Indistinguishability: Counting, Constraints, and Contractions

By

Benjamin Morgen Young

A dissertation submitted in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

(Computer Sciences)

at the

UNIVERSITY OF WISCONSIN-MADISON

2026

Date of final oral examination: 06/04/2026

The dissertation is approved by the following members of the Final Oral Committee:
Jin-Yi Cai, Professor, Computer Sciences
Eric Bach, Emeritus Professor, Computer Sciences
Rishab Goyal, Assistant Professor, Computer Sciences
Paul Terwilliger, Professor, Mathematics

Abstract

This thesis studies indistinguishability for counting problems on graphs in the Holant, #CSP, and graph homomorphism frameworks. Valiant's Holant theorem, a powerful tool for algorithms and reductions for counting problems, states that if two sets $\mathcal{F}$ and $\mathcal{G}$ of tensors (a.k.a. constraint functions or signatures) are related by a holographic transformation, then $\mathcal{F}$ and $\mathcal{G}$ are Holant-indistinguishable, meaning every tensor network with tensors from $\mathcal{F}$ contracts to the same value when we replace the tensors from $\mathcal{F}$ with the corresponding tensors from $\mathcal{G}$. Our results are partial converses of the Holant theorem: if $\mathcal{F}$ and $\mathcal{G}$ are indistinguishable in the context of some counting problem, then there is a certain type, depending on the counting problem, of holographic transformation between them. These include isomorphism, orthogonal transformation, sequences of invertible transformations mapping $\mathcal{F}$ arbitrarily close to $\mathcal{G}$, and, for planar Holant and #CSP, quantum isomorphism and quantum orthogonal transformation, with connections to quantum information theory.

Our proofs are largely based on the algebra of quantum Holant gadgets, and develop new connections between counting problems and other fields, including quantum algebra and classical and geometric invariant theory. Along the way, we prove several other results related to counting indistinguishability, including the undecidability of vertex separation by planar homomorphism counting gadgets, a combinatorial perspective on simultaneously orthogonally decomposable (odeco) signature sets, a characterization of vanishing Holant signature sets on all domains, and the first characterization of homomorphism-indistinguishability over graphs of bounded degree.

Acknowledgments

First and foremost, I want to thank my advisor Jin-Yi Cai. His guidance, encouragement, insight, and vigilant attention to detail have been the driving forces behind my development as a researcher and teacher.

Thank you to my labmates Ashwin, Jin Soo, Zhuxiao, Austen, and Hugh for years of enlightening and entertaining discussion around the department. Also thank you to David Roberson, Tim Seppelt, Radu Curticapean, and my committee members Eric Bach, Rishab Goyal, and Paul Terwilliger for their valuable feedback on my work.

Finally, thank you to Yuna and to my family and friends for always supporting me and bringing joy to my life here in Madison.

To Yuna and my family.

Contents

Chapter 1

Introduction

Let $F, G \in \mathbb{C}^{q \times q}$ be two matrices. If F and G are similar, then $\operatorname{tr}(F^k) = \operatorname{tr}(G^k)$ for every k – that is, F and G are *indistinguishable* by the trace-of-kth-power function function $\operatorname{tr}((\cdot)^k)$. Conversely, if $\operatorname{tr}(F^k) = \operatorname{tr}(G^k)$ for every k, then we may only conclude that F and G have the same multiset of eigenvalues; F and G are not necessarily similar. What other conditions are equivalent to $\operatorname{tr}((\cdot)^k)$-indistinguishability, and what other assumptions on F and G suffice to obtain similarity? In this thesis, we consider questions that are vast generalizations of this example.

Holant Problems and the Holant theorem. Holant problems, first introduced in [CLX11], are a framework for counting problems on graphs. Holant captures information about graphs expressible by local constraints, including matchings, vertex and edge colorings, Eulerian orientations, and homomorphisms to a fixed graph X. The input to a Holant problem is a *signature grid* Ω, a graph with a constraint – also called a *signature* – assigned to each vertex v, applied to the edges incident to v. Given a coloring of the edges of Ω, the constraints detect whether the coloring is locally valid at each vertex. For example, if the edges are colored red or black and each constraint checks that exactly one of its incident edges is red, then the red edges in a valid coloring form a perfect matching in Ω (see Figure 1.1). If instead the edges are colored by $1, \ldots, q$ and each constraint checks that its incident edges have pairwise distinct colors, then a valid coloring is a proper edge coloring of Ω. Viewing each constraint as a function evaluating to 1 when satisfied and 0 when unsatisfied, the *Holant value* – the number of perfect matchings or proper edge colorings – is the sum over all edge colorings of the product of the constraint evaluations (as the product corresponding to a coloring is 1 if and only if this coloring satisfies all constraints). This sum-of-product perspective generalizes to signatures that may not be 0-1-valued; for any set $\mathcal{F}$ of complex-valued signatures, define the computational problem $\text{Holant}(\mathcal{F})$ as follows: given a signature grid Ω with vertices assigned signatures from $\mathcal{F}$, compute the Holant value of Ω. Equivalently, a signature over q edge colors is a tensor over the vector space $\mathbb{C}^q$; with this perspective, the Holant value of Ω is the contraction of Ω as a tensor network.

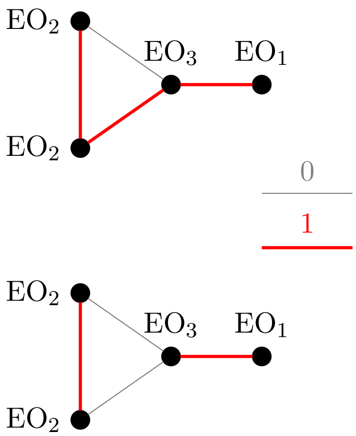

$$\text{EO}_2(1,0)\cdot\text{EO}_2(1,1)\cdot\text{EO}_3(0,1,1)\cdot\text{EO}_1(1) = 1\cdot 0\cdot 0\cdot 1$$

$$\text{EO}_2(1,0)\cdot\text{EO}_2(1,0)\cdot\text{EO}_3(0,0,1)\cdot\text{EO}_1(1) = 1\cdot 1\cdot 1\cdot 1$$

Figure 1.1: Define the *exact-one* signature $\text{EO}_n : \{0,1\}^n \to \{0,1\}$ by $\text{EO}_n(x_1,\ldots,x_n) = 1$ iff exactly one of $x_1,\ldots,x_n$ is 1. Then the Holant value of a signature grid Ω in which each vertex of degree d is assigned EO_d is the number of perfect matchings in Ω. This figure shows two edge colorings in Ω and the corresponding product of constraint evaluations. Each edge is either thin and gray (assigned 0) or thick and red (assigned 1). The red edges in the lower coloring form a perfect matching, so this coloring contributes 1.

While Holant is very expressive, it is also restrictive enough to prove sweeping dichotomy theorems. These classify $\text{Holant}(\mathcal{F})$ as either computable in polynomial time or #P-hard for *any* set $\mathcal{F}$ [CLX11; CLX08; HL16; CGW16; Cai+15; LW18; SC20; CLX13; CI25]. While most existing work focuses on domain size $q = 2$ (the *Boolean domain*) or $q = 3$, the techniques employed in this thesis apply to all q.

Valiant's *Holant theorem* [Val08; Val06], the genesis for Holant problems, states that, if two sets $\mathcal{F}$ and $\mathcal{G}$ of signatures are related by a *holographic transformation* – essentially a basis change by a $T \in \text{GL}_q$ – then $\mathcal{F}$ and $\mathcal{G}$ are *Holant-indistinguishable*, meaning that every signature grid Ω has the same Holant value whether its vertices are assigned signatures from $\mathcal{F}$ or the corresponding transformed signatures in $\mathcal{G}$. This implies that $\text{Holant}(\mathcal{F})$ and $\text{Holant}(\mathcal{G})$ have the same complexity, leading to the notions of *holographic reductions* between Holant problems and *holographic* *algorithms*. Later work [CC07; CL09; CL11] formalized the Holant theorem and holographic reductions in terms of covariant and contravariant tensors. We take this perspective in this thesis. The Holant theorem originated in the context of planar Holant, which restricts to planar signature grids. Valiant [Val06; Val08] gave polynomial-time holographic algorithms for planar Holant($\mathcal{F}$) for certain $\mathcal{F}$ transformable to the signatures of *matchgates*, for which the Holant value is efficiently computable using the FKT algorithm for counting the number of perfect matchings in a planar graph [TF61; Kas61]. Succeeding work on the complexity of planar Holant led to a *trichotomy* classifying symmetric complex-valued signatures on the Boolean domain into exactly three classes: (1) polynomial-time solvable; (2) #P-hard for general instances but solvable in polynomial-time over planar instances; and (3) #P-hard over planar instances [Cai+15].

Counting CSP and graph homomorphisms An (unweighted) counting constraint satisfaction problem #CSP($\mathcal{F}$) is parameterized by a set $\mathcal{F}$ of 0-1-valued *constraint functions* on one or more inputs from a finite domain. The input to the problem is a #CSP *instance*, consisting of a multiset C of constraints and set V of variables, with each constraint applying a constraint function in $\mathcal{F}$ to a subset of V. The output is the number of variable assignments that satisfy all of the constraints. More generally, we consider #CSP($\mathcal{F}$) for a set $\mathcal{F}$ of $\mathbb{C}$-valued constraint functions. We can model #CSP($\mathcal{F}$) as a Holant problem by constructing a signature grid from the constraint-variable incidence graph of a #CSP($\mathcal{F}$) instance, with constraint vertices assigned the corresponding signatures in $\mathcal{F}$ and variable vertices assigned auxiliary signatures forcing each variable to take the same value in all constraints in which it appears. Thus we consider #CSP to be a "special case" of Holant. Unweighted #CSP itself generalizes the problem of counting homomorphisms (edge-preserving maps) from a graph K to a graph X. Counting graph homomorphisms is a central problem in counting complexity, both in its own right and as a special case of #CSP. Both settings have seen significant dichotomy theorems classifying the partition function as either tractable or #P-hard to compute, depending on X or $\mathcal{F}$, respectively. Graph homomorphism dichotomies have been established for unweighted graphs [DG00], nonnegative-real-weighted graphs [BG05; CC19], real-weighted graphs [Gol+10], and complex-weighted graphs [CCL13]. For #CSP, dichotomies have been established for sets of 0-1 valued constraint functions [Bul13; DR13b], nonnegative-real-valued constraint functions [CCL16], and complex-valued constraint functions [CC17b]. As with

Holant, when we restrict to planar #CSP instances (for which the constraint-variable incidence graph is planar), a further trichotomy is known for the Boolean domain [CF22]. Furthermore, Valiant's holographic algorithm with matchgates [Val08] is universal in that *every* #CSP problem that is solvable in polynomial time only in the planar setting is solvable by this one algorithmic strategy. The same trichotomy is known for counting (weighted) homomorphisms from planar graphs to graphs with up to four vertices [CM23; CM24a].

Vertex and edge coloring models. This thesis uses and generalizes ideas from the theory of vertex coloring models and edge coloring models, two well-studied classes of Holant problems. De la Harpe and Jones [HJ93] defined vertex and edge coloring models as extensions of statistical mechanics models (e.g. the Ising model), calling them "spin models" and "vertex models", respectively. A vertex coloring model (also called a spin system) is defined by a graph X with edge and possibly vertex weights. Given an input graph K, one aims to compute the *partition function*, the number of (weighted) graph homomorphisms from K to X. An edge coloring model is defined by a set $\mathcal{F}$ of symmetric signatures containing exactly one signature of each arity, and the problem of computing its partition function is equivalent to Holant($\mathcal{F}$) (this restriction on $\mathcal{F}$ ensures that edge coloring models take ordinary graphs, rather than signature grids, as input).

One thread of prior work on vertex and edge coloring models characterizes which scalar-valued functions on graphs are expressible as vertex coloring models [FLS07; Sch09] or as edge coloring models [Sze07; Sch08a; Dra+12; Reg13a]. Another, related, line of works compute the rank of *connection matrices* for vertex coloring models [Lov06] and edge coloring models [Reg12; DR13a]. See [Reg13b] for an overview of many of the above results. Following Freedman, Lovász, and Schrijver [FLS07], these works use *quantum labeled graphs*, algebras of formal linear combinations of graphs with labeled vertices. This thesis follows many of these works in their application of techniques from invariant theory, either of the symmetric group in the case of vertex coloring models [Sch09], or of the orthogonal group $O(q)$ in the case of edge-coloring models [Sze07; Sch08a; Dra+12; Reg12; DR13a; Reg13a].

Homomorphism indistinguishability. Since graph homomorphism counting is a Holant problem, this thesis encompasses many results (and proves some new ones) in the area of *homomor-*

phism indistinguishability. Lovász [Lov67] showed that two graphs X and Y are isomorphic if and only if they admit the same number of homomorphisms from – that is, are homomorphism-indistinguishable over – all graphs. This result was later improved to X and Y with edge and vertex weights [Lov06; Sch09; CG21]. Another line of work aims to determine the relaxations of isomorphism which must relate any X and Y indistinguishable under homomorphism counts from restricted classes of graphs [Dvo10; DGR18; MR20; Kar+25; RS23; GRS25; RS24]. Given the rich complexity theory of counting problems on planar graphs – including counting homomorphisms from planar graphs – discussed above, we specifically highlight the work of Mančinska and Roberson [MR20], who showed that X and Y are indistinguishable by homomorphism counts from planar graphs if and only if X and Y are *quantum isomorphic*, a relaxation of isomorphism originally defined in terms of a quantum nonlocal game. Mančinska and Roberson's proof makes use of quantum bi-labeled graphs, a version of the quantum labeled graphs discussed in the previous paragraph. Both labeled graphs and bi-labeled graphs are a special case of Holant *gadgets*, the primary reduction tool – along with holographic transformation – in the study of Holant complexity. Mančinska and Roberson also apply a *duality theorem* from the theory of quantum permutation groups [Wor88] leading to a characterization of the expressive power of quantum bi-labeled graphs. We apply variants of this duality theorem throughout this thesis.

Our results. We prove indistinguishability theorems for #CSP, Holant, and their planar and bipartite variants. Such a theorem has the form of a restricted converse of the Holant theorem: if two signature sets $\mathcal{F}$ and $\mathcal{G}$ are #CSP- or Holant-indistinguishable, then there is a holographic transformation between them. The form of this holographic transformation depends on the problem in question. For #CSP, it is isomorphism (Chapter 3); for planar #CSP, it is quantum isomorphism (Chapter 7); for Holant, it is orthogonal transformation (Chapter 8); for planar Holant, it is quantum orthogonal transformation (Chapter 9); for bipartite Holant (the setting of Valiant's original Holant theorem), the converse does not hold – we do not obtain an invertible transformation – but we obtain a sequence of invertible transformations mapping $\mathcal{F}$ and $\mathcal{G}$ arbitrarily close to each other (Chapter 10).

