have that $(*)$ $\lambda(f, \mathbf{x}, s) > 0$ if and only if $s \in \mathcal{H}(\overline{\text{Set}(\mathbf{x})})$. Hence, λ satisfies (\dagger) . It remains to prove that λ is feasible for $\text{SA}_k(\mathbb{A}, \mathbb{B})$. Conditions (SA4) and (SA2) follow from the definition. For condition (SA3), recall first that the function $c^* : \text{tup}(\mathbb{A}) \rightarrow \mathbb{Q}_{\geq 0}$, whose existence is guaranteed by the fact that \mathbb{A} is a core, satisfies that

$$\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) c^*(f, \mathbf{x}) < \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) c^*(f, g(\mathbf{x}))$$

for every nonsurjective mapping $g : A \rightarrow A$. In particular, we have the following: $\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) c^*(f, \mathbf{x}) < \infty$. It follows that $f^{\mathbb{A}}(\mathbf{x}) = \infty$ implies $c^*(f, \mathbf{x}) = 0$ for every $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$. By contradiction, suppose that condition (SA3) does not hold. Then there is $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$ and $s : \text{Set}(\mathbf{x}) \rightarrow B$ such that $f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{B}}(s(\mathbf{x})) = \infty$, but $\lambda(f, \mathbf{x}, s) > 0$. By $(*)$, we know that $s \in \mathcal{H}(\overline{\text{Set}(\mathbf{x})})$. In particular, s is a partial homomorphism from \mathbf{A} to \mathbf{B} . On the other hand, note that $f^{\mathbb{B}}(s(\mathbf{x})) < \infty$ by construction of \mathbb{B} , and then $f^{\mathbb{A}}(\mathbf{x}) = \infty$. By a previous remark, we have that $c^*(f, \mathbf{x}) = 0$. Additionally, note that $\mathbf{x} \in R_f^{\mathbf{A}}$, and since s is a partial homomorphism, $s(\mathbf{x}) \in R_f^{\mathbf{B}}$. By definition of \mathbb{B} , it follows that $f^{\mathbb{B}}(s(\mathbf{x})) = c^*(f, \pi(s(\mathbf{x}))) = c^*(f, \mathbf{x}) = 0$. Hence, $f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{B}}(s(\mathbf{x})) = 0$, a contradiction.

For condition (SA1), let $(f, \mathbf{x}), (p, \mathbf{y}) \in \text{tup}(\mathbb{A}_k)_{>0}$ and define $X := \text{Set}(\mathbf{x})$ and $Y := \text{Set}(\mathbf{y})$. Suppose that $|X| \leq k$ and $X \subseteq Y$. Let $s : X \rightarrow B$. We need to show that

$$\lambda(f, \mathbf{x}, s) = \sum_{\substack{r: Y \rightarrow B \\ r|_X = s}} \lambda(p, \mathbf{y}, r) = \sum_{r \in \mathcal{H}(Y)^{X, s}} \lambda(p, \mathbf{y}, r),$$

where the last equality holds due to (*). Using again (*) and Claim 5.16, item 2, the required equality holds directly when $s \notin \mathcal{H}(X)$. Then we assume that $s \in \mathcal{H}(X)$. For any $Z \in \mathcal{X}$ such that $X \subseteq Z \subseteq \bar{Y}$, we have that

$$(5.5) \quad \begin{aligned} \Pr_{h \sim U(\mathcal{H}(\bar{Y}))} [h|_X = s] &= \sum_{\substack{r: Z \rightarrow B \\ r|_X = s}} \Pr_{h \sim U(\mathcal{H}(\bar{Y}))} [h|_X = s, h|_Z = r] \\ &= \sum_{\substack{r: Z \rightarrow B \\ r|_X = s}} \Pr_{h \sim U(\mathcal{H}(\bar{Y}))} [h|_Z = r]. \end{aligned}$$

By applying (5.5) with $Z = \bar{X}$, we obtain

$$(5.6) \quad \Pr_{h \sim U(\mathcal{H}(\bar{Y}))} [h|_X = s] = \sum_{\substack{r: \bar{X} \rightarrow B \\ r|_X = s}} \Pr_{h \sim U(\mathcal{H}(\bar{Y}))} [h|_{\bar{X}} = r] = \sum_{r \in \mathcal{H}(\bar{X})^{X,s}} \Pr_{h \sim U(\mathcal{H}(\bar{Y}))} [h|_{\bar{X}} = r],$$

where the last equality holds due to Claim 5.16, item 2. By definition, we have $\Pr_{h \sim U(\mathcal{H}(\bar{Y}))} [h|_{\bar{X}} = r] = |\mathcal{H}(\bar{Y})^{\bar{X},r}| / |\mathcal{H}(\bar{Y})|$. Since $r \in \mathcal{H}(\bar{X})$, we can apply Claim 5.18 and obtain that $\Pr_{h \sim U(\mathcal{H}(\bar{Y}))} [h|_{\bar{X}} = r] = 1/|\mathcal{H}(\bar{X})|$. Using this in (5.6), we obtain

$$\Pr_{h \sim U(\mathcal{H}(\bar{Y}))} [h|_X = s] = \sum_{r \in \mathcal{H}(\bar{X})^{X,s}} 1/|\mathcal{H}(\bar{X})| = \Pr_{h \sim U(\mathcal{H}(\bar{X}))} [h|_X = s] = \lambda(f, \mathbf{x}, s).$$

Hence,

$$\begin{aligned} \lambda(f, \mathbf{x}, s) &= \Pr_{h \sim U(\mathcal{H}(\bar{Y}))} [h|_X = s] \\ &= \sum_{\substack{r: Y \rightarrow B \\ r|_X = s}} \Pr_{h \sim U(\mathcal{H}(\bar{Y}))} [h|_Y = r] \quad (\text{By (5.5) with } Z = Y) \\ &= \sum_{\substack{r: Y \rightarrow B \\ r|_X = s}} \lambda(p, \mathbf{y}, r). \end{aligned}$$

Therefore, λ is a feasible solution of $\text{SA}_k(\mathbb{A}, \mathbb{B})$.

Finally, we remark that, as in Proposition 4.6, we can define \mathbb{B} to be finite-valued by replacing ∞ with a sufficiently large N (taking $N = 1 + (M^* / \min\{f^{\mathbb{A}}(\mathbf{x}) : (f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{>0}\})$ will do). This finishes the proof of Theorem 5.12. \square

THEOREM 5.19. *Let \mathbb{A} be a valued σ -structure and let $k \geq 1$. Let \mathbb{A}' be the core of \mathbb{A} . If $\text{tw}_{ms}(\mathbb{A}') \geq k$, then the Sherali–Adams relaxation of level k is not always tight for \mathbb{A} .*

Proof. We can apply Theorem 5.12 to \mathbb{A}' and obtain \mathbb{B} such that $\text{opt}_k^{\text{SA}}(\mathbb{A}', \mathbb{B}) < \text{opt}(\mathbb{A}', \mathbb{B})$. Since $\mathbb{A}' \equiv \mathbb{A}$, we know that $\text{opt}(\mathbb{A}, \mathbb{B}) = \text{opt}(\mathbb{A}', \mathbb{B})$ (by the definition of equivalence and cores) and $\text{opt}_k^{\text{SA}}(\mathbb{A}, \mathbb{B}) = \text{opt}_k^{\text{SA}}(\mathbb{A}', \mathbb{B})$ (by Proposition 5.3). Hence, $\text{opt}_k^{\text{SA}}(\mathbb{A}, \mathbb{B}) < \text{opt}(\mathbb{A}, \mathbb{B})$, and the result follows. \square

5.5. Necessity of bounded overlap. In this section we provide the last missing piece in the proof of Theorem 5.4.

THEOREM 5.20. *Let \mathbb{A} be a valued σ -structure and let $k \geq 1$. Suppose that \mathbb{A} is a core and that the overlap of \mathbb{A} is at least $k + 1$. Then there exists a valued σ -structure \mathbb{B} such that $\text{opt}_k^{\text{SA}}(\mathbb{A}, \mathbb{B}) < \text{opt}(\mathbb{A}, \mathbb{B})$.*

Proof. Given a tuple \mathbf{t} and an index $i \in \{1, \dots, |\mathbf{t}|\}$, we denote by $\mathbf{t}[i]$ the i th entry of \mathbf{t} . By extension, given a set I of indices we use $\mathbf{t}[I]$ to denote the tuple obtained from \mathbf{t} after discarding all entries $\mathbf{t}[j]$ with $j \notin I$. Given a tuple \mathbf{t} , we use $\|\mathbf{t}\|$ to denote the number of distinct elements appearing in \mathbf{t} , i.e., $\|\mathbf{t}\| := |\text{Set}(\mathbf{t})|$.

Because the overlap of \mathbb{A} is at least $k+1$, there exist $(p, \mathbf{x}), (q, \mathbf{y}) \in \text{tup}(\mathbb{A})$ such that

- (i) $p^{\mathbb{A}}(\mathbf{x}), q^{\mathbb{A}}(\mathbf{y}) > 0$,
- (ii) $|\text{Set}(\mathbf{x}) \cap \text{Set}(\mathbf{y})| > k$, and
- (iii) $(p, \mathbf{x}) \neq (q, \mathbf{y})$.

Let I_x, I_y be sets of indices of size $n = |\text{Set}(\mathbf{x}) \cap \text{Set}(\mathbf{y})|$ such that $\text{Set}(\mathbf{x}[I_x]) = \text{Set}(\mathbf{y}[I_y]) = \text{Set}(\mathbf{x}) \cap \text{Set}(\mathbf{y})$. (There may be more than one possible choice for I_x and I_y .) We first define a relational structure \mathbf{B} over the signature $\text{rel}(\sigma) = \{R_f \mid f \in \sigma, \text{ar}(R_f) = \text{ar}(f)\}$ and universe $B = A \times \{0, 1\}$. Let $\pi : B \rightarrow A$ and $\eta : B \rightarrow \{0, 1\}$ be the projections onto the first and second coordinates, respectively. For every symbol $R_f \in \text{rel}(\sigma)$, we add to $R_f^{\mathbf{B}}$ every tuple $\mathbf{t} \in B^{\text{ar}(R_f)}$ such that $\pi(\mathbf{t}) \in R_f^{\text{Pos}(\mathbb{A})}$. Then, we remove from $R_p^{\mathbf{B}}$ all the tuples \mathbf{t} such that $\pi(\mathbf{t}) = \mathbf{x}$ and $\bigoplus_{i \in I_x} \eta(\mathbf{t}[i]) = 1$ and we remove from $R_q^{\mathbf{B}}$ all the tuples \mathbf{t} such that $\pi(\mathbf{t}) = \mathbf{y}$ and $\bigoplus_{i \in I_y} \eta(\mathbf{t}[i]) = 0$. This completes the definition of \mathbf{B} .

By construction, π is a homomorphism from \mathbf{B} to $\text{Pos}(\mathbb{A})$. As the next claim shows, whenever h is a homomorphism from $\text{Pos}(\mathbb{A})$ to \mathbf{B} , the mapping $\pi \circ h$ cannot be surjective.

CLAIM 5.21. *If $h : A \rightarrow B$ is a homomorphism from $\text{Pos}(\mathbb{A})$ to \mathbf{B} , then $(\pi \circ h)(A) \neq A$.*

Proof. Toward a contradiction, suppose that $(\pi \circ h)(A) = A$. In that case, the map $\pi \circ h$ is bijective and has an inverse l . The map l is an isomorphism from $\text{Pos}(\mathbb{A})$ to itself, so in particular we have that $l(\mathbf{x}) \in R_p^{\text{Pos}(\mathbb{A})}$ and $l(\mathbf{y}) \in R_q^{\text{Pos}(\mathbb{A})}$. It follows that $h(l(\mathbf{x})) \in R_p^{\mathbf{B}}$, $h(l(\mathbf{y})) \in R_q^{\mathbf{B}}$, $\pi(h(l(\mathbf{x}))) = \mathbf{x}$ and $\pi(h(l(\mathbf{y}))) = \mathbf{y}$. Now, if we define

$$b_S := \bigoplus_{v \in \text{Set}(\mathbf{x}) \cap \text{Set}(\mathbf{y})} \eta(h(l(v)))$$

then by construction of \mathbf{B} we have both $b_S = \bigoplus_{i \in I_x} \eta(h(l(\mathbf{x}))[i]) = 0$ and also $b_S = \bigoplus_{i \in I_y} \eta(h(l(\mathbf{y}))[i]) = 1$, a contradiction. \square

Now, since \mathbb{A} is a core, by Proposition 3.13 there exists a function $c^* : \text{tup}(\mathbb{A}) \rightarrow \mathbb{Q}_{\geq 0}$ such that for every nonsurjective mapping $g : A \rightarrow A$,

$$M^* := \sum_{(f, \mathbf{z}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{z}) c^*(f, \mathbf{z}) < \sum_{(f, \mathbf{z}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{z}) c^*(f, g(\mathbf{z})).$$

In particular, M^* is finite so $f^{\mathbb{A}}(\mathbf{z}) = \infty$ implies $c^*(f, \mathbf{z}) = 0$ for all $(f, \mathbf{z}) \in \text{tup}(\mathbb{A})$. Then, we define a valued σ -structure \mathbb{B} over the universe B such that for all $f \in \sigma$ and $\mathbf{t} \in B^{\text{ar}(f)}$,

$$f^{\mathbb{B}}(\mathbf{t}) := \begin{cases} c^*(f, \pi(\mathbf{t})) & \text{if } \mathbf{t} \in R_f^{\mathbf{B}}, \\ \infty & \text{otherwise.} \end{cases}$$

CLAIM 5.22. $\text{opt}(\mathbb{A}, \mathbb{B}) > M^*$.

Proof. Let h be a homomorphism from $\text{Pos}(\mathbb{A})$ to \mathbf{B} (if it is not, then $\text{cost}_{\mathbb{A} \rightarrow \mathbb{B}}(h) = \infty$). By Claim 5.21, we have $(\pi \circ h)(A) \neq A$, i.e., $\pi \circ h$ is not surjective. By definition of c^* , we have

$$\text{cost}_{\mathbb{A} \rightarrow \mathbb{B}}(h) = \sum_{(f, \mathbf{z}) \in \text{tup}(\mathbb{A})_{>0}} f^{\mathbb{A}}(\mathbf{z}) f^{\mathbb{B}}(h(\mathbf{z})) = \sum_{(f, \mathbf{z}) \in \text{tup}(\mathbb{A})_{>0}} f^{\mathbb{A}}(\mathbf{z}) c^*(f, \pi(h(\mathbf{z}))) > M^*$$

and the claim follows. \square

The next step is to exhibit a solution to $\text{SA}_k(\mathbb{A}, \mathbb{B})$ of cost exactly M^* . Recall that the signature σ_k of the modified instance $(\mathbb{A}_k, \mathbb{B}_k)$ used to define $\text{SA}_k(\mathbb{A}, \mathbb{B})$ contains an additional function symbol ρ_k of arity k .

CLAIM 5.23. $\text{opt}_k^{\text{SA}}(\mathbb{A}, \mathbb{B}) \leq M^*$.

Proof. Let $(f, \mathbf{z}) \in \text{tup}(\mathbb{A}_k)_{>0}$ and $s : \text{Set}(\mathbf{z}) \rightarrow B_k$. If $\pi(s(\mathbf{z})) \neq \mathbf{z}$ then we set $\lambda(f, \mathbf{z}, s) := 0$. If instead $\pi(s(\mathbf{z})) = \mathbf{z}$, then we set $\lambda(f, \mathbf{z}, s)$ according to the following rules:

- If $(p, \mathbf{x}) \neq (f, \mathbf{z}) \neq (q, \mathbf{y})$, then $\lambda(f, \mathbf{z}, s) := 1/2^{\|\mathbf{z}\|}$.
- If $(f, \mathbf{z}) = (p, \mathbf{x})$, then $\lambda(p, \mathbf{x}, s) := 1/2^{\|\mathbf{x}\|-1}$ if $\bigoplus_{i \in I_x} \eta(s(\mathbf{x}))[i] = 0$ and $\lambda(p, \mathbf{x}, s) := 0$ otherwise.
- If $(f, \mathbf{z}) = (q, \mathbf{y})$, then $\lambda(q, \mathbf{y}, s) := 1/2^{\|\mathbf{y}\|-1}$ if $\bigoplus_{i \in I_y} \eta(s(\mathbf{y}))[i] = 1$ and $\lambda(q, \mathbf{y}, s) := 0$ otherwise.

