

The Complexity of Weighted Boolean #CSP with Mixed Signs^{*†}

Andrei Bulatov,[‡] Martin Dyer,[§] Leslie Ann Goldberg,[¶]
Markus Jalsenius[¶] and David Richerby[§]

December 24, 2008

Abstract

We give a complexity dichotomy for the problem of computing the partition function of a weighted Boolean constraint satisfaction problem. Such a problem is parameterized by a set Γ of rational functions, each of which assigns a weight to each variable assignment. Our dichotomy extends previous work in which the weight functions were restricted to being non-negative. We represent a weight function as a product of the form $(-1)^s g$, where the polynomial s determines the sign of the weight and the non-negative function g determines its magnitude. We show that the problem of computing the partition function (the sum of the weights of all possible variable assignments) is computable in polynomial time if either every function in Γ can be defined by a “pure affine” magnitude with a quadratic sign polynomial or every function can be defined by a magnitude of “product type” with a linear sign polynomial. In all other cases, computing the partition function is $\text{FP}^{\#P}$ -complete.

1 Introduction

The principal result of this paper is a dichotomy theorem for the complexity of computing the partition function of a weighted Boolean constraint satisfaction problem. This problem has a set of functions Γ that are used to assign a *weight* to any configuration. (A configuration is an assignment of values to the *variables* in the instance.) These functions generalise constraint relations in the classical constraint satisfaction problem (CSP). In the classical setting, all functions in Γ have codomain $\{0, 1\}$. The problem we consider here is, given an instance of a weighted CSP, compute its *partition function*, that is, the sum of weights of all configurations. Computing the partition function generalises

* Subsequent to the completion of this work, the authors learned that J.-Y. Cai, P. Lu and M. Xia have recently (and independently) obtained a dichotomy for Boolean #CSP with *complex* weights, using holographic techniques.

† Funded in part by the EPSRC grant “The Complexity of Counting in Constraint Satisfaction Problems”.

‡ School of Computing Science, Simon Fraser University, Burnaby, Canada.

§ School of Computing, University of Leeds, Leeds, LS2 9JT, UK.

¶ Department of Computer Science, University of Liverpool, Liverpool, L69 3BX, UK.

the problem of counting the number of satisfying solutions of a CSP. We will denote by $\#\text{CSP}(\Gamma)$ the problem of computing the partition function of weighted CSP instances which use functions from the set Γ .

The term “partition function” originates in statistical physics, and certain problems from statistical physics may be expressed as weighted CSPs. For example, the *Potts model* [21] can be expressed a weighted CSP, whereas only the “hard core” version can be expressed as a classical CSP. The two possible hard core versions of the Potts model correspond to graph colouring, in the *antiferromagnetic* case, and the trivial problem of colouring each component of a graph with a single colour, in the *ferromagnetic* case.

Here, we extend the work of Dyer, Goldberg and Jerrum [7], who gave a dichotomy theorem for the complexity of $\#\text{CSP}(\Gamma)$ when every function in Γ is restricted to have non-negative values. They defined two classes of functions, those that are “pure affine” and those “of product type”, and showed that $\#\text{CSP}(\Gamma)$ is in FP if and only if all functions in Γ are either pure affine or of product type. Otherwise $\#\text{CSP}(\Gamma)$ is complete for $\text{FP}^{\#\text{P}}$. The existence of algorithms for testing the properties of being purely affine or of product type means that the dichotomy is effectively decidable.

The contribution of this paper is a dichotomy theorem for $\#\text{CSP}(\Gamma)$ where Γ is allowed to contain function which give values of either sign. This extension is of particular interest because functions having mixed signs can cause cancellations in the partition function, and hence may make it easier to compute. Many natural problems can be expressed as weighted $\#\text{CSP}$ problems with functions of mixed signs. (See [11] for some examples.) We give a dichotomy theorem showing that $\#\text{CSP}(\Gamma)$ is either in FP or is complete for $\text{FP}^{\#\text{P}}$. As in the non-negative case, there are algorithms for determining whether $\#\text{CSP}(\Gamma)$ is in FP or $\text{FP}^{\#\text{P}}$ -complete, so the dichotomy is decidable.

1.1 Constraint Satisfaction

Constraint Satisfaction provides a general framework for modelling decision problems, and has many practical applications, particularly in Artificial Intelligence. (See, for example [18].) Decisions are modelled by *variables*, which are subject to *constraints*, which model the logical and resource restrictions. Many interesting problems can be modelled in this way, from satisfiability problems to scheduling problems to graph-theoretic problems. Consequently, studying the computational complexity of constraint satisfaction problems has become a major and active area [5, 14].

A Constraint Satisfaction Problem (CSP) has a finite *domain*, which we may denote by $\{0, 1, \dots, q - 1\}$ for some positive integer q . In this paper we are interested only in the *Boolean* case, where $q = 2$. A *constraint language* Γ with domain $\{0, 1, \dots, q - 1\}$ is a set of relations on $\{0, 1, \dots, q - 1\}$. For example, let $q = 2$, and consider the relation $R = \{(1, 0, 0), (1, 0, 1), (0, 1, 0), (0, 1, 1)\}$. This is a 3-ary relation on the domain $\{0, 1\}$, having four tuples.

Given a constraint language Γ , an *instance* of the $\text{CSP}(\Gamma)$ is a set of *variables* $V = \{v_1, \dots, v_n\}$ and a set of *constraints*. Each constraint has a *scope*, which is a tuple of variables (for example, (v_3, v_7, v_2)) and a relation from Γ of the same arity, which constrains the variables in the scope. A *configuration* σ is a function from V to $\{0, 1, \dots, q - 1\}$. The configuration σ is *satisfying* if the scope of every constraint is mapped to a tuple that is in the corresponding relation. In our example above, a con-

figuration σ satisfies the constraint with scope (v_3, v_7, v_2) and relation R if and only if it maps either v_3 or v_7 , but not both, to the value 1. Given an instance I of a CSP with constraint language Γ , the *decision problem* $\text{CSP}(\Gamma)$ is to determine whether any configuration satisfies I . The *counting problem* $\#\text{CSP}(\Gamma)$ is to determine the *number* of (distinct) satisfying configurations of I .

Varying the constraint language Γ defines the classes CSP and $\#\text{CSP}$ of decision and counting problems. These contain problems of very different computational complexity. For example, if $\Gamma = \{R_1, R_2, R_3\}$ where R_1, R_2 and R_3 are the three binary relations defined by $R_1 = \{(0, 1), (1, 0), (1, 1)\}$, $R_2 = \{(0, 0), (0, 1), (1, 1)\}$ and $R_3 = \{(0, 0), (0, 1), (1, 0)\}$, then $\text{CSP}(\Gamma)$ is the classical 2-Satisfiability problem, which is in P . On the other hand, there is a similar constraint language Γ' such that 3-Satisfiability (which is NP -complete) can be expressed as $\text{CSP}(\Gamma')$. It is possible that the counting problem is harder than the decision problem. If Γ is the constraint language defining 2-Satisfiability, then $\#\text{CSP}(\Gamma)$ contains the problem of counting independent sets in graph, which is $\#\text{P}$ -complete [20] even for 3-regular graphs [13].

Any problem in CSP is in NP , but not every problem in NP can be expressed in CSP . For example, the question “Is the graph G Hamiltonian?” cannot be expressed in CSP , because the property of being Hamiltonian cannot be captured by constraints of fixed size. This is a limitation of the class CSP , but it also has an advantage. If $\text{P} \neq \text{NP}$, there are problems which are neither in P nor NP -complete [15]. But, for smaller classes of decision problems, the situation may be more straightforward. A dichotomy theorem may be possible, partitioning all problems in the class into those which are in P and those which are NP -complete, with no problems of intermediate complexity. It has been conjectured, in a seminal paper of Feder and Vardi [10], that there is a dichotomy theorem for CSP . Although much progress has been made towards proving this, it remains unproven to date.

In the Boolean case, the status of CSP was resolved by Schaefer [19]. Schaefer stated a theorem for the domain $\{0, 1\}$, giving four conditions on the constraint language Γ . If any of the conditions holds then $\text{CSP}(\Gamma)$ is in P , otherwise $\text{CSP}(\Gamma)$ is NP -complete. For details, the interested reader is referred to the paper [19] or to Theorem 6.2 of the textbook [5]. An interesting feature is that Schaefer’s conditions are all algorithmically checkable. Thus, given a constraint language Γ with domain $\{0, 1\}$, we can determine whether $\text{CSP}(\Gamma)$ is in P or NP -complete.

While the conjectured dichotomy for CSP remains open, Bulatov [1] has recently made a major breakthrough for $\#\text{CSP}$. He has shown that there is a dichotomy between FP and $\#\text{P}$ -complete, for the whole of $\#\text{CSP}$. However, his proof [1] sheds very little light on which constraint languages Γ render $\#\text{CSP}(\Gamma)$ in FP , and which Γ render it $\#\text{P}$ -hard. The difficulty is that, while Γ itself is of fixed size, the criterion of [1] involves finding a defect in any of a potentially infinite class of structures built on Γ . Whether the criterion is algorithmically checkable is an open question.