After preliminaries in Chapter 2, we prove the #CSP indistinguishability theorem for any characteristic-0 field $\mathbb{K}$ in Chapter 3. This chapter, in contrast with the rest of the thesis, applies an entirely constructive and explicit technique using only the invertibility of a Vandermonde matrix, following Cai and Govorov's proof of the graph homomorphism special case [CG21]. In Chapter 5, we prove the special case of the #CSP indistinguishability theorem for $\mathbb{K} = \mathbb{C}$ via a proof technique in the style of Mančinska and Roberson [MR20], using quantum gadgets and duality. The remaining indistinguishability proofs – with the exception of the planar Holant converse in Chapter 9 and the first bipartite Holant converse in Section 10.2, which use even more abstract and nonconstructive theorems – follow, at a very high level, the same approach, applying variants of the same duality theorem [Sch08b; DM23]. Upon moving from #CSP to Holant (motivated by a conjecture of Xia [Xia10] on the converse of the Holant theorem), we incorporate a novel technique: induction on the domain size. See Figure 4.3 and Table 4.1 for an overview of the indistinguishability results in this thesis.

Chapter 2

Counting Problem Preliminaries

Throughout this thesis, $\mathbb{K}$ is a field of characteristic 0. Occasionally, we also assume $\mathbb{K}$ is algebraically closed. For $n \in \mathbb{N}$, write $[n] = \{0, 1, \ldots, n-1\}$. Let $A \sqcup B$ denote the disjoint union of sets A and B. Abbreviate tuples $(x_0, \ldots, x_{n-1})$ by boldface letter $\mathbf{x}$. A *multigraph* $X = (V(X), E(X))$ is defined by a finite set $V(X)$ of vertices and a multiset $E(X)$ of edges, which are one- or two-element subsets of $V(X)$. That is, edges are undirected and loops are allowed. A *graph* X has no loops or multiedges (so $E(X)$ is a set of two-element subsets of $V(X)$). We will often identify the vertex set of a q-vertex graph $V(X)$ with $[q] := \{0, 1, \ldots, q-1\}$, and consider the *adjacency matrix* $A_X \in \{0,1\}^{q \times q}$ of X, defined by $(A_X)_{xy} = 1 \iff \{x, y\} \in E(X)$.

2.1 Tensors and signatures

We work with the finite-dimensional vector space $\mathbb{K}^q$, with standard basis vectors $e_1, \ldots, e_q$, and its dual space $(\mathbb{K}^q)^*$ with standard basis vectors $e_1^*, \ldots, e_q^*$. For $\ell, r \geq 0$, $(\mathbb{K}^q)^{\otimes \ell} \otimes ((\mathbb{K}^q)^*)^{\otimes r}$ is, as usual, a $q^{\ell+r}$-dimensional vector space with basis vectors $e_{x_1} \otimes \ldots \otimes e_{x_\ell} \otimes e_{y_1}^* \otimes \ldots \otimes e_{y_r}^*$ for $(x_1, \ldots, x_\ell, y_1, \ldots, y_r) \in [q]^{\ell+r}$.

Definition 2.1.1 ($\mathcal{V}$)**.** The *mixed tensor algebra* over $\mathbb{K}^q$ is

$$\mathcal{V} = \mathcal{V}(\mathbb{K}^q) := \bigcup_{\ell, r \geq 0} {}_\ell\mathcal{V}_r, \text{ where } {}_\ell\mathcal{V}_r = (\mathbb{K}^q)^{\otimes \ell} \otimes ((\mathbb{K}^q)^*)^{\otimes r}.$$

$\mathcal{V}(\mathbb{K}^q)$ is bigraded $\mathbb{K}$-vector space, meaning each grade ${}_\ell\mathcal{V}_r$ is a $\mathbb{K}$-vector space. Tensors in $\bigcup_{n \geq 1} {}_n\mathcal{V}_0 \subset \mathcal{V}(\mathbb{K}^q)$ are *contravariant* or *left-facing*, and tensors in $\bigcup_{n \geq 1} {}_0\mathcal{V}_n$ are *covariant* or *right-*

facing. Tensors in ${}_\ell\mathcal{V}_r$ are said to have *arity* $\ell + r$, with *left arity* ℓ and *right arity* r. Note that ${}_0\mathcal{V}_0 = \mathbb{K}$.

We will also view a tensor

$$\sum_{x_1,\ldots,x_\ell=1}^{q} \sum_{y_1,\ldots,y_r=1}^{q} a_{x_1\ldots x_\ell, y_1\ldots, y_r} e_{x_1} \otimes \ldots \otimes e_{x_\ell} \otimes e^*_{y_1} \otimes \ldots \otimes e^*_{y_r} \in {}_\ell\mathcal{V}_r \tag{2.1.1}$$

in *flattened* form: a matrix in $\mathbb{K}^{q^\ell \times q^r}$ (indexed by $[q]^\ell \times [q]^r$ lexicographically) whose entry at $x_1 \ldots x_\ell, y_1 \ldots y_r$ is $a_{x_1\ldots x_\ell, y_1\ldots, y_r}$. Then contravariant and covariant tensors are column and row vectors, respectively. Use $I = I_q$ to denote the identity matrix $\sum_{x=1}^q e_i \otimes e_i^*$. We depict a tensor in ${}_\ell\mathcal{V}_r$ diagrammatically as a node with contravariant and covariant inputs along left- and right-facing *dangling edges*, respectively, as shown in Figure 2.1.

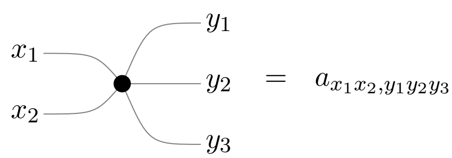

Figure 2.1: Diagrammatic representation of the tensor in (2.1.1) with $\ell = 2$ and $r = 3$.

Definition 2.1.2 (Signature)**.** A *signature* is a function $F : V(F)^n \to \mathbb{K}$ on $n = \text{arity}(F)$ inputs from a finite *domain* $V(F)$. We usually have $V(F) = [q]$.

A signature is *symmetric* if its value is invariant under reordering of its inputs.

Associate a signature $F : [q]^n \to \mathbb{K}$ with the tensors $F^{\ell,r} \in {}_\ell\mathcal{V}(\mathbb{K}^q)_r$ for every $\ell, r \geq 0$ with $\ell + r = n$, where

$$F^{\ell,r} := \sum_{x_1,\ldots,x_\ell=1}^{q} \sum_{y_1,\ldots,y_r=1}^{q} F(x_1, \ldots, x_\ell, y_r, \ldots, y_1) e_{x_1} \otimes \ldots \otimes e_{x_\ell} \otimes e^*_{y_1} \otimes \ldots \otimes e^*_{y_r}. \tag{2.1.2}$$

Conversely, the coefficients of any tensor $F' \in {}_\ell\mathcal{V}_r$ determine a unique n-ary signature F such that $F^{\ell,r} = F'$. Note that the order of inputs to F from the covariant (right) factors is reversed, so that the signature associated with a tensor takes its inputs in cyclic counterclockwise order, while the tensor's coefficients are defined by top-down input order on both sides. This convention ensures that the reshaping of $F^{\ell,r}$ to $F^{\ell',r'}$ with $\ell' + r' = \ell + r$ is realized diagrammatically by *pivoting* dangling edges between left and right while preserving their cyclic order around the central vertex (see Figure 2.2). In particular, pivoting is a planar operation – it does not entail crossing any dangling edges. This will become relevant in Chapter 7.

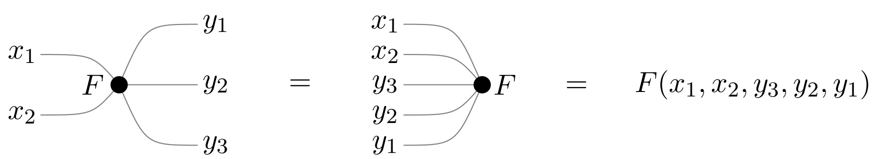

Figure 2.2: The diagrammatic representation of two reshapings $F^{2,3}$ and $F^{5,0}$ of a signature F. Since we don't think of F as having a defined shape, the central node itself is labeled F. Cyclically pivoting dangling edges between left and right does not change F.

We use the symbol F to refer either to a specific tensor in ${}_\ell\mathcal{V}_r$ with shape (ℓ, r), or to the signature (shapeless tensor) underlying such a tensor and encoded by its coefficients, depending on context. Conventionally, we view signatures as contravariant/left-facing tensors (in such a tensor, no input order is reversed):

Definition 2.1.3 ($\mathcal{S}, {}_n\mathcal{S}$)**.** Define

$$\mathcal{S} = \mathcal{S}(\mathbb{K}^q) := \bigcup_{n \geq 0} {}_n\mathcal{V}(\mathbb{K}^q)_0$$

to be the set of all left-facing tensors, thought of as shapeless signatures. Given $\mathcal{F} \subset \mathcal{S}$, define ${}_n\mathcal{F} := \mathcal{F} \cap {}_n\mathcal{V}_0$ and ${}_{\leq n}\mathcal{F} := \mathcal{F} \cap \bigcup_{i \leq n} {}_n\mathcal{V}_0$ to be the subspaces of signatures of arity n and at most n, respectively.

We frequently work with the following set of *equality* signatures:

Definition 2.1.4 ($\mathcal{EQ}$)**.** Define $\mathcal{EQ} := \{=_n \mid n \geq 1\}$, where $=_n$ is the n-ary signature defined by

$$(=_n)(x_1, \ldots, x_n) = \begin{cases} 1 & x_1 = \ldots = x_n \\ 0 & \text{otherwise.} \end{cases}$$

Definition 2.1.5 ($\bullet$)**.** Given $F, G \in {}_n\mathcal{S}$, define the entrywise product $F \bullet G \in {}_n\mathcal{S}$ by, for $\mathbf{x} \in [q]^n$,

$$(F \bullet G)(\mathbf{x}) = F(\mathbf{x})G(\mathbf{x}).$$

2.2 Holant problems

Given a signature set $\mathcal{F} \subset \mathcal{S}$, a *signature grid* (or $\mathcal{F}$-grid) Ω is a multigraph along with an assignment of an n-ary $F_v \in \mathcal{F}$ to each degree-n vertex v in Ω, with an ordering of the n edges $\delta(v)$ incident to v to serve as the n inputs to F.

For technical reasons, we also allow Ω to contain *vertexless loops* $\bigcirc$ (a loop with one edge and no vertex). The goal of Holant($\mathcal{F}$) is to compute the *Holant value*

$$\text{Holant}_\Omega(\mathcal{F}) = \sum_{\sigma:E(\Omega)\to[q]} \prod_{v\in V(\Omega)} F_v(\sigma(\delta(v))) \tag{2.2.1}$$

of Ω, where $F_v(\sigma(\delta(v)))$ is the evaluation of F_v on the n domain elements assigned to the incident edges of v (recall the example in Figure 1.1 in the introduction. There, we didn't need to specify an edge ordering around each vertex because the signatures EO_n are symmetric). The Holant value of a disconnected signature grid is the product of the Holant values of its connected components. In particular, each vertexless loop in Ω contributes a factor q.

We also consider the well-studied bipartite version of Holant.

Definition 2.2.1 (Holant($\mathcal{F} \mid \mathcal{F}'$))**.** Given two signature sets $\mathcal{F}, \mathcal{F}' \subset \mathcal{S}(\mathbb{K}^q)$, define Holant($\mathcal{F} \mid \mathcal{F}'$) to be the restriction of Holant($\mathcal{F} \cup \mathcal{F}'$) to bipartite signature grids in which vertices from the two bipartitions are assigned signatures from $\mathcal{F}$ and $\mathcal{F}'$, respectively.

2.2.1 Gadgets

Definition 2.2.2 (Holant gadget)**.** For signature set $\mathcal{F}$, an Holant $\mathcal{F}$ *gadget* $\mathbf{K}$ is a Holant($\mathcal{F}$) signature grid equipped with an ordered set of *dangling edges* with zero or one endpoints. An (ℓ, r) gadget has $\ell \geq 0$ left dangling edges $e_0, \ldots, e_{\ell-1}$ and $r \geq 0$ right dangling edges $e'_0, \ldots, e'_{r-1}$ drawn extending out to the left and right of the gadget, respectively, in top-to-bottom order of the dangling ends. Include the dangling edges in the set $E(\mathbf{K})$ of edges of $\mathbf{K}$. See Figure 2.2 for examples of gadgets.

Two-sided dangling edges, also called *wires*, can have two left ends, two right ends, or one left and one right end (see Figure 2.1; in particular, $e_0, \ldots, e_{\ell-1}, e'_0, \ldots, e'_{r-1}$ may not be distinct when this list contains both ends of a wire). The first two types are called *pivots*, and are important enough to merit the diagrammatic shorthands $\supset$ and $\subset$. The third type is called a *crossing wire.*

See Figure 2.2 for examples of gadgets. The following definition ensures that, if F is the signature of a Holant($\mathcal{F}$) gadget $\mathbf{K}_F$, then any $(\mathcal{F} \cup \{F\})$-grid corresponds to a Holant($\mathcal{F}$) signature grid with the same Holant value constructed by replacing every vertex assigned F with $\mathbf{K}_F$ (with appropriately ordered dangling edges).

Definition 2.2.3 ($M(\mathbf{K})$)**.** Every (ℓ, r) Holant($\mathcal{F}$) gadget $\mathbf{K}$ defines a tensor $M(\mathbf{K}) \in {}_\ell\mathcal{V}_r$, thought of as a matrix in $\mathbb{K}^{q^\ell \times q^r}$, by letting the coefficient of $e_{x_1} \otimes \ldots \otimes e_{x_\ell} \otimes e^*_{y_1} \otimes \ldots \otimes e^*_{y_r}$ in $M(\mathbf{K})$ be the Holant value of $\mathbf{K}$ when the ℓ left and r right dangling edges are fixed to values $x_1, \ldots, x_\ell$ and $y_1, \ldots, y_r$ (summing only over assignments σ to the internal edges). That is,

$$M(\mathbf{K})_{\mathbf{x},\mathbf{y}} = \sum_{\substack{\sigma: E(\mathbf{K}) \to [q] \\ \forall i \in [\ell]: \sigma(e_i) = x_i \\ \forall j \in [r]: \sigma(e'_j) = y_j}} \prod_{v \in V} F_v(\sigma(\delta(v))) \quad \text{for } \mathbf{x} \in [q]^\ell \text{ and } \mathbf{y} \in [q]^r.$$

Call $M(\mathbf{K})$ the tensor or matrix of $\mathbf{K}$, and the underlying signature of $M(\mathbf{K})$ the signature of $\mathbf{K}$.

As in Figure 2.2, the signature of a gadget depends only on the cyclic order, and not the left/right orientations, of its dangling edges. The signature of a wire is $(=_2)$, as the inputs to its two ends must match. In particular, $\supset = (=_2)^{2,0}$, $\subset = (=_2)^{0,2}$, and the crossing wire has tensor $I = (=_2)^{1,1}$.

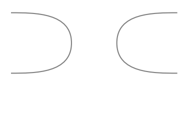

Figure 2.1: Left wire/pivot $=_2$ in ${}_2\mathcal{V}_0$ (top left), right wire/pivot ${}_0\mathcal{V}_2$ (top right) and crossing wire $I \in {}_1\mathcal{V}_1$ (bottom), respectively.