Because $\lambda(f, \mathbf{z}, s)$ may only be nonzero if $\pi(s(\mathbf{z})) = \mathbf{z}$ and $f^{\mathbb{A}}(\mathbf{z}) = \infty$ implies $f^{\mathbb{B}}(s(\mathbf{z})) = c^*(f, \pi(s(\mathbf{z}))) = c^*(f, \mathbf{z}) = 0$ for all such (s, \mathbf{z}) , λ satisfies the condition (SA3). By construction, λ satisfies the conditions (SA2) and (SA4) as well.

For any $(f, \mathbf{z}) \in \text{tup}(\mathbb{A}_k)$ we use $\text{id}(f^{\mathbb{B}}, \mathbf{z})$ to denote the set of all assignments $r : \text{Set}(\mathbf{z}) \rightarrow B$ such that $\pi(r(\mathbf{z})) = \mathbf{z}$ and either $f = \rho_k$ or $r(\mathbf{z}) \in R_f^{\mathbb{B}}$. Then, the cost of λ is

$$\begin{aligned} & \sum_{\substack{(f, \mathbf{z}) \in \text{tup}(\mathbb{A}_k)_{>0}, s : \text{Set}(\mathbf{z}) \rightarrow B_k, \\ f^{\mathbb{A}_k}(\mathbf{z}) \times f^{\mathbb{B}_k}(s(\mathbf{z})) < \infty}} \lambda(f, \mathbf{z}, s) f^{\mathbb{A}_k}(\mathbf{z}) f^{\mathbb{B}_k}(s(\mathbf{z})) \\ &= \sum_{(f, \mathbf{z}) \in \text{tup}(\mathbb{A}_k)_{>0}, s \in \text{id}(f^{\mathbb{B}}, \mathbf{z})} \lambda(f, \mathbf{z}, s) f^{\mathbb{A}_k}(\mathbf{z}) f^{\mathbb{B}_k}(s(\mathbf{z})) \\ &= \sum_{(f, \mathbf{z}) \in \text{tup}(\mathbb{A})_{>0}, s \in \text{id}(f^{\mathbb{B}}, \mathbf{z})} \lambda(f, \mathbf{z}, s) f^{\mathbb{A}}(\mathbf{z}) f^{\mathbb{B}}(s(\mathbf{z})) \\ &= \sum_{(f, \mathbf{z}) \in \text{tup}(\mathbb{A})_{>0}, s \in \text{id}(f^{\mathbb{B}}, \mathbf{z})} \lambda(f, \mathbf{z}, s) f^{\mathbb{A}}(\mathbf{z}) c^*(f, \pi(s(\mathbf{z}))) \\ &= \sum_{(f, \mathbf{z}) \in \text{tup}(\mathbb{A})_{>0}, s \in \text{id}(f^{\mathbb{B}}, \mathbf{z})} \lambda(f, \mathbf{z}, s) f^{\mathbb{A}}(\mathbf{z}) c^*(f, \mathbf{z}) \\ &= \sum_{(f, \mathbf{z}) \in \text{tup}(\mathbb{A})_{>0}} f^{\mathbb{A}}(\mathbf{z}) c^*(f, \mathbf{z}) \\ &= M^*. \end{aligned}$$

At this point, to prove that λ is indeed a solution to $\text{SA}_k(\mathbb{A}, \mathbb{B})$ of cost M^* we need only show that it satisfies the condition (SA1). Let $(f, \mathbf{z}), (g, \mathbf{w}) \in \text{tup}(\mathbb{A}_k)_{>0}$ such that $\text{Set}(\mathbf{z}) \subseteq \text{Set}(\mathbf{w})$ and $\|\mathbf{z}\| \leq k$. Let $s : \text{Set}(\mathbf{z}) \rightarrow B_k$. If $\pi(s(\mathbf{z})) \neq \mathbf{z}$ then all the lambdas involved in the equation (SA1) for the triple $((f, \mathbf{z}), (g, \mathbf{w}), s)$ are zero, so in this case the condition holds. Now, assume that $\pi(s(\mathbf{z})) = \mathbf{z}$ and observe that $(p, \mathbf{x}) \neq (f, \mathbf{z}) \neq (q, \mathbf{y})$ because $\|\mathbf{x}\|, \|\mathbf{y}\| > k$. We consider two cases.

If $(p, \mathbf{x}) \neq (g, \mathbf{w}) \neq (q, \mathbf{y})$ then there are exactly $2^{\|\mathbf{w}\| - \|\mathbf{z}\|}$ mappings r from $\text{Set}(\mathbf{w})$ to B_k such that $\pi(r(w)) = w$ for all $w \in \text{Set}(\mathbf{w})$ and $\eta(r(z)) = \eta(s(z))$ for all $z \in \text{Set}(\mathbf{z})$.

Thus we have

$$\lambda(f, \mathbf{z}, s) = 1/2^{\|\mathbf{z}\|} = 2^{\|\mathbf{w}\| - \|\mathbf{z}\|} 1/2^{\|\mathbf{w}\|} = \sum_{r: \text{Set}(\mathbf{w}) \rightarrow B_k, r|_{\text{Set}(\mathbf{z})=s} \lambda(g, \mathbf{w}, r)$$

and (SA1) holds for the triple $((f, \mathbf{z}), (g, \mathbf{w}), s)$.

In the second case, we have either $(g, \mathbf{w}) = (p, \mathbf{x})$ or $(g, \mathbf{w}) = (q, \mathbf{y})$. Let $Z' := \text{Set}(\mathbf{z}) \cap \text{Set}(\mathbf{x}) \cap \text{Set}(\mathbf{y})$ and $Z'' := \text{Set}(\mathbf{z}) \setminus Z'$. Since Z' is a strict subset of $\text{Set}(\mathbf{x}) \cap \text{Set}(\mathbf{y})$, the mapping s has exactly $2^{|\text{Set}(\mathbf{x}) \cap \text{Set}(\mathbf{y})| - |Z'| - 1}$ distinct extensions $r : \text{Set}(\mathbf{z}) \cup (\text{Set}(\mathbf{x}) \cap \text{Set}(\mathbf{y})) \rightarrow B_k$ such that $\pi(r(w)) = w$ for all $w \in \text{Set}(\mathbf{z}) \cup (\text{Set}(\mathbf{x}) \cap \text{Set}(\mathbf{y}))$ and $\bigoplus_{v \in \text{Set}(\mathbf{x}) \cap \text{Set}(\mathbf{y})} \eta(r(v)) = 0$ (if $(g, \mathbf{w}) = (p, \mathbf{x})$) or $\bigoplus_{v \in \text{Set}(\mathbf{x}) \cap \text{Set}(\mathbf{y})} \eta(r(v)) = 1$ (if $(g, \mathbf{w}) = (q, \mathbf{y})$). In turn, each of these mappings r has $2^{\|\mathbf{w}\| - |\text{Set}(\mathbf{x}) \cap \text{Set}(\mathbf{y})| - |Z''|}$ distinct extensions defined on the whole of $\text{Set}(\mathbf{w})$ that satisfy $\pi(r(w)) = w$ for all $w \in \text{Set}(\mathbf{w})$. Putting everything together, s has $2^{\|\mathbf{w}\| - (|Z'| + |Z''|) - 1} = 2^{\|\mathbf{w}\| - \|\mathbf{z}\| - 1}$ extensions to $\text{Set}(\mathbf{w})$ with the correct bit sum. Thus we have

$$\lambda(f, \mathbf{z}, s) = 1/2^{\|\mathbf{z}\|} = 2^{\|\mathbf{w}\| - \|\mathbf{z}\| - 1} 1/2^{\|\mathbf{w}\| - 1} = \sum_{r: \text{Set}(\mathbf{w}) \rightarrow B_k, r|_{\text{Set}(\mathbf{z})=s} \lambda(g, \mathbf{w}, r)$$

and (SA1) holds in this last case as well, which concludes the proof of the claim. \square

The proof of Theorem 5.20 is now established by Claims 5.22 and 5.23. As in Theorem 5.12, we can make \mathbb{B} to be finite-valued by replacing ∞ with a sufficiently large number N . \square

THEOREM 5.24. *Let \mathbb{A} be a valued σ -structure and let $k \geq 1$. Let \mathbb{A}' be the core of \mathbb{A} . If the overlap of \mathbb{A}' is at least $k + 1$, then the Sherali–Adams relaxation of level k is not always tight for \mathbb{A} .*

Proof. We can apply Theorem 5.20 to \mathbb{A}' and obtain \mathbb{B} such that $\text{opt}_k^{\text{SA}}(\mathbb{A}', \mathbb{B}) < \text{opt}(\mathbb{A}', \mathbb{B})$. Since $\mathbb{A}' \equiv \mathbb{A}$, we know that $\text{opt}(\mathbb{A}, \mathbb{B}) = \text{opt}(\mathbb{A}', \mathbb{B})$ (by the definition of equivalence and cores) and $\text{opt}_k^{\text{SA}}(\mathbb{A}, \mathbb{B}) = \text{opt}_k^{\text{SA}}(\mathbb{A}', \mathbb{B})$ (by Proposition 5.3). Hence, $\text{opt}_k^{\text{SA}}(\mathbb{A}, \mathbb{B}) < \text{opt}(\mathbb{A}, \mathbb{B})$, and the result follows. \square

6. Search VCSP(\mathcal{C} , $-$). If a class \mathcal{C} of valued structures has bounded treewidth modulo equivalence then the Sherali–Adams LP hierarchy can be used to solve in polynomial time $\text{VCSP}(\mathcal{C}, -)$, that is, to compute the minimum cost of a mapping from $\mathbb{A} \in \mathcal{C}$ to some arbitrary valued structure \mathbb{B} . However, it may be the case that computing a mapping of that cost is NP-hard even though we know that one exists. In this section we will focus on the search version of the VCSP, which explicitly asks for a minimum-cost mapping and will be denoted by SVCSP .

If \mathcal{C} is a class of valued structures, we denote by $\text{CORE COMPUTATION}(\mathcal{C})$ the problem that takes as input some $\mathbb{A} \in \mathcal{C}$ and asks us to compute a mapping $g : A \rightarrow A$ such that $g(\mathbb{A})$ is a core of \mathbb{A} and there exists an IFH ω from \mathbb{A} to \mathbb{A} such that $g \in \text{supp}(\omega)$.

Building on our results from section 5 and adapting techniques from [50], we will prove our third main result.

THEOREM 6.1 (search classification). *Assume $\text{FPT} \neq \text{W}[1]$. Let \mathcal{C} be a recursively enumerable class of valued structures of bounded arity. Then, the following are equivalent:*

1. $\text{SVCSP}(\mathcal{C}, -)$ is in PTIME .
2. \mathcal{C} is of bounded treewidth modulo equivalence and $\text{CORE COMPUTATION}(\mathcal{C})$ is in PTIME .

Remark 6.2. Given a class \mathcal{C} of relational structures, let $\text{SCSP}(\mathcal{C}, -)$ denote the search variant of $\text{CSP}(\mathcal{C}, -)$; i.e., given \mathbf{A} and \mathbf{B} with $\mathbf{A} \in \mathcal{C}$, the task is to return a homomorphism from \mathbf{A} to \mathbf{B} if one exists. When applied to (bounded-arity, recursively enumerable) classes of relational structures, Theorem 6.1 states that $\text{SCSP}(\mathcal{C}, -)$ is in PTIME if and only if \mathcal{C} has bounded treewidth modulo homomorphic equivalence and computing a homomorphism from any given $\mathbf{A} \in \mathcal{C}$ to one of its cores is in PTIME. This result is folklore and can be easily derived from Grohe [32] and Dalmau, Kolaitis, and Vardi [16]. We provide here a brief sketch of the argument since our proof of Theorem 6.1 for classes of valued structures follows roughly the same strategy, although the technical details are significantly more involved.

Let \mathcal{C} be a class of relational structures and \mathcal{C}' denote the class of all cores of structures in \mathcal{C} . If computing a homomorphism g from any given $\mathbf{A} \in \mathcal{C}$ to one of its cores is in PTIME, then $\text{SCSP}(\mathcal{C}, -)$ is polynomial-time reducible to $\text{SCSP}(\mathcal{C}', -)$ by simply replacing each instance (\mathbf{A}, \mathbf{B}) of $\text{SCSP}(\mathcal{C}, -)$ with $(g(\mathbf{A}), \mathbf{B})$. (If a homomorphism h from $g(\mathbf{A})$ to \mathbf{B} is returned, then $h \circ g$ is a homomorphism from \mathbf{A} to \mathbf{B} . Otherwise, no homomorphism from \mathbf{A} to \mathbf{B} exists.) Furthermore, if \mathcal{C} has bounded treewidth modulo homomorphic equivalence then $\text{SCSP}(\mathcal{C}', -)$ is in PTIME; one can, for example, compute an optimal tree decomposition of any $\mathbf{A}' \in \mathcal{C}'$ in linear time (using Bodlaender's algorithm [6]) and then use standard dynamic programming techniques. It follows that $\text{SCSP}(\mathcal{C}, -)$ is in PTIME for any class \mathcal{C} satisfying both conditions. For the converse implication, if $\text{SCSP}(\mathcal{C}, -)$ is in PTIME then, by Grohe's result [32] (and under our assumptions), \mathcal{C} is of bounded treewidth modulo homomorphic equivalence. It only remains to prove that computing a homomorphism from $\mathbf{A} \in \mathcal{C}$ to one of its cores is in PTIME; for this task we will use a simple argument from Chen and Mengel [14]. Observe that a relational structure \mathbf{C} is not a core if and only if there exists a relational structure \mathbf{C}_a obtained by removing one element a from the universe of \mathbf{C} (as well as all tuples containing a) that is homomorphically equivalent to \mathbf{C} . It follows that an algorithm that starts with the pair of structures (\mathbf{A}, \mathbf{A}) and greedily removes elements from the universe of the right-hand-side structure while maintaining homomorphic equivalence will eventually terminate with $(\mathbf{A}, \mathbf{A}')$, where \mathbf{A}' is the core of \mathbf{A} . Recall that $\mathbf{A} \in \mathcal{C}$ and \mathcal{C} has bounded treewidth modulo homomorphic equivalence. Thus the homomorphism tests required by the algorithm outlined above can be done in polynomial time [16]. It then suffices to run the assumed algorithm for $\text{SCSP}(\mathcal{C}, -)$ on the pair $(\mathbf{A}, \mathbf{A}')$ to compute a homomorphism from \mathbf{A} to one of its cores.

If \mathcal{C} is a class of valued structures, the problem $\text{REDUCTION STEP}(\mathcal{C})$ takes as input some $\mathbb{A} \in \mathcal{C}$ and a mapping $g : A \rightarrow A$ that belongs to the support of some IFH from \mathbb{A} to \mathbb{A} . The goal is to compute a mapping $g^+ : A \rightarrow A$ such that $g^+(A) \subsetneq g(A)$ and g^+ belongs to the support of some IFH from \mathbb{A} to \mathbb{A} , or assert that no such mapping exist. The relevance of this problem to CORE COMPUTATION is highlighted by the following proposition.