In the Boolean case, which is our focus here, a dichotomy theorem for $\#\text{CSP}$ had already been established by Creignou and Hermann [4], and the dichotomy is decidable. Before stating their theorem we introduce the following definition. A Boolean relation R is said to be *affine* if the set of tuples $x \in R$ is the set of solutions to a system of linear equations over $\text{GF}(2)$. A constraint language Γ is affine if every relation $R \in \Gamma$ is affine. The dichotomy theorem by Creignou and Hermann [4] states that $\#\text{CSP}(\Gamma)$

is in FP if Γ is affine, otherwise $\#\text{CSP}(\Gamma)$ is $\#\text{P}$ -complete. There is an algorithm for determining whether a Boolean constraint language Γ is affine, so there is an algorithm for determining whether $\#\text{CSP}(\Gamma)$ is in FP or $\#\text{P}$ -complete. In addition to Creignou and Hermann’s dichotomy, Dyer, Goldberg and Jerrum [6] have given an approximation trichotomy for Boolean $\#\text{CSP}$. Let $\#\text{BIS}$ denote the problem of counting number of independent sets in a bipartite graph and let $\#\text{SAT}$ denote the problem of counting satisfying assignments to a Boolean formula in conjunctive normal form. Dyer, et al. [6] have shown that if Γ is not affine (hence $\#\text{CSP}(\Gamma)$ is $\#\text{P}$ -complete) then there is an approximation-preserving reduction between $\#\text{CSP}(\Gamma)$ and either $\#\text{BIS}$ or $\#\text{SAT}$.

1.2 Weighted $\#\text{CSP}$

The counting problem $\#\text{CSP}(\Gamma)$ can be naturally extended by replacing the relations in Γ by *functions*. We refer to the corresponding class of problems as *weighted $\#\text{CSP}$* . The functions are used to assign weights to configurations. The *partition function* computes the sum of the weights over all configurations. We give a formal definition below. The partition function of a weighted $\#\text{CSP}$ generalises the number of satisfying solutions of a classical $\#\text{CSP}$. The classical setting may be recovered by restricting the codomain of every function to $\{0, 1\}$.

In weighted $\#\text{CSP}$, a *constraint language* over a finite domain \mathcal{D} is a finite collection of functions $\Gamma \subseteq \{f_i: \mathcal{D}^{r_i} \rightarrow \mathbb{Q} \mid i \in \{1, 2, \dots\}\}$. The natural number r_i is called the *arity* of the function f_i ; we refer to functions of arity one, two and three as *unary*, *binary* and *ternary*, respectively. In this paper we consider exclusively the Boolean domain, where $\mathcal{D} = \{0, 1\}$.

An *instance* of a weighted constraint satisfaction problem over a constraint language Γ is a pair $I = (V, C)$, where $V = \{v_1, \dots, v_n\}$ is a set of *variables* and C is a finite set of *constraints*. Each constraint is of the form $f(v_{i_1}, \dots, v_{i_r})$, where f is an r -ary function in the set Γ . To keep notation simple, we will often use x_1, x_2, \dots as “metavariables”, standing for variables in V .

A *configuration* of an instance (V, C) is a function $\sigma: V \rightarrow \mathcal{D}$, assigning a value from the domain to each variable. The *weight* of a configuration σ is

$$W(\sigma) := \prod_{f(x_1, \dots, x_r) \in C} f(\sigma(x_1), \dots, \sigma(x_r)).$$

We are interested in computing the *partition function* of an instance I . This is the sum $Z(I)$ of the weights of all possible configurations:

$$Z(I) := \sum_{\sigma: V \rightarrow \mathcal{D}} W(\sigma).$$

The *weighted constraint satisfaction problem* is the problem of computing $Z(I)$ given an instance I . Since this paper is exclusively about weighted CSPs, we will drop the word weighted and write $\#\text{CSP}(\Gamma)$ for the weighted constraint satisfaction problem over the constraint language Γ and $\#\text{CSP}$ for the union of $\#\text{CSP}(\Gamma)$ over all rational-weighted constraint languages. For constraint languages with only a single function f , we write $\#\text{CSP}(f)$, rather than $\#\text{CSP}(\{f\})$.

A constraint $f(x_1, \dots, x_r)$ is *satisfied* by a configuration σ if $f(\sigma(x_1), \dots, \sigma(x_r)) \neq 0$. Therefore, the weight of a configuration is zero if it does not satisfy some constraint. If we restrict to constraint languages where every function has codomain $\{0, 1\}$, the weight of every configuration is either zero or one and $Z(I)$ is just the number of satisfying configurations for I . This corresponds precisely to the counting constraint satisfaction problem.

1.3 Related work

Bulatov's counting dichotomy [1] can be extended to weighted $\#\text{CSP}$ as long as the range of every function $f \in \Gamma$ is $\mathbb{Q}^{\geq 0}$ (the set of non-negative rationals) [2]. However, it is not known whether it extends to weighted $\#\text{CSP}$ with functions of mixed signs. Furthermore, in the current state of knowledge, there is no algorithm to determine whether $\#\text{CSP}(\Gamma)$ is in FP or $\text{FP}^{\#\text{P}}$ -complete, given a constraint language Γ . For the special case of *graph homomorphisms*, an effective dichotomy is known for functions of mixed signs [11]. There is also a dichotomy theorem by Dyer et al. for Boolean weighted $\#\text{CSP}$ for functions that are non-negative [7]. This is expressed in terms of two classes of functions: *pure affine* and *product type*. The definitions of these two classes are given in Section 2.1. Then Dyer et al. [7] gave the following theorem.

Theorem 1 (Theorem 4, [7]). *Suppose Γ is a constraint language such that the range of every function $f \in \Gamma$ is a set of non-negative rationals. If every function in Γ is pure affine then $\#\text{CSP}(\Gamma)$ is in FP . If every function in Γ is of product type then $\#\text{CSP}(\Gamma)$ is in FP . Otherwise, $\#\text{CSP}(\Gamma)$ is $\text{FP}^{\#\text{P}}$.*

There exist algorithms for testing whether a Boolean constraint language Γ is pure affine or of product type. This means that the dichotomy is effectively decidable.

The contribution of this paper (Theorem 2 below) extends Theorem 1 to constraint languages Γ containing arbitrary rational-valued functions. This is an interesting extension since functions with negative values can cause cancellations and may make the partition function easier to compute¹.

The case of mixed signs has been considered previously by Goldberg, Grohe, Jerrum and Thurley [11], in the case of one symmetric binary function on an arbitrary finite domain. Their theorem generalises that of Bulatov and Grohe [3] for the non-negative case. Goldberg et al. [3] give two examples, which can also be expressed as Boolean weighted $\#\text{CSP}$, and fall within the scope of this paper. The first appeared as an open problem in [3]. The complexity of these problems can be deduced from [11] and from the results of this paper.

Example 1. The first example in [11] is the function $f : \{0, 1\} \rightarrow \{-1, 1\}$, where

$$\begin{array}{ll} f(0, 0) = 1 & f(0, 1) = 1 \\ f(1, 0) = 1 & f(1, 1) = -1. \end{array}$$

An instance I of $\#\text{CSP}(f)$ can be represented by a graph $G = (V, E)$ with n vertices². The set of variables in I is V and for each edge $(u, v) \in E$ we have the constraint $f(u, v)$

¹In a related context, recall the sharp distinction in complexity between computing the permanent and the determinant of a matrix.

²In fact G can be a multigraph here, and may have self-loops, but the argument is unaffected.

in I . Then $\frac{1}{2}Z(G) + 2^{n-1}$ is the number of induced subgraphs of G with an even number of edges [11]. Hence up to a simple transformation, the partition function $Z(G)$ counts induced subgraphs with an even number of edges. To see this, observe that for every configuration σ , the term $\prod_{(u,v) \in E} f(u,v)$ is 1 if the subgraph of G induced by $\sigma^{-1}(1)$ has an even number of edges and -1 otherwise. In terms of our Theorem 2 below, $f(x,y) = (-1)^{xy}$, so this problem is in FP, and an algorithm for computing $Z(G)$ follows from Lemma 2 below.

Example 2. The second example in [11] is $\#\text{CSP}(f)$, where

$$\begin{aligned} f(0,0) &= 1 & f(0,1) &= -1 \\ f(1,0) &= -1 & f(1,1) &= 1. \end{aligned}$$

In terms of Theorem 2 below, $f(x,y) = (-1)^{x+y}$, so this problem is also in FP. This can easily be shown directly. Let $G = (V, E)$ be a graph with n vertices. Note that $Z(G)$ is unchanged by removing any circuit from G . Thus we may reduce G to a forest F , which will have edges if and only if G was not Eulerian. If F has no edges, then $Z(G) = Z(F) = 2^n$. Otherwise F has at least one leaf vertex v , and then we have $Z(G) = Z(F) = Z(F \setminus v) - Z(F \setminus v) = 0$. Thus $Z(G) = 0$ unless G is Eulerian, in which case $Z(G) = 2^n$, and hence the problem is trivially in FP.

1.4 Complexity

Since our weights are arbitrary rationals, the partition function Z is not, in general, an integer function. As such, $\#\text{CSP}(\Gamma)$ is not, in general, in the class $\#\text{P}$. However, it is easy to see that, for every constraint language Γ , there is a partition function Z' in $\#\text{P}$ and an FP-computable integer function K such that, for all instances I of $\#\text{CSP}(\Gamma)$, $Z(I) = Z'(I)/K(I)$. This is achieved by “clearing denominators”, see [12].