Definition 2.2.4 (Gadget $\circ, \otimes, \top, \dagger$)**.** Define the following four operations on Holant($\mathcal{F}$) gadgets:

- For (ℓ, r) gadget $\mathbf{K}$ and (r, d) gadget $\mathbf{J}$, construct the (ℓ, d) gadget $\mathbf{K} \circ \mathbf{J}$ from the disjoint union of $\mathbf{K}$ and $\mathbf{J}$, with $\mathbf{K}$ placed to the left of $\mathbf{J}$, by connecting the ith right dangling edge of $\mathbf{K}$ with the i left dangling edge of $\mathbf{J}$ for $i \in [r]$. The left and right dangling edges of $\mathbf{K} \circ \mathbf{J}$ are the left dangling edges of $\mathbf{K}$ and the right dangling edges of $\mathbf{J}$, respectively.

- For (ℓ, r) gadget $\mathbf{K}$ and (ℓ', r') gadget $\mathbf{J}$, construct the $(\ell + \ell', r + r')$ gadget $\mathbf{K} \otimes \mathbf{J}$ as the disjoint union of $\mathbf{K}$ and $\mathbf{J}$, placing $\mathbf{K}$ above $\mathbf{J}$. The left/right dangling edges of $\mathbf{K} \otimes \mathbf{J}$ are the concatenation of the left/right dangling edges of $\mathbf{K}$, then $\mathbf{J}$.

- For (ℓ, r) gadget $\mathbf{K}$, construct the (r, ℓ) gadget $\mathbf{K}^\top$ by exchanging the roles of left and right dangling edges, preserving the order of both (visually, horizontally reflect $\mathbf{K}$).

- If $\mathcal{F} \subset \mathcal{S}(\mathbb{C}^q)$, then, for (ℓ, r) gadget $\mathbf{K}$, construct the (r, ℓ) gadget $\mathbf{K}^\dagger$ from $\mathbf{K}^\top$ by replacing every signature F assigned to a vertex in $\mathbf{K}^\top$ with its entrywise conjugate $\overline{F}$.

See Figure 2.2. Now we use gadget operations to define the corresponding tensor operations.

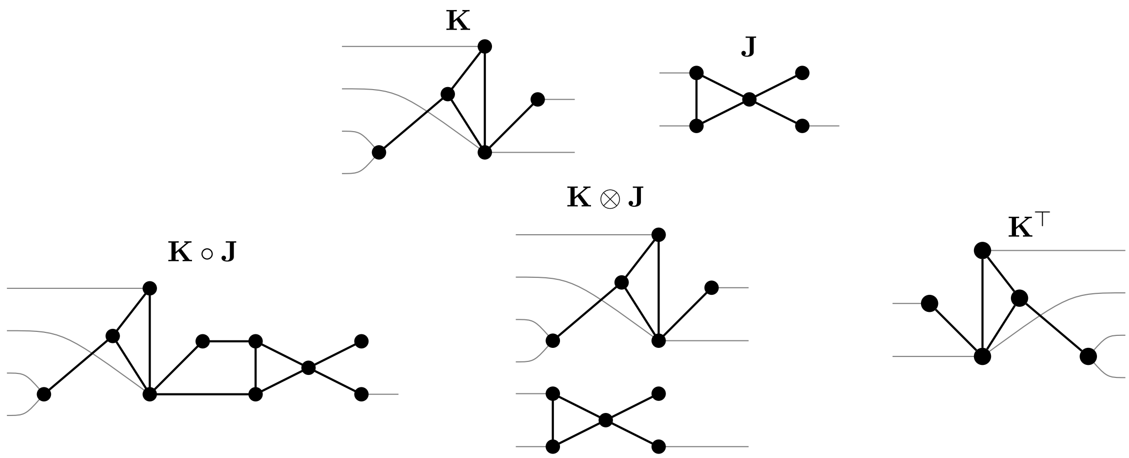

Figure 2.2: Operations on the $(4, 2)$ gadget **K** and the $(2, 1)$ gadget **J**. Dangling edges are drawn thinner than internal edges. See Definition 2.2.4.

Definition 2.2.5 ($\mathbf{V}(F)$)**.** The diagrammatic representation in Figure 2.1 of a tensor $F^{\ell,r} \in {}_\ell\mathcal{V}_r$ is a gadget $\mathbf{V}(F)$ consisting a single vertex assigned F incident to ℓ left and r right dangling edges satisfying $M(\mathbf{V}(F)) = F^{\ell,r}$.

Definition 2.2.6 (Tensor $\circ, \otimes, \top, \dagger$)**.** Define the following operations on tensors.

- Given $F \in {}_\ell\mathcal{V}_r$ and $G \in {}_r\mathcal{V}_d$, define $F \circ G \in {}_\ell\mathcal{V}_r$ as the tensor of $\mathbf{V}(F) \circ \mathbf{V}(G)$.
- Given $F \in {}_\ell\mathcal{V}_r$ and $G \in {}_{\ell'}\mathcal{V}_{r'}$, define $F \otimes G \in {}_{\ell+\ell'}\mathcal{V}_{r+r'}$ as the tensor of $\mathbf{V}(F) \otimes \mathbf{V}(G)$.
- Given $F \in {}_\ell\mathcal{V}_r$, define $F^\top \in {}_r\mathcal{V}_\ell$ as the tensor of $\mathbf{V}(F)^\top$.
- Given $F \in {}_\ell\mathcal{V}(\mathbb{C}^q)_r$, define $F^\dagger \in {}_r\mathcal{V}(\mathbb{C}^q)_\ell$ as the tensor of $\mathbf{V}(F)^\dagger$.

As one might expect, the gadget operations $\circ, \otimes, \top, \dagger$ induce the respective matrix operations – composition, Kronecker product, transpose, conjugate transpose – on the flattened forms of their tensors (see e.g. [CC17a, Section 1.3]). In particular, $\otimes$ and $\dagger$ (also denoted $*$) are the usual operations on $\mathcal{V}$ – a $\mathbb{K}$ algebra with product $\otimes$ – as in e.g. [Sch08b; DM23].

Proposition 2.2.1. *For gadgets* **K** *and* **J** *with matrices* $M(\mathbf{K})$ *and* $M(\mathbf{J})$*, we have* $M(\mathbf{K} \circ \mathbf{J}) = M(\mathbf{K}) \circ M(\mathbf{J})$*,* $M(\mathbf{K} \otimes \mathbf{J}) = M(\mathbf{K}) \otimes M(\mathbf{J})$*,* $M(\mathbf{K}^\top) = M(\mathbf{K})^\top$*, and* $M(\mathbf{K}^\dagger) = M(\mathbf{K})^\dagger$*.*

Definition 2.2.7 ($\langle\cdot,\cdot\rangle, \|\cdot\|$)**.** Define the standard bilinear form $\langle\cdot,\cdot\rangle : {}_\ell\mathcal{V}_r \times {}_r\mathcal{V}_\ell \to \mathbb{K}$ by $\langle F_1, F_2\rangle = \mathrm{tr}(F_1 \circ F_2)$, where tr denotes matrix trace.

Equivalently, if $F_1 \in {}_\ell\mathcal{V}_r$ and $F_2 \in {}_r\mathcal{V}_\ell$ are the tensors of Holant($\mathcal{F}$) gadgets $\mathbf{K}_1$ and $\mathbf{K}_2$, then $\langle F_1, F_2\rangle$ is the Holant value of the Holant($\mathcal{F}$) grid constructed from $\mathbf{K}_1 \circ \mathbf{K}_2$ by also connecting the ith left dangling edge of $\mathbf{K}_1$ with the ith right dangling edge of $\mathbf{K}_2$ for $i \in [\ell]$ (see Figure 2.3).

For $F \in \mathcal{V}(\mathbb{C}^q)$, define the Frobenius norm $\|F\| := \langle F, F^\dagger\rangle$, which satisfies $\|F\| \geq 0$ and $\|F\| = 0 \iff F = 0$.

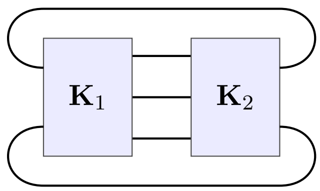

Figure 2.3: The grid Ω such that Holant(Ω) $= \langle F_1, F_2\rangle$ for $(2,3)$-$\mathcal{F}$-gadget $\mathbf{K}_1$ with tensor F_1 and $(3,2)$-$\mathcal{F}$-gadget $\mathbf{K}_2$ with tensor F_2.

Definition 2.2.8. Gadget $\mathbf{K}$ is *planar* if its underlying multigraph with dangling edges has a plane embedding (called a *plane gadget*) such that no edges (dangling or not) cross, and the dangling ends of the left and right dangling edges are in top-down order in the outer face: there are vertical lines L_ℓ and L_r in the plane lying entirely to the left and right of $\mathbf{K}$, except the dangling ends of $e_0, \ldots, e_{\ell-1}$ lie on L_ℓ in top-down order and the dangling ends of $e'_0, \ldots, e'_{r-1}$ lie on L_r in top-down order.

As is intuitively clear from Figure 2.2, the gadget operations preserve planarity. See also [MR20, Lemmas 5.12-5.14], the same result for operations on bilabeled graphs, which are equivalent to Holant($\mathcal{EQ}, A_X$) gadgets (the setting of counting homomorphisms to an unweighted graph X).

Proposition 2.2.2. *The gadget operations* $\circ, \otimes, \top, \dagger$ *preserve planarity:*

- *If* $\mathbf{K}$ *and* $\mathbf{J}$ *are planar, then* $\mathbf{K} \circ \mathbf{J}$ *is planar.*
- *If* $\mathbf{K}$ *and* $\mathbf{J}$ *are planar, then* $\mathbf{K} \otimes \mathbf{J}$ *is planar.*
- *If* $\mathbf{K}$ *is planar, then* $\mathbf{K}^\top$ *is planar.*
- *If* $\mathbf{K}$ *is planar, then* $\mathbf{K}^\dagger$ *is planar.*

Proof. For $\circ$, identify the right line L_r of $\mathbf{K}$ with the left line L_ℓ of $\mathbf{J}$. Then, since composition matches the dangling edges in top-down order, move the left dangling edges of $\mathbf{K}$ and right dangling edges of $\mathbf{J}$ up and down locally around this merged line without crossing to connect the dangling ends. Do not otherwise modify the plane embeddings of $\mathbf{K}$ and $\mathbf{J}$. The other operations $\otimes, \top, \dagger$ also preserve planarity of the underlying multigraphs, so it suffices to show that the new dangling edge orders require no crossings. For $\top$ and $\dagger$, horizontally reflect the underlying multigraph of $\mathbf{K}$ (preserving its planarity) and exchange the positions of L_ℓ and L_r. Let the left and right lines of $\mathbf{K} \circ \mathbf{J}$ be the left line of $\mathbf{K}$ and right line of $\mathbf{J}$, respectively, shifted out horizontally if necessary. For $\otimes$, after appropriate vertical shifts of $\mathbf{K}$ and $\mathbf{J}$ and appropriate horizontal extensions of the dangling edges of $\mathbf{K}$ or $\mathbf{J}$ (with no other modifications to either planar embedding), identify the right lines of $\mathbf{K}$ and $\mathbf{J}$ and the left lines of $\mathbf{K}$ and $\mathbf{J}$, with the dangling ends of $\mathbf{K}$'s dangling edges above those of $\mathbf{J}$ on both sides. □

2.3 Counting constraint satisfaction problems

Definition 2.3.1 (#CSP, Z)**.** A counting constraint satisfaction problem #CSP($\mathcal{F}$) is parameterized by a set $\mathcal{F}$ signatures (usually called *constraint functions* in the CSP literature). A *#CSP($\mathcal{F}$) instance* $I = (V, C)$ is defined by a set V of variables and a multiset C of *constraints*. Each constraint $(F, v_1, \dots, v_{n_F})$ consists of a constraint function $F \in \mathcal{F}$ and an ordered tuple of variables to which F is applied.

The *partition function* Z, on input #CSP($\mathcal{F}$) instance $I = (V, C)$, outputs

$$Z(I) = \sum_{\phi: V \to V(\mathcal{F})} \prod_{(F, v_1, \dots, v_{n_F}) \in C} F(\phi(v_1), \dots, \phi(v_{n_F})).$$

A common example of #CSP is computing the number of homomorphisms from graph K to X, where a homomorphism $K \to X$ is a map $\phi : V(K) \to V(X)$ that preserves adjacency: $uv \in E(K) \implies \phi(u)\phi(v) \in E(X)$. Given K and X, consider a #CSP(A_X) instance I_K where the vertices of K are variables and each edge of K is a constraint applying binary signature A_X (whose matrix form is the adjacency matrix of X) to the edge's two endpoints. Then

$$Z(I_K) = \sum_{\phi: V(K) \to V(X)} \prod_{uv \in E(K)} (A_X)_{uv} = \hom(K, X), \tag{2.3.1}$$

the number of homomorphisms $K \to X$. We also consider (2.3.1) for $\#\mathrm{CSP}(A_X)$ instances I_K for an arbitrary matrix $A_X \in \mathbb{K}^{V(X)\times V(X)}$, considered to be the adjacency matrix of a $\mathbb{K}$-edge-weighted graph X, and define the number of $\mathbb{K}$-weighted homomorphisms $\hom(K, X) = Z(I_K)$.

Let $\mathcal{F}$ be a set of constraint functions. To each $\#\mathrm{CSP}(\mathcal{F})$ instance $I = (V, C)$ we associate a biparite signature grid $\Omega(I)$ in the context of $\mathrm{Holant}(\mathcal{F} \mid \mathcal{EQ})$, constructed as follows. For each variable $v \in V$, create a vertex u_v – called an *equality vertex* – assigned $=_{n_v}$, where n_v is the total number of appearances of v in the constraints in C. Then, for every constraint $c = (F, v_1, \dots, v_{n_F}) \in C$, create a vertex w_c assigned F, called a *constraint vertex*, adjacent to the equality vertices $u_{v_1}, \dots, u_{v_{n_F}}$ (ordering the edges incident to w_c accordingly). Any edge assignment σ must assign all edges incident to an equality vertex the same value (or else the term corresponding to σ is 0), so we can view σ as #CSP variable assignment. Hence $Z(I) = \mathrm{Holant}_{\Omega(I)}(\mathcal{F} \mid \mathcal{EQ})$. Conversely, given a $\mathrm{Holant}(\mathcal{F} \mid \mathcal{EQ})$-grid, construct a $\#\mathrm{CSP}(\mathcal{F})$ instance by creating a variable or constraint for every vertex assigned a signature in $\mathcal{EQ}$ or $\mathcal{F}$, respectively, and placing variables in constraints according to adjacency, respecting the order of edges incident to each $\mathcal{F}$ vertex. These maps are inverses, so $\#\mathrm{CSP}(\mathcal{F})$ instances are in bijection with $\mathrm{Holant}(\mathcal{F} \mid \mathcal{EQ})$ grids.