PROPOSITION 6.3. *Let \mathbb{A} be a valued structure and $g : A \rightarrow A$ be a mapping that belongs to the support of some IFH from \mathbb{A} to \mathbb{A} . Then, $g(\mathbb{A})$ is not the core of \mathbb{A} if and only if there exists a mapping $g^+ : A \rightarrow A$ such that $g^+(A) \subsetneq g(A)$ and g^+ belongs to the support of some IFH from \mathbb{A} to \mathbb{A} .*

Proof. If there exists such a mapping g^+ then by Proposition A.1(3) in Appendix A we have $\mathbb{A} \equiv g(\mathbb{A}) \equiv g^+(\mathbb{A})$. Then, the cores of $g(\mathbb{A})$ and $g^+(\mathbb{A})$ are equivalent and by Proposition 3.10 they are isomorphic. By Proposition 3.12 the universe of the core of $g^+(\mathbb{A})$ has size at most $|g^+(A)| < |g(A)|$, so $g(\mathbb{A})$ is not a core.

For the converse implication, suppose that $g(\mathbb{A})$ is not a core. Let ω be an IFH from \mathbb{A} to \mathbb{A} such that $g \in \text{supp}(\omega)$. By the definition of a core, there exists a nonsurjective IFH from $g(\mathbb{A})$ to $g(\mathbb{A})$. Let g^* be a mapping in its support such that $g^*(g(\mathbb{A})) \subsetneq g(\mathbb{A})$. By Proposition A.1(3), we have $\mathbb{A} \equiv g(\mathbb{A}) \equiv g^*(g(\mathbb{A}))$. Then, by Proposition 3.6 there exist an IFH ω_1 from \mathbb{A} to $g^*(g(\mathbb{A}))$ and an IFH ω_2 from $g^*(g(\mathbb{A}))$ to \mathbb{A} . Let g_1, g_2 be arbitrary mappings in the support of ω_1, ω_2 , respectively. If we define $\omega^+ := \omega \circ \omega_2 \circ \omega_1$, where

$$\omega^+(h) = (\omega \circ \omega_2 \circ \omega_1)(h) = \sum_{\substack{h_1: A \rightarrow g^*(g(\mathbb{A})), \\ h_2: g^*(g(\mathbb{A})) \rightarrow A, \\ h_3: A \rightarrow A: \\ h_3 \circ h_2 \circ h_1 = h}} \omega(h_3) \omega_2(h_2) \omega_1(h_1)$$

then ω^+ is an IFH from \mathbb{A} to \mathbb{A} . Let $g^+ := g \circ g_2 \circ g_1$. Because $|g^+(A)| \leq |g_1(A)| \leq |g^*(g(\mathbb{A}))| < |g(\mathbb{A})|$, we have $g^+(A) \subsetneq g(\mathbb{A})$. Moreover,

$$\omega^+(g^+) \geq \omega(g) \omega_2(g_2) \omega_1(g_1) > 0$$

so g^+ belongs to the support of at least one IFH from \mathbb{A} to \mathbb{A} , which concludes the proof. \square

LEMMA 6.4. *Let \mathcal{C} be a class of valued structures. If REDUCTION STEP(\mathcal{C}) is in PTIME, then CORE COMPUTATION(\mathcal{C}) is in PTIME.*

Proof. Suppose that REDUCTION STEP(\mathcal{C}) is in PTIME, and let $\mathbb{A} \in \mathcal{C}$. We initialize a variable $g : A \rightarrow A$ to the identity mapping on A and invoke the polynomial-time algorithm R for REDUCTION STEP(\mathcal{C}) on input (\mathbb{A}, g) . If R asserts that no mapping $g^+ : A \rightarrow A$ such that $g^+(A) \subsetneq g(A)$ and g^+ belongs to the support of some IFH from \mathbb{A} to \mathbb{A} exists, then from Proposition 6.3 we deduce that $g(\mathbb{A})$ is a core and we are done. Otherwise, we set $g := g^+$ and repeat the procedure until R finds that $g(\mathbb{A})$ is a core (via Proposition 6.3). In this case, by Proposition A.1(3) from Appendix A, it holds that $g(\mathbb{A}) \equiv \mathbb{A}$, so $g(\mathbb{A})$ is the core of \mathbb{A} and we return g . The procedure terminates after at most $|A|$ calls to R. \square

We start by proving the implication $1 \Rightarrow 2$ of Theorem 6.1. If SVCSP($\mathcal{C}, -$) is in PTIME, then, by Theorem 4.1 (and under the assumption that FPT \neq W[1]), \mathcal{C} is of bounded treewidth modulo equivalence; the nontrivial part is to show that CORE COMPUTATION(\mathcal{C}) is in PTIME. To achieve this, we adapt an algorithm from [50, Proposition 4.7] originally used to determine in polynomial time the complexity of core finite-valued constraint languages (here, “core” refers to the notion for right-hand-side valued structures, which differs from our own; see [50, Definition 2.6] for a precise definition). The central idea is to show that REDUCTION STEP(\mathcal{C}) can be solved by the *ellipsoid algorithm* using a separation oracle that makes polynomially many calls to the assumed polynomial-time algorithm for SVCSP($\mathcal{C}, -$).

Before we proceed with the main proof we need the following definitions and result from combinatorial optimization, as well as two minor technical lemmas.

DEFINITION 6.5 (see [33]). *Let $A \in \mathbb{Q}^{m \times n}$, $\bar{b} \in \mathbb{Q}^m$, and $P = \{\bar{x} \in \mathbb{Q}^n : A\bar{x} \leq \bar{b}\}$. A strong separation oracle for P is an algorithm that, given on input a vector $\bar{y} \in \mathbb{Q}^n$, either concludes that $\bar{y} \in P$ or computes a vector $\bar{a} \in \mathbb{Q}^n$ such that $\bar{a}^T \bar{y} > \bar{a}^T \bar{x}$ for all $\bar{x} \in P$.*

DEFINITION 6.6 (see [33]). *Let $A \in \mathbb{Q}^{m \times n}$, $\bar{b} \in \mathbb{Q}^m$, $P = \{\bar{x} \in \mathbb{Q}^n : A\bar{x} \leq \bar{b}\}$, $\bar{c} \in \mathbb{Q}^n$, and SEP be a strong separation oracle for P . A basic optimum dual solution with oracle inequalities is a set of inequalities $\bar{a}_1^T \bar{x} \leq \alpha_1, \dots, \bar{a}_k^T \bar{x} \leq \alpha_k$ valid for*

P , where $\bar{a}_1, \dots, \bar{a}_k$ are linearly independent outputs of SEP, and dual variables $\lambda_1, \dots, \lambda_k \in \mathbb{Q}_{\geq 0}$ such that $\lambda_1 \bar{a}_1 + \dots + \lambda_k \bar{a}_k = \bar{c}$ and $\lambda_1 \alpha_1 + \dots + \lambda_k \alpha_k = \max_{\bar{x} \in P} \bar{c}^T \bar{x}$.

LEMMA 6.7 (see [33, Lemma 6.5.15]). Let $A \in \mathbb{Q}^{m \times n}$, $\bar{b} \in \mathbb{Q}^m$, $P = \{\bar{x} \in \mathbb{Q}^n : A\bar{x} \leq \bar{b}\}$, and $\bar{c} \in \mathbb{Q}^n$. Suppose that the bit sizes of the coefficients of A and \bar{b} are bounded by ϕ . Given a strong separation oracle SEP for P where every output has encoding size at most ϕ , one can, in time polynomial in n , ϕ and the encoding size of \bar{c} , and with polynomially many oracle queries to SEP, either

- find a basic optimum dual solution with oracle inequalities, or
- assert that the dual problem is unbounded or has no solution.

LEMMA 6.8. There exists a polynomially computable function ε_Δ which maps any two valued σ -structures \mathbb{A}, \mathbb{B} to a positive rational number $\varepsilon_\Delta(\mathbb{A}, \mathbb{B})$ such that for any two mappings $h_1, h_2 : A \rightarrow B$ satisfying $\text{cost}(h_1) < \text{cost}(h_2) < \infty$, we have $\text{cost}(h_2) - \text{cost}(h_1) > \varepsilon_\Delta(\mathbb{A}, \mathbb{B})$.

Proof. We assume without loss of generality that every nonnegative rational number p/q is encoded as a sequence of two nonnegative integers p and q . Let $q_{\mathbb{A}}$ (resp., $q_{\mathbb{B}}$) denote the product of all denominators of the values $f^{\mathbb{A}}(\mathbf{x})$ with $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}$ (resp., values $f^{\mathbb{B}}(\mathbf{x})$ with $(f, \mathbf{x}) \in \text{tup}(\mathbb{B})_{<\infty}$). The cost of any mapping $h : A \rightarrow B$ is an integer multiple of $1/(q_{\mathbb{A}}q_{\mathbb{B}})$, so for any two mappings $h_1, h_2 : A \rightarrow B$ satisfying $\text{cost}(h_1) < \text{cost}(h_2) < \infty$ we have that $\text{cost}(h_2) - \text{cost}(h_1) \geq 1/(q_{\mathbb{A}}q_{\mathbb{B}})$. We can then choose $\varepsilon_\Delta(\mathbb{A}, \mathbb{B}) := 1/2 \cdot 1/(q_{\mathbb{A}}q_{\mathbb{B}})$, which is positive and polynomially computable. \square

LEMMA 6.9. There exists a polynomially computable function ε_Ω which maps any two valued σ -structures \mathbb{A}, \mathbb{B} to a positive rational number $\varepsilon_\Omega(\mathbb{A}, \mathbb{B})$ such that for any subset \mathcal{H}_S of B^A , the following statements are equivalent:

- There exists an IFH ω from \mathbb{A} to \mathbb{B} such that $\sum_{h \in \mathcal{H}_S} \omega(h) > 0$.
- There exists an IFH ω from \mathbb{A} to \mathbb{B} such that $\sum_{h \in \mathcal{H}_S} \omega(h) \geq \varepsilon_\Omega(\mathbb{A}, \mathbb{B})$.

Proof. Let ϕ be a nondecreasing polynomial such that every feasible linear program with encoding size n has an optimum value with encoding size at most $\phi(n)$. (Since linear programming is solvable in polynomial time, such a polynomial exists and depends on the encoding scheme chosen; for more details we refer the reader to [46].) Let f_{\min} be the polynomially computable function that maps each natural number $n \geq 2$ to the smallest positive rational number with encoding size at most n . We define the set

$$\mathcal{H}_{<\infty} := \{h \in B^A : f^{\mathbb{B}}(\mathbf{x}) < \infty \Rightarrow f^{\mathbb{A}}(h^{-1}(\mathbf{x})) < \infty, \text{ for all } (f, \mathbf{x}) \in \text{tup}(\mathbb{B})\}$$

and observe that the statement (i) is true if and only if the linear program

$$(6.1) \quad \begin{aligned} & \max \left(\sum_{h \in \mathcal{H}_S \cap \mathcal{H}_{<\infty}} \omega(h) \right) \\ & \sum_{h \in \mathcal{H}_{<\infty}} \omega(h) f^{\mathbb{A}}(h^{-1}(\mathbf{x})) \leq f^{\mathbb{B}}(\mathbf{x}) \quad \forall (f, \mathbf{x}) \in \text{tup}(\mathbb{B})_{<\infty}, \\ & \sum_{h \in \mathcal{H}_{<\infty}} \omega(h) = 1, \\ & \omega(h) \geq 0 \end{aligned} \quad \forall h \in \mathcal{H}_{<\infty},$$

is feasible and its optimum value is positive. This program has polynomially many inequalities; hence if it is feasible then there exists an optimum solution ω^* such

that $\text{supp}(\omega^*)$ has polynomial size. The restriction of the linear program (6.1) to the variables in $\text{supp}(\omega^*)$ has encoding size at most $p(|\mathbb{A}| + |\mathbb{B}|)$ for some polynomial p and has the same optimum value. Now, we define

$$\varepsilon_\Omega(\mathbb{A}, \mathbb{B}) := f_{\min}(\phi(p(|\mathbb{A}| + |\mathbb{B}|)))$$

and we observe that if the linear program (6.1) is feasible and its optimum value is positive, then it is at least $\varepsilon_\Omega(\mathbb{A}, \mathbb{B})$. This establishes the implication (i) \Rightarrow (ii) for the function ε_Ω . The implication (ii) \Rightarrow (i) is trivial and given \mathbb{A}, \mathbb{B} the function $\varepsilon_\Omega(\mathbb{A}, \mathbb{B})$ is polynomially computable, so the claim follows. \square

The following lemma is the main technical part in establishing $1 \Rightarrow 2$ in Theorem 6.1.

LEMMA 6.10. *Let \mathcal{C} be a class of valued structures. If $\text{SVCSP}(\mathcal{C}, -)$ is in PTIME then $\text{REDUCTION STEP}(\mathcal{C})$ is in PTIME as well.*

Proof. Let (\mathbb{A}, g) be an input to $\text{REDUCTION STEP}(\mathcal{C})$. We assume that $|g(A)| > 1$; otherwise the problem is trivial. If \mathbb{A} is $\{0, \infty\}$ -valued then we can solve the instance by following the argument described in Remark 6.2, so we also assume that $\max_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} (f^{\mathbb{A}}(\mathbf{x})) > 0$. We define the set

$$\mathcal{H}_{<\infty} := \{h \in A^A : f^{\mathbb{A}}(\mathbf{x}) < \infty \Rightarrow f^{\mathbb{A}}(h^{-1}(\mathbf{x})) < \infty \ \forall (f, \mathbf{x}) \in \text{tup}(\mathbb{A})\}$$

and we recall that for every IFH ω from \mathbb{A} to \mathbb{A} , every mapping $g' \in \text{supp}(\omega)$ belongs to $\mathcal{H}_{<\infty}$. We denote by \mathcal{H}^* the set of all mappings h in $\mathcal{H}_{<\infty}$ such that $h(A) \subsetneq g(A)$. Let us consider the following linear program:

$$(6.2) \quad \begin{aligned} & \min 0, \\ & \sum_{h \in \mathcal{H}_{<\infty}} \omega(h) f^{\mathbb{A}}(h^{-1}(\mathbf{x})) \leq f^{\mathbb{A}}(\mathbf{x}) \quad \forall (f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}, \\ & \sum_{h \in \mathcal{H}_{<\infty}} \omega(h) = 1, \\ & \sum_{h \in \mathcal{H}^*} \omega(h) \geq \varepsilon_\Omega(\mathbb{A}, \mathbb{A}), \\ & \omega(h) \geq 0 \quad \forall h \in \mathcal{H}_{<\infty}. \end{aligned}$$

By Lemma 6.9, this program is not feasible if and only if $g(\mathbb{A})$ is a core; otherwise a solution gives a mapping $g \in \mathcal{H}^*$ with the desired property. This program has exponentially many variables, and hence we will solve its dual instead:

$$(6.3) \quad \begin{aligned} & \max \delta_1 + \varepsilon_\Omega(\mathbb{A}, \mathbb{A}) \delta_2 \\ & \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) (f^{\mathbb{A}}(\mathbf{x}) - f^{\mathbb{A}}(h^{-1}(\mathbf{x}))) + \delta_1 \leq 0 \quad \forall h \in \mathcal{H}_{<\infty} \setminus \mathcal{H}^*, \\ & \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) (f^{\mathbb{A}}(\mathbf{x}) - f^{\mathbb{A}}(h^{-1}(\mathbf{x}))) + \delta_1 + \delta_2 \leq 0 \quad \forall h \in \mathcal{H}^*, \\ & z(f, \mathbf{x}) \geq 0 \quad \forall (f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}, \\ & \delta_2 \geq 0. \end{aligned}$$

We now describe a strong separation oracle for the associated polyhedron P . Given a vector $(z, \delta_1, \delta_2) \in \mathbb{Q}^{|\text{tup}(\mathbb{A})_{<\infty}| \times \mathbb{Q}^2}$, we first check if there exists $(f^*, \mathbf{x}^*) \in \text{tup}(\mathbb{A})_{<\infty}$

such that $z(f^*, \mathbf{x}^*) < 0$; if it is the case then $a(f^*, \mathbf{x}^*) = -1$ and 0 otherwise defines a separating hyperplane. Similarly, if $\delta_2 < 0$ then we can set $a(\delta_2) = -1$ and 0 otherwise to obtain a separation.