Following [12], we write $\#\text{P}_{\mathbb{Q}}$ for the class of functions of the form f/g , where $f \in \#\text{P}$ and $g \in \text{FP}$. It is immediate that

$$\#\text{CSP} \subseteq \#\text{P}_{\mathbb{Q}} \subseteq \text{FP}^{\#\text{P}}.$$

Proposition 1. *Every $\#\text{CSP}$ problem that is $\#\text{P}$ -hard is $\text{FP}^{\#\text{P}}$ -complete.*

Proof. If $Z(I) \in \#\text{CSP}$ is $\#\text{P}$ -hard, we can use an oracle for $Z(I)$ to construct an oracle for $Z'(I)$, as described above. With this oracle, we can compute any problem in $\text{FP}^{\#\text{P}}$ in a polynomial number of steps. Therefore, $Z(I)$ is $\text{FP}^{\#\text{P}}$ -complete. \square

Let Γ be a constraint language. We say that Γ *simulates* a function $f \notin \Gamma$ if, given an instance I of $\#\text{CSP}(\Gamma \cup \{f\})$, we can construct, in polynomial time, an instance I' of $\#\text{CSP}(\Gamma)$ such that $Z(I) = K(I)Z(I')$ for some FP-computable function K . This generalizes parsimonious reductions [17]; clearly, if Γ simulates f then $\#\text{CSP}(\Gamma \cup \{f\}) \leq_{\text{T}} \#\text{CSP}(\Gamma)$, where \leq_{T} denotes polynomial-time Turing reducibility. If $\#\text{CSP}(\Gamma) \leq_{\text{T}} \#\text{CSP}(\Gamma') \leq_{\text{T}} \#\text{CSP}(\Gamma)$, we write $\#\text{CSP}(\Gamma) \equiv_{\text{T}} \#\text{CSP}(\Gamma')$.

1.5 Organization of the paper

Our paper is organized as follows. In Section 2, we define notation and the classes of functions we will use throughout the paper. In Section 3, we state our dichotomy result and prove the polynomial-time cases. The following sections prove that all other cases are $\text{FP}^{\#\text{P}}$ -complete. We give useful tools for proving hardness in Sections 4. In Sections 5 and 6, we show, respectively, that any constraint language containing a pure affine function of degree greater than 2 is $\#\text{P}$ -hard and that any language with a function of product type of degree greater than 1 can be made $\#\text{P}$ -hard by adding a simple function. Finally, we complete the proof of the dichotomy in Section 7, showing that the simple function is definable in terms of the functions already present.

2 Some notation

All sets and other objects referred to in this paper are finite unless it is stated otherwise. We write \bar{a} for a tuple of elements (a_1, \dots, a_r) for some r and, for natural numbers $m \leq n$, write $[m, n]$ for the set $\{m, m+1, \dots, n\}$.

The *support* of a function $f: X^r \rightarrow \mathbb{Q}$ is the r -ary relation $\{\bar{a} \mid f(\bar{a}) \neq 0\}$. For a tuple $\bar{a} \in X^r$, we write $f(\bar{a})$ for the tuple $(f(a_1), \dots, f(a_r))$.

We write $\mathcal{F}_{\mathbb{B}}$ for the set of all functions, of all positive arities, from the set $\{0, 1\}$ to \mathbb{Q} , the rationals, and $\mathcal{F}_{\mathbb{B}}^{\geq 0}$ for the subset of $\mathcal{F}_{\mathbb{B}}$ consisting of all functions with non-negative ranges. We write \mathcal{P}_k for the set of multivariate polynomials in variables x_1, \dots, x_k over $\text{GF}(2)$. We sometimes write $p(x_1, \dots, x_k)$ for a polynomial $p \in \mathcal{P}_k$ or other function, to emphasize that p is a function of those variables.

A function $f(x_1, \dots, x_k)$ *depends on* a variable x_i if there are constants $c_1, \dots, c_{i-1}, c_{i+1}, \dots, c_k \in \{0, 1\}$ such that

$$f(c_1, \dots, c_{i-1}, 0, c_{i+1}, \dots, c_k) \neq f(c_1, \dots, c_{i-1}, 1, c_{i+1}, \dots, c_k).$$

2.1 Classes of functions

In this section, we define the classes of functions we will use throughout the paper. Our definitions of pure affine functions and functions of product type are those used by Dyer et al. [7] but multiplied by a term $(-1)^s$ for some polynomial s , which determines the sign.

Recall that a relation over $\{0, 1\}$ is *affine* if it is the solution set of a set of linear equations over $\text{GF}(2)$. We say that a function $f \in \mathcal{F}_{\mathbb{B}}$ is affine if it has affine support.

Definition 1. A k -ary function $f \in \mathcal{F}_{\mathbb{B}}$ is *pure affine* if there is a constant $w \in \mathbb{Q}^{>0}$, an affine function $g \in \mathcal{P}_k$ and a polynomial $s \in \mathcal{P}_k$ such that

$$f(\bar{x}) = w(-1)^{s(\bar{x})}g(\bar{x}). \tag{1}$$

Note that the range of f in (1) is included in $\{-w, 0, w\}$. The polynomial g is uniquely defined, up to the identities $x \oplus x = 0$ and $x^2 = x$. However, because the value of f does not depend on the value of s for values of its inputs where $g(\bar{x}) = 0$, there may be several distinct polynomials s for which the identity (1) holds. If s is of minimal degree d such

that (1) holds, we say that s is *degree-minimized* with respect to g and that f is *pure affine of degree d* . For the purposes of this paper, we consider the constant zero and one polynomials to have degree zero.

We write $\chi_ =$ and χ_{\neq} for the binary equality and disequality functions, respectively, defined as

$$\chi_=(x, y) = x \oplus y \oplus 1 \qquad \chi_{\neq}(x, y) = x \oplus y.$$

Definition 2. A k -ary function $f \in \mathcal{F}_{\mathbb{B}}$ is of *product type* if there are unary functions $U_1(x_1), \dots, U_k(x_k): \{0, 1\} \rightarrow \mathbb{Q}^{>0}$, a polynomial $g \in \mathcal{P}_k$ that is a product of binary functions of the form $\chi_ =$ and χ_{\neq} , and a polynomial $s \in \mathcal{P}_k$ such that

$$f(\bar{x}) = (-1)^{s(\bar{x})} U_1(x_1) \cdots U_k(x_k) g(\bar{x}). \quad (2)$$

The function f is of *product type of degree d* if s is of degree d and is degree-minimized with respect to $U_1(x_1) \cdots U_k(x_k) g(\bar{x})$.

Let $f \in \mathcal{F}_{\mathbb{B}}$ be of product type and let $(-1)^{s(\bar{x})} U_1(x_1) \cdots U_k(x_k) g(\bar{x})$ be an expression of f as in the definition. We call a variable x_i in the representation *determined* if exactly one of the terms in g is an equality or disequality involving x_i , $U_i(0) = U_i(1) = 1$ and s does not depend on x_i .

It is convenient to impose certain restrictions on expressions for functions of product type. We say that the expression for f is *normalized* if the following conditions are met:

- at least one variable in every equality and disequality term in g is determined,
- if, for some i , $U_i(0) = 0$ or $U_i(1) = 0$, then g and s do not depend on x_i , and
- s is degree-minimized with respect to $U_1(x_1) \cdots U_k(x_k) g(\bar{x})$.

Lemma 1. *Every function $f \in \mathcal{F}_{\mathbb{B}}$ that is of product type is defined by a normalized expression.*

Proof. Let $(-1)^{s(\bar{x})} U_1(x_1) \cdots U_k(x_k) g(\bar{x})$ be a non-normalized expression defining f .

Suppose g contains a term $\chi_=(x_i, x_j)$ where neither x_i nor x_j is determined. First, substitute x_i for x_j in every other term of g and in s . Replace U_j with the function that maps both 0 and 1 to 1 and U_i with the function $U_i(x_i)U_j(x_i)$. The variable x_j is now determined in the resulting expression, which still defines f .

Suppose g contains a term $\chi_{\neq}(x_i, x_j)$ where neither x_i nor x_j is determined. We proceed as above but substitute $x_i \oplus 1$ for x_j . Having done so, there may be terms $\chi_=(x_\ell, x_i \oplus 1)$ and $\chi_{\neq}(x_\ell, x_i \oplus 1)$; replace these with $\chi_{\neq}(x_\ell, x_i)$ and $\chi_=(x_\ell, x_i)$, respectively, and similarly for the terms with the parameters the other way round.

Suppose that $U_i(c) = 0$ for some $c \in \{0, 1\}$ but g or s depends on x_i . Since f is zero if $x_i = c$, we may replace x_i with $c \oplus 1$ throughout g and s . Performing such a replacement in a term of g results in that term becoming a unary function, which can be incorporated into the corresponding U_j .