For example, the Holant signature grid $\Omega(I_K)$ constructed from the $\#\mathrm{CSP}(A_X)$ instance K for counting graph homomorphisms from K to X starts with underlying graph K, with K's vertices assigned the appropriate equality signature from $\mathcal{EQ}$, and subdivides each of K's edges by placing degree-2 constraint vertices, assigned signature A_X, to the labeled equality vertices. See Figure 2.1.

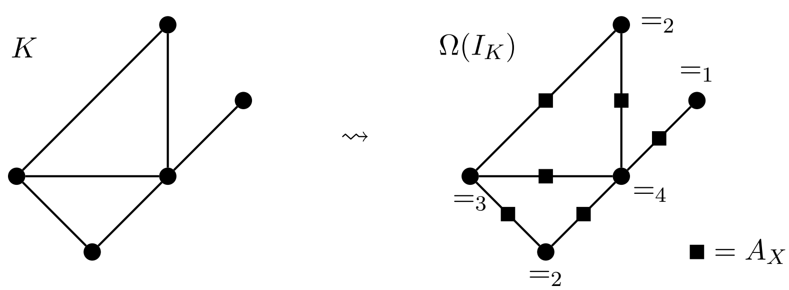

Figure 2.1: A graph K and the corresponding $\mathrm{Holant}(A_X \mid \mathcal{EQ})$ grid $\Omega(I_K)$ for computing the number of homomorphisms from K to X. Square vertices are assigned signature A_X.

Given a not-necessarily-bipartite $\mathrm{Holant}(\mathcal{F} \cup \mathcal{EQ})$ grid or gadget, we may construct a bipartite $\mathrm{Holant}(\mathcal{F} \mid \mathcal{EQ})$ grid or gadget with the same value or signature by inserting a dummy degree-2 constraint vertex assigned $(=_2) \in \mathcal{EQ}$ between adjacent $\mathcal{F}$ vertices and contracting the edge between adjacent $\mathcal{EQ}$ vertices assigned $=_a$ and $=_b$, combining them into a single vertex assigned $=_{a+b-2}$. Informally,

$$\#\mathrm{CSP}(\mathcal{F}) \equiv \mathrm{Holant}(\mathcal{F} \mid \mathcal{EQ}) \equiv \mathrm{Holant}(\mathcal{F} \cup \mathcal{EQ}). \tag{2.3.2}$$

2.4 Transformations and indistinguishability

Throughout, we consider pairs of multisubsets $\mathcal{F}, \mathcal{G} \subset \mathcal{V}(\mathbb{K}^q)$, called *tensor sets* or *signature sets*.

Definition 2.4.1 (bijective, $\leftrightsquigarrow$)**.** Multisubsets $\mathcal{F} \subset \mathcal{V}(\mathbb{K}^q)$ and $\mathcal{G} \subset \mathcal{V}(\mathbb{K}^p)$ are *bijective* if there is a left- and right-arity-preserving multiset bijection (for signature sets, just an arity-preserving multiset bijection) $\leftrightsquigarrow$ between $\mathcal{F}$ and $\mathcal{G}$. Call $\mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}$ *corresponding* tensors.

Definition 2.4.2 ($\Omega_{\mathcal{F}\to\mathcal{G}}$, Holant-indistinguishable)**.** For a pair $\mathcal{F}, \mathcal{G}$ of bijective signature sets and a $\mathrm{Holant}(\mathcal{F})$ signature grid Ω, define the $\mathrm{Holant}(\mathcal{G})$ signature grid $\Omega_{\mathcal{F}\to\mathcal{G}}$ by replacing every $F \in \mathcal{F}$ assigned to a vertex in Ω with the corresponding $G \in \mathcal{G}$.

Say $\mathcal{F}$ and $\mathcal{G}$ are *Holant-indistinguishable* if $\mathrm{Holant}_{\mathcal{F}}(\Omega) = \mathrm{Holant}_{\mathcal{G}}(\Omega_{\mathcal{F}\to\mathcal{G}})$ for every $\mathrm{Holant}(\mathcal{F})$ signature grid Ω.

Observe that bijective $\mathcal{F}$ and $\mathcal{G}$ are not required to have the same domain size (we may have $p \neq q$). However, if $\mathcal{F}$ and $\mathcal{G}$ are Holant-indistinguishable (or Bi-Holant-indistinguishable; see below), then they must have the same domain size because the Holant value of the vertexless loop $\bigcirc$ equals the domain size, and $\bigcirc_{\mathcal{F}\to\mathcal{G}} = \bigcirc$.

Define $\mathrm{GL}_q := \mathrm{GL}_q(\mathbb{K})$ to be general linear group of invertible $q \times q$ matrices over $\mathbf{K}$. Define $O_q(\mathbb{K}) \subset \mathrm{GL}_q(\mathbb{K})$ to be the group of orthogonal matrices: those H satisfying $H^\top H = HH^\top = I$. We usually consider the well-behaved *real* orthogonal group $O_q := O_q(\mathbb{R})$. For $\mathbb{K} = \mathbb{C}$, define $U_q \subset \mathrm{GL}_q(\mathbb{C})$ to be the group of complex unitary matrices: those U satisfying $H^\dagger H = HH^\dagger = I$, where $\dagger$ denotes conjugate transpose. This thesis studies the actions of these classical linear algebraic groups on signatures and tensors.

Definition 2.4.3 ($T \cdot F$, $T\mathcal{F}$, $T(\mathcal{F} \mid \mathcal{F}')$)**.** For $T \in \mathrm{GL}_q$ and $F \in {}_\ell\mathcal{V}_r$, define $T \cdot F := T^{\otimes \ell} \circ F \circ (T^{-1})^{\otimes r} \in {}_\ell\mathcal{V}_r$. Then for $\mathcal{F} \subset \mathcal{V}$, define $T\mathcal{F} = \{T \cdot F \mid F \in \mathcal{F}\}$, a set bijective to $\mathcal{F}$ via $T \cdot F \leftrightsquigarrow F$.

In particular, T acts on a signature $F \in {}_n\mathcal{S} = {}_n\mathcal{V}_0$ by $T \cdot F = T^{\otimes n} \circ F$ (a matrix-column-vector product). However, for the purpose of transformation, we think of $\mathcal{F}'$ in a bipartite Holant set $\mathcal{F} \mid \mathcal{F}'$ as a set of covariant tensors (or row vectors), so $T(\mathcal{F} \mid \mathcal{F}') = (T\,\mathcal{F} \mid \mathcal{F}'\,T^{-1})$, where $\mathcal{F}'\,T^{-1} = \{F' \circ (T^{-1})^{\otimes\,\mathrm{arity}(F')} : F' \in \mathcal{F}\}$.

As the definition suggests, it is useful to think of T itself as a binary signature that can be assigned to a vertex. This gives the gadget perspective in Figure 2.1.

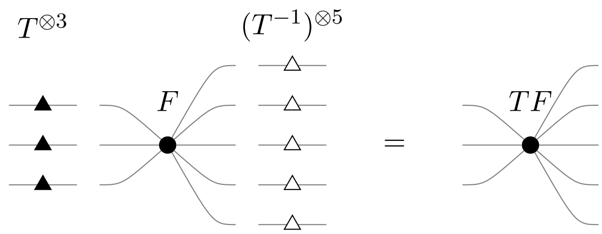

Figure 2.1: The transformation $T^{\otimes 3}F(T^{-1})^{\otimes 5} = T \cdot F$ for $F \in {}_3\mathcal{V}_5$.

The following theorem, a powerful reduction tool introduced by Valiant [Val08], is the genesis for Holant problems.

Theorem 2.4.1 (The Holant Theorem [Val08])**.** *If $\mathcal{F} \,|\, \mathcal{F}' = T(\mathcal{G} \,|\, \mathcal{G}')$ for $T \in \mathrm{GL}_q$, then $\mathcal{F} \,|\, \mathcal{F}'$ and $\mathcal{G} \,|\, \mathcal{G}'$ are Holant-indistinguishable.*

Theorem 2.4.1 follows from the fact that left/right contractions are GL_q-equivariant for the action of GL_q in Definition 2.4.3. See Figure 2.2

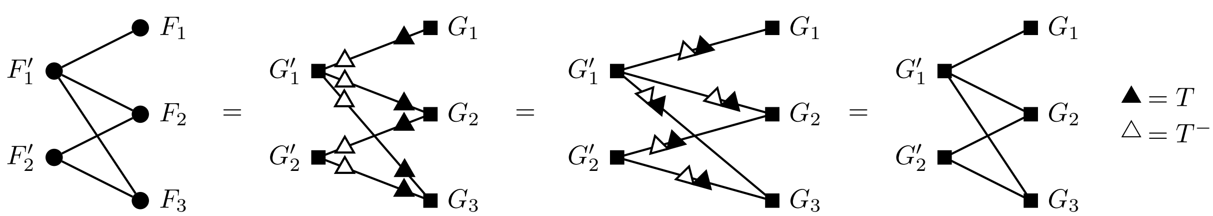

Figure 2.2: An example of Theorem 2.4.1, with $F_i' = G_i'(T^{-1})^{\otimes n_i}$ and $F_i = T^{\otimes n_i} G_i$.

The next proposition follows from direct calculation, or the diagrammatic calculus in Figure 2.3

Proposition 2.4.1. *A matrix T is orthogonal if and only if $T \cdot \supset = \supset$ or $T \cdot \subset = \subset$.*

Figure 2.3: Illustrating Proposition 2.4.1. Pivoting the lower dangling edge between left and right flips the T on this edge, transposing it.

Convert an $\mathcal{F}$-grid Ω into a $(\mathcal{F} \mid =_2)$-grid by placing a degree-2 vertex assigned $=_2$ on each edge. The resulting grid is bipartite between $=_2$ and $\mathcal{F}$ and, since $=_2$ acts identically to an edge, does not change the Holant value. Conversely, given an $(\mathcal{F} \mid =_2)$-grid Ω, replace each vertex assigned $=_2$ with an edge. This connects arbitrary inputs of signatures in $\mathcal{F}$, but this is allowed in Holant($\mathcal{F}$).

$$\text{Holant}(\mathcal{F}) \equiv \text{Holant}(\mathcal{F} \mid =_2). \tag{2.4.1}$$

Recall that $=_2$ is the underlying signature of $\subset$. We will sometimes abuse notation and write the latter problem as Holant($\mathcal{F} \mid \subset$). Now Proposition 2.4.1 gives the following well-known special case of the Holant theorem

Corollary 2.4.1 (The Orthogonal Holant Theorem)**.** *If $\mathcal{G} = H\,\mathcal{F}$ for $H \in O_q(\mathbb{K})$, then $\mathcal{F}$ and $\mathcal{G}$ are Holant-indistinguishable.*

2.5 Overview: counting indistinguishability theorems

Say graphs F and G are *homomorphism indistinguishable* over a class $\mathfrak{G}$ of simple graphs if $\hom(X, F) = \hom(X, G)$ for every $X \in \mathfrak{G}$. Almost 60 years aro, Lovász proved the first *counting indistinguishability theorem*, showing that homomorphism counts from all graphs determine a graph up to isomorphism.

Theorem 2.5.1 (Lovász [Lov67])**.** *Graphs X and Y are homomorphism-indistingushable over the class of all graphs if and only if they are isomorphic.*

The past several years have seen a resurgence in the study of homomorphism indistinguishability, with new results characterizing when two graphs are homomorphism-indistinguishable over smaller graph classes $\mathfrak{G}$. These characterizations include quantum isomorphism for planar graphs [MR20; Kar+25], fractional isomorphism for trees [DGR18], orthogonal transformation for cycles, pseudo-stochastic transformation for paths [DGR18; GRS25], indistinguishability by bounded-variable counting logic for bounded-treewidth graphs [Dvo10], isomorphism itself for 2-degenerate graphs [Dvo10], and more [Sep24; GRS25; RS24].

This thesis extends the concept of counting indistinguishability from graph homomorphisms to #CSP, and, more generally, to Holant problems. In general, the more expressive a counting problem, the more powerful its indistinguishability theorem. Since graph homomorphisms are expressible in the #CSP framework and #CSP is expressible in the Holant framework, many of the graph homomorphism results in the previous paragraph are special cases of the results in this thesis, which progress roughly from least to most general. In Chapter 3, we prove the indistinguishability theorem for #CSP: constraint function sets $\mathcal{F}$ and $\mathcal{G}$ are #CSP-indistinguishable if and only if they are isomorphic [1]. Lovász's theorem is the special case $\mathcal{F} = \{A_X\}$ and $\mathcal{G} = \{A_Y\}$ for graphs X and Y. In Chapter 5, we give a cleaner alternate proof of this #CSP indistinguishability theorem for $\mathbb{K} = \mathbb{C}$. We then employ this alternate proof pattern in Chapter 7 to prove a generalization from graph homomorphisms to #CSP of the aforementioned quantum isomorphism result of Mančinska and Roberson [MR20]: $\mathcal{F}$ and $\mathcal{G}$ are quantum isomorphic if and only if they are Pl-#CSP-indistinguishable (where Pl-#CSP restricts to instances with planar constraint-variable incidence graphs). Since #CSP is expressible using Pl-#CSP by adding to every constraint function set a special constraint $\times$ that allows edge crossings, this planar/quantum #CSP indistinguishability theorem generalizes the classical version in Chapter 5. In Chapter 8, we prove the indistinguishability theorem for Holant, the converse of the orthogonal Holant theorem (Corollary 2.4.1): there is an orthogonal transformation between (real-valued) $\mathcal{F}$ and $\mathcal{G}$ if and only if $\mathcal{F}$ and $\mathcal{G}$ are Holant-indistinguishable. This time, adding $\mathcal{EQ}$ to all signature sets returns us to the #CSP setting, and an orthogonal transformation preserving $\mathcal{EQ}$ must be an isomorphism, so we again recover the #CSP indistinguishability theorem. In Chapter 9, we prove a planar/quantum generalization (adding $\times$ again expresses the classical case) of the Holant indistinguishability theorem: $\mathcal{F}$ and $\mathcal{G}$ are Pl-Holant-indistinguishable (where Pl-Holant restricts to planar signature grids) if and only if there is a quantum orthogonal transformation between them. Finally, in Chapter 10, we prove two *near*-converses of Valiant's bipartite Holant theorem (Theorem 2.4.1), as the true converse is known to be false. Since $\text{Holant}(\mathcal{F}) \equiv \text{Holant}(\mathcal{F} \mid \subset)$, we again recover the converse of the orthogonal Holant theorem by adding $\subset$ to both signature sets. We leave as an open problem the formulation and proof of a combined planar *and* bipartite Holant indistinguishability theorem that would unite the two main threads of this thesis.

[1]This result holds over any field $\mathbb{K}$ with characteristic zero, so at its most general is not subsumed by results in later chapters, since these hold only over $\mathbb{R}$ or $\mathbb{C}$, or require that $\mathbb{K}$ be algebraically closed.

Chapter 3

#CSP and Isomorphism: Interpolation

This chapter is based on the second half of [You25a].