Now, let \mathbb{B}_0 be a valued σ -structure with universe $B_0 = A$ such that for all $(f, \mathbf{x}) \in \text{tup}(\mathbb{B}_0)$, $f^{\mathbb{B}_0}(\mathbf{x}) = z(f, \mathbf{x})$ if $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}$ and $f^{\mathbb{B}_0}(\mathbf{x}) = 0$ otherwise. We will be interested in computing a mapping $h \in B_0^A \cap \mathcal{H}_{<\infty}$ with minimum cost. However, invoking the assumed algorithm for $\text{SVCSP}(\mathcal{C}, -)$ on the instance $(\mathbb{A}, \mathbb{B}_0)$ is not sufficient for this task because the returned mapping might not belong to $\mathcal{H}_{<\infty}$. In order to solve this problem, we define

$$\varepsilon := \frac{\varepsilon_{\Delta}(\mathbb{A}, \mathbb{B}_0)}{|\text{tup}(\mathbb{A})| \cdot \max_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} (f^{\mathbb{A}}(\mathbf{x}))},$$

where the function ε_{Δ} is as in Lemma 6.8, and we let \mathbb{B} be the valued σ -structure with universe $B = A$ such that for all $(f, \mathbf{x}) \in \text{tup}(\mathbb{B})$, $f^{\mathbb{B}}(\mathbf{x}) = z(f, \mathbf{x}) + \varepsilon$ if $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}$ and $f^{\mathbb{B}}(\mathbf{x}) = 0$ otherwise. Because $\varepsilon > 0$, the set of finite-cost mappings $A \rightarrow B$ is precisely $\mathcal{H}_{<\infty}$ and cannot be empty because of the identity mapping. Since (\mathbb{A}, \mathbb{B}) is an instance of $\text{SVCSP}(\mathcal{C}, -)$, we can find a mapping h^* from A to B of minimum cost in polynomial time. Then, we have

$$\begin{aligned} h^* &\in \operatorname{argmin}_{h \in B^A} \left(\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{B})} f^{\mathbb{B}}(\mathbf{x}) f^{\mathbb{A}}(h^{-1}(\mathbf{x})) \right) \\ &= \operatorname{argmin}_{h \in \mathcal{H}_{<\infty}} \left(\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{B})} f^{\mathbb{B}}(\mathbf{x}) f^{\mathbb{A}}(h^{-1}(\mathbf{x})) \right) \\ &= \operatorname{argmin}_{h \in \mathcal{H}_{<\infty}} \left(\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) f^{\mathbb{A}}(h^{-1}(\mathbf{x})) \right. \\ &\quad \left. + \varepsilon \left(\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} f^{\mathbb{A}}(h^{-1}(\mathbf{x})) \right) \right). \end{aligned}$$

Recall that we are looking for a mapping in $B_0^A \cap \mathcal{H}_{<\infty}$ with minimum cost with respect to the instance $(\mathbb{A}, \mathbb{B}_0)$. The mapping h^* we have just computed realizes the minimum of a slightly different objective function, but the next claim shows that this is not an issue.

CLAIM 6.11. *Let \mathcal{H}_S be a nonempty subset of $\mathcal{H}_{<\infty}$ and $h' \in \mathcal{H}_S$. If*

$$\begin{aligned} h' &\in \operatorname{argmin}_{h \in \mathcal{H}_S} \left(\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) f^{\mathbb{A}}(h^{-1}(\mathbf{x})) \right. \\ &\quad \left. + \varepsilon \left(\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} f^{\mathbb{A}}(h^{-1}(\mathbf{x})) \right) \right), \end{aligned}$$

then

$$h' \in \operatorname{argmin}_{h \in \mathcal{H}_S} \left(\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) f^{\mathbb{A}}(h^{-1}(\mathbf{x})) \right).$$

Proof. We prove the statement by contraposition. Let $h' \in \mathcal{H}_S$ and suppose that there exists $h_1 \in \mathcal{H}_S$ such that

$$\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) f^{\mathbb{A}}(h_1^{-1}(\mathbf{x})) < \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) f^{\mathbb{A}}(h'^{-1}(\mathbf{x})).$$

These two quantities are the respective costs of h_1 and h' with respect to the instance $(\mathbb{A}, \mathbb{B}_0)$, and both are finite. By Lemma 6.8 we have

$$\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) f^{\mathbb{A}}(h_1^{-1}(\mathbf{x})) + \varepsilon_{\Delta}(\mathbb{A}, \mathbb{B}_0) < \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) f^{\mathbb{A}}(h'^{-1}(\mathbf{x}))$$

and since

$$\frac{\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} f^{\mathbb{A}}(h_1^{-1}(\mathbf{x}))}{|\text{tup}(\mathbb{A})| \cdot \max_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} (f^{\mathbb{A}}(\mathbf{x}))} \leq 1$$

it follows from the definition of ε that

$$\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) f^{\mathbb{A}}(h_1^{-1}(\mathbf{x})) + \varepsilon \left(\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} f^{\mathbb{A}}(h_1^{-1}(\mathbf{x})) \right)$$

is strictly smaller than $\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) f^{\mathbb{A}}(h'^{-1}(\mathbf{x}))$ and a fortiori strictly smaller than $\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) f^{\mathbb{A}}(h'^{-1}(\mathbf{x})) + \varepsilon(\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} f^{\mathbb{A}}(h'^{-1}(\mathbf{x})))$, which concludes the proof. \square

As a consequence of Claim 6.11 (with $\mathcal{H}_S = \mathcal{H}_{<\infty}$) we have that h^* realizes the maximum of the function

$$\zeta(h) = \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) (f^{\mathbb{A}}(\mathbf{x}) - f^{\mathbb{A}}(h^{-1}(\mathbf{x}))) + \delta_1$$

over $\mathcal{H}_{<\infty}$. If this maximum is positive, then the vector a such that $a(f, \mathbf{x}) = f^{\mathbb{A}}(\mathbf{x}) - f^{\mathbb{A}}(h^{*-1}(\mathbf{x}))$, $a(\delta_1) = 1$ and $a(\delta_2) = 0$ defines a separating hyperplane (this is true even if $h^* \in \mathcal{H}^*$ because δ_2 is nonnegative). In this case we output a together with the mapping h^* . (A separation oracle is supposed to only output the vector a , but later on we will need to know to which mapping it corresponds.) Otherwise, only two possibilities remain: either $(z, \delta_1, \delta_2) \in P$, or the maximum of $\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) (f^{\mathbb{A}}(\mathbf{x}) - f^{\mathbb{A}}(h^{-1}(\mathbf{x}))) + \delta_1 + \delta_2$ over \mathcal{H}^* is positive.

To verify the latter condition, for every $a \in g(A)$ we construct a valued σ -structure \mathbb{B}_a with universe $B_a = g(A) \setminus \{a\}$ such that for all $(f, \mathbf{x}) \in \text{tup}(\mathbb{B}_a)$, $f^{\mathbb{B}_a}(\mathbf{x}) = z(f, \mathbf{x}) + \varepsilon$ if $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}$ and $f^{\mathbb{B}_a}(\mathbf{x}) = 0$ otherwise. Then, for each $a \in g(A)$ we use the algorithm for $\text{SVCSP}(\mathcal{C}, -)$ to compute in polynomial time a minimum-cost mapping h_a^* for the instance $(\mathbb{A}, \mathbb{B}_a)$. If the cost of h_a^* is infinite for each $a \in g(A)$ then \mathcal{H}^* is empty; it follows that $g(\mathbb{A})$ is a core and we can stop. Otherwise, for each $a \in g(A)$ such that the cost of h_a^* is finite we have

$$\begin{aligned}
h_a^* &\in \operatorname{argmin}_{h \in B_a^A} \left(\sum_{(f, \mathbf{x}) \in \operatorname{tup}(\mathbb{B}_a)} f^{\mathbb{B}_a}(\mathbf{x}) f^{\mathbb{A}}(h^{-1}(\mathbf{x})) \right) \\
&= \operatorname{argmin}_{h \in \mathcal{H}_{<\infty}: h(A) \subseteq g(A) \setminus \{a\}} \left(\sum_{(f, \mathbf{x}) \in \operatorname{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) f^{\mathbb{A}}(h^{-1}(\mathbf{x})) \right. \\
&\quad \left. + \varepsilon \left(\sum_{(f, \mathbf{x}) \in \operatorname{tup}(\mathbb{A})_{<\infty}} f^{\mathbb{A}}(h^{-1}(\mathbf{x})) \right) \right)
\end{aligned}$$

and by Claim 6.11,

$$h_a^* \in \operatorname{argmin}_{h \in \mathcal{H}_{<\infty}: h(A) \subseteq g(A) \setminus \{a\}} \left(\sum_{(f, \mathbf{x}) \in \operatorname{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) f^{\mathbb{A}}(h^{-1}(\mathbf{x})) \right).$$

Therefore, if we pick $h^{**} \in \operatorname{argmin}_{h_a^*: a \in g(A)} (\sum_{(f, \mathbf{x}) \in \operatorname{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) f^{\mathbb{A}}(h^{-1}(\mathbf{x})))$ then we have

$$h^{**} \in \operatorname{argmin}_{h \in \mathcal{H}_{<\infty}: h(A) \subseteq g(A)} \left(\sum_{(f, \mathbf{x}) \in \operatorname{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) f^{\mathbb{A}}(h^{-1}(\mathbf{x})) \right)$$

and either $\sum_{(f, \mathbf{x}) \in \operatorname{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) (f^{\mathbb{A}}(\mathbf{x}) - f^{\mathbb{A}}(h^{**^{-1}}(\mathbf{x}))) + \delta_1 + \delta_2 > 0$ and the vector a such that $a(f, \mathbf{x}) = f^{\mathbb{A}}(\mathbf{x}) - f^{\mathbb{A}}(h^{**^{-1}}(\mathbf{x}))$, $a(\delta_1) = 1$ and $a(\delta_2) = 1$ defines a separating hyperplane (which we output together with the mapping h^{**}) or $(z, \delta_1, \delta_2) \in P$. This concludes the description of our strong separation oracle.

We now apply Lemma 6.7 to the linear program (6.3). Its dual (6.2) is bounded, and if the ellipsoid algorithm returns that it is not feasible then $g(\mathbb{A})$ is a core. Otherwise, the algorithm will return a set of polynomially many valid inequalities of the form

$$\begin{aligned}
\sum_{(f, \mathbf{x}) \in \operatorname{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) (f^{\mathbb{A}}(\mathbf{x}) - f^{\mathbb{A}}(h^{-1}(\mathbf{x}))) + \delta_1 &\leq \alpha'_h \\
\forall h \in \mathcal{H}' \subseteq (\mathcal{H}_{<\infty} \setminus \mathcal{H}^*), \\
\sum_{(f, \mathbf{x}) \in \operatorname{tup}(\mathbb{A})_{<\infty}} z(f, \mathbf{x}) (f^{\mathbb{A}}(\mathbf{x}) - f^{\mathbb{A}}(h^{-1}(\mathbf{x}))) + \delta_1 + \delta_2 &\leq \alpha''_h \\
\forall h \in \mathcal{H}'' \subseteq \mathcal{H}^* \\
-z(f, \mathbf{x}) &\leq \alpha_{f, \mathbf{x}} \\
\forall (f, \mathbf{x}) \in T \subseteq \operatorname{tup}(\mathbb{A})_{<\infty} \\
-\delta_2 &\leq \alpha_2,
\end{aligned}$$

where each mapping h appearing in the inequalities is explicitly known (because we modified the output of the strong separation oracle), and dual variables that satisfy

$$\begin{aligned}
\sum_{h \in \mathcal{H}'} \lambda'_h (f^{\mathbb{A}}(\mathbf{x}) - f^{\mathbb{A}}(h^{-1}(\mathbf{x}))) + \sum_{h \in \mathcal{H}''} \lambda''_h (f^{\mathbb{A}}(\mathbf{x}) - f^{\mathbb{A}}(h^{-1}(\mathbf{x}))) - \lambda_{f, \mathbf{x}} &= 0 \\
\forall (f, \mathbf{x}) \in \operatorname{tup}(\mathbb{A})_{<\infty},
\end{aligned}$$

$$\begin{aligned} \sum_{h \in \mathcal{H}'} \lambda'_h + \sum_{h \in \mathcal{H}''} \lambda''_h &= 1, \\ \sum_{h \in \mathcal{H}''} \lambda''_h - \lambda_2 &= \varepsilon_\Omega(\mathbb{A}, \mathbb{A}), \end{aligned}$$

where we set $\lambda_{f, \mathbf{x}} := 0$ if $(f, \mathbf{x}) \notin T$. Now, we define

$$\omega(h) := \begin{cases} \lambda'_h & \text{if } h \in \mathcal{H}', \\ \lambda''_h & \text{if } h \in \mathcal{H}'', \\ 0 & \text{otherwise} \end{cases}$$

and we deduce from the above system that

$$(6.4) \quad \sum_{h \in \mathcal{H}_{<\infty}} \omega(h) (f^{\mathbb{A}}(\mathbf{x}) - f^{\mathbb{A}}(h^{-1}(\mathbf{x}))) - \lambda_{f, \mathbf{x}} = 0 \quad \forall (f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty},$$

$$(6.5) \quad \sum_{h \in \mathcal{H}_{<\infty}} \omega(h) = 1,$$

$$(6.6) \quad \sum_{h \in \mathcal{H}^*} \omega(h) \geq \varepsilon_\Omega(\mathbb{A}, \mathbb{A}) > 0.$$

Then, from (6.4), (6.5), and the nonnegativity of the dual variables we can deduce that for all $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}$,

$$\sum_{h \in \mathcal{H}_{<\infty}} \omega(h) f^{\mathbb{A}}(h^{-1}(\mathbf{x})) \leq f^{\mathbb{A}}(\mathbf{x})$$

and hence ω (complemented with $\omega(h) = 0$ for all $h \notin \mathcal{H}_{<\infty}$) is an IFH with a support of polynomial size. Finally, we search $\text{supp}(\omega)$ for a mapping g^+ such that $g^+(A) \subsetneq g(A)$, which is guaranteed to exist by the definition of ω and (6.6). \square

The next lemma is the last missing ingredient in the proof of Theorem 6.1.

LEMMA 6.12. *Let $k \geq 1$ and \mathcal{C} be a class of valued σ -structures such that for every $\mathbb{A} \in \mathcal{C}$, $\text{tw}_{ms}(\mathbb{A}) \leq k - 1$ and the overlap of \mathbb{A} is at most k . Then, $\text{SVCSP}(\mathcal{C}, -)$ is in PTIME.*

Proof. Given a pair (\mathbb{A}, \mathbb{B}) of valued structures over some signature σ and two elements $a \in A$, $b \in B$, we say that a is *fixed* to b if there exists a unary symbol $f_{ab} \in \sigma$ such that $f_{ab}^{\mathbb{A}}(\mathbf{x}) = \infty$ if $\mathbf{x} = (a)$ and 0 otherwise, and $f_{ab}^{\mathbb{B}}(\mathbf{x}) = 0$ if $\mathbf{x} = (b)$ and ∞ otherwise. Clearly, if a is fixed to b then any finite-cost mapping from \mathbb{A} to \mathbb{B} must map a to b .