Finally, if s is not degree-minimized, replace it with a polynomial in the appropriate variables that is. \square

We say that a k -ary function $f: \{0, 1\}^k \rightarrow \mathbb{Q}$ is *positive pure affine* or of *positive product type* if it can be written according to Definition 1 or Definition 2, respectively, but choosing the sign polynomial s to be identically zero. Thus, positive pure affine and positive product type correspond exactly to the definitions of pure affine and product type used by Dyer et al. for functions $\{0, 1\}^k \rightarrow \mathbb{Q}^{\geq 0}$ [7]. Observe that, if a function $f: \{0, 1\}^k \rightarrow \mathbb{Q}^{\geq 0}$ is pure affine (respectively, of product type) then is positive pure affine (respectively, of positive product type). This is because we must have $s(\bar{x}) = 0$ whenever $f(\bar{x}) \neq 0$ and, when $f(\bar{x}) = 0$, we can set $s(\bar{x}) = 0$ without altering the value of f . Thus, all properties of the functions that Dyer et al. call “pure affine” or “of product type” in [7] carry over to non-negative functions that we call pure affine and of product type, respectively.

3 The dichotomy

We now give our main result, a complexity dichotomy for Boolean $\#\text{CSP}$ with rational weights. In this section, we prove the tractability of the polynomial-time cases and comment on our definitions of the classes of pure affine and product-type functions. Proving $\text{FP}^{\#\text{P}}$ -completeness of the remaining cases requires considerably more work and is the subject of the remainder of the paper.

Theorem 2. *Let $\Gamma \subseteq \mathcal{F}_{\mathbb{B}}$. If every function in Γ is pure affine of degree at most 2, then $\#\text{CSP}(\Gamma)$ is in FP . If every function in Γ is of product type of degree at most 1, then $\#\text{CSP}(\Gamma)$ is in FP . Otherwise, $\#\text{CSP}(\Gamma)$ is $\text{FP}^{\#\text{P}}$ -complete.*

Proof. The two polynomial-time cases are covered by Lemmas 2 and 4 in this section. If we are not in one of these cases, then Γ must contain functions f and g (not necessarily distinct) such that f is not pure affine of degree at most 2 and g is not of product type of degree at most 1. $\text{FP}^{\#\text{P}}$ -completeness follows from Lemmas 5 and 22. \square

Following from the observations at the end of the previous section, if we have $\Gamma \subseteq \mathcal{F}_{\mathbb{B}}^{\geq 0}$, then Theorem 2 is equivalent to Theorem 4 of Dyer et al. [7].

It is worth pointing out that we cannot simply dispense with the “sign polynomial” in the definitions of pure affine and product type and, instead, allow the constants and unary functions to take negative values. Temporarily call a function $f: \{0, 1\}^k \rightarrow \mathbb{Q}$ *weakly pure affine* if there is a constant $w \in \mathbb{Q}$ and an affine polynomial $g \in \mathcal{P}_k$ such that $f(\bar{x}) = wg(\bar{x})$ and of *weak product type* if there are unary functions $U_i: \{0, 1\} \rightarrow \mathbb{Q}$ and a product g of equalities and disequalities such that $f(\bar{x}) = U_1(x_1) \cdots U_k(x_k)g(\bar{x})$. It is not hard to see that every function that is weakly pure affine or of weak product type is pure affine or of product type, respectively. However, the converse does not hold. The function $f(x, y) = (-1)^{xy}$ of Example 1 above is not weakly pure affine (there is no rational w such that its codomain is $\{0, w\}$) and not of weak product type (it is nowhere zero so there can be no non-trivial equality or disequality terms and the sign cannot be expressed as a combination of unary functions). However, it is trivially pure affine and of product type (of degree two in both cases).

Lemma 2. *Let $\Gamma \subseteq \mathcal{F}_{\mathbb{B}}$. If every function in Γ is pure affine of degree at most 2, then $\#\text{CSP}(\Gamma) \in \text{FP}$.*

Proof. Let $\Gamma = \{f_1, \dots, f_m\}$, where each $f_i = w_i(-1)^{s_i}g_i$ and let $\Gamma' = \{f'_1, \dots, f'_m\}$, where each $f'_i = f_i/w_i = (-1)^{s_i}g_i$. Note that the range of each f'_i is included in $\{-1, 0, 1\}$.

Let I be an instance of $\#\text{CSP}(\Gamma)$ and, for each $i \in [1, m]$, let k_i be the number of constraints in I that involve the function f_i . Let I' be the instance of $\#\text{CSP}(\Gamma')$ made by replacing each constraint $f(\bar{x})$ in I with $f'(\bar{x})$. We have

$$Z(I) = Z(I') \prod_{1 \leq i \leq m} w_i^{k_i},$$

so it suffices to show that we can compute $Z(I')$ in a polynomial number of steps.

If there are k constraints and n variables in I' , $Z(I')$ is a sum of terms of the form

$$\prod_{1 \leq j \leq k} (-1)^{s_{i_j}(\bar{x}_j)} g_{i_j}(\bar{x}_j) = (-1)^{s(v_1, \dots, v_n)} \prod_{1 \leq j \leq k} g_{i_j}(\bar{x}_j),$$

where $s(\bar{v}) = \sum_{1 \leq j \leq k} s_{i_j}(\bar{x}_j)$.

We can write $Z(I') = N^+ - N^-$, where N^+ is the number of configurations of the variables of I with weight 1 and N^- is the number with weight -1 . Now, N^+ is the number of solutions of the simultaneous equations

$$g_{i_1}(\bar{x}_1) = \dots = g_{i_k}(\bar{x}_k) = 1$$

over $\text{GF}(2)$ that have $s(\bar{v}) = 1$ and N^- is the number of solutions with $s(\bar{v}) = 0$. Since Γ' is pure affine of degree at most 2, each g_i is linear and s is quadratic. Lemma 3 below shows that the number of solutions to such a system of equations can be computed in polynomial time. \square

Lemma 3. *There is a polynomial-time algorithm for the following problem: given a multivariate quadratic polynomial q over $\text{GF}(2)$ and k multivariate linear polynomials ℓ_1, \dots, ℓ_k over $\text{GF}(2)$, determine the number of solutions that satisfy $q = 0, \ell_1 = 0, \dots, \ell_k = 0$ simultaneously.*

Proof. Suppose q and ℓ_1, \dots, ℓ_k are in indeterminates x_1, \dots, x_n . Suppose without loss of generality that ℓ_k depends on x_n . The polynomial ℓ_k evaluates to 0 if and only if $x_n = h(x_1, \dots, x_{n-1}) = \ell_k \oplus x_n$, where h is a linear polynomial in x_1, \dots, x_{n-1} . Substitute h for x_n in q and $\ell_1, \dots, \ell_{k-1}$ to obtain q' and $\ell'_1, \dots, \ell'_{k-1}$, respectively. The number of solutions that satisfy $q = 0, \ell_1 = 0, \dots, \ell_k = 0$ is the same as the number of solutions that satisfy $q' = 0, \ell'_1 = 0, \dots, \ell'_{k-1} = 0$, which may be found recursively. We recur until the system of equations contains one quadratic equation and no linear equations, or only linear equations. The number of solutions to a quadratic polynomial equation over $\text{GF}(2)$ can be computed in polynomial time [9, 16]. The number of solutions of a system of linear equations over $\text{GF}(2)$ can be computed by Gaussian elimination in polynomial time. \square

The case where every function in Γ is of product type of degree at most 1 is essentially the same as the corresponding case for non-negative functions [7] but we give a full proof for completeness.

Lemma 4. *Let $\Gamma \subseteq \mathcal{F}_B$. If every function in Γ is of product type of degree at most 1, then $\#\text{CSP}(\Gamma) \in \text{FP}$.*

Proof. Observe that, since each function $f \in \Gamma$ is of product type of degree at most 1, each can be written in the form

$$\begin{aligned} f(\bar{x}) &= (-1)^{x_{i_1} + \dots + x_{i_\ell} + c} U_1(x_1) \cdots U_k(x_k) g(\bar{x}) \\ &= (-1)^{x_{i_1}} \cdots (-1)^{x_{i_\ell}} (-1)^c U_1(x_1) \cdots U_k(x_k) g(\bar{x}), \end{aligned}$$

for some $c \in \{0, 1\}$. Thus, we can, instead, write $f(\bar{x}) = U'_1(x_1) \cdots U'_k(x_k) \cdot h(\bar{x})$ where each U'_i is a function $\{0, 1\} \rightarrow \mathbb{Q}$ instead of $\{0, 1\} \rightarrow \mathbb{Q}^{\geq 0}$. The remainder of the proof is the same as the corresponding case for non-negative functions.

Let I be an instance of $\#\text{CSP}(\Gamma)$, with variables V . Let \approx be the finest equivalence relation over V such that $v_i \approx v_j$ if $i = j$ or some constraint in I requires that either $v_i = v_j$ or $v_i \neq v_j$. We process each equivalence class in turn, independently of the others.

Let $S \subseteq V$ be an equivalence class of \approx . If there is no assignment to the variables in S that satisfies the equalities and disequalities in I 's constraints, then $Z(I) = 0$ and we are done. Otherwise, S must have a partition into sets S_0 and S_1 so that each variable in S_0 must have the same value and each variable in S_1 (which may be empty) must have the opposite value. The variables in S contribute one weight, say α , to $Z(I)$ if the variables in S_0 are set to 0 and another weight, say β , if they are set to 1. Thus, we can write $Z(I) = (\alpha + \beta)Z'(I)$, where $Z'(I)$ is the partition function $Z(I)$ with all terms involving the variables in S deleted. We may then proceed to factor out the next equivalence class. \square

4 Useful reductions

In this section, we give several reductions that are useful for proving hardness of weighted Boolean CSPs.