3.1 Introduction

Extending the notion of graph isomorphism, say two signatures (also called constraint functions) F and G of the same arity n are isomorphic if there is a bijection $\sigma : V(F) \to V(G)$ such that $F(x_1, \ldots, x_n) = G(\sigma(x_1), \ldots, \sigma(x_n))$ for all $x_1, \ldots, x_n \in V(F)$. Bijective sets $\mathcal{F}$ and $\mathcal{G}$ of constraint functions are isomorphic ($\mathcal{F} \cong \mathcal{G}$) if there is a common isomorphism between each $F \in \mathcal{F}$ and the corresponding $G \in \mathcal{G}$. Some similar concepts exist: Böhler et al. [Böh+02; Böh+04] study "constraint isomorphism" between Boolean #CSP *instances* (rather than constraint functions), involving permuting variables (rather than domain elements). Viewing an n-ary constraint function F as a tensor in ${}_n\mathcal{S}(\mathbb{K}^q)$, the notion of tensor isomorphism [GQ23] is a relaxation of constraint function isomorphism from a bijection $V(F) \to V(G)$ to possibly distinct invertible linear transformations $T \in \mathrm{GL}_q$ on each of the n dimensions: $F = (T_1 \otimes \ldots \otimes T_n) \circ G$.

The notion of graph homomorphism naturally extends to directed edge-weighted graphs X by using the weighted adjacency matrix A_X in (2.3.1). Lovász, Freedman, and Schrijver [FLS07; Lov06; LS09] studied the problem of counting homomorphisms to graphs with a real weight assigned to each edge and a nonnegative real weight assigned to each vertex. Lovász [Lov06] extended to these weighted graphs his result, proved forty years prior [Lov67], that two graphs X and Y are isomorphic if and only if they are homomorphism-indistinguishable over all graphs. Lovász's proof for real-weighted graphs [Lov06] used *graph algebras* of formal $\mathbb{C}$-linear combinations of k-*labeled graphs* (which we generalize to k-labeled #CSP instances in Definition 3.2.1). Still using k-labeled graphs, but applying invariant theory and algebraic geometry, Schrijver [Sch09] proved an analogous indistinguishability-to-isomorphism theorem for graphs with complex edge weights but without vertex weights. Using similar techniques, Regts [Reg13b] extended this result graphs with arbitrary vertex and edge weights, provided that no nonempty subset of vertex weights sums to zero. Finally, Cai and Govorov [CG21] extended the previous results to graphs with any vertex and edge weights from a field $\mathbb{K}$ of characteristic 0, and provide a counterexample to the existence of such a theorem for graphs with weights from any field of nonzero characteristic. Cai and Govorov overcame the algebraic approaches' technical difficulties of vertex weights summing to 0 by applying a simple, direct *Vandermonde interpolation* technique, dependent only on the fact that a Vandermonde matrix with distinct roots is nonsingular. It is remarkable that such a simple tool unifies all previous homomorphism indistinguishability theorems; in this chapter we further demonstrate its power by using it to extend Cai and Govorov's results to #CSP.

This chapter's results. Extending the notion of homomorphism-indistinguishablity, say bijective $\mathcal{F}$ and $\mathcal{G}$ are *#CSP-indistinguishable* if the partition function value of every #CSP($\mathcal{F}$) instance is preserved when we replace every constraint function in $\mathcal{F}$ with the corresponding function in $\mathcal{G}$. The main result of this chapter is the following theorem.

Theorem (Theorem 3.2.1, informal)**.** *For a field $\mathbb{K}$ of characteristic 0, sets $\mathcal{F}$ and $\mathcal{G}$ of $\mathbb{K}$-valued constraint functions are isomorphic if and only if $\mathcal{F}$ and $\mathcal{G}$ are #CSP-indistinguishable.*

Bulatov et al. [Bul+12, Theorem 4 and Remark 5] show that the problem of computing the partition function of a #CSP($\mathcal{F}$) instance reduces to computing the partition function of a (much larger) #CSP($\mathcal{F}'$) instance, where $\mathcal{F}'$ is a set of binary (arity-2) constraint functions. This reduction is isomorphism-preserving – that is, if $\mathcal{F} \cong \mathcal{G}$, then the corresponding binary $\mathcal{F}' \cong \mathcal{G}'$. Hence to prove Theorem 3.2.1 it suffices to extend Cai and Govorov's homomorphism indistinguishability theorem from a single binary constraint function (i.e. graph homomorphism) to a set of binary constraint functions. However, hypothetical proofs of an indistinguishability theorem for a set of binary constraint functions would extend quite naturally to the interpolation proof presented below for a set of arbitrary-arity constraint functions. Furthermore, the #CSP instances constructed by our proof are much smaller, simpler and more uniform than those produced by the reduction in [Bul+12], which is important because these #CSP instances serve as witnesses for constraint function nonisomorphism, and can be enumerated to make complexity dichotomies *effective* (see below). Thus we choose to use our interpolation technique to tackle the arbitrary-arity case directly.

In Chapter 5, we will prove the $\mathbb{K} = \mathbb{C}$ case of Theorem 3.2.1 using a different, less technical, approach. The Vandermonde interpolation-based proof in this chapter, in addition to supporting any field of characteristic zero, actually proves a more general result (Theorem 3.4.1) applying to k-labeled #CSP instances. It is also *constructive* (see Section 3.5): if $\mathcal{F}$ and $\mathcal{G}$ are finite and not isomorphic, then the proof provides a finite, explicit list of #CSP instances which must contain an instance on which $\mathcal{F}$ and $\mathcal{G}$ are distinguishable. Cai and Govorov use their constructive interpolation proof to make the graph homomorphism dichotomy of Cai, Chen, and Lu [CCL13] *effective*, meaning there is an algorithm that decides whether the problem is #P-hard (i.e. the dichotomy is decidable) and, if so, constructs a reduction from a #P-hard problem, rather than simply asserting such a reduction exists. Notably, the current complex-weighted #CSP dichotomy [CC17b] is not even known to be decidable; our constructive proof could similarly play a role in a decidable or effective #CSP dichotomy.

3.2 Preliminaries

3.2.1 Labeled #CSP instances

A k-labeled graph is a graph in which k distinct vertices are labeled by $[k]$ [FLS07; Lov06; CG21]. In the context of counting homomorphisms from graphs K to a fixed graph X, we study k-labeled 'right hand side' graphs K. Recall that, to model $\hom(K, X)$ as a #CSP instance I_K, the vertices of K become variables of I_K. Therefore the definitions in this subsection generalize k-labeled graphs and their associated definitions.

Definition 3.2.1 (k-labeled #CSP instance (product), $\mathcal{PLI}[\mathcal{F}; k]$, $\mathcal{PLI}^{\text{simp}}[\mathcal{F}; k]$)**.** A *$k$-labeled*

#CSP instance $\mathbf{K} = (V, C)$ is a #CSP instance in which k distinct variables are labeled by $[k]$. Define the *product* $\mathbf{K}_1 \mathbf{K}_2$ of two k-labeled #CSP($\mathcal{F}$) instances $\mathbf{K}_1 = (V_1, C_1), \mathbf{K}_2 = (V_2, C_2)$ as follows. For $i \in [k]$, let $u_i \in V_1, v_i \in V_2$ be the variables labeled i in V_1 and V_2, respectively. Define a new variable set V by starting with $V_1 \sqcup V_2$, then for each $i \in [k]$ merging u_i and v_i into a new variable w_i, and label w_i by i. Then define a new constraint multiset C by starting with $C_1 \sqcup C_2$ (multiset union), and for every $i \in [k]$ replacing every occurrence of u_i or v_i in each constraint with w_i. Then take $\mathbf{K}_1 \mathbf{K}_2 = (V, C)$.

Define $\mathcal{PLI}[\mathcal{F}; k]$ to be the set of k-labeled #CSP($\mathcal{F}$) instances. Let $U_k = (V, \varnothing) \in \mathcal{PLI}[\mathcal{F}; k]$, where V contains exactly k variables (all labeled). The k-labeled instance product is commutative and associative and has identity U_k, so $\mathcal{PLI}[\mathcal{F}; k]$ forms a commutative monoid under this product. Let $\mathcal{PLI}^{\text{simp}}[\mathcal{F}; k]$ denote the submonoid of $\mathcal{PLI}[\mathcal{F}; k]$ consisting of *simple* instances – those where the variables in any constraint $c \in C$ are distinct and the multiplicity of every constraint in C is 1 up to permutation of the order of its variables.

Definition 3.2.2 (Z^ψ)**.** For $\mathbf{K} = (V, C) \in \mathcal{PLI}[\mathcal{F}; k]$ and a map $\psi : [k] \to V(\mathcal{F})$ fixing, or *pinning*, the values of the labeled variables, define

$$Z^\psi(\mathbf{K}) = \sum_{\phi : V \to V(\mathcal{F}) \text{ extends } \psi} \prod_{(F, v_1, \ldots, v_{n_F}) \in C} F(\phi(v_1), \ldots, \phi(v_{n_F})),$$

where ϕ extends ψ means ϕ assigns value $\psi(i)$ to the variable labeled i. In this way, every $\mathbf{K} \in \mathcal{PLI}[\mathcal{F}; k]$ defines a k-ary signature $K \in \mathcal{S}(\mathbb{K}^q)$ defined by, for $\mathbf{x} \in [q]^k$,

$$K(\mathbf{x}) = Z^{[k] \mapsto \mathbf{x}}(\mathbf{K}).$$

Just as a #CSP($\mathcal{F}$) instance I correpsonds to a Holant($\mathcal{F} \cup \mathcal{EQ}$) grid $\Omega(I)$, a k-labeled #CSP($\mathcal{F}$) instance $\mathbf{K}$ corresponds to a k-ary Holant($\mathcal{F} \cup \mathcal{EQ}$) gadget $\mathbf{K}'$ with dangling edges incident to the vertices corresponding to the labeled variables of $\mathbf{K}$. Comparing Definition 2.2.3 and Definition 3.2.2, we find that $\mathbf{K}$ and $\mathbf{K}'$ have the same signature. However, not every Holant($\mathcal{F} \cup \mathcal{EQ}$) gadget has an equivalent k-labeled #CSP($\mathcal{F}$) instance, because a single vertex could have multiple dangling edges, and each variable is allowed at most one label. We will readdress this issue in Proposition 5.1.1.

The k-labeled instance product $\mathbf{K}_1 \mathbf{K}_2$ merges the labeled variables, and the unlabeled variables of $\mathbf{K}_1$ and $\mathbf{K}_2$ both still appear in constraints from $\mathbf{K}_1$ and $\mathbf{K}_2$. The unlabeled variables of $\mathbf{K}_1$ take values independently of the unlabeled variables of $\mathbf{K}_2$ (i.e. they appear in no constraints with each other). Hence the k-labeled instance product induces an entrywise product on the corresponding signatures:

$$Z^{\psi}(\mathbf{K}_1\,\mathbf{K}_2) = Z^{\psi}(\mathbf{K}_1)Z^{\psi}(\mathbf{K}_2). \tag{3.2.1}$$

Definition 3.2.3 ($(\cdot)_{\mathcal{F}\to\mathcal{G}}$)**.** For bijective constraint function sets $\mathcal{F}$ and $\mathcal{G}$ and any k-labeled #CSP($\mathcal{F}$) instance $\mathbf{K}$ (possibly with $k = 0$), define a k-labeled #CSP($\mathcal{G}$) instance $\mathbf{K}_{\mathcal{F}\to\mathcal{G}}$ by replacing each constraint $(F, v_1, \dots, v_{n_F})$ of $\mathbf{K}$ with $(G, v_1, \dots, v_{n_F})$, for $G \leftrightsquigarrow F$.

Definition 3.2.4 (#CSP-indistinguishable)**.** For pinning maps $\varphi : [k] \to V(\mathcal{F})$ and $\psi : [k] \to V(\mathcal{G})$, say $(\mathcal{F}, \varphi)$ and $(\mathcal{G}, \psi)$ are *#CSP-indistinguishable* if $Z^{\varphi}(K) = Z^{\psi}(K_{\mathcal{F}\to\mathcal{G}})$ for every $K \in \mathcal{PLI}[\mathcal{F}; k]$.

If $k = 0$, simply say that $\mathcal{F}$ and $\mathcal{G}$ are #CSP-indistinguishable.

Say $(\mathcal{F}, \varphi)$ and $(\mathcal{G}, \psi)$ are *simple-*#CSP*-indistinguishable* if $Z^{\varphi}(K) = Z^{\psi}(K_{\mathcal{F}\to\mathcal{G}})$ for every $K \in \mathcal{PLI}^{\text{simp}}[\mathcal{F}; k]$.

Definition 3.2.5 ($\cong$, Aut)**.** Let F, G be constraint functions of common arity n and $|V(\mathcal{F})| = |V(\mathcal{G})|$. A bijection $\sigma : V(F) \to V(G)$ is a *isomorphism* from F to G if $F(\mathbf{x}) = G(\sigma(\mathbf{x}))$ for all $\mathbf{x} \in V(F)^n$, where $\sigma(\mathbf{x}) = (\sigma(x_0), \dots, \sigma(x_{n-1}))$.

Say bijective constraint function sets $\mathcal{F}$ and $\mathcal{G}$ are isomorphic ($\mathcal{F} \cong \mathcal{G}$) if there is a single $\sigma : V(\mathcal{F}) \to V(\mathcal{G})$ which is an isomorphism between every pair $\mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}$.

For $\varphi : [k] \to V(\mathcal{F})$ and $\psi : [k] \to V(\mathcal{G})$, say $(\mathcal{F}, \varphi) \cong (\mathcal{G}, \psi)$ if there is an isomorphism $\sigma : V(\mathcal{F}) \to V(\mathcal{G})$ satisfying $\psi = \sigma \circ \varphi$.

Define $\text{Aut}(\mathcal{F})$ to be the group of all isomorphisms from $\mathcal{F}$ to itself.

We emphasize that every corresponding pair of functions in $\mathcal{F}$ and $\mathcal{G}$ must be isomorphic via the same σ. If $\mathcal{F} \cong \mathcal{G}$, then, since an isomorphism is just a relabeling of the domain elements, $\mathcal{F}$ and $\mathcal{G}$ are #CSP-indistinguishable. The main theorem of this chapter is the converse of this fact, a generalization of the fact that graphs $X \cong Y$ iff X and Y are homomorphism-indistinguishable:

Theorem 3.2.1. *Let $\mathbb{K}$ be a field of characteristic 0, and let $\mathcal{F}, \mathcal{G}$ be sets of $\mathbb{K}$-valued signatures. Then $\mathcal{F} \cong \mathcal{G}$ if and only if $\mathcal{F}$ and $\mathcal{G}$ are* #CSP*-indistinguishable.*

3.2.2 Domain weights

As usual, for sets A and B, A^B denotes the set of functions from B to A. For a set B with an implicit linear order and $a_b \in A$ for every $b \in B$, $(a_b)_{b\in B}$ denotes a tuple of elements of A indexed and ordered by B. Let $\mathbb{S}_q$ be the symmetric group of permutations on $[q]$.