Now, let (\mathbb{A}, \mathbb{B}) be an instance of $\text{SVCSP}(\mathcal{C}, -)$. Suppose that $\text{opt}(\mathbb{A}, \mathbb{B})$ is finite and that the overlap of \mathbb{A} is positive (otherwise the problem is trivial). If every element $a \in A$ is fixed to some element $g(a) \in B$, then the mapping g is the only finite-cost mapping from \mathbb{A} to \mathbb{B} and we can output g . Otherwise, we pick an element $a \in A$ that is not fixed to any element of B . For each $b \in B$, we create one new instance $(\mathbb{A}_{ab}, \mathbb{B}_{ab})$ of SVCSP by adding to (\mathbb{A}, \mathbb{B}) a new symbol f_{ab} that fixes a to b . The overlap and treewidth modulo scopes of \mathbb{A}_{ab} are no greater than those of \mathbb{A} , so we can use Theorem 5.8 to compute $\text{opt}(\mathbb{A}_{ab}, \mathbb{B}_{ab})$ in polynomial time.

Observe that there necessarily exists a value $b^* \in B$ such that $\text{opt}(\mathbb{A}, \mathbb{B}) = \text{opt}(\mathbb{A}_{ab^*}, \mathbb{B}_{ab^*})$ (just take $b^* = h(a)$ for any minimum-cost mapping h from \mathbb{A} to \mathbb{B}). The instance $(\mathbb{A}_{ab^*}, \mathbb{B}_{ab^*})$ has the same optimum cost as (\mathbb{A}, \mathbb{B}) , but one extra element

of A is fixed. We can then repeat the operation until all elements of the left-hand-side structure are fixed, in which case the unique finite-cost mapping (which can be found in polynomial time by inspecting the signature) is a minimum-cost mapping from \mathbb{A} to \mathbb{B} . \square

Proof of Theorem 6.1. If 1 holds, then, by Theorem 4.1, \mathcal{C} has bounded treewidth modulo equivalence. Furthermore, by Lemma 6.10, $\text{REDUCTION STEP}(\mathcal{C})$ is in PTIME and, by Lemma 6.4, $\text{CORE COMPUTATION}(\mathcal{C})$ is in PTIME as well. For the converse implication, assume that 2 holds and let (\mathbb{A}, \mathbb{B}) be an instance of $\text{SVCSP}(\mathcal{C}, -)$. Let k denote the maximum treewidth of the core of a structure in \mathcal{C} . We compute in polynomial time the core $g(\mathbb{A})$ of \mathbb{A} and the associated mapping $g : A \rightarrow A$. Because $\text{tw}(g(\mathbb{A})) \leq k$ implies both $\text{tw}_{ms}(g(\mathbb{A})) \leq k$ and an upper bound of $k + 1$ for the overlap of $g(\mathbb{A})$, we can use Lemma 6.12 to compute in polynomial time a minimum-cost solution $h : g(A) \rightarrow B$ to $(g(\mathbb{A}), \mathbb{B})$. Since $\text{cost}_{g(\mathbb{A}) \rightarrow \mathbb{B}}(h) = \text{cost}_{\mathbb{A} \rightarrow \mathbb{B}}(h \circ g)$ and $\text{opt}(\mathbb{A}, \mathbb{B}) = \text{opt}(g(\mathbb{A}), \mathbb{B})$, we conclude that $h \circ g$ is a minimum-cost mapping from A to B and the theorem follows. \square

7. Related problems. In this section we provide a quick overview of the complexity of deciding the various natural questions on valued structures that arise from our characterizations. We also highlight some interesting implications of our results in the context of database theory.

We establish tight complexity bounds for the following problems. We note that while hardness mostly follows directly from existing results on relational structures, the technical machinery of section 3 is required in order to derive precise upper bounds.

- **IMPROVEMENT:** given two valued structures \mathbb{A}, \mathbb{B} , is it true that $\mathbb{A} \preceq \mathbb{B}$?
- **EQUIVALENCE:** given two valued structures \mathbb{A}, \mathbb{B} , is it true that $\mathbb{A} \equiv \mathbb{B}$?
- **CORE RECOGNITION:** given a valued structure \mathbb{A} , is \mathbb{A} a core?
- **CORE TREewidth:** given a valued structure \mathbb{A} and $k \geq 1$, is the treewidth of the core of \mathbb{A} at most k ?
- **SHERALI-ADAMS TIGHTNESS:** given a valued structure \mathbb{A} and $k \geq 1$, is the Sherali-Adams relaxation of level k tight for \mathbb{A} ?

PROPOSITION 7.1. *IMPROVEMENT and EQUIVALENCE are NP-complete.*

Proof. We first prove that **IMPROVEMENT** is in NP, which implies that **EQUIVALENCE** is in NP as well. By Proposition 3.6, an instance (\mathbb{A}, \mathbb{B}) of **IMPROVEMENT** is a yes-instance if and only if there exists an IFH from \mathbb{A} to \mathbb{B} , or equivalently if $\mathcal{G}_{<\infty} := \{g \in B^A \mid f^{\mathbb{A}}(g^{-1}(\mathbf{x})) < \infty \text{ for all } (f, \mathbf{x}) \in \text{tup}(\mathbb{B})_{<\infty}\}$ (where $\text{tup}(\mathbb{B})_{<\infty} := \{(f, \mathbf{x}) \in \text{tup}(\mathbb{B}) \mid f^{\mathbb{B}}(\mathbf{x}) < \infty\}$) is not empty and the system

$$\begin{aligned} \sum_{g \in \mathcal{G}_{<\infty}} \omega(g) f^{\mathbb{A}}(g^{-1}(\mathbf{x})) &\leq f^{\mathbb{B}}(\mathbf{x}) && \forall (f, \mathbf{x}) \in \text{tup}(\mathbb{B})_{<\infty}, \\ \sum_{g \in \mathcal{G}_{<\infty}} \omega(g) &\leq 1, \\ - \sum_{g \in \mathcal{G}_{<\infty}} \omega(g) &\leq -1, \\ - \omega(g) &\leq 0 && \forall g \in \mathcal{G}_{<\infty} \end{aligned}$$

has a rational solution ω . Since the number of inequalities is polynomial in $|\mathbb{A}|$ and $|\mathbb{B}|$, this system has a solution if and only if it has one with a polynomial number of nonzero variables. Such a subset of nonzero variables is an NP certificate: the

corresponding restriction of the system has polynomial size and its satisfiability can be checked in polynomial time.

For hardness, we note that in the special case of $\{0, \infty\}$ -valued structures (that is, relational structures) the IMPROVEMENT and EQUIVALENCE problems correspond respectively to HOMOMORPHISM and HOMOMORPHIC EQUIVALENCE, which are well-known to be NP-complete even in the bounded arity case [13]. \square

PROPOSITION 7.2. CORE RECOGNITION is coNP-complete.

Proof. We start by establishing membership in coNP. Let (\mathbb{A}) be an instance of CORE RECOGNITION. By the definition of a core, (\mathbb{A}) is a no-instance if and only if there exists a nonsurjective IFH from \mathbb{A} to \mathbb{A} . This is true if and only if the optimum of the linear program

$$\begin{aligned} & \min \left(- \sum_{g \in \mathcal{G}^*} \omega(g) \right), \\ & \sum_{g \in \mathcal{G}_{<\infty}} \omega(g) f^{\mathbb{A}}(g^{-1}(\mathbf{x})) \leq f^{\mathbb{A}}(\mathbf{x}) \quad \forall (f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}, \\ & \sum_{g \in \mathcal{G}_{<\infty}} \omega(g) \leq 1, \\ & - \sum_{g \in \mathcal{G}_{<\infty}} \omega(g) \leq -1, \\ & - \omega(g) \leq 0 \quad \forall g \in \mathcal{G}_{<\infty} \end{aligned}$$

is strictly negative, where $\text{tup}(\mathbb{A})_{<\infty} := \{(f, \mathbf{x}) \in \text{tup}(\mathbb{A}) \mid f^{\mathbb{A}}(\mathbf{x}) < \infty\}$, $\mathcal{G}_{<\infty} := \{g \in A^A \mid f^{\mathbb{A}}(g^{-1}(\mathbf{x})) < \infty \text{ for all } (f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{<\infty}\}$, and \mathcal{G}^* is the restriction of $\mathcal{G}_{<\infty}$ to nonsurjective mappings. Again, the number of inequalities in this system is polynomial in $|\mathbb{A}|$ so there exists a solution of minimum cost with a polynomial number of nonzero variables. Such a subset of variables is a coNP certificate.

On $\{0, \infty\}$ -valued structures with a single binary symmetric function symbol, CORE RECOGNITION coincides with the problem of deciding if a graph is a core in the usual sense (that is, the problem of deciding if all of its endomorphisms are surjective). This problem is coNP-complete [34], so CORE RECOGNITION is coNP-complete as well. \square

PROPOSITION 7.3. CORE TREewidth is NP-complete even for fixed $k \geq 1$, and SHERALI-ADAMS TIGHTNESS is NP-complete even for fixed $k \geq 1$.

Proof. First, we prove that these problems belong to NP when k is part of the input. For CORE TREewidth, the certificate for a yes-instance (\mathbb{A}, k) is a valued structure \mathbb{B} , a polynomially sized certificate that $\mathbb{B} \equiv \mathbb{A}$ (which exists because EQUIVALENCE is in NP by Proposition 7.1), and a tree decomposition of $G(\text{Pos}(\mathbb{B}))$ of width at most k . Correctness follows from Proposition 3.16. For SHERALI-ADAMS TIGHTNESS, the certificate for a yes-instance (\mathbb{A}, k) is a valued structure \mathbb{B} whose overlap is at most k , a polynomially sized certificate that $\mathbb{B} \equiv \mathbb{A}$, and a tree decomposition of $G(\text{Pos}(\mathbb{B}))$ of width modulo scopes at most $k - 1$. For correctness, by Theorem 5.4 it is sufficient to prove that this certificate exists if and only if the core \mathbb{A}' of \mathbb{A} has treewidth modulo scopes at most $k - 1$ and overlap at most k . One implication is immediate: if \mathbb{A}' has treewidth modulo scopes at most $k - 1$ and overlap at most k then we can take $\mathbb{B} := \mathbb{A}'$. For the converse implication, if this certificate exists then by Theorem 5.8

the Sherali–Adams relaxation of level k is always tight for \mathbb{B} . Then, by Theorem 5.4 the core \mathbb{B}' of \mathbb{B} has treewidth modulo scopes at most $k - 1$ and overlap at most k . Furthermore, $\mathbb{A} \equiv \mathbb{B}$ so by Proposition 3.10 \mathbb{A}' and \mathbb{B}' are isomorphic, and finally \mathbb{A}' has treewidth modulo scopes at most $k - 1$ and overlap at most k .

As before, we derive hardness from the $\{0, \infty\}$ -valued case. Determining whether a graph has a core of treewidth at most k is NP-complete for all fixed $k \geq 1$ [16], so CORE TREewidth is NP-complete even for fixed $k \geq 1$ and arity at most 2. For fixed $k \geq 2$ and on valued structures of arity at most 2, SHERALI–ADAMS TIGHTNESS is equivalent to CORE TREewidth with $k' = k - 1$, and hence it is NP-complete. For the case $k = 1$ and arity at most 2 (i.e., for directed graphs), SHERALI–ADAMS TIGHTNESS is equivalent to deciding whether the core of a directed graph is a disjoint union of *oriented trees*, i.e., simple directed graphs whose underlying undirected graphs are trees. It follows from the proof of [16, Theorem 13] that this problem is NP-complete. \square

7.1. Application to database theory. It is well-known that the evaluation (containment) problem for CQs (i.e., first-order queries using only conjunction and existential quantification) is equivalent to the homomorphism problem and hence equivalent to CSPs [13, 37]. This observation has been fundamental in providing principled techniques for the static analysis and optimization of CQs. Indeed, in their seminal work [13], Chandra and Merlin exploited this connection to show that the containment and equivalence problem for CQs is NP-complete. They also provided tools for minimizing CQs with strong theoretical guarantees. In terms of homomorphisms, minimizing a CQ corresponds essentially to computing the (relational) core of a relational structure.

The situation is less clear in the context of annotated databases [31]. In this framework, the tuples of the database are annotated with values from a particular semiring \mathcal{K} , and the semantics of a CQ is a value from \mathcal{K} . For instance, the Boolean semiring $(\{0, 1\}, \vee, \wedge, 0, 1)$ gives us the usual semantics of CQs, and the natural semiring $(\mathbb{N}, +, \times, 0, 1)$ corresponds to the so-called bag semantics of CQs. Another semiring considered in the literature is the *tropical* semiring $(\mathbb{Q}_{\geq 0}, \min, +, \infty, 0)$, which provides a *minimum-cost* semantics [31]. Unfortunately, the homomorphism machinery cannot be applied directly to the study of containment and equivalence in the semiring setting. While there are some works in this direction (see, e.g., [40, 30]), several basic problems remain open. In particular, the precise complexity of containment/equivalence of CQs over the tropical semiring is open (it was shown in [40] to be NP-hard and in Π_2^P , the second level of the polynomial-time hierarchy). Our first observation is that these two problems are actually NP-complete. Indeed, it is well-known that VCSP is equivalent to CQ evaluation over the tropical semiring. Moreover, containment and equivalence of CQs over the tropical semiring correspond to improvement and (valued) equivalence of valued structures. By applying Proposition 7.1, we directly obtain NP-completeness of these problems.

Our second observation is that our notion of (valued) core provides a notion of minimization of CQs over the tropical semiring with theoretical guarantees. Indeed, as the following proposition shows, the core of a valued structure is always an equivalent valued structure with a minimal number of elements, or in terms of CQs, with a minimal number of variables.

PROPOSITION 7.4. *Let \mathbb{A} and \mathbb{B} be valued σ -structures. Then the following are equivalent:*

1. \mathbb{B} is the core of \mathbb{A} .
2. \mathbb{B} is a minimal (with respect to the size of the universe) valued structure equivalent to \mathbb{A} .

Proof. For $1 \Rightarrow 2$, suppose that \mathbb{B} is the core of \mathbb{A} . By contradiction, assume that (2) is false, i.e., there is a valued σ -structure \mathbb{B}' such that $|B'| < |B|$ and $\mathbb{B}' \equiv \mathbb{A}$. In particular, $\mathbb{B}' \equiv \mathbb{B}$ and then by Proposition 3.10 the core \mathbb{B}'' of \mathbb{B}' is isomorphic to \mathbb{B} . By Proposition 3.12, we have that $|B''| \leq |B'| < |B|$, a contradiction. For $2 \Rightarrow 1$, suppose by contradiction that \mathbb{B} is not the core of \mathbb{A} . Since $\mathbb{B} \equiv \mathbb{A}$ by hypothesis, the only possibility is that \mathbb{B} is not a core. Hence, there is an IFH ω and a nonsurjective mapping $g : B \rightarrow B$ such that $g \in \text{supp}(\omega)$. By Proposition A.1(3) from Appendix A, $g(\mathbb{B}) \equiv \mathbb{B} \equiv \mathbb{A}$. Since $|g(B)| < |B|$, this is a contradiction. \square

Proposition 3.12 also gives an algorithm to compute the core of a CQ over the tropical semiring. (In fact, a PSPACE algorithm.) Finally, it is worth mentioning that our classification result from Theorem 4.1 can be interpreted as a characterization of the classes of CQs over the tropical semiring that can be evaluated in PTIME.

Appendix A. Missing proofs from section 3. We prove several propositions and show how they establish results from section 3.

PROPOSITION A.1.

- (1) For every valued σ -structures \mathbb{A} and \mathbb{B} with $A = B$, $\mathbb{A} \leq \mathbb{B}$ implies $\mathbb{A} \preceq \mathbb{B}$.
- (2) Let \mathbb{A} be a valued σ -structure, $g \in B^A$, and ω a probability distribution on B^A . Then, $\mathbb{A} \preceq g(\mathbb{A})$ and $\mathbb{A} \preceq \omega(\mathbb{A})$.
- (3) If ω is an IFH from \mathbb{A} to \mathbb{A} then $\mathbb{A} \equiv \omega(\mathbb{A}) \equiv g(\mathbb{A})$ for every $g \in \text{supp}(\omega)$.

Proof.