4.1 Pinning

Let δ_0 and δ_1 be the unary functions defined as

$$\begin{array}{ll} \delta_0(0) = 1 & \delta_1(0) = 0 \\ \delta_0(1) = 0 & \delta_1(1) = 1 \end{array}$$

These functions are referred to as *pinning* functions, since a constraint $\delta_c(x)$ “forces” the variable x to take value c by giving weight zero to any configuration with $x \neq c$. The proof of the following lemma is identical to the proof of [7, Lemma 8], except that the condition “ $f(x) > f(\bar{x}) \geq 0$ ” in the first sentence of Case 2 needs to be replaced with “ $f(x) \neq f(\bar{x})$ ”.

Lemma 5. *For every $\Gamma \subseteq \mathcal{F}_B$, $\#\text{CSP}(\Gamma \cup \{\delta_0, \delta_1\}) \leq_T \#\text{CSP}(\Gamma)$.*

4.2 Arity reduction

Given a k -ary function $f \in \mathcal{F}_B$ and $i \in [1, k]$, the function obtained by *projecting out the i th variable* is

$$g(x_1, \dots, x_{k-1}) = \sum_{y \in \{0, 1\}} f(x_1, \dots, x_{i-1}, y, x_i, \dots, x_{k-1}).$$

The following is a special case of [7, Lemma 6]. Although that Lemma is stated only for classes of non-negative rational functions, the proof does not rely on this.

Lemma 6. *Let $\Gamma \subseteq \mathcal{F}_B$, let $f \in \Gamma$ and let g be defined by projecting out a variable of f . $\#\text{CSP}(\Gamma \cup \{g\}) \leq_T \#\text{CSP}(\Gamma)$.*

Given a ternary function $f \in \mathcal{F}_B$, the *contraction* of f is the function

$$g(x_1, x_2) = \sum_{y, z \in \{0,1\}} f(x_1, y, z) f(y, z, x_2).$$

(In principle, we could define contractions in terms of any sequence of function arguments but we only make use of the version defined here.)

Lemma 7. *Let $\Gamma \subseteq \mathcal{F}_B$ and let g be the contraction of some $f \in \Gamma$. $\#\text{CSP}(\Gamma \cup \{g\}) \leq_T \#\text{CSP}(\Gamma)$.*

Proof. Replace each constraint C of the form $g(x, y)$ with the two constraints $f(x, x_C, y_C)$ and $f(x_C, y_C, y)$, where x_C and y_C are new variables used only in these two constraints. \square

4.3 Arithmetic techniques

For a constant $q \in \mathbb{Q}$ and a function $f \in \mathcal{F}_B$, write qf for the function that maps \bar{x} to $qf(\bar{x})$.

Lemma 8. *Let $f \in \Gamma \subseteq \mathcal{F}_B$ and let $q \neq 0$ be rational. $\#\text{CSP}(\Gamma \cup \{qf\}) \leq_T \#\text{CSP}(\Gamma)$.*

Proof. Let I be an instance of $\#\text{CSP}(\Gamma \cup \{qf\})$ and let I' be the instance of $\#\text{CSP}(\Gamma)$ made by replacing every constraint $qf(\bar{x})$ in I with $f(\bar{x})$. $Z(I) = q^m Z(I')$, where m is the number of qf -constraints in I . \square

Given a constraint language Γ , let Γ^2 be the constraint language that replaces every function $f(\bar{x})$ with the function $(f(\bar{x}))^2$. The following lemma is immediate from the observation that an instance of $\#\text{CSP}(\Gamma^2)$ can be converted to one of Γ with the same partition function just by including an extra copy of each constraint.

Lemma 9. $\#\text{CSP}(\Gamma^2) \leq_T \#\text{CSP}(\Gamma)$.

Further, if $\Gamma \subseteq \mathcal{F}_B$, then $\Gamma^2 \subseteq \mathcal{F}_B^{\geq 0}$. This fact and the following lemma allow us to reuse results on those functions from [7].

Lemma 10. *$f \in \mathcal{F}_B$ is pure affine (respectively, of product type) if, and only if, $f^2 \in \mathcal{F}_B^{\geq 0}$ is pure affine (respectively, of product type).*

Proof. Let $f \in \mathcal{F}_B$ be k -ary. It is clear that, if f is pure affine (respectively, of product type), then so is f^2 ; we show the converse. We assume that f^2 is not identically zero as this case is trivial.

First, suppose $f^2(\bar{x})$ is pure affine and equal to $w^2 g(\bar{x})$ as in Definition 1. There is a polynomial $s(\bar{x})$ that assigns the correct sign to each input such that $f(\bar{x}) = w(-1)^{s(\bar{x})} g(\bar{x})$. Therefore f is pure affine.

Now, suppose $f^2(\bar{x}) = U_1(x_1) \cdots U_k(x_k)g(\bar{x})$, as in Definition 2, is of product type. For each $i \in [1, k]$, let $U'_i(x_i) = \sqrt{U_i(x_i)}$. The functions U'_i are not necessarily rational but we certainly have

$$f(\bar{x}) = (-1)^{s(\bar{x})}U'_1(x_1) \cdots U'_k(x_k)g(\bar{x}), \quad (3)$$

for some suitable polynomial s , as before. By the arguments of Lemma 1, which do not depend on the rationality of the functions U'_i , we may assume that this is a normalized expression for f , except for the possible irrationality of the U'_i . We now describe how the functions U'_i can be replaced by rational functions, keeping the expression for f normalized. The function f is not identically zero so there is a tuple $\bar{a} \in \{0, 1\}^k$ such that $f(\bar{a}) \neq 0$. Since $f \in \mathcal{F}_B$, $f(\bar{a})$ is rational. For $i \in [1, k]$, let $U''_i(x_i) = U'_i(x_i)/U'_i(a_i)$. Then

$$f(\bar{x}) = |f(\bar{a})|(-1)^{s(\bar{x})}U''_1(x_1) \cdots U''_k(x_k)g(\bar{x}). \quad (4)$$

The expression for f in Equation (4) is not necessarily normalized because of the factor $|f(\bar{a})|$, however, as we will see next, the functions U''_i are rational. Once we have established this fact we will see that the factor $|f(\bar{a})|$ (which is rational) can be included in one of the unary functions U''_i , giving us a normalized expression for f . Note that for $i \in [1, k]$, $U''_i(a_i) = 1$, which is rational. Therefore we need to show that $U''_i(a_i \oplus 1)$ is rational. Observe that for each $i \in [1, k]$ such that x_i is a determined variable in the expression for f in Equation (3), $U''_i(0) = U''_i(1) = 1$. We now show that for each $i \in [1, k]$ such that x_i is not a determined variable, $U''_i(a_i \oplus 1)$ is rational. Suppose $U''_i(a_i \oplus 1) \neq 0$. Let $\bar{a}_i \in \{0, 1\}^k$ be the tuple obtained from \bar{a} by replacing a_i with $a_i \oplus 1$ and replacing a_j with $a_j \oplus 1$ for every determined variable x_j that occur together with x_i in an equality or disequality function of g . Thus, $g(\bar{a}_i) = 1$ and $|f(\bar{a}_i)| = |f(\bar{a})|U''_i(a_i \oplus 1) > 0$. Since $f \in \mathcal{F}_B$ and $f(\bar{a})$ is rational, $U''_i(a_i \oplus 1)$ is rational. Lastly we notice that the factor $|f(\bar{a})|$ in Equation (4) can be absorbed by any of the unary functions U''_i such that x_i is not a determined variable in the expression for f in Equation (3). If all variables are determined then we note that $|f(\bar{a})| = 1$ and we can disregard it completely. We finally conclude that $f(\bar{x})$ is of product type. \square

4.4 Matrix techniques

Given a $k \times k$ rational matrix, $A = (A_{ij})$, and a directed multigraph, $G = (V, E)$, that may have loops, let

$$Z_A(G) = \sum_{\sigma: V \rightarrow [1, k]} \prod_{(x, y) \in E} A_{\sigma(x)\sigma(y)}.$$

The problem of computing $Z_A(G)$ for a given input graph G is denoted by $\text{EVAL}(A)$. Bulatov and Grohe have given the complexity of $\text{EVAL}(A)$ for any symmetric matrix A with non-negative entries [3]. Here we only need the following special case.

Lemma 11. *Let A be a symmetric 2×2 matrix with non-negative rational entries. If A has rank 2 and at most one entry of A is zero then $\text{EVAL}(A)$ is $\#\text{P}$ -hard.*

For any $k \times k$ rational matrix A , $\text{EVAL}(A)$ is just the same thing as $\#\text{CSP}(f)$ for an appropriate binary function f over a domain of size k . In particular, then, 2×2 matrices correspond to binary Boolean functions.

Lemma 12. *Let $f \in \mathcal{F}_B$ be a binary function and let A be the matrix*

$$A = \begin{pmatrix} f(0,0) & f(0,1) \\ f(1,0) & f(1,1) \end{pmatrix}.$$

Then $\text{EVAL}(A) \equiv_T \#\text{CSP}(f)$.