Definition 3.2.6 (#CSP($\mathcal{F}, \alpha$), Z_α^ψ)**.** The problem #CSP($\mathcal{F}, \alpha$) is parameterized by a set $\mathcal{F}$ of constraint functions and a vector $\alpha \in (\mathbb{K} \setminus \{0\})^{V(\mathcal{F})}$ of *domain weights*. The partition function Z_α^ψ for α and $\psi : [k] \to V(\mathcal{F})$ is defined on k-labeled #CSP($\mathcal{F}$) instances $\mathbf{K} = (V, C)$ by

$$Z_\alpha^\psi(\mathbf{K}) = \sum_{\phi:V\to V(\mathcal{F}) \text{ extends } \psi} \frac{\alpha_\phi}{\alpha_\psi} \prod_{(F,v_1,\dots,v_{n_F})\in C} F(\phi(v_1), \dots, \phi(v_{n_F})),$$

where

$$\alpha_\phi = \prod_{v\in V} \alpha_{\phi(v)} \text{ and } \alpha_\psi = \prod_{i\in[k]} \alpha_{\psi(i)}.$$

In particular, $Z_{\mathbf{1}}^\psi = Z^\psi$, where $\mathbf{1}$ is the all-ones vector. One can model domain weights in the ordinary #CSP setting by applying a unary constraint function to each variable. However, explicit domain weights will prove useful for removing twin domain elements in Corollary 3.4.1, so we consider them separately.

Call $(\mathcal{F}, \alpha, \varphi)$ and $(\mathcal{G}, \beta, \psi)$ (simple)-#CSP-indistinguishable if $Z_\alpha^\varphi(\mathbf{K}) = Z_\beta^\psi(\mathbf{K}_{\mathcal{F}\to\mathcal{G}})$ for every $\mathbf{K} \in \mathcal{PLI}[\mathcal{F}; k]$ (resp. $\mathcal{PLI}^{\text{simp}}[\mathcal{F}; k]$), and say σ is an isomorphism of $(\mathcal{F}, \alpha, \varphi)$ and $(\mathcal{G}, \beta, \psi)$ if σ is an isomorphism of $(\mathcal{F}, \varphi)$ and $(\mathcal{G}, \psi)$ and $\alpha_i = \beta_{\sigma(i)}$ for all $i \in V(\mathcal{F})$.

For $\mathcal{F} = \{F_j \mid j \in T\}$ indexed by a possibly infinite set T, define the following set, which represents all 'configurations' of the remaining arguments of an application of a function in $\mathcal{F}$ when given a single distinguished argument at position r:

$$\mathcal{J}(\mathcal{F}) := \{(j, \mathbf{x}, r) \mid j \in T, \mathbf{x} \in V(\mathcal{F})^{n_j-1}, r \in [n_j]\}, \tag{3.2.2}$$

where n_j is the arity of $F_j \in \mathcal{F}$. If $n_j = 1$ (F_j is unary), then say $V(\mathcal{F})^{n_j-1} = V(\mathcal{F})^0 = \{()\}$ (the set containing the empty tuple). For a length-n tuple $\mathbf{x}$, index $r \in [n+1]$ and any y, define $\mathbf{x}^{r \hookleftarrow y}$ as the length-$n+1$ tuple $(x_0, \dots, x_{r-1}, y, x_r, \dots, x_{n-1})$ created by inserting y at position r. Domain elements $i, i' \in V(\mathcal{F})$ are *twins* if

$$F_j\left(\mathbf{x}^{r \hookleftarrow i}\right) = F_j\left(\mathbf{x}^{r \hookleftarrow i'}\right) \text{ for every } (j, \mathbf{x}, r) \in \mathcal{J}(\mathcal{F}).$$

If $n_j = 1$ and $\mathbf{x} = ()$, then $F_j\left(\mathbf{x}^{r \hookleftarrow i}\right) = F_j(i)$. If every $F \in \mathcal{F}$ is symmetric, meaning F is invariant under permutations of the order of its inputs, then say $\mathcal{F}$ is symmetric, and $i, i' \in V(\mathcal{F})$ are twins if $F_j(i, \mathbf{x}) = F_j(i', \mathbf{x})$ for every $j \in T$ and $\mathbf{x} \in V(\mathcal{F})^{n_j - 1}$, where we abbreviate $F_j(i, \mathbf{x}) = F_j\left(\mathbf{x}^{0 \hookleftarrow i}\right)$. $\mathcal{F}$ is *twin-free* if no two domain elements are twins. Equivalently, $\mathcal{F}$ is twin-free iff the tuples

$$\left(F_j\left(\mathbf{x}^{r \hookleftarrow i}\right)\right)_{(j, \mathbf{x}, r) \in \mathcal{J}(\mathcal{F})}$$

are pairwise distinct for $i \in V(\mathcal{F})$. If $\mathcal{F}$ is symmetric, then $\mathcal{F}$ is twin free if and only if the tuples $(F_j(i, \mathbf{x}))_{j \in T, \mathbf{x} \in V(\mathcal{F})^{n_j - 1}}$ are pairwise distinct for $i \in V(\mathcal{F})$.

3.2.3 Vandermonde interpolation

Next, we introduce the iterated Vandermonde interpolation technique of Cai and Govorov [CG21]. The following simple proposition follows from the fact that a square Vandermonde matrix with distinct rows is always nonsingular.

Proposition 3.2.1 ([CG21, Lemma 4.1])**.** *Let I be a finite index set, and a_i, $b_i \in \mathbb{K}$ for all $i \in I$. Let $I = \bigsqcup_{\ell \in [s]} I_\ell$ be the partition of I into equivalence classes defined by relation $\sim$, where $i \sim i'$ iff $b_i = b_{i'}$. If*

$$\sum_{i \in I} a_i b_i^j = 0$$

for every choice of $j \in [|I|]$, then $\sum_{i \in I_\ell} a_i = 0$ for every $\ell \in [s]$.

Iterated applications of Proposition 3.2.1 yield the following.

Lemma 3.2.1 ([CG21, Corollary 4.2])**.** *Let I and J be finite index sets, and a_i, $b_{i,j} \in \mathbb{K}$ for all $i \in I$, $j \in J$. Further, let $I = \bigsqcup_{\ell \in [s]} I_\ell$ be the partition of I into equivalence classes defined by relation $\sim$, where $i \sim i'$ iff $b_{i,j} = b_{i',j}$ for all $j \in J$. If*

$$\sum_{i \in I} a_i \prod_{j \in J} b_{i,j}^{p_j} = 0$$

for all choices of $(p_j)_{j \in J}$ where each $0 \leq p_j < |I|$, then $\sum_{i \in I_\ell} a_i = 0$ for every $\ell \in [s]$.

Often, I (and J) will be the set of all m-tuples whose entries range over a fixed finite set. In this case, we have the following corollary.

Corollary 3.2.1. *Let I and J be finite index sets and $m \geq 1$. Let $a_{\mathbf{i}} \in \mathbb{K}$ for $\mathbf{i} \in I^m$ and $b_{i,j} \in \mathbb{K}$ for $i \in I, j \in J$. Define $\sim$ as in Lemma 3.2.1. Let $I^m = \bigsqcup_{k \in [s_m]} I_k^{(m)}$ be a partition of I^m into equivalence classes defined by relation $\sim_m$, where $\mathbf{i} \sim_m \mathbf{i}'$ if $i_h \sim i'_h$ for all $h \in [m]$. If*

$$\sum_{\mathbf{i} \in I^m} a_{\mathbf{i}} \prod_{j \in J, h \in [m]} b_{i_h,j}^{p_{h,j}} = 0 \tag{3.2.3}$$

for every choice of $(p_{h,j})_{j \in J, h \in [m]}$, where each $0 \leq p_{h,j} < |I|$, then, for every $k \in [s_m]$,

$$\sum_{\mathbf{i} \in I_k^{(m)}} a_{\mathbf{i}} = 0.$$

Proof. Apply induction on m. The case $m = 1$ is exactly Lemma 3.2.1. Otherwise, assume Corollary 3.2.1 holds for $m - 1$ and any choice of $a_{(i_0,\ldots,i_{m-2})}$ for each $(i_0, \ldots, i_{m-2}) \in I^{m-1}$. Separating the sum in (3.2.3) over i_{m-1} gives

$$\sum_{i_{m-1} \in I} \left(\sum_{(i_0,\ldots,i_{m-2}) \in I^{m-1}} a_{\mathbf{i}} \prod_{j \in J, h \in [m-1]} b_{i_h,j}^{p_{h,j}} \right) \left(\prod_{j \in J} b_{i_{m-1},j}^{p_{m-1,j}} \right) = 0. \tag{3.2.4}$$

Now apply Lemma 3.2.1 with $i := i_{m-1}$ and

$$a_{i_{m-1}} := \sum_{(i_0,\ldots,i_{m-2}) \in I^{m-1}} a_{\mathbf{i}} \prod_{j \in J, h \in [m-1]} b_{i_h,j}^{p_{h,j}} \quad \text{and} \quad p_j := p_{m-1,j}.$$

For every $\ell \in [s]$ (with s and $(I_\ell)_{\ell \in [s]}$ from Lemma 3.2.1), we obtain

$$\begin{aligned} 0 = \sum_{i_{m-1} \in I_\ell} a_{i_{m-1}} &= \sum_{i_{m-1} \in I_\ell} \left(\sum_{(i_0,\ldots,i_{m-2}) \in I^{m-1}} a_{\mathbf{i}} \prod_{j \in J, h \in [m-1]} b_{i_h,j}^{p_{h,j}} \right) \\ &= \sum_{(i_0,\ldots,i_{m-2}) \in I^{m-1}} \left(\sum_{i_{m-1} \in I_\ell} a_{\mathbf{i}} \right) \prod_{j \in J, h \in [m-1]} b_{i_h,j}^{p_{h,j}}. \end{aligned}$$

Now, inductively applying Corollary 3.2.1 with, for each $(i_0, \ldots, i_{m-2}) \in I^{m-1}$,

$$a_{(i_0,\ldots,i_{m-2})} := \sum_{i_{m-1} \in I_\ell} a_{\mathbf{i}},$$

we conclude that, for every $k' \in [s_{m-1}]$,

$$\sum_{(i_0,\ldots,i_{m-2}) \in I_{k'}^{(m-1)}} \sum_{i_{m-1} \in I_\ell} a_{\mathbf{i}} = 0. \tag{3.2.5}$$

By definition, $(i_0, \ldots, i_{m-2}) \sim_{m-1} (i'_0, \ldots, i'_{m-2}) \wedge i_{m-1} \sim i'_{m-1} \iff \mathbf{i} \sim_m \mathbf{i}'$. Hence, for fixed k' and ℓ, the double sum in (3.2.5) is a sum over a single equivalence class $I_k^{(m)}$, and since (3.2.5) holds for every choice of $k' \in [s_{m-1}]$ and $\ell \in [s]$, we obtain the desired $\sum_{\mathbf{i} \in I_k^{(m)}} a_{\mathbf{i}} = 0$ for every $k \in [s_m]$. $\square$

3.2.4 The unary case

First, we must separately address the case where $\mathcal{F}$ and $\mathcal{G}$ contain only unary constraint functions, where Lemma 3.4.1 below does not apply. The proof uses simple versions of the arguments we will apply throughout this section.

Lemma 3.2.2. *Let $(\mathcal{F}, \alpha)$ and $(\mathcal{G}, \beta)$ be finite bijective domain-weighted sets of unary constraint functions with $|V(\mathcal{F})| \geq |V(\mathcal{G})|$. Assume $\mathcal{F}$ is twin-free. Let $k \geq 0$ and $\varphi : [k] \to V(\mathcal{F})$ and $\psi : [k] \to V(\mathcal{G})$. If $(\mathcal{F}, \alpha, \varphi)$ and $(\mathcal{G}, \beta, \psi)$ are simple-#*CSP*-indistinguishable, then $|V(\mathcal{F})| = |V(\mathcal{G})|$ and $(\mathcal{F}, \alpha, \varphi) \cong (\mathcal{G}, \beta, \psi)$.*

Proof. Write $\mathcal{F} = \{F_i \mid i \in [t]\}$ and $\mathcal{G} = \{G_i \mid i \in [t]\}$. For $\mathbf{p} \in [|V(\mathcal{F}) \sqcup V(\mathcal{G})|]^t$, let $\mathbf{K}^{\mathbf{p}} \in \mathcal{PLI}^{\text{simp}}[\mathcal{F}; k]$ be the instance with an unlabeled variable v, k unused labeled variables, and p_j copies of the constraint (F_j, v), for $j \in [t]$. Then $Z_\alpha^\varphi(\mathbf{K}^{\mathbf{p}}) = Z_\beta^\psi(\mathbf{K}^{\mathbf{p}}_{\mathcal{F}\to\mathcal{G}})$ is equivalent to

$$\sum_{i \in V(\mathcal{F})} \alpha_i \prod_{j \in [t]} F_j(i)^{p_j} + \sum_{i \in V(\mathcal{G})} (-\beta_i) \prod_{j \in [t]} G_j(i)^{p_j} = 0. \tag{3.2.6}$$

Considering (3.2.6) for every $\mathbf{p} \in [|V(\mathcal{F}) \sqcup V(\mathcal{G})|]^t$, we may apply Lemma 3.2.1 with $I := V(\mathcal{F}) \sqcup V(\mathcal{G})$, $J := [t]$, $a_i := \alpha_i$ or β_i, and $b_{i,j} := F_j(i)$ or $G_j(i)$ for $i \in V(\mathcal{F})$ or $i \in V(\mathcal{G})$, respectively. As $\mathcal{F}$ is twin-free, the tuples $(b_{i,j})_{j\in[t]} = (F_j(i))_{j\in[t]}$ are distinct for distinct $i \in V(\mathcal{F})$. Therefore no equivalence class I_ℓ from Lemma 3.2.1 contains more than one element of $V(\mathcal{F})$. However, every $\alpha_i \neq 0$ by definition, so no equivalence class consists of only a single element of $V(\mathcal{F})$. Thus there is a function $\sigma : V(\mathcal{F}) \to V(\mathcal{G})$ such that $i \sim \sigma(i)$ – or equivalently $\forall j \in [t] : F_j(i) = b_{i,j} = b_{\sigma(i),j} = G_j(\sigma(i))$ – for every $i \in V(\mathcal{F})$. Furthermore, since no two elements of $V(\mathcal{F})$ are in the same equivalence class, σ is injective, hence bijective, as $|V(\mathcal{F})| \geq |V(\mathcal{G})|$. Therefore $|V(\mathcal{F})| = |V(\mathcal{G})|$, and Lemma 3.2.1 concludes that $\alpha_i - \beta_{\sigma(i)} = 0$ for every $i \in V(\mathcal{F})$, so σ is an isomorphism of $(\mathcal{F}, \alpha)$ and $(\mathcal{G}, \beta)$.

Now, for every $c \in [k]$ and $F \in \mathcal{F}$, let $\mathbf{K}^{F,c} \in \mathcal{PLI}^{\text{simp}}[\mathcal{F}; k]$ be an instance with k labeled variables, no unlabeled variables, and a single constrant applying F to the variable labeled c. Then

$$F(\varphi(c)) = Z_\alpha^\varphi(\mathbf{K}^{F,c}) = Z_\beta^\psi(\mathbf{K}^{F,c}_{\mathcal{F}\to\mathcal{G}}) = G(\psi(c)) = F(\sigma^{-1}(\psi(c))),$$

where $G \leftrightsquigarrow F$. By twin-freeness of $\mathcal{F}$, we conclude that $\varphi(c) = \sigma^{-1}(\psi(c))$. Therefore $\psi = \sigma \circ \varphi$. $\square$

3.3 The symmetric ternary case

For clarity of exposition, we first prove the key lemma for the special case in which all constraint functions are symmetric and ternary.