- (1) $\mathbb{A} \leq \mathbb{B}$ implies, for every \mathbb{C} and $g : A \rightarrow C$, $\text{cost}_{\mathbb{A} \rightarrow \mathbb{C}}(g) \leq \text{cost}_{\mathbb{B} \rightarrow \mathbb{C}}(g)$. Thus, $\text{opt}(\mathbb{A}, \mathbb{C}) \leq \text{opt}(\mathbb{B}, \mathbb{C})$ for every \mathbb{C} , and so $\mathbb{A} \preceq \mathbb{B}$.
- (2) Since $\text{cost}_{g(\mathbb{A}) \rightarrow \mathbb{C}}(h) = \text{cost}_{\mathbb{A} \rightarrow \mathbb{C}}(h \circ g)$, for every \mathbb{C} and $h : B \rightarrow C$, we obtain $\mathbb{A} \preceq g(\mathbb{A})$.
Then, since $\text{opt}(\sum_{g \in B^A} \omega(g)g(\mathbb{A}), \mathbb{C}) \geq \sum_{g \in B^A} \omega(g)\text{opt}(g(\mathbb{A}), \mathbb{C})$, we get $\mathbb{A} \preceq \omega(\mathbb{A})$.
- (3) From (2), $\mathbb{A} \preceq \omega(\mathbb{A})$. From the definition of IFH (Definition 3.5), we have $\omega(\mathbb{A}) \leq \mathbb{A}$. By (1), $\omega(\mathbb{A}) \preceq \mathbb{A}$. Together, $\mathbb{A} \equiv \omega(\mathbb{A})$.
By (2) again, $\mathbb{A} \preceq g(\mathbb{A})$ for every $g \in \text{supp}(\omega)$. In particular, $\text{opt}(\mathbb{A}, \mathbb{C}) \leq \text{opt}(g(\mathbb{A}), \mathbb{C})$ for every \mathbb{C} . Assume for contradiction that $\text{opt}(g(\mathbb{A}), \mathbb{C}) > \text{opt}(\mathbb{A}, \mathbb{C})$ for at least one $g \in \text{supp}(\omega)$ and some \mathbb{C} . Then,

$$\begin{aligned} \text{opt}(\omega(\mathbb{A}), \mathbb{C}) &= \text{opt}\left(\sum_{g \in A^A} \omega(g)g(\mathbb{A}), \mathbb{C}\right) \geq \sum_{g \in A^A} \omega(g)\text{opt}(g(\mathbb{A}), \mathbb{C}) \\ &> \sum_{g \in A^A} \omega(g)\text{opt}(\mathbb{A}, \mathbb{C}) = \text{opt}(\mathbb{A}, \mathbb{C}), \end{aligned}$$

which contradicts the already established $\mathbb{A} \equiv \omega(\mathbb{A})$. Hence $\mathbb{A} \equiv g(\mathbb{A})$. \square

PROPOSITION A.2. Let \mathbb{A} be a core valued σ -structure and ω be an IFH from \mathbb{A} to itself. Then, for every $g \in \text{supp}(\omega)$ and every $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$, we have $f^{\mathbb{A}}(\mathbf{x}) = f^{\mathbb{A}}(g(\mathbf{x}))$.

Proof. By contradiction, take $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$ and $g \in \text{supp}(\omega)$ such that $f^{\mathbb{A}}(\mathbf{x}) \neq f^{\mathbb{A}}(g(\mathbf{x}))$ with minimal $c := f^{\mathbb{A}}(\mathbf{x})$. By minimality, for any $h \in \text{supp}(\omega)$, we have $f^{\mathbb{A}}(\mathbf{y}) = f^{\mathbb{A}}(h(\mathbf{y}))$ if $f^{\mathbb{A}}(\mathbf{y}) < c$. Since h is a bijection (\mathbb{A} is a core), and by counting, we obtain $f^{\mathbb{A}}(\mathbf{y}) \geq c$ if and only if $f^{\mathbb{A}}(h(\mathbf{y})) \geq c$. If $c = \infty$, then $f^{\mathbb{A}}(g(\mathbf{x}))$ needs to be $> \infty$, a contradiction.

If $c < \infty$, we have $f^{\mathbb{A}}(\mathbf{x}) < f^{\mathbb{A}}(g(\mathbf{x}))$. By counting, this implies that there is \mathbf{w} such that $f^{\mathbb{A}}(g(\mathbf{w})) = c$ and $f^{\mathbb{A}}(\mathbf{w}) > c$. Then $\sum_{h \in \text{supp}(\omega); h(\mathbf{y})=g(\mathbf{w})} \omega(h) f^{\mathbb{A}}(\mathbf{y}) > c = f^{\mathbb{A}}(g(\mathbf{w}))$ since $f^{\mathbb{A}}(\mathbf{y}) \geq c$ and at least for one \mathbf{y} we have $f^{\mathbb{A}}(\mathbf{y}) > c$. This contradicts ω being an IFH. \square

COROLLARY A.3. *Let \mathbb{A} and \mathbb{B} be two core valued σ -structures. Assume that there is an IFH ω from \mathbb{A} to \mathbb{B} and an IFH ω' from \mathbb{B} to \mathbb{A} . Then, \mathbb{A} and \mathbb{B} are isomorphic.*

Proof. By composition, $\omega' \circ \omega$ is an IFH from \mathbb{A} to itself. Since \mathbb{A} is a core, the support of $\omega' \circ \omega$ contains only bijections. Similarly, since \mathbb{B} is a core the support of $\omega \circ \omega'$ contains only bijections. These two facts imply that $|A| = |B|$ and the supports of ω and ω' also contain only bijections. Assume for contradiction that there is $g \in \text{supp}(\omega)$ such that for some $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$ we have $f^{\mathbb{A}}(\mathbf{x}) \neq f^{\mathbb{B}}(g(\mathbf{x}))$.

Case 1: $f^{\mathbb{A}}(\mathbf{x}) < f^{\mathbb{B}}(g(\mathbf{x}))$. Since ω' is an IFH, there must exist $g' \in \text{supp}(\omega')$ such that for some \mathbf{y} we have $g'(\mathbf{y}) = \mathbf{x}$ and $f^{\mathbb{B}}(\mathbf{y}) \leq f^{\mathbb{A}}(\mathbf{x})$. But then, $g \circ g' \in \text{supp}(\omega \circ \omega')$, which contradicts Proposition A.2.

Case 2: $f^{\mathbb{A}}(\mathbf{x}) > f^{\mathbb{B}}(g(\mathbf{x}))$. In this case, we consider any $g' \in \text{supp}(\omega')$ and $(f, g(\mathbf{x}))$. As $g' \circ g \in \text{supp}(\omega' \circ \omega)$, Proposition A.2 guarantees that $f^{\mathbb{B}}(g(\mathbf{x})) < f^{\mathbb{A}}(g'(g(\mathbf{x})))$. We can then proceed as in Case 1, interchanging the roles of \mathbb{A} and \mathbb{B} . \square

We will need the following variant of Farkas' lemma, known as Motzkin's transposition theorem. [50, Lemma 2.8] shows how it can be derived from Farkas' lemma [46, Corollary 7.1k].

LEMMA A.4. *For any $A \in \mathbb{Q}^{m \times n}$ and $B \in \mathbb{Q}^{p \times n}$ exactly one of the following holds:*

- $A\bar{y} > 0, B\bar{y} \geq 0$ for some $\bar{y} \in \mathbb{Q}_{\geq 0}^n$, or
- $A^T \bar{z}_1 + B^T \bar{z}_2 \leq 0$ for some $0 \neq \bar{z}_1 \in \mathbb{Q}_{\geq 0}^m, \bar{z}_2 \in \mathbb{Q}_{\geq 0}^p$.

PROPOSITION A.5. *Let \mathbb{A} and \mathbb{B} be two valued σ -structures and let $\mathcal{G} \subseteq B^A$. Then, either there is an IFH ω from \mathbb{A} to \mathbb{B} such that $\mathcal{G} \cap \text{supp}(\omega) \neq \emptyset$, or there is a finite-valued σ -structure \mathbb{C} with $C = B$ such that $\text{cost}_{\mathbb{A} \rightarrow \mathbb{C}}(g) > \text{cost}_{\mathbb{B} \rightarrow \mathbb{C}}(\text{id})$ for every $g \in \mathcal{G}$. Moreover, deciding whether such a ω exists and if so then computing it, as well as computing \mathbb{C} if no such ω exists, can be done effectively from \mathbb{A} , \mathbb{B} , and \mathcal{G} .*

Proof. Let $\mathcal{M} \subseteq B^A$ be the set of all mappings g such that for every $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$, $f^{\mathbb{B}}(g(\mathbf{x})) = \infty$ whenever $f^{\mathbb{A}}(\mathbf{x}) = \infty$.

Case 1: Assume $\mathcal{G} \cap \mathcal{M} = \emptyset$, i.e., for every $g \in \mathcal{G}$ there is some $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$ satisfying $f^{\mathbb{A}}(\mathbf{x}) = \infty$ and $f^{\mathbb{B}}(g(\mathbf{x})) < \infty$. Define \mathbb{C} by setting $f^{\mathbb{C}}(\mathbf{x}) = 0$ if $f^{\mathbb{B}}(\mathbf{x}) = \infty$ and $f^{\mathbb{C}}(\mathbf{x}) = 1$ otherwise. By the definition of \mathbb{C} , we have $\text{cost}_{\mathbb{B} \rightarrow \mathbb{C}}(\text{id}) < \infty$. By the definition of \mathbb{C} again and the assumption, $\text{cost}_{\mathbb{A} \rightarrow \mathbb{C}}(g) = \infty$ for every $g \in \mathcal{G}$. Hence the desired inequality holds.

Case 2: Assume there is $g \in \mathcal{G} \cap \mathcal{M}$ such that for every $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$, $f^{\mathbb{B}}(g(\mathbf{x})) = \infty$ whenever $0 < f^{\mathbb{A}}(\mathbf{x}) < \infty$. Define ω by $\omega(g) = 1$; it is an IFH from \mathbb{A} to \mathbb{B} , since for every $(f, \mathbf{y}) \in \text{tup}(\mathbb{B})$, $\sum_{\mathbf{x}: g(\mathbf{x})=\mathbf{y}} f^{\mathbb{A}}(\mathbf{x}) \leq f^{\mathbb{B}}(\mathbf{y})$.

Case 3: Assume that neither Case 1 nor Case 2 applies. We will use Lemma A.4. We set $n := |\mathcal{M}| + 1$ and $m := 1$. Thus the matrix A is just a row vector, where we have a coordinate for every map $g \in \mathcal{M}$ and one more (say, the last), coordinate. We define $A(g) = 1$ if $g \in \mathcal{M} \cap \mathcal{G}$ and $A(g) = 0$ otherwise. Note that since Case 1 does not apply, then $\mathcal{M} \cap \mathcal{G} \neq \emptyset$. We set A to be 0 in the last coordinate. Condition $A\bar{y} > 0$ means that $\bar{y}(g) > 0$ for some $g \in \mathcal{M} \cap \mathcal{G}$. Let $\text{tup}(\mathbb{B})_{<\infty} := \{(f, \mathbf{y}) \in \text{tup}(\mathbb{B}) : f^{\mathbb{B}}(\mathbf{y}) < \infty\}$. We set $p := |\text{tup}(\mathbb{B})_{<\infty}| + 1$ and define the matrix B as follows. We have a row in B for every $(f, \mathbf{y}) \in \text{tup}(\mathbb{B})_{<\infty}$. The entry of B with row (f, \mathbf{y}) and column $g \in \mathcal{M}$ is

$-\sum_{\mathbf{x}:g(\mathbf{x})=\mathbf{y}} f^{\mathbb{A}}(\mathbf{x})$. In the last column of B we put $f^{\mathbb{B}}(\mathbf{y})$. Finally, we add one more row to B of the form $(1, \dots, 1, -1)$. Note that B is actually a rational matrix.

We now distinguish the two cases of Lemma A.4.

Case 3a: In the first case of Lemma A.4, we will construct an IFH ω from \mathbb{A} to \mathbb{B} . Let c be the value of the last coordinate of \bar{y} .

If $c = 1$, we define ω as follows: $\omega(g)$ is the g th coordinate of \bar{y} . Using all, except for the last one, rows of $B\bar{y}$, we get $\omega(\mathbb{A}) \leq \mathbb{B}$. Using $A\bar{y} > 0$, $\mathcal{G} \cap \text{supp}(\omega) \neq \emptyset$. Finally, using the last row of $B\bar{y}$, we get that $d := \sum_{g \in B^A} \omega(g) \geq c = 1$. If $d \neq 1$, to make ω a probability distribution, we scale it by $1/d$.

If $c > 0$, we can scale \bar{y} by $1/c$ and are in the $c = 1$ case.

If $c = 0$, let $g \in \mathcal{M} \cap \mathcal{G}$ be a coordinate in \bar{y} with $\bar{y}(g) > 0$. Since Case 2 does not apply, there is $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$ such that $0 < f^{\mathbb{A}}(\mathbf{x}) < \infty$ and $f^{\mathbb{B}}(g(\mathbf{x})) < \infty$. We claim that in $B\bar{y}$ the entry corresponding to $(f, g(\mathbf{x})) \in \text{tup}(\mathbb{B})_{<\infty}$ is negative, which contradicts $B\bar{y} \geq 0$. Indeed, by the definition of B , the entry is equal to $\sum_{h \in \mathcal{M}} \bar{y}(h) (-\sum_{\mathbf{y}:h(\mathbf{y})=g(\mathbf{x})} f^{\mathbb{A}}(\mathbf{y}))$. This is negative as every term in the sum is ≤ 0 and there is one term < 0 (take $h = g$ and $\mathbf{y} = \mathbf{x}$).

Case 3b: In the second case of Lemma A.4, we construct a valued σ -structure \mathbb{C} as required. Let c be the last entry of \bar{z}_2 given by Lemma A.4. Define \mathbb{C} by $f^{\mathbb{C}}(\mathbf{x})$ to be $\bar{z}_2((f, \mathbf{x})) + \varepsilon$ for every $(f, \mathbf{x}) \in \text{tup}(\mathbb{B})_{<\infty}$, where ε is a positive number to be defined later. For $(f, \mathbf{x}) \in \text{tup}(\mathbb{B}) \setminus \text{tup}(\mathbb{B})_{<\infty}$ we set $f^{\mathbb{C}}(\mathbf{x}) = 0$. The last row of the inequality $A^T \bar{z}_1 + B^T \bar{z}_2 \leq 0$ implies that $\text{cost}_{\mathbb{B} \rightarrow \mathbb{C}}(\text{id}) = (\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{B})_{<\infty}} f^{\mathbb{B}}(\mathbf{x}) \bar{z}_2((f, \mathbf{x}))) + \varepsilon N \leq c + \varepsilon N$, where $N := \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{B})_{<\infty}} f^{\mathbb{B}}(\mathbf{x})$. Since $\text{cost}_{\mathbb{B} \rightarrow \mathbb{C}}(\text{id}) < \infty$, the required inequality holds for every $g \in \mathcal{G} \setminus \mathcal{M}$, as we have $\text{cost}_{\mathbb{A} \rightarrow \mathbb{C}}(g) = \infty$ (here we use the fact that $f^{\mathbb{C}}(\mathbf{x}) > 0$ whenever $f^{\mathbb{B}}(\mathbf{x}) < \infty$). On the other hand, for every $g \in \mathcal{G} \cap \mathcal{M}$ we have $\text{cost}_{\mathbb{A} \rightarrow \mathbb{C}}(g) \geq \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{B})_{<\infty}} \bar{z}_2((f, \mathbf{x})) \sum_{\mathbf{y}:g(\mathbf{y})=\mathbf{x}} f^{\mathbb{A}}(\mathbf{y})$. Using the row corresponding to g in the inequality $A^T \bar{z}_1 + B^T \bar{z}_2 \leq 0$, we obtain $\text{cost}_{\mathbb{A} \rightarrow \mathbb{C}}(g) \geq \bar{z}_1 + c$. We can choose ε so $\bar{z}_1 > \varepsilon N$ and obtain the required inequality.