Proof. Given an instance graph $G = (V, E)$ of $\text{EVAL}(A)$, let I be the instance of $\#\text{CSP}(f)$ with variables V that has a constraint $f(x, y)$ for every edge (x, y) in E . Thus $Z_A(G) = Z(I)$. \square

While Lemma 12 applies to all rational functions f , Lemma 11 can only be used if the resulting matrix is both symmetric and non-negative. The following lemma, essentially due to Dyer and Greenhill [8] will allow us to transform the matrix corresponding to a function f into a symmetric, non-negative matrix. For a matrix $A = (A_{ij})$, we write $A^{(2)}$ for the matrix (A_{ij}^2) .

Lemma 13. *For any rational square matrix A , the problems $\text{EVAL}(A^{(2)})$, $\text{EVAL}(AA^T)$, $\text{EVAL}(A^T A)$ and $\text{EVAL}(A^2)$ are polynomial-time Turing-reducible to $\text{EVAL}(A)$.*

Proof. For any graph G ,

- $Z_{A^{(2)}}(G) = Z_A(G_1)$, where G_1 is the multigraph formed by replacing each edge of G with two parallel edges;
- $Z_{AA^T}(G) = Z_A(G_2)$, where G_2 is the graph obtained by introducing a new vertex v_e for each edge $e = (x, y) \in G$ and replacing e with the edges (x, v_e) and (y, v_e) ;
- $Z_{A^T A}(G) = Z_A(G_3)$, where G_3 is made in the same way as G_2 but replacing e with (v_e, x) and (v_e, y) ;
- $Z_{A^2}(G) = Z_A(G_4)$, where G_4 is made in the same way as G_2 but replacing e with (x, v_e) and (v_e, y) . \square

5 High-degree pure affine functions

We have seen that there is a polynomial-time algorithm for $\#\text{CSP}(\Gamma)$ if every function in Γ is pure affine of degree at most two. We now show that computing partition functions of pure affine functions of higher degree is $\#\text{P}$ -hard. The main result of this section is the following lemma.

Lemma 14. *If $f \in \mathcal{F}_B$ is pure affine of degree at least three, then $\#\text{CSP}(f)$ is $\#\text{P}$ -hard.*

We first consider the restricted case $f(x, y, z) = (-1)^{s(x,y,z)}$, for ternary functions g of degree exactly 3 and then show that the case $f(\bar{x}) = (-1)^{s(\bar{x})}$ of degree-3 functions of arbitrary arity greater than three follows. Finally, we prove Lemma 14.

Lemma 15. *Let $f(x, y, z) = (-1)^{s(x,y,z)}$ where $s \in \mathcal{P}_3$ is of degree 3. $\#\text{CSP}(f)$ is $\#\text{P}$ -hard.*

$s(x, y, z)$	Method	A	A'
xyz	Project out z	$\begin{pmatrix} 2 & 2 \\ 2 & 0 \end{pmatrix}$	$\begin{pmatrix} 64 & 16 \\ 16 & 16 \end{pmatrix}$
$xyz + x$	Project out z	$\begin{pmatrix} 2 & 2 \\ -2 & 0 \end{pmatrix}$	$\begin{pmatrix} 64 & 16 \\ 16 & 16 \end{pmatrix}$
$xyz + x + y$	Project out z	$\begin{pmatrix} 2 & -2 \\ -2 & 0 \end{pmatrix}$	$\begin{pmatrix} 64 & 16 \\ 16 & 16 \end{pmatrix}$
$xyz + x + y + z$	Contract	$\begin{pmatrix} 4 & -2 \\ -2 & 4 \end{pmatrix}$	$\begin{pmatrix} 400 & 256 \\ 256 & 400 \end{pmatrix}$
$xyz + xy$	Project out z	$\begin{pmatrix} 2 & 2 \\ 2 & 0 \end{pmatrix}$	$\begin{pmatrix} 64 & 16 \\ 16 & 16 \end{pmatrix}$
$xyz + xy + x$	Project out z	$\begin{pmatrix} 2 & 2 \\ -2 & 0 \end{pmatrix}$	$\begin{pmatrix} 64 & 16 \\ 16 & 16 \end{pmatrix}$
$xyz + xy + x + y$	Project out z	$\begin{pmatrix} 2 & -2 \\ -2 & 0 \end{pmatrix}$	$\begin{pmatrix} 64 & 16 \\ 16 & 16 \end{pmatrix}$
$xyz + xy + z$	Project out y	$\begin{pmatrix} 2 & -2 \\ 0 & -2 \end{pmatrix}$	$\begin{pmatrix} 64 & 16 \\ 16 & 16 \end{pmatrix}$
$xyz + xy + x + z$	Project out y	$\begin{pmatrix} 2 & -2 \\ 0 & 2 \end{pmatrix}$	$\begin{pmatrix} 64 & 16 \\ 16 & 16 \end{pmatrix}$
$xyz + xy + x + y + z$	Contract	$\begin{pmatrix} 2 & -4 \\ 0 & 2 \end{pmatrix}$	$\begin{pmatrix} 400 & 64 \\ 64 & 16 \end{pmatrix}$
$xyz + xy + xz$	Project out z	$\begin{pmatrix} 2 & 2 \\ 0 & -2 \end{pmatrix}$	$\begin{pmatrix} 64 & 16 \\ 16 & 16 \end{pmatrix}$
$xyz + xy + xz + x$	Project out z	$\begin{pmatrix} 2 & 2 \\ 0 & 2 \end{pmatrix}$	$\begin{pmatrix} 64 & 16 \\ 16 & 16 \end{pmatrix}$
$xyz + xy + xz + y$	Project out z	$\begin{pmatrix} 2 & -2 \\ 0 & 2 \end{pmatrix}$	$\begin{pmatrix} 64 & 16 \\ 16 & 16 \end{pmatrix}$
$xyz + xy + xz + x + y$	Project out z	$\begin{pmatrix} 2 & -2 \\ 0 & -2 \end{pmatrix}$	$\begin{pmatrix} 64 & 16 \\ 16 & 16 \end{pmatrix}$
$xyz + xy + xz + y + z$	Contract	$\begin{pmatrix} 2 & -4 \\ 0 & 2 \end{pmatrix}$	$\begin{pmatrix} 400 & 64 \\ 64 & 16 \end{pmatrix}$
$xyz + xy + xz + x + y + z$	Project out x	$\begin{pmatrix} 0 & -2 \\ -2 & 2 \end{pmatrix}$	$\begin{pmatrix} 16 & 16 \\ 16 & 64 \end{pmatrix}$
$xyz + xy + xz + yz$	Contract	$\begin{pmatrix} 4 & -2 \\ -2 & 4 \end{pmatrix}$	$\begin{pmatrix} 400 & 256 \\ 256 & 400 \end{pmatrix}$
$xyz + xy + xz + yz + x$	Project out x	$\begin{pmatrix} 0 & 2 \\ 2 & -2 \end{pmatrix}$	$\begin{pmatrix} 16 & 16 \\ 16 & 64 \end{pmatrix}$
$xyz + xy + xz + yz + x + y$	Project out y	$\begin{pmatrix} 0 & 2 \\ -2 & 2 \end{pmatrix}$	$\begin{pmatrix} 16 & 16 \\ 16 & 64 \end{pmatrix}$
$xyz + xy + xz + yz + x + y + z$	Project out z	$\begin{pmatrix} 0 & -2 \\ -2 & -2 \end{pmatrix}$	$\begin{pmatrix} 16 & 16 \\ 16 & 64 \end{pmatrix}$

Table 1: The twenty ternary degree-3 polynomials considered in Lemma 15, with the methods used to prove them hard and the corresponding matrices.

Proof. Since s is of degree 3, it must contain the term xyz . Note that $\#\text{CSP}((-1)^{s(x,y,z)})$ is equivalent to $\#\text{CSP}((-1)^{s(x,y,z)+1})$ under Turing reductions, since $Z(I) = (-1)^m Z(I')$ where I and I' are instances of the two problems with the same m constraints. Therefore, we may assume that s does not contain the constant term 1.

Given this assumption, the terms of s are xyz and some subset of the terms xy , yz , zx , x , y and z . By symmetry between the variables, there are twenty cases to consider, listed in Table 1. Each case is proven $\#\text{P}$ -hard by either projecting out a variable or contracting, as detailed in the table.

For each polynomial s listed in the table, let $f(x, y, z) = (-1)^{s(x,y,z)}$. Note that f has the same value when s is evaluated over \mathbb{Z} as it does when s is evaluated over $\text{GF}(2)$, so we need not distinguish between $+$ and \oplus . The operation given (projecting out a variable or contracting) produces a new function f' in two variables which we will call x and y . By Lemma 6 (projection), or Lemma 7 (contraction), $\#\text{CSP}(f') \equiv_{\text{T}} \#\text{CSP}(f)$. Further, by Lemma 12, $\text{EVAL}(A) \leq_{\text{T}} \#\text{CSP}(f')$, where

$$A = \begin{pmatrix} f'(0,0) & f'(0,1) \\ f'(1,0) & f'(1,1) \end{pmatrix}$$

is given in the table. Let $A' = (AA^{\text{T}})^{(2)}$. For any rational matrix A , the corresponding

A' is symmetric and non-negative and, by Lemma 13, $\text{EVAL}(A') \leq_T \text{EVAL}(A)$. All of the matrices A' given in the table have rank 2 and at most one zero entry so, by Lemma 11, $\text{EVAL}(A')$ is $\#P$ -hard. \square

Lemma 16. *Let $\bar{x} = x_1 \dots x_k$ for some $k > 3$ and let $f(\bar{x}) = (-1)^{s(\bar{x})}$ for some $s \in \mathcal{P}_k$ of degree at least 3. $\#CSP(f)$ is $\#P$ -hard.*

Proof. Renaming variables if necessary, we may assume that one of the terms of least degree greater than or equal to three in s is $x_1 \dots x_\ell$ for some ℓ with $3 \leq \ell \leq k$. Let $c_4 = \dots = c_\ell = 1$ and $c_{\ell+1} = \dots = c_k = 0$ and let $s'(x_1, x_2, x_3) = s(x_1, x_2, x_3, c_4, \dots, c_k)$ and $f'(x_1, x_2, x_3) = (-1)^{s'(x_1, x_2, x_3)}$.