Proposition 3.3.1. *Let $\mathcal{F} = \{F_j \mid j \in [t]\}$ and $\mathcal{G} = \{G_j \mid j \in [t]\}$ be finite bijective constraint function sets with $|V(\mathcal{F})| \geq |V(\mathcal{G})|$ such that every $F \in \mathcal{F}$ and $G \in \mathcal{G}$ are symmetric and have arity 3, and assume $\mathcal{F}$ is twin-free. Let α, β be the domain weights associated with $\mathcal{F}$ and $\mathcal{G}$, respectively.*

Let $\theta : [2k] \to V(\mathcal{F})$ and $\psi : [2k] \to V(\mathcal{G})$ for $k \geq 0$, and for every $x, y \in V(\mathcal{F})$, let

$$I_{xy} = \{a \in [k] \mid (\theta(a), \theta(a+k)) = (x, y)\} \subset [k]$$

Assume θ is well-balanced *– that is, for every $x, y \in V(\mathcal{F})$, $|I_{xy}| \geq 2|V(\mathcal{F})|^3$. If $Z_\alpha^\theta(\mathbf{K}) = Z_\beta^\psi(\mathbf{K}_{\mathcal{F}\to\mathcal{G}})$ for every $\mathbf{K} \in \mathcal{PLI}^{simp}[\mathcal{F}; 2k]$, then $|V(\mathcal{F})| = |V(\mathcal{G})|$ and $(\mathcal{F}, \alpha, \theta) \cong (\mathcal{G}, \beta, \psi)$.*

Proof. By the pigeonhole principle, since $|I_{xy}| \geq 2|V(\mathcal{F})|^3 \geq 2|V(\mathcal{F})| \cdot |V(\mathcal{G})|^2$ for every $x, y \in V(\mathcal{F})$, there exists a function $s : V(\mathcal{F})^2 \to V(\mathcal{G})^2$ such that for every $x, y \in V(\mathcal{F})$, the set

$$J_{xy} := I_{xy} \cap \{a \in [k] \mid (\psi(a), \psi(a+k)) = s(x, y)\}$$

satisfies $|J_{xy}| \geq 2|V(\mathcal{F})|$. Consider the variable set

$$V_1 = \{v, u_0, \ldots, u_{2k-1}\},$$

where each u_ℓ is labeled ℓ. For each choice of $\mathbf{p} = (p_{xyj})_{x,y \in V(\mathcal{F}), j \in [t]} \in [2|V(\mathcal{F})|]^{V(\mathcal{F})^2 \times [t]}$, construct a $2k$-labeled #CSP($\mathcal{F}$) instance $\mathbf{K}^\mathbf{p} \in \mathcal{PLI}^{\text{simp}}[\mathcal{F}; 2k]$ as follows. For each $x, y \in V(\mathcal{F})$ and $j \in [t]$, choose an arbitrary $P_{xyj} \subset J_{xy}$ with $|P_{xyj}| = p_{xyj}$, and define

$$C^\mathbf{p}(v) := \bigcup_{x,y \in V(\mathcal{F})} \bigcup_{j \in [t]} \bigcup_{a \in P_{xyj}} (F_j, v, u_a, u_{a+k}). \tag{3.3.1}$$

Then define $\mathbf{K}^\mathbf{p} := (V_1, C^\mathbf{p}(v))$. If $a \in P_{xyj} \subset J_{xy}$, then by definition $(\theta(a), \theta(a+k)) = (x, y)$ and $(\psi(a), \psi(a+k)) = s(x, y)$. Hence the variables (u_a, u_{a+k}) take values (x, y) and $s(x, y)$ under θ and ψ, respectively, independent of the choice of a in P_{xyj}. Therefore $Z_\alpha^\theta(\mathbf{K}^\mathbf{p}) = Z_\beta^\psi(\mathbf{K}^\mathbf{p}_{\mathcal{F}\to\mathcal{G}})$ is equivalent to

$$\sum_{i \in V(\mathcal{F})} \alpha_i \prod_{x,y \in V(\mathcal{F}), j \in [t]} F_j(i, x, y)^{p_{xyj}} + \sum_{i \in V(\mathcal{G})} (-\beta_i) \prod_{x,y \in V(\mathcal{G}), j \in [t]} G_j(i, s(x, y))^{p_{xyj}} = 0, \tag{3.3.2}$$

where we write $G_j(i, s(x,y))$ to mean $G_j(i, s(x,y)_0, s(x,y)_1)$. The sums over i correspond to the choice of assignment for the only free variable v in $\mathbf{K}^{\mathbf{p}}$ and $\mathbf{K}^{\mathbf{p}}_{\mathcal{F}\to\mathcal{G}}$, respectively. Overall, the LHS of (3.3.2) is a sum over $V(\mathcal{F})\sqcup V(\mathcal{G})$, with $|V(\mathcal{F})\sqcup V(\mathcal{G})| = |V(\mathcal{F})|+|V(\mathcal{G})| \le 2|V(\mathcal{F})|$. By constructing $\mathbf{K}^{\mathbf{p}}$ and considering (3.3.2) for every $\mathbf{p} \in [2|V(\mathcal{F})|]^{V(\mathcal{F})^2\times[t]}$, we may apply Lemma 3.2.1 with

$$I := V(\mathcal{F}) \sqcup V(\mathcal{G}),\ \ J := V(\mathcal{F})^2 \times [t],$$
$$a_i := \begin{cases} \alpha_i & i \in V(\mathcal{F}) \\ -\beta_i & i \in V(\mathcal{G}) \end{cases},\ \ b_{i,xyj} := \begin{cases} F_j(i,x,y) & i \in V(\mathcal{F}) \\ G_j(i, s(x,y)) & i \in V(\mathcal{G}) \end{cases}.$$

Since $\mathcal{F}$ is twin-free, the tuples $(F_j(i,x,y))_{x,y\in V(\mathcal{F}), j\in[t]}$ are pairwise distinct for $i \in V(\mathcal{F})$. Hence no equivalence class contains more than one element of $V(\mathcal{F})$. However, every $\alpha_i \neq 0$, so no equivalence class consists of only a single element of $V(\mathcal{F})$. Thus there is a function $\sigma : V(\mathcal{F}) \to V(\mathcal{G})$ such that $i \sim \sigma(i)$ for every $i \in V(\mathcal{F})$ – that is,

$$(F_j(i,x,y))_{x,y\in V(\mathcal{F}),j\in[t]} = (G_j(\sigma(i), s(x,y)))_{x,y\in V(\mathcal{F}),j\in[t]} \text{ for every } i \in V(\mathcal{F}). \tag{3.3.3}$$

Since no two elements of $V(\mathcal{F})$ are in the same equivalence class, σ is injective, hence bijective, as $|V(\mathcal{F})| \ge |V(\mathcal{G})|$. Therefore $|V(\mathcal{F})| = |V(\mathcal{G})|$, and Lemma 3.2.1 gives

$$\alpha_i = \beta_{\sigma(i)} \text{ for } i \in V(\mathcal{F}). \tag{3.3.4}$$

Next, we improve σ to an isomorphism between $(\mathcal{F}, \alpha)$ and $(\mathcal{G}, \beta)$. Define another family of #CSP($\mathcal{F}$) instances as follows. First, construct a $2k$-labeled variable set

$$V_2 := V_1 \cup \{v', v''\}$$

with two new variables v', v''. Fix $\mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}$. For $\mathbf{p}, \mathbf{p}', \mathbf{p}'' \in [|V(\mathcal{F})|]^{V(\mathcal{F})^2\times[t]}$, define (recall (3.3.1))

$$C^{F,\mathbf{p},\mathbf{p}',\mathbf{p}''} = \{(F, v, v', v'')\} \cup C^{\mathbf{p}}(v) \cup C^{\mathbf{p}'}(v') \cup C^{\mathbf{p}''}(v''),$$

a set of constraints on V_2, and define $\mathbf{K}^{F,\mathbf{p},\mathbf{p}',\mathbf{p}''} = (V_2, C^{F,\mathbf{p},\mathbf{p}',\mathbf{p}''}) \in \mathcal{PLI}^{\text{simp}}[\mathcal{F}; 2k]$. Now

$Z^{\theta}_{\alpha}(\mathbf{K}^{F,\mathbf{p},\mathbf{p}',\mathbf{p}''}) = Z^{\psi}_{\beta}(\mathbf{K}^{F,\mathbf{p},\mathbf{p}',\mathbf{p}''}_{\mathcal{F}\to\mathcal{G}})$ is equivalent to

$$
\begin{aligned}
0 = \sum_{i,i',i''\in V(\mathcal{F})} & \alpha_i\alpha_{i'}\alpha_{i''}F(i,i',i'') \\
& \cdot \prod_{x,y\in V(\mathcal{F}),j\in[t]} F_j(i,x,y)^{p_{xyj}}F_j(i',x,y)^{p'_{xyj}}F_j(i'',x,y)^{p''_{xyj}} \\
+ \sum_{i,i',i''\in V(\mathcal{G})} & -\beta_i\beta_{i'}\beta_{i''}G(i,i',i'') \\
& \cdot \prod_{x,y\in V(\mathcal{G}),j\in[t]} G_j(i,s(x,y))^{p_{xyj}}G_j(i',s(x,y))^{p'_{xyj}}G_j(i'',s(x,y))^{p''_{xyj}}
\end{aligned}
$$

(where $F \leftrightsquigarrow G \in \mathcal{G}$). Applying (3.3.4) and (3.3.3) then gives

$$
\begin{aligned}
0 = \sum_{i,i',i''\in V(\mathcal{F})} & \alpha_i\alpha_{i'}\alpha_{i''}\big(F(i,i',i'') - G(\sigma(i),\sigma(i'),\sigma(i''))\big) \\
& \cdot \prod_{x,y,j} F_j(i,x,y)^{p_{xyj}}F_j(i',x,y)^{p'_{xyj}}F_j(i'',x,y)^{p''_{xyj}}. \qquad (3.3.5)
\end{aligned}
$$

Constructing $\mathbf{K}^{F,\mathbf{p},\mathbf{p}',\mathbf{p}''}$ and considering (3.3.5) for all choices of $\mathbf{p},\mathbf{p}',\mathbf{p}'' \in [|V(\mathcal{F})|]^{V(\mathcal{F})^2\times[t]}$, we may apply Corollary 3.2.1 with

$$
\begin{aligned}
& m := 3,\ J = V(\mathcal{F})^2\times[t],\ a_{i,i',i''} := \alpha_i\alpha_{i'}\alpha_{i''}\big(F(i,i',i'') - G(\sigma(i),\sigma(i'),\sigma(i''))\big), \\
& b_{i_h,xyj} := F_j(i_h,x,y),\ p_{1,xyj} := p_{xyj},\ p_{2,xyj} := p'_{xyj},\ p_{3,xyj} := p''_{xyj}.
\end{aligned}
$$

Again, by twin-freeness, the tuples $(F_j(i,x,y))_{xyj}$ are pairwise distinct for $i \in V(\mathcal{F})$, so, for distinct (i,i',i''), the tuples

$$(F_j(i,x,y), F_j(i',x,y), F_j(i'',x,y))_{xyj}$$

are distinct. Now Corollary 3.2.1 gives $\alpha_i\alpha_{i'}\alpha_{i''}(F(i,i',i'') - G(\sigma(i),\sigma(i'),\sigma(i''))) = 0$ for all i,i',i''. Since each $\alpha_i \neq 0$ and our choice of F was arbitrary, this implies

$$F(i,i',i'') = G(\sigma(i),\sigma(i'),\sigma(i'')) \text{ for every } i,i',i'' \in V(\mathcal{F}) \text{ and every pair } F \leftrightsquigarrow G. \qquad (3.3.6)$$

Combined with (3.3.4), (3.3.6) implies that σ is a isomorphism between $(\mathcal{F},\alpha)$ and $(\mathcal{G},\beta)$.

To conclude that $(\mathcal{F},\alpha,\theta) \cong (\mathcal{G},\beta,\psi)$, it remains to show that $\psi = \sigma\circ\theta$. Again let $\mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}$. Define a third family of #CSP($\mathcal{F}$) instances. Define a $2k$-labeled variable set

$$V_3 := V_1 \cup \{v'\} = V_2 \setminus \{v''\}.$$

Fix $c \in [2k]$ and, for $\mathbf{p}, \mathbf{p}' \in [|V(\mathcal{F})|]^{V(\mathcal{F})^2 \times [t]}$, define the following set of constraints on V_3:

$$C^{F,\mathbf{p},\mathbf{p}'} := \{(F, u_c, v, v')\} \cup C^{\mathbf{p}}(v) \cup C^{\mathbf{p}'}(v')$$

and let $\mathbf{K}^{F,\mathbf{p},\mathbf{p}'} = (V_3, C^{F,\mathbf{p},\mathbf{p}'}) \in \mathcal{PLI}^{\text{simp}}[\mathcal{F}; 2k]$. Now $Z^{\theta}_{\alpha}(\mathbf{K}^{F,\mathbf{p},\mathbf{p}'}) = Z^{\psi}_{\beta}(\mathbf{K}^{F,\mathbf{p},\mathbf{p}'}_{\mathcal{F}\to\mathcal{G}})$ is equivalent to

$$\begin{aligned} 0 = &\sum_{i,i' \in V(\mathcal{F})} \alpha_i \alpha_{i'} F(\theta(c), i, i') \prod_{x,y,j} F_j(i, x, y)^{p_{xyj}} F_j(i', x, y)^{p'_{xyj}} \\ &+ \sum_{i,i' \in V(\mathcal{G})} -\beta_i \beta_{i'} G(\psi(c), i, i') \prod_{x,y,j} G_j(i, s(x, y))^{p_{xyj}} G_j(i', s(x, y))^{p'_{xyj}}. \end{aligned}$$

Applying (3.3.4) and (3.3.3) then gives

$$\sum_{i,i' \in V(\mathcal{F})} \alpha_i \alpha_{i'} \big(F(\theta(c), i, i') - G(\psi(c), \sigma(i), \sigma(i'))\big) \prod_{x,y,j} F_j(i, x, y)^{p_{xyj}} F_j(i', x, y)^{p'_{xyj}} = 0.$$

As above, the tuples $(F_j(i, x, y), F_j(i', x, y))_{xyj}$ are distinct for distinct (i, i'), so by a similar application of Corollary 3.2.1 with $m := 2$, we have $F(\theta(c), i, i') = G(\psi(c), \sigma(i), \sigma(i'))$ for all $i, i' \in V(\mathcal{F})$. This holds for any pair $F \leftrightsquigarrow G$, so, by (3.3.6),

$$F_j(\theta(c), i, i') = G_j(\psi(c), \sigma(i), \sigma(i')) = F_j(\sigma^{-1}(\psi(c)), i, i')$$

for all $i, i' \in V(\mathcal{F})$ and $j \in [t]$. Since $\mathcal{F}$ is twin-free, we have $\theta(c) = \sigma^{-1}(\psi(c))$, hence $\sigma(\theta(c)) = \psi(c)$. We chose $c \in [2k]$ arbitrarily, so $\psi = \sigma \circ \theta$. □

3.4 The general case

We now extend Proposition 3.3.1 to general finite sets of arbitrary arity, non-necessarily-symmetric constraint functions, containing at least one non-unary constraint function. The general proof requires more sophisticated indexing but is not fundamentally different from the proof of Proposition 3.3.1. Instead of constructing instances from one, three, or two tuples $\mathbf{p}$ indexed by $V(\mathcal{F})^2 \times [t]$ in the three steps of the proof Proposition 3.3.1, we use one, n, and $n-1$ tuples $\mathbf{p}$ indexed by $\mathcal{J}(\mathcal{F})$ (recall (3.2.2)), respectively, where n is the maximum arity among the functions in $\mathcal{F}$. This accounts for possible asymmetry and distinct arities of functions in $\mathcal{F}$.