Computability of ω and \mathbb{C} follows directly from the previous arguments and from solving the corresponding system of linear inequalities (from Lemma A.4) whenever necessary. \square

We now show how the previous propositions imply the results from section 3.

Propositions A.1(1) and (2) imply the “if direction” of Proposition 3.6. For the “only if direction,” we can use Proposition A.5 with $\mathcal{G} = B^A$. We obtain that if there is no IFH from \mathbb{A} to \mathbb{B} , then there is a structure \mathbb{C} with $\text{opt}(\mathbb{A}, \mathbb{C}) > \text{opt}(\mathbb{B}, \mathbb{C})$, i.e., $\mathbb{A} \not\preceq \mathbb{B}$ does not hold.

Proposition 3.10 follows from Corollary A.3 and Proposition 3.6.

The existence of a core for every valued structure \mathbb{A} in Proposition 3.12 follows from the fact that either \mathbb{A} is a core or, by Proposition A.1(3), $\mathbb{A} \equiv g(\mathbb{A})$, where $g : A \rightarrow A$ is a nonsurjective mapping. We can iterate this argument until we find a core of \mathbb{A} , whose universe has size $\leq |A|$. Uniqueness of cores follows from Corollary A.3. The computability of the core can be obtained by invoking Proposition A.5 with $\mathbb{A} = \mathbb{B}$ and \mathcal{G} being the set of nonsurjective functions on A . In this case, Proposition A.5 gives us that if \mathbb{A} is not a core, then we can compute a nonsurjective IFH ω from \mathbb{A} to itself, and in particular, we can compute a nonsurjective $g : A \rightarrow A$ such that $\mathbb{A} \equiv g(\mathbb{A})$. Hence, we can turn the previous existence argument into an algorithm for computing the core.

For the “only if direction” of Proposition 3.13, we again use Proposition A.5 with $\mathbb{A} = \mathbb{B}$ and \mathcal{G} being the set of nonsurjective functions on A . We obtain that if \mathbb{A} is a core, then there exists the required mapping c (which is the finite-valued structure \mathbb{C})

and it can be effectively computed. For the “if direction”, and towards a contradiction, suppose \mathbb{A} is not a core and assume there exists a mapping c as in the statement of the proposition. Take a nonsurjective IFH ω from \mathbb{A} to itself and a nonsurjective $g \in \text{supp}(\omega)$. Let \mathbb{C} be the finite-valued structure with $C = A$ and $f^{\mathbb{C}}(\mathbf{x}) = c(f, \mathbf{x})$. By Proposition A.1(3), $\mathbb{A} \equiv g(\mathbb{A})$. Let h be a minimum-cost mapping from $g(\mathbb{A})$ to \mathbb{C} . Since $\text{cost}_{g(\mathbb{A}) \rightarrow \mathbb{C}}(h) = \text{cost}_{\mathbb{A} \rightarrow \mathbb{C}}(h \circ g)$ and $\text{opt}(g(\mathbb{A}), \mathbb{C}) = \text{opt}(\mathbb{A}, \mathbb{C})$, we obtain that $h \circ g$ is a minimum-cost mapping from \mathbb{A} to \mathbb{C} . However, by the properties of c (and hence \mathbb{C}) and the fact that $h \circ g : A \rightarrow A$ is nonsurjective, we obtain that $\text{cost}_{\mathbb{A} \rightarrow \mathbb{C}}(h \circ g) > \text{cost}_{\mathbb{A} \rightarrow \mathbb{C}}(\text{id})$. This is a contradiction.

Appendix B. Proof of Example 4.4. We first recall the construction from Example 4.4. Consider the signature $\sigma = \{f, \mu\}$, where f and μ are binary and unary function symbols, respectively. For $n \geq 1$, let \mathbb{A}_n be the valued σ -structure with universe $A_n = \{1, \dots, n\} \times \{1, \dots, n\}$ such that (i) $f^{\mathbb{A}_n}((i, j), (i', j')) = \infty$ if $i \leq i'$, $j \leq j'$, and $(i' - i) + (j' - j) = 1$; otherwise $f^{\mathbb{A}_n}((i, j), (i', j')) = 0$, and (ii) $\mu^{\mathbb{A}_n}((i, j)) = 1$, for all $(i, j) \in A_n$. Also, for $n \geq 1$, let \mathbb{A}'_n be the valued σ -structure with universe $A'_n = \{1, \dots, 2n - 1\}$ such that (i) $f^{\mathbb{A}'_n}(i, j) = \infty$ if $j = i + 1$; otherwise $f^{\mathbb{A}'_n}(i, j) = 0$, and (ii) $\mu^{\mathbb{A}'_n}(i) = i$, for $1 \leq i \leq n$, and $\mu^{\mathbb{A}'_n}(i) = 2n - i$, for $n + 1 \leq i \leq 2n - 1$.

We prove the following.

PROPOSITION B.1. *For every $n \geq 1$, \mathbb{A}'_n is the core of \mathbb{A}_n .*

Proof. Let $n \geq 1$. First, $\text{Pos}(\mathbb{A}'_n)$ is a relational core, so \mathbb{A}'_n is a core. We define $g : A_n \rightarrow A'_n$ as $g((i, j)) := i + j - 1$; then for all $k \in A'_n$ we have

$$|g^{-1}(k)| = |\{(i, j) \in A_n \mid i + j - 1 = k\}| = \begin{cases} k & \text{if } 1 \leq k \leq n, \\ 2n - k & \text{if } n + 1 \leq k \leq 2n - 1 \end{cases}$$

and hence for every $k \in A'_n$ we have $\mu^{\mathbb{A}_n}(g^{-1}(k)) = |g^{-1}(k)| \leq \mu^{\mathbb{A}'_n}(k)$. Furthermore, $f^{\mathbb{A}_n}((i, j), (i', j')) = \infty$ implies $g(i', j') - g(i, j) = (i' - i) + (j' - j) = 1$, and in turn $f^{\mathbb{A}'_n}(g(i, j), g(i', j')) = \infty$. In particular, for all $k_1, k_2 \in A'_n$ it holds that $f^{\mathbb{A}_n}(g^{-1}((k_1, k_2))) \leq f^{\mathbb{A}'_n}((k_1, k_2))$, which completes the proof that the distribution $\omega(g) := 1$ is an IFH from \mathbb{A}_n to \mathbb{A}'_n .

We now turn to the more delicate task of constructing an IFH ω' from \mathbb{A}'_n to \mathbb{A}_n . Note that for any (hypothetical) mapping $g' \in \text{supp}(\omega')$, if $f^{\mathbb{A}'_n}(\mathbf{x}) = \infty$ then $f^{\mathbb{A}_n}(g'(\mathbf{x})) = \infty$. This implies that for any such mapping the sequence of vertices $g'(1), \dots, g'(2n - 1)$ is a directed path from $(1, 1)$ to (n, n) in $\text{Pos}(\mathbb{A}_n)$ and if $g'(m) = (i, j)$ then $i + j - 1 = m$ (i.e., (i, j) lies on the m th diagonal). Additionally, the definitions of $\mu^{\mathbb{A}_n}$ and $\mu^{\mathbb{A}'_n}$ imply that ω' is an IFH if and only if the marginal for each vertex on the m th diagonal is $1/m$ if $m \leq n$ and $1/(2n - m)$ if $n \leq m \leq 2n - 1$. Formally, for every i, j ,

$$(B.1) \quad \sum_{g': g'(i+j-1)=(i,j)} \omega'(g') = \begin{cases} 1/(i+j-1) & \text{if } i+j-1 \leq n, \\ 1/(2n-i-j+1) & \text{otherwise.} \end{cases}$$

Now, let $B_n := \{(i, j) \in A_n : i + j - 1 \leq n\}$ and \mathbb{B}_n be the restriction of \mathbb{A}_n to B_n . Similarly, let $B'_n := \{1, \dots, n\}$ and \mathbb{B}'_n be the restriction of \mathbb{A}'_n to B'_n . Note that it suffices to find an IFH ω'_B from \mathbb{B}'_n to \mathbb{B}_n satisfying (B.1) (the “otherwise” part of the equation is void for $\mathbb{B}_n, \mathbb{B}'_n$): given any such ω'_B we can define $\omega'(g') := \omega'_B(g'|_{B'_n})$ if for all $m \leq n$, $g'(m) = (i, j)$ implies that $g'(2n - m) = (n - j, n - i)$, and 0 otherwise (i.e., every mapping g' in the support of ω' is a mirror image of itself via the diagonal).

Recall that, in order for ω' to be an IFH, any $g'_B \in \text{supp}(\omega'_B)$ must define a path in $\text{Pos}(\mathbb{B}_n)$ of length n from $(1, 1)$ to a vertex on the last diagonal (i.e., a vertex (i, j) such that $i + j - 1 = n$). Each vertex (i, j) in B_n , except for those on the last diagonal, has two outgoing edges: to $(i + 1, j)$ and to $(i, j + 1)$. We choose the first edge with probability $i/(i + j)$ and the second with probability $j/(i + j)$. The probability $\omega'_B(g'_B)$ is the product of the probabilities of the edges in the path defined by g'_B .

It remains to prove that (B.1) holds for ω'_B . Note that the left-hand side of (B.1) is the probability that (i, j) is on a path chosen according to distribution ω'_B , i.e., the probability that the initial part of the path $g'_B(1), \dots, g'_B(i + j - 1)$ ends in the vertex (i, j) . We now prove that (B.1) holds by induction on the diagonal of \mathbb{B}_n (i.e., for increasing values of $i + j - 1$). Choose (i, j) . If $i = j = 1$, the left-hand side of (B.1) is equal to 1 and the claim holds. Otherwise, we have two cases. If $i \neq 1 \neq j$, then the probability that (i, j) is on the path is equal to the probability that $(i - 1, j)$ is and the edge $((i - 1, j), (i, j))$ was chosen, plus the probability that $(i, j - 1)$ is and the edge $((i, j - 1), (i, j))$ was chosen. By induction, this probability is equal to $1/(i + j - 2) \times (j - 1)/(i + j - 1) + 1/(i + j - 2) \times (i - 1)/(i + j - 1)$, which is equal to $1/(i + j - 1)$ as required. In the case where $i = 1$ or $j = 1$, one of the summands is 0, but the required equality still holds. This concludes the proof that ω' is an IFH from \mathbb{A}'_n to \mathbb{A}_n , and hence \mathbb{A}'_n is the core of \mathbb{A}_n , as claimed.

As a final remark, we note that if we alter the definitions of \mathbb{A}_n and \mathbb{A}'_n so that the edges of \mathbb{A}'_n have weight 1 (instead of ∞) and the weight of each edge e of \mathbb{A}'_n is the marginal of ω' on e , then the distributions ω and ω' are still IFHs between \mathbb{A}_n and \mathbb{A}'_n . Since \mathbb{A}'_n is still a core, this new construction shows that bounded treewidth modulo (valued) equivalence is a strictly more general property than bounded treewidth *even for finite-valued structures*. \square

Appendix C. Proof of Proposition 5.2.

We prove the following.

PROPOSITION C.1. *Let \mathbb{A}, \mathbb{B} be valued σ -structures and $k \geq 1$. If there exists an IFH from \mathbb{A} to \mathbb{B} , then $\mathbb{A} \preceq_k^{SA} \mathbb{B}$.*

Proof. Let \mathbb{C} be an arbitrary valued σ -structure, ω be an IFH from \mathbb{A} to \mathbb{B} , and λ be a solution to $\text{SA}_k(\mathbb{B}, \mathbb{C})$ of minimum cost. We can write the cost of λ as a sum over all $(f, \mathbf{x}) \in \text{tup}(\mathbb{B}_k)_{>0}$ and $s : \text{Set}(\mathbf{x}) \rightarrow C_k$, as constraint (SA3) ensures that $\lambda(f, \mathbf{x}, s) = 0$, whenever $f^{\mathbb{B}_k}(\mathbf{x}) \times f^{\mathbb{C}_k}(s(\mathbf{x})) = \infty$. Then, we have

$$\begin{aligned} \text{opt}_k^{SA}(\mathbb{B}, \mathbb{C}) &= \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{B}_k)_{>0}, s : \text{Set}(\mathbf{x}) \rightarrow C_k} \lambda(f, \mathbf{x}, s) f^{\mathbb{B}_k}(\mathbf{x}) f^{\mathbb{C}_k}(s(\mathbf{x})) \\ &= \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{B}), s : \text{Set}(\mathbf{x}) \rightarrow C} \lambda(f, \mathbf{x}, s) f^{\mathbb{B}}(\mathbf{x}) f^{\mathbb{C}}(s(\mathbf{x})) \\ &\geq \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{B}), s : \text{Set}(\mathbf{x}) \rightarrow C} \left(\sum_{g \in \text{supp}(\omega)} \omega(g) f^{\mathbb{A}}(g^{-1}(\mathbf{x})) \right) \lambda(f, \mathbf{x}, s) f^{\mathbb{C}}(s(\mathbf{x})) \\ &= \sum_{g \in \text{supp}(\omega)} \omega(g) \left(\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{B}), s : \text{Set}(\mathbf{x}) \rightarrow C} \lambda(f, \mathbf{x}, s) f^{\mathbb{A}}(g^{-1}(\mathbf{x})) f^{\mathbb{C}}(s(\mathbf{x})) \right) \\ &= \sum_{g \in \text{supp}(\omega)} \omega(g) \left(\sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{A}), s : \text{Set}(g(\mathbf{y})) \rightarrow C} \lambda(f, g(\mathbf{y}), s) f^{\mathbb{A}}(\mathbf{y}) f^{\mathbb{C}}(s(g(\mathbf{y}))) \right) \end{aligned}$$

and hence there exists $g \in \text{supp}(\omega)$ such that

$$(C.1) \quad \begin{aligned} \text{opt}_k^{\text{SA}}(\mathbb{B}, \mathbb{C}) &\geq \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{A}), s: \text{Set}(g(\mathbf{y})) \rightarrow C} \lambda(f, g(\mathbf{y}), s) f^{\mathbb{A}}(\mathbf{y}) f^{\mathbb{C}}(s(g(\mathbf{y}))) \\ &= \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{A}_k)_{>0}, s: \text{Set}(g(\mathbf{y})) \rightarrow C_k} \lambda(f, g(\mathbf{y}), s) f^{\mathbb{A}_k}(\mathbf{y}) f^{\mathbb{C}_k}(s(g(\mathbf{y}))). \end{aligned}$$

Since $g \in \text{supp}(\omega)$, we have that g is a homomorphism from $\text{Pos}(\mathbb{A})$ to $\text{Pos}(\mathbb{B})$ (see Remark 3.7). It follows that $(f, g(\mathbf{y})) \in \text{tup}(\mathbb{B}_k)_{>0}$ for every $(f, \mathbf{y}) \in \text{tup}(\mathbb{A}_k)_{>0}$. Hence, for any $(f, \mathbf{y}) \in \text{tup}(\mathbb{A}_k)_{>0}$ and $r: \text{Set}(\mathbf{y}) \rightarrow C_k$, we can define

$$\lambda'(f, \mathbf{y}, r) = \begin{cases} \lambda(f, g(\mathbf{y}), s) & \text{if there exists } s: \text{Set}(g(\mathbf{y})) \rightarrow C_k \text{ such that } s \circ g = r, \\ 0 & \text{otherwise.} \end{cases}$$

Equation (C.1) then becomes

$$\text{opt}_k^{\text{SA}}(\mathbb{B}, \mathbb{C}) \geq \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{A}_k)_{>0}, r: \text{Set}(\mathbf{y}) \rightarrow C_k} \lambda'(f, \mathbf{y}, r) f^{\mathbb{A}_k}(\mathbf{y}) f^{\mathbb{C}_k}(r(\mathbf{y})).$$