The degree of s' is 3 since it has only three variables and includes exactly one term $x_1 x_2 x_3 1 \dots 1 = x_1 x_2 x_3$. Therefore, $\#CSP(f')$ is $\#P$ -hard by the previous lemma.

It remains to show that $\#CSP(f') \leq_T \#CSP(f)$. To see this, let I' be any instance of $\#CSP(f')$. We create an instance I'' of $\#CSP(\{f, \delta_0, \delta_1\})$ such that $Z(I'') = Z(I')$ as follows, and the result is then immediate from Lemma 5. Let z_4, \dots, z_k be new variables. Let I'' have the constraints $\delta_1(z_4), \dots, \delta_1(z_\ell), \delta_0(z_{\ell+1}), \dots, \delta_0(z_k)$ and, for each constraint $f'(y_1, y_2, y_3)$ in I' , the constraint $f(y_1, y_2, y_3, z_4, \dots, z_k)$. \square

We now prove the main result of this section, namely that $\#CSP(f)$ is $\#P$ -hard if f is pure affine of degree at least three.

Proof of Lemma 14. Let $f(x_1, \dots, x_k)$ be pure affine of degree at least three. Thus, we may write $f(\bar{x}) = w(-1)^{s(\bar{x})}g(\bar{x})$, where $w > 0$, $g \in \mathcal{P}_k$ is affine and $s \in \mathcal{P}_k$ is degree-minimized with respect to g and has degree at least three. By Lemma 8, we may assume that $w = 1$.

Since g is affine, we may write

$$f(\bar{x}) = (-1)^{s(\bar{x})} \prod_{i \in [1, m]} g_i(\bar{x}),$$

where each $g_i \in \mathcal{P}_k$ is linear. We show that f is $\#P$ -hard by induction on m . The base case, $m = 0$ is Lemma 16.

For the inductive step $m > 0$, we may assume without loss of generality that g_m depends on x_k . If $f(\bar{x}) \neq 0$, we must have $g_m(\bar{x}) = 1$ and, therefore, $x_k = g_m(\bar{x}) \oplus x_k \oplus 1$. Note that $g_m(\bar{x}) \oplus x_k \oplus 1$ does not depend on x_k . Let $g'_1, \dots, g'_{m-1}, s' \in \mathcal{P}_{k-1}$ be the polynomials that result from substituting $g_m(\bar{x}) \oplus x_k \oplus 1$ for x_k in g_1, \dots, g_{m-1} and s , respectively.

Since $s'(\bar{x}) = s(\bar{x})$ whenever $f(\bar{x}) \neq 0$, we have $f(\bar{x}) = (-1)^{s'(\bar{x})}g'(\bar{x})$. Because s is degree-minimized with respect to g , s' must have the same degree as s .

Let

$$f'(x_1, \dots, x_{k-1}) = (-1)^{s'(\bar{x})} \prod_{i \in [1, m-1]} g'_i(\bar{x}).$$

Suppose that s' is not degree-minimized with respect to $g'(\bar{x}) = \prod_i g'_i(\bar{x})$. Then there is another polynomial s'' of strictly lower degree such that $f'(\bar{x}) = (-1)^{s''(\bar{x})}g'(\bar{x})$. But then, $f(\bar{x}) = (-1)^{s''(\bar{x})}g_m(\bar{x})g'(\bar{x}) = (-1)^{s''(\bar{x})}g(\bar{x})$, contradicting degree-minimality of s . Therefore, s' is degree-minimized with respect to g' . Further, s and s' have the same degree, so $\#CSP(f')$ is $\#P$ -hard by the inductive hypothesis.

It remains to show that $\#\text{CSP}(f') \leq_{\text{T}} \#\text{CSP}(f)$. Let I' be an instance of $\#\text{CSP}(f')$ and let I be the instance of $\#\text{CSP}(f)$ that has a constraint $f(x_1, \dots, x_{k-1}, x_C)$ for every constraint $C = f'(x_1, \dots, x_{k-1})$ in I' . Then $Z(I) = Z(I')$ and we are done. \square

6 Higher-degree product-type functions

We now construct the machinery for the remaining hard case: functions of product type of degree two or more that are not pure affine of degree two.

For any $\lambda \in \mathbb{Q}$, we write $\Theta_\lambda(x)$ for the function $\Theta_\lambda(0) = 1$, $\Theta_\lambda(1) = \lambda$. (In [7], these functions are written U_λ but we wish to avoid the potential for confusion with the functions U_1, \dots, U_k used to define a k -ary function of product type.)

The main result of this section is the following lemma.

Lemma 17. *Let $f \in \mathcal{F}_{\text{B}}$ be of product type of degree at least two. Then, $\#\text{CSP}(\{f, \Theta_\lambda\})$ is $\#\text{P}$ -hard for any positive rational $\lambda \neq 1$.*

If f is both of product type of degree two and pure affine of degree two, then $\#\text{CSP}(f)$ is computable in polynomial time by Lemma 2. In the following section, we will show that, for all other functions of product type of degree two or more, we have $\#\text{CSP}(\{f, \Theta_\lambda\}) \equiv_{\text{T}} \#\text{CSP}(f)$ so Lemma 17 is sufficient for our needs, even though it appears, at first sight, to be weaker than the desired result.

As in the previous section, we first consider simplified cases.

Lemma 18. *Let f be of product type of degree at least two. There are non-zero rationals α and β such that $\#\text{CSP}(f') \leq_{\text{T}} \#\text{CSP}(f)$, where $f'(x, y) = (-1)^{xy} \Theta_\alpha(x) \Theta_\beta(y)$.*

Proof. Let $(-1)^{s(\bar{x})} U_1(x_1) \cdots U_k(x_k) g(\bar{x})$ be a normalized expression defining f .

We may assume, renaming variables if necessary, that

$$s(\bar{x}) = x_1 x_2 p(x_3, \dots, x_k) + q(x_1, \dots, x_k),$$

where p and q are polynomials in the stated variables, p is not identically zero and q contains no term that has $x_1 x_2$ as a factor. Let X be the set of variables on which s depends. Because $x_1 \in X$, there must be an assignment $\sigma: X \rightarrow \{0, 1\}$ such that $s(0, \sigma(x_2), \dots, \sigma(x_k)) \neq s(1, \sigma(x_2), \dots, \sigma(x_k))$. We may assume that $s(\sigma(\bar{x})) = 1$.

Now let Y be the set of variables on which g depends. Since the expression is normalized, at least one variable in each term $\chi_{=(x_i, x_j)}$ or $\chi_{\neq}(x_i, x_j)$ is determined and no determined variable appears in s . Therefore, we can extend σ to an assignment $\sigma': X \cup Y \rightarrow \{0, 1\}$ such that $s(\sigma'(\bar{x})) = g(\sigma'(\bar{x})) = 1$.

Further, for every i with $x_i \in X \cup Y$, $U_i(0)$ and $U_i(1)$ are both non-zero. For each $x_i \notin (X \cup Y)$, we must have $U_i(0) \neq 0$ or $U_i(1) \neq 0$ or both; otherwise, f is identically zero (and, thus, of product type of degree zero). Therefore, we can extend σ' to an assignment $\sigma'': \{x_1, \dots, x_k\} \rightarrow \{0, 1\}$ such that $f(\sigma''(\bar{x})) \neq 0$.

Finally, suppose that $f'(x, y) = f(x, y, \sigma''(x_3), \dots, \sigma''(x_k))$. Then clearly $\#\text{CSP}(f') \leq_{\text{T}} \#\text{CSP}(\{f, \delta_0, \delta_1\})$ so, by Lemma 5, $\#\text{CSP}(f') \leq_{\text{T}} \#\text{CSP}(f)$.