All that remains to do is to show that λ' is a feasible solution to $\text{SA}_k(\mathbb{A}, \mathbb{C})$. The fact that the condition (SA4) is satisfied is immediate. Also, it follows from $\omega(g) > 0$ that $f^{\mathbb{A}_k}(\mathbf{y}) = \infty$ implies $f^{\mathbb{B}_k}(g(\mathbf{y})) = \infty$ and $f^{\mathbb{A}_k}(\mathbf{y}) > 0$ implies $f^{\mathbb{B}_k}(g(\mathbf{y})) > 0$; thus $f^{\mathbb{A}_k}(\mathbf{y}) \times f^{\mathbb{C}_k}(r(\mathbf{y})) = \infty$ implies that $f^{\mathbb{B}_k}(g(\mathbf{y})) \times f^{\mathbb{C}_k}(r(\mathbf{y})) = \infty$ and the condition (SA3) is satisfied for all r that can be written as $s \circ g$ for some s , as λ satisfies (SA3). The definition of λ' ensures that it also holds for all other mappings r . For the condition (SA2), observe that for all $(f, \mathbf{y}) \in \text{tup}(\mathbb{A}_k)_{>0}$ we have

$$\sum_{r: \text{Set}(\mathbf{y}) \rightarrow C_k} \lambda'(f, \mathbf{y}, r) = \sum_{s: \text{Set}(g(\mathbf{y})) \rightarrow C_k} \lambda'(f, \mathbf{y}, s \circ g) = \sum_{s: \text{Set}(g(\mathbf{y})) \rightarrow C_k} \lambda(f, g(\mathbf{y}), s) = 1.$$

That only leaves the condition (SA1). Let $(f, \mathbf{x}), (p, \mathbf{y}) \in \text{tup}(\mathbb{A}_k)_{>0}$ such that $\text{Set}(\mathbf{x}) \subseteq \text{Set}(\mathbf{y})$ and $|\text{Set}(\mathbf{x})| \leq k$, and let $t: \text{Set}(\mathbf{x}) \rightarrow C_k$ be any mapping. If there does not exist a mapping $z: \text{Set}(g(\mathbf{x})) \rightarrow C_k$ such that $t = z \circ g$, then

$$\lambda'(f, \mathbf{x}, t) = 0 = \sum_{r: \text{Set}(\mathbf{y}) \rightarrow C_k, r|_{\text{Set}(\mathbf{x})} = t} \lambda'(p, \mathbf{y}, r)$$

and the relevant constraints are satisfied, so let us assume that such a mapping z exists. In this case, we have

$$\begin{aligned} \sum_{r: \text{Set}(\mathbf{y}) \rightarrow C_k, r|_{\text{Set}(\mathbf{x})} = t} \lambda'(p, \mathbf{y}, r) &= \sum_{s: \text{Set}(g(\mathbf{y})) \rightarrow C_k, s \circ g|_{\text{Set}(\mathbf{x})} = t} \lambda'(p, \mathbf{y}, s \circ g) \\ &= \sum_{s: \text{Set}(g(\mathbf{y})) \rightarrow C_k, s \circ g|_{\text{Set}(\mathbf{x})} = t} \lambda(p, g(\mathbf{y}), s) \\ &= \sum_{s: \text{Set}(g(\mathbf{y})) \rightarrow C_k, s|_{\text{Set}(g(\mathbf{x}))} = z} \lambda(p, g(\mathbf{y}), s) \\ &= \lambda(f, g(\mathbf{x}), z) \\ &= \lambda'(f, \mathbf{x}, t) \end{aligned}$$

and again the condition (SA1) is satisfied. (The second equality holds by definition of λ' and the penultimate equality holds by applying (SA1) to λ , $(f, g(\mathbf{x}))$ and $(p, g(\mathbf{y}))$.) Therefore λ' is a feasible solution to $\text{SA}_k(\mathbb{A}, \mathbb{C})$, and finally $\mathbb{A} \preceq_k^{\text{SA}} \mathbb{B}$. \square

Acknowledgments. We would like to thank the anonymous referees of both the conference [12] and this full version of the paper. In particular, one of the journal referees suggested a more concise way of presenting some of the results, which shortened the paper.

REFERENCES

- [1] A. ATSERIAS, A. A. BULATOV, AND V. DALMAU, *On the power of k -consistency*, in Proceedings of the 34th International Colloquium on Automata, Languages and Programming, Lecture Notes in Comput. Sci. 4596, Springer, New York, 2007, pp. 279–290, https://doi.org/10.1007/978-3-540-73420-8_26.
- [2] L. BARTO AND M. KOZIK, *Constraint satisfaction problems solvable by local consistency methods*, J. ACM, 61 (2014), 3, <https://doi.org/10.1145/2556646>.
- [3] L. BARTO, M. KOZIK, AND T. NIVEN, *The CSP dichotomy holds for digraphs with no sources and no sinks (a positive answer to a conjecture of Bang-Jensen and Hell)*, SIAM J. Comput., 38 (2009), pp. 1782–1802, <https://doi.org/10.1137/070708093>.
- [4] U. BERTELÉ AND F. BRIOSCHI, *Nonserial Dynamic Programming*, Academic Press, New York, 1972.
- [5] D. BIENSTOCK AND N. ÖZBAY, *Tree-width and the Sherali-Adams operator*, Discrete Optim., 1 (2004), pp. 13–21, <https://doi.org/10.1016/j.disopt.2004.03.002>.
- [6] H. L. BODLAENDER, *A linear-time algorithm for finding tree-decompositions of small treewidth*, SIAM J. Comput., 25 (1996), pp. 1305–1317, <https://doi.org/10.1137/S0097539793251219>.
- [7] A. BULATOV, *A dichotomy theorem for constraint satisfaction problems on a 3-element set*, J. ACM, 53 (2006), pp. 66–120, <https://doi.org/10.1145/1120582.1120584>.
- [8] A. BULATOV, *A dichotomy theorem for nonuniform CSPs*, in Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science, IEEE, 2017, pp. 319–330.
- [9] A. BULATOV, V. DALMAU, M. GROHE, AND D. MARX, *Enumerating homomorphisms*, J. Comput. System Sci., 78 (2012), pp. 638–650.
- [10] A. BULATOV, P. JEAVONS, AND A. KROKHIN, *Classifying the complexity of constraints using finite algebras*, SIAM J. Comput., 34 (2005), pp. 720–742, <https://doi.org/10.1137/S0097539700376676>.
- [11] A. A. BULATOV, *Complexity of conservative constraint satisfaction problems*, ACM Trans. Comput. Logic, 12 (2011), 24, <https://doi.org/10.1145/1970398.1970400>.
- [12] C. CARBONNEL, M. ROMERO, AND S. ŽIVNÝ, *The complexity of general-valued CSPs seen from the other side*, in Proceedings of the 59th Annual IEEE Symposium on Foundations of Computer Science, IEEE, 2018, pp. 236–246, <https://doi.org/10.1109/FOCS.2018.00031>.
- [13] A. K. CHANDRA AND P. M. MERLIN, *Optimal implementation of conjunctive queries in relational data bases*, in Proceedings of the 9th Annual ACM Symposium on Theory of Computing ACM, 1977, pp. 77–90, <https://doi.org/10.1145/800105.803397>.
- [14] H. CHEN AND S. MENGEL, *A trichotomy in the complexity of counting answers to conjunctive queries*, in Proceedings of the 18th International Conference on Database Theory, 2015, pp. 110–126, <https://doi.org/10.4230/LIPIcs.ICDT.2015.110>.
- [15] D. A. COHEN, M. C. COOPER, P. G. JEAVONS, AND A. A. KROKHIN, *The complexity of soft constraint satisfaction*, Artificial Intelligence, 170 (2006), pp. 983–1016, <https://doi.org/10.1016/j.artint.2006.04.002>.
- [16] V. DALMAU, P. G. KOLAITIS, AND M. Y. VARDI, *Constraint satisfaction, bounded treewidth, and finite-variable logics*, in Proceedings of the 8th International Conference on Principles and Practice of Constraint Programming, Lecture Notes in Comput. Sci. 2470, Springer, New York, 2002, pp. 310–326, https://doi.org/10.1007/3-540-46135-3_21.
- [17] R. DECHTER, *Constraint Processing*, Morgan Kaufmann, Burlington, MA, 2003.
- [18] R. DIESTEL, *Graph Theory*, 4th ed., Springer, New York, 2010.
- [19] R. G. DOWNEY AND M. R. FELLOWS, *Fixed-parameter tractability and completeness I: Basic results*, SIAM J. Comput., 24 (1995), pp. 873–921, <https://doi.org/10.1137/S0097539792228228>.
- [20] R. G. DOWNEY AND M. R. FELLOWS, *Fixed-parameter tractability and completeness II: On completeness for $W[1]$* , Theoret. Comput. Sci., 141 (1995), pp. 109–131.
- [21] T. FÄRNQVIST, *Constraint optimization problems and bounded tree-width revisited*, in Proceedings of the 9th International Conference on Integration of Artificial Intelligence and Operations Research Techniques in Constraint Programming, Lecture Notes in Comput. Sci. 7298, Springer, New York, 2012, pp. 163–197, https://doi.org/10.1007/978-3-642-29828-8_11.

- [22] T. FÄRNQVIST, *Exploiting Structure in CSP-related Problems*, Ph.D. thesis, Department of Computer Science and Information Science, Linköping University, 2013.
- [23] T. FÄRNQVIST AND P. JONSSON, *Bounded tree-width and CSP-related problems*, in Proceedings of the 18th International Symposium on Algorithms and Computation, Lecture Notes in Comput. Sci. 4835, Springer, New York, 2007, pp. 632–643, https://doi.org/10.1007/978-3-540-77120-3_55.
- [24] T. FEDER AND M. Y. VARDI, *The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory*, SIAM J. Comput., 28 (1998), pp. 57–104, <https://doi.org/10.1137/S0097539794266766>.
- [25] J. FLUM AND M. GROHE, *Parameterized Complexity Theory*, Springer, New York, 2006.
- [26] E. C. FREUDER, *A sufficient condition for backtrack-free search*, J. ACM, 29 (1982), pp. 24–32.
- [27] E. C. FREUDER, *Complexity of K-tree structured constraint satisfaction problems*, in Proceedings of the 8th National Conference on Artificial Intelligence, 1990, pp. 4–9.
- [28] G. GOTTLÖB, G. GRECO, N. LEONE, AND F. SCARCELLO, *Hypertree decompositions: Questions and answers*, in Proceedings of the 35th SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, 2016, pp. 57–74.
- [29] G. GOTTLÖB, G. GRECO, AND F. SCARCELLO, *Tractable optimization problems through hypergraph-based structural restrictions*, in Proceedings of the 36th International Colloquium on Automata, Languages and Programming, Part II, Lecture Notes in Comput. Sci. 5556, Springer, New York, 2009, pp. 16–30, https://doi.org/10.1007/978-3-642-02930-1_2.
- [30] T. J. GREEN, *Containment of conjunctive queries on annotated relations*, Theory Comput. Syst., 49 (2011), pp. 429–459.
- [31] T. J. GREEN AND V. TANNEN, *The semiring framework for database provenance*, in Proceedings of the 36th SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, 2017, pp. 93–99.
- [32] M. GROHE, *The complexity of homomorphism and constraint satisfaction problems seen from the other side*, J. ACM, 54 (2007), 1, <https://doi.org/10.1145/1206035.1206036>.
- [33] M. GRÖTSCHTEL, L. LOVÁSZ, AND A. SCHRIJVER, *Geometric Algorithms and Combinatorial Optimization*, Algorithms Combin. 2, Springer, New York, 1988.
- [34] P. HELL AND J. NEŠETŘIL, *The core of a graph*, Discrete Math., 109 (1992), pp. 117–126, [https://doi.org/https://doi.org/10.1016/0012-365X\(92\)90282-K](https://doi.org/https://doi.org/10.1016/0012-365X(92)90282-K).
- [35] P. HELL AND J. NEŠETŘIL, *Graphs and Homomorphisms*, Oxford University Press, New York, 2004.
- [36] P. M. IDZIAK, P. MARKOVIC, R. MCKENZIE, M. VALERIOTE, AND R. WILLARD, *Tractability and learnability arising from algebras with few subpowers*, SIAM J. Comput., 39 (2010), pp. 3023–3037, <https://doi.org/10.1137/090775646>.
- [37] P. G. KOLAITIS AND M. Y. VARDI, *Conjunctive-query containment and constraint satisfaction*, in Proceedings of the 17th SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, 1998, pp. 205–213, <https://doi.org/10.1145/275487.275511>.
- [38] P. KOLMAN AND M. KOUTECKÝ, *Extended formulation for CSP that is compact for instances of bounded treewidth*, Electron. J. Combin., 22 (2015), P4.30.
- [39] V. KOLMOGOROV, A. A. KROKHIN, AND M. ROLÍNEK, *The complexity of general-valued CSPs*, SIAM J. Comput., 46 (2017), pp. 1087–1110, <https://doi.org/10.1137/16M1091836>.
- [40] E. V. KOSTYLEV, J. L. REUTTER, AND A. Z. SALAMON, *Classification of annotation semirings over containment of conjunctive queries*, ACM Trans. Database Syst., 39 (2014), 1, <https://doi.org/10.1145/2556524>.
- [41] M. KOZIK AND J. OCHREMIK, *Algebraic properties of valued constraint satisfaction problem*, in Proceedings of the 42nd International Colloquium on Automata, Languages and Programming, Lecture Notes in Comput. Sci. 9134, Springer, New York, 2015, pp. 846–858, https://doi.org/10.1007/978-3-662-47672-7_69.
- [42] D. MARX, *Tractable hypergraph properties for constraint satisfaction and conjunctive queries*, J. ACM, 60 (2013), 42, <https://doi.org/10.1145/2535926>.
- [43] U. MONTANARI, *Networks of constraints: Fundamental properties and applications to picture processing*, Inform. Sci., 7 (1974), pp. 95–132, [https://doi.org/10.1016/0020-0255\(74\)90008-5](https://doi.org/10.1016/0020-0255(74)90008-5).
- [44] N. ROBERTSON AND P. D. SEYMOUR, *Graph minors. III. Planar tree-width*, J. Combin. Theory Ser. B, 36 (1984), pp. 49–64, [https://doi.org/10.1016/0095-8956\(84\)90013-3](https://doi.org/10.1016/0095-8956(84)90013-3).
- [45] N. ROBERTSON AND P. D. SEYMOUR, *Graph minors. V. Excluding a planar graph*, J. Combin. Theory Ser. B, 41 (1986), pp. 92–114, [https://doi.org/10.1016/0095-8956\(86\)90030-4](https://doi.org/10.1016/0095-8956(86)90030-4).
- [46] A. SCHRIJVER, *Theory of Linear and Integer Programming*, John Wiley & Sons, New York, 1986.
- [47] P. D. SEYMOUR AND R. THOMAS, *Graph searching and a min-max theorem for tree-width*, J. Combin. Theory Ser. B, 58 (1993), pp. 22–33.

- [48] H. D. SHERALI AND W. P. ADAMS, *A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems*, SIAM J. Discrete Math., 3 (1990), pp. 411–430.
- [49] J. THAPPER AND S. ŽIVNÝ, *The power of linear programming for valued CSPs*, in Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science, IEEE, 2012, pp. 669–678, <https://doi.org/10.1109/FOCS.2012.25>.
- [50] J. THAPPER AND S. ŽIVNÝ, *The complexity of finite-valued CSPs*, J. ACM, 63 (2016), 37, <https://doi.org/10.1145/2974019>.
- [51] J. THAPPER AND S. ŽIVNÝ, *The power of Sherali-Adams relaxations for general-valued CSPs*, SIAM J. Comput., 46 (2017), pp. 1241–1279, <https://doi.org/10.1137/16M1079245>.
- [52] D. ZHUK, *A proof of the CSP dichotomy conjecture*, J. ACM, 67 (2020), 30, <https://doi.org/10.1145/3402029>.