There are constants $w \in \mathbb{Q}$ and $a, b, c \in \{0, 1\}$ such that

$$\begin{aligned} f'(x, y) &= w(-1)^{xy+ax+by+c} U_1(x) U_2(y) \\ &= w(-1)^c (-1)^{xy} U_1'(x) U_2'(y), \end{aligned}$$

where $U_1'(x) = (-1)^{ax}U_1(x)$ and $U_2'(y) = (-1)^{by}U_2(y)$. Putting $\alpha = U_1'(1)/U_1'(0)$ and $\beta = U_2'(1)/U_2'(0)$, we have

$$f'(x, y) = wU_1'(0)U_2'(0)(-1)^c(-1)^{xy}\Theta_\alpha(x)\Theta_\beta(y)$$

and, by Lemma 8, we can discard the constant factor $wU_1'(0)U_2'(0)(-1)^c$. \square

Lemma 19. *If $f(x, y) = (-1)^{xy}\Theta_\alpha(x)\Theta_\beta(y)$, where $\alpha \in \mathbb{Q} \setminus \{-1, 0, 1\}$, $\beta \in \mathbb{Q} \setminus \{0\}$, then $\#\text{CSP}(f)$ is $\#\text{P}$ -hard.*

Proof. Let

$$A = \begin{pmatrix} f(0, 0) & f(0, 1) \\ f(1, 0) & f(1, 1) \end{pmatrix} = \begin{pmatrix} 1 & \beta \\ \alpha & -\alpha\beta \end{pmatrix}$$

and let

$$B = (A^T A)^{(2)} = \begin{pmatrix} 1 + \alpha^2 & \beta(1 - \alpha^2) \\ \beta(1 - \alpha^2) & \beta^2(1 + \alpha^2) \end{pmatrix}^{(2)}.$$

Since $\beta \neq 0$ and $\alpha^2 \neq 1$, every entry of B is positive. We have

$$|B| = (1 + \alpha^2)^4\beta^4 - (1 - \alpha^2)^4\beta^4 = 8\alpha^2\beta^4(1 + \alpha^4) > 0.$$

Therefore, B has rank two and hence $\text{EVAL}(B)$ is $\#\text{P}$ -hard by Lemma 11. By Lemmas 12 and 13, $\text{EVAL}(B) \leq_T \text{EVAL}(A) \equiv_T \#\text{CSP}(f)$. \square

We now prove Lemma 17, namely that, if f is of product type of degree at least two, then $\#\text{CSP}(\{f, \Theta_\lambda\})$ is $\#\text{P}$ -hard for any positive rational $\lambda \neq 1$.

Proof of Lemma 17. Let $f \in \mathcal{F}_B$ be of product type of degree at least 2 and let $\lambda \in \mathbb{Q}^{>0} \setminus \{1\}$. By Lemma 18 there are non-zero rational constants α and β such that $\#\text{CSP}(g) \leq_T \#\text{CSP}(f)$, where

$$g(x, y) = (-1)^{xy}\Theta_\alpha(x)\Theta_\beta(y).$$

If at most one of α and β is 1 or -1 , then $\#\text{CSP}(g)$ is $\#\text{P}$ -hard by Lemma 19 and we are done. Otherwise, we have $\alpha, \beta \in \{-1, 1\}$. Let

$$g'(x, y) = (-1)^{xy}\Theta_{\alpha\lambda}(x)\Theta_\beta(y).$$

$\#\text{CSP}(g')$ is $\#\text{P}$ -hard by Lemma 19, since $\alpha\lambda \notin \{-1, 0, 1\}$. It just remains to show that $\#\text{CSP}(g') \leq_T \#\text{CSP}(\{g, \Theta_\lambda\})$ but this is easy: given an instance of $\#\text{CSP}(g')$, replace every constraint $g'(x, y)$ by the pair of constraints $g(x, y)$ and $\Theta_\lambda(x)$. \square

7 Proving the dichotomy

We now have all the tools we need to prove the remaining side of the dichotomy, namely that, unless either every $f \in \Gamma$ is pure affine of degree at most two or every f is of product type of degree at most one, then $\#\text{CSP}(\Gamma)$ is $\#\text{P}$ -hard.

Lemma 20. *If $f \in \mathcal{F}_B$ does not have affine support, then $\#\text{CSP}(f)$ is $\#\text{P}$ -hard.*

Proof. $f^2 \in \mathcal{F}_B^{\geq 0}$ has the same support as f . By [7, Lemma 11], $\#\text{CSP}(f^2)$ is $\#\text{P}$ -hard and $\#\text{CSP}(f^2) \leq_T \#\text{CSP}(f)$ by Lemma 9. \square

Lemma 21. *If $f \in \mathcal{F}_B$ is not of product type of degree at most one then the problem $\#\text{CSP}(\{f, \delta_0, \delta_1, \Theta_\lambda\})$ is $\#\text{P}$ -hard for any positive rational $\lambda \neq 1$.*

Proof. If f is not of product type then, by Lemma 10, $f^2 \in \mathcal{F}_B^{\geq 0}$ is also not of product type. By [7, Lemma 15], $\#\text{CSP}(\{f^2, \delta_0, \delta_1, \Theta_\lambda\})$ is $\#\text{P}$ -hard for any positive, rational $\lambda \neq 1$ and the result follows by Lemma 9.

If f is of product type but of degree two or more, the result follows from Lemma 17. \square

The next lemma corresponds to [7, Lemma 16] and its proof is based on the same idea as the proof there. The only difference is a slight adjustment to deal with mixed signs.

Lemma 22. *If $f \in \mathcal{F}_B$ is not pure affine of degree at most two and $g \in \mathcal{F}_B$ is not of product type of degree at most one, then $\#\text{CSP}(\{f, g, \delta_0, \delta_1\})$ is $\#\text{P}$ -hard.*

Proof. Suppose f is not pure affine of degree at most two. If f does not even have affine support, we are done by Lemma 20 and, if f is pure affine of degree three or higher, we are done by Lemma 14. So we may assume that f is not pure affine. By Lemma 10, f^2 is also not pure affine and, by Lemma 9, it suffices to show that $\#\text{CSP}(\{f^2, g, \delta_0, \delta_1\})$ is $\#\text{P}$ -hard.

Since f^2 has affine support but is not pure affine, there must be at least two positive values in its range. The proof now proceeds exactly as that of Lemma 16 in [7]. By using pinning and projection, we extract from f^2 a unary function Θ_λ for some positive rational $\lambda \neq 1$. The function Θ_λ is simulated by f^2 and we show hardness of $\#\text{CSP}(\{f^2, g, \delta_0, \delta_1\})$ by reduction from $\#\text{CSP}(\{g, \delta_0, \delta_1, \Theta_\lambda\})$, which is $\#\text{P}$ -hard by Lemma 21. We do not repeat the details here; refer to the proof in [7], starting with the second paragraph and noting that the function g referred to there is the function f^2 here. \square

References

- [1] A. Bulatov. The complexity of the counting constraint satisfaction problem. In *Automata, Languages and Programming, 35th International Colloquium (ICALP 2008) Part 1*, volume 5125, pages 646–661. Springer, 2008.
- [2] A. Bulatov, M. Dyer, L. A. Goldberg and M. Jerrum. Personal communication.
- [3] A. Bulatov and M. Grohe. The complexity of partition functions. *Theoretical Computer Science*, 348(2):148–186, 2005.
- [4] N. Creignou and M. Hermann. Complexity of generalized satisfiability counting problems. *Information and Computation*, 125(1):1–12, 1996.
- [5] N. Creignou, S. Khanna and M. Sudan. *Complexity classifications of Boolean constraint satisfaction problems*. SIAM Press, 2001.
- [6] M. Dyer, L. A. Goldberg and M. Jerrum. An approximation trichotomy for Boolean $\#\text{CSP}$. <http://arxiv.org/abs/0710.4272>, 2007.

- [7] M. Dyer, L. A. Goldberg and M. Jerrum. The complexity of weighted Boolean #CSP. *SIAM Journal of Computing*, 2008. To appear.
- [8] M. Dyer and C. Greenhill. The complexity of counting graph homomorphisms. *Random Structures and Algorithms*, 17(3–4):260–289, 2000.
- [9] A. Ehrenfeucht and M. Karpinski. The computational complexity of (XOR, AND)-counting problems. Technical Report 8543-CS, University of Bonn, 1990. Available at <http://citeseer.ist.psu.edu/ehrenfeucht90computational.html>.
- [10] T. Feder and M. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory. *SIAM Journal on Computing*, 28(1):57–104, 1999.
- [11] L. A. Goldberg, M. Grohe, M. Jerrum and M. Thurley. A complexity dichotomy for partition functions with mixed signs. To appear in STACS 2009.
- [12] L. A. Goldberg and M. Jerrum. Inapproximability of the Tutte polynomial. *Information and Computation*, 207(7):908–929, 2008.
- [13] C. Greenhill. The complexity of counting colourings and independent sets in sparse graphs and hypergraphs. *Computational Complexity*, 9(1):52–72, 2000.
- [14] P. Hell and J. Nešetřil. *Graph homomorphisms*. Oxford University Press, 2004.
- [15] R. Ladner. On the structure of polynomial time reducibility. *Journal of the Association for Computing Machinery*, 22(1):155–171, 1975.
- [16] R. Lide and H. Niederreiter. *Finite fields*, volume 20 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 2nd edition, 1997.
- [17] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- [18] F. Rossi, P. van Beek and T. Walsh (Eds.). *Handbook of constraint programming*. Elsevier, 2006.
- [19] T. Schaefer. The complexity of satisfiability problems. In *Tenth Annual ACM Symposium on Theory of Computing*, pages 216–226. ACM Press, 1978.
- [20] L. Valiant. The complexity of enumeration and reliability problems. *SIAM Journal on Computing*, 8(3):410–421, 1979.
- [21] D. Welsh. *Complexity: knots, colourings and counting*. Cambridge University Press, 1993.