Definition 3.4.1. Let $(\mathcal{F}, \alpha)$ be a domain-weighted constraint function set. Say $\theta : [k] \to V(\mathcal{F})$ is an *isomorphism pinning* for $(\mathcal{F}, \alpha)$ if, for any bijective domain-weighted constraint function set

$(\mathcal{G}, \beta)$ with $|V(\mathcal{G})| \leq |V(\mathcal{F})|$ and any $\psi : [k] \to V(\mathcal{G})$, if $(\mathcal{F}, \alpha, \theta)$ and $(\mathcal{G}, \beta, \psi)$ are simple-#CSP-indistinguishable, then $|V(\mathcal{F})| = |V(\mathcal{G})|$ and $(\mathcal{F}, \alpha, \theta) \cong (\mathcal{G}, \beta, \psi)$.

Lemma 3.4.1. *Let $\mathcal{F} = \{F_j \mid j \in [t]\}$ be twin-free, with domain weights α. Let $n = \max_{j\in[t]} n_j$ be the maximum arity among all functions in $\mathcal{F}$, and assume $n \geq 2$. Suppose $\theta : [(n-1)k] \to V(\mathcal{F})$, and for every $\mathbf{x} \in V(\mathcal{F})^{n-1}$, let*

$$I_{\mathbf{x}} = \{a \in [k] \mid (\theta(a + dk))_{d\in[n-1]} = \mathbf{x}\}.$$

If θ is well-balanced *– that is, $|I_{\mathbf{x}}| \geq 2|V(\mathcal{F})|^n$ for every $\mathbf{x} \in V(\mathcal{F})^{n-1}$ – then θ is an isomorphism pinning for $(\mathcal{F}, \alpha)$.*

Proof. Let $(\mathcal{G}, \beta)$ be bijective to $\mathcal{F}$ with $|V(\mathcal{G})| \leq |V(\mathcal{F})|$, and suppose $\psi : [(n-1)k] \to V(\mathcal{G})$ satisfies $Z^{\theta}_{\alpha}(\mathbf{K}) = Z^{\psi}_{\beta}(\mathbf{K}_{\mathcal{F}\to\mathcal{G}})$ for every $\mathbf{K} \in \mathcal{PLI}^{\text{simp}}[\mathcal{F}; (n-1)k]$. Since each $|I_{\mathbf{x}}| \geq 2|V(\mathcal{F})|^n \geq 2|V(\mathcal{F})| \cdot |V(\mathcal{G})|^{n-1}$, by the pigeonhole principle there exists a function $s : V(\mathcal{F})^{n-1} \to V(\mathcal{G})^{n-1}$ such that for every $\mathbf{x} \in V(\mathcal{F})^{n-1}$, the set

$$J_{\mathbf{x}} := I_{\mathbf{x}} \cap \{a \in [k] \mid (\psi(a + dk))_{d\in[n-1]} = s(\mathbf{x})\}$$

satisfies $|J_{\mathbf{x}}| \geq 2|V(\mathcal{F})|$. For every $\mathbf{x} \in V(\mathcal{F})^{n_j-1}$ with $n_j < n$, choose an arbitrary $\mathbf{x}' \in V(\mathcal{F})^{n-1}$ extending $\mathbf{x}$ (i.e. such that $x'_d = x_d$ for every $d \in [n_j - 1]$) and define $J_{\mathbf{x}} := J_{\mathbf{x}'}$.

Consider the $(n-1)k$-labeled variable set

$$V_1 = \{v\} \cup \{u^{(d)}_a\}_{a\in[k],d\in[n-1]},$$

where $u^{(d)}_a$ is labeled $a + dk$. For each choice of $\mathbf{p} = (p_{j\,\mathbf{x}\,r})_{(j,\mathbf{x},r)\in\mathcal{J}(\mathcal{F})} \in [2|V(\mathcal{F})|]^{\mathcal{J}(\mathcal{F})}$ (recall (3.2.2)), construct a set $C^{\mathbf{p}}(v)$ of constraints on V_1 as follows. Choose an arbitrary $P_{j\,\mathbf{x}\,r} \subset J_{\mathbf{x}}$ with $|P_{j\,\mathbf{x}\,r}| = p_{j\,\mathbf{x}\,r}$ for each $(j, \mathbf{x}, r) \in \mathcal{J}(\mathcal{F})$ and define

$$C^{\mathbf{p}}(v) = \bigcup_{(j,\mathbf{x},r)\in\mathcal{J}(\mathcal{F})} \bigcup_{a\in P_{j\,\mathbf{x}\,r}} \left(F_j, (u^{(d)}_a)^{r \hookleftarrow v}_{d\in[n_j-1]}\right). \tag{3.4.1}$$

Then define an $(n-1)k$-labeled #CSP($\mathcal{F}$) instance $\mathbf{K}^{\mathbf{p}} = (V_1, C^{\mathbf{p}}(v)) \in \mathcal{PLI}^{\text{simp}}[\mathcal{F}; (n-1)k]$. If $a \in P_{j\,\mathbf{x}\,r} \subset J_{\mathbf{x}}$ then $\theta(a + dk) = x_d$ and $\psi(a + dk) = s(\mathbf{x})_d$ for all $d \in [n_j - 1]$. Hence the variable $u^{(d)}_a$ takes value x_d and $s(\mathbf{x})_d$ under θ and ψ, respectively, for $d \in [n_j - 1]$. These values are independent of the choice of a within $P_{j\,\mathbf{x}\,r}$. Therefore $Z^{\theta}_{\alpha}(\mathbf{K}^{\mathbf{p}}) = Z^{\psi}_{\beta}(\mathbf{K}^{\mathbf{p}}_{\mathcal{F}\to\mathcal{G}})$ is equivalent to

$$\sum_{i\in V(\mathcal{F})} \alpha_i \prod_{(j,\mathbf{x},r)\in\mathcal{J}(\mathcal{F})} F_j\big(\mathbf{x}^{r \hookleftarrow i}\big)^{p_{j\,\mathbf{x}\,r}} + \sum_{i\in V(\mathcal{G})} (-\beta_i) \prod_{(j,\mathbf{x},r)\in\mathcal{J}(\mathcal{F})} G_j\big(s(\mathbf{x})^{r \hookleftarrow i}\big)^{p_{j\,\mathbf{x}\,r}} = 0. \tag{3.4.2}$$

The sums over i correspond to the choice of assignment for the free variable v in $\mathbf{K}^{\mathbf{p}}$ and $\mathbf{K}^{\mathbf{p}}_{\mathcal{F}\to\mathcal{G}}$, respectively. Overall, the LHS of (3.4.2) is a sum over $V(\mathcal{F}) \sqcup V(\mathcal{G})$, with $|V(\mathcal{F}) \sqcup V(\mathcal{G})| = |V(\mathcal{F})| + |V(\mathcal{G})| \leq 2|V(\mathcal{F})|$. Therefore, by constructing $\mathbf{K}^{\mathbf{p}}$ and considering (3.4.2) for every choice of $\mathbf{p} \in [2|V(\mathcal{F})|]^{\mathcal{J}(\mathcal{F})}$, we may apply Lemma 3.2.1 with

$$I := V(\mathcal{F}) \sqcup V(\mathcal{G}),\; J := \mathcal{J}(\mathcal{F}),\; a_i := \begin{cases} \alpha_i & i \in V(\mathcal{F}) \\ -\beta_i & i \in V(\mathcal{G}) \end{cases},$$

$$b_{i,j\,\mathbf{x}\,r} := \begin{cases} F_j\big(\mathbf{x}^{r \hookleftarrow i}\big) & i \in V(\mathcal{F}) \\ G_j\big(s(\mathbf{x})^{r \hookleftarrow i}\big) & i \in V(\mathcal{G}) \end{cases}.$$

Since $\mathcal{F}$ is twin-free, the tuples $\big(F_j\big(\mathbf{x}^{r \hookleftarrow i}\big)\big)_{(j,\mathbf{x},r)\in\mathcal{J}(\mathcal{F})}$ are pairwise distinct for $i \in V(\mathcal{F})$. Hence no equivalence class contains more than one element of $V(\mathcal{F})$. However, every $\alpha_i \neq 0$, so no equivalence class consists of only a single element of $V(\mathcal{F})$. Thus there is a function $\sigma : V(\mathcal{F}) \to V(\mathcal{G})$ such that $i \sim \sigma(i)$ for every $i \in V(\mathcal{F})$ – that is,

$$\big(F_j\big(\mathbf{x}^{r \hookleftarrow i}\big)\big)_{(j,\mathbf{x},r)\in\mathcal{J}(\mathcal{F})} = \Big(G_j\big(s(\mathbf{x})^{r \hookleftarrow \sigma(i)}\big)\Big)_{(j,\mathbf{x},r)\in\mathcal{J}(\mathcal{F})} \text{ for every } i \in V(\mathcal{F}). \tag{3.4.3}$$

Since no two elements of $V(\mathcal{F})$ are in the same equivalence class, σ is injective, hence bijective, as $|V(\mathcal{F})| \geq |V(\mathcal{G})|$. Therefore $|V(\mathcal{F})| = |V(\mathcal{G})|$, and Lemma 3.2.1 gives

$$\alpha_i = \beta_{\sigma(i)} \text{ for } i \in V(\mathcal{F}). \tag{3.4.4}$$

Next, define another family of #CSP($\mathcal{F}$) instances as follows. Fix $\mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}$ with common arity $n_F \leq n$, and let

$$V_2 := V_1 \setminus \{v\} \cup \{v_h \mid h \in [n_F]\}.$$

For $\mathbf{p}^{(0)}, \ldots, \mathbf{p}^{(n_F-1)} \in [|V(\mathcal{F})|]^{\mathcal{J}(\mathcal{F})}$, define the following set of constraints on V_2 (recall (3.4.1)):

$$C^{F,\mathbf{p}^{(0)},\ldots,\mathbf{p}^{(n_F-1)}} := \{(F, v_0, \ldots, v_{n_F-1})\} \cup \bigcup_{h\in[n_F]} C^{\mathbf{p}^{(h)}}(v_h). \tag{3.4.5}$$

Define the labeled instance $\mathbf{K}^{F,\mathbf{p}^{(0)},\ldots,\mathbf{p}^{(n_F-1)}} := (V_2, C^{F,\mathbf{p}^{(0)},\ldots,\mathbf{p}^{(n_F-1)}}) \in \mathcal{PLI}^{\text{simp}}[\mathcal{F}; (n-1)k]$. Then the assumption $Z_\alpha^\theta\big(\mathbf{K}^{F,\mathbf{p}^{(0)},\dots,\mathbf{p}^{(n_F-1)}}\big) = Z_\beta^\psi\big(\mathbf{K}_{\mathcal{F}\to\mathcal{G}}^{F,\mathbf{p}^{(0)},\dots,\mathbf{p}^{(n_F-1)}}\big)$ is equivalent to

$$\begin{aligned} 0 = &\sum_{\mathbf{i}\in V(\mathcal{F})^{n_F}} \left(\prod_{h\in[n_F]} \alpha_{i_h}\right) F(\mathbf{i}) \prod_{(j,\mathbf{x},r)\in\mathcal{J}(\mathcal{F}),h\in[n_F]} F_j\big(\mathbf{x}^{r\hookleftarrow i_h}\big)^{p^{(h)}_{j\,\mathbf{x}\,r}} \\ &+ \sum_{\mathbf{i}\in V(\mathcal{G})^{n_F}} \left(-\prod_{h\in[n_F]} \beta_{i_h}\right) G(\mathbf{i}) \prod_{(j,\mathbf{x},r)\in\mathcal{J}(\mathcal{F}),h\in[n_F]} G_j\big(s(\mathbf{x})^{r\hookleftarrow i_h}\big)^{p^{(h)}_{j\,\mathbf{x}\,r}}, \end{aligned}$$

where the sums over $\mathbf{i}$ corresponds to the choice of assignment for the free variables $v_0,\dots,v_{n_F-1}$. Applying (3.4.4) and (3.4.3) then gives

$$\sum_{\mathbf{i}\in V(\mathcal{F})^{n_F}} \left(\prod_{h\in[n_F]} \alpha_{i_h}\right) \big(F(\mathbf{i}) - G(\sigma(\mathbf{i}))\big) \prod_{(j,\mathbf{x},r)\in\mathcal{J}(\mathcal{F}),h\in[n_F]} F_j\big(\mathbf{x}^{r\hookleftarrow i_h}\big)^{p^{(h)}_{j\,\mathbf{x}\,r}} = 0. \tag{3.4.6}$$

Considering $\mathbf{K}^{F,\mathbf{p}^{(0)},\dots,\mathbf{p}^{(n_F-1)}}$ and (3.4.6) for every $\mathbf{p}^{(0)},\dots,\mathbf{p}^{(n_F-1)} \in [|V(\mathcal{F})|]^{\mathcal{J}(\mathcal{F})}$, we may apply Corollary 3.2.1 with

$$m := n_F,\ J = \mathcal{J}(\mathcal{F}),\ a_{\mathbf{i}} := \left(\prod_{h\in[n_F]} \alpha_{i_h}\right) \big(F(\mathbf{i}) - G(\sigma(\mathbf{i}))\big),$$
$$b_{i,j\,\mathbf{x}\,r} := F_j\big(\mathbf{x}^{r\hookleftarrow i}\big),\ p_{h,j\,\mathbf{x}\,r} := p^{(h)}_{j\,\mathbf{x}\,r}.$$

Again, by twin-freeness, the tuples $\big(F_j\big(\mathbf{x}^{r\hookleftarrow i}\big)\big)_{(j,\mathbf{x},r)\in\mathcal{J}(\mathcal{F})}$ are distinct for distinct $i \in V(\mathcal{F})$, so the larger tuples $\big(F_j\big(\mathbf{x}^{r\hookleftarrow i_h}\big)\big)_{h\in[n_F],(j,\mathbf{x},r)\in\mathcal{J}(\mathcal{F})}$ are distinct for distinct $\mathbf{i} \in V(\mathcal{F})^{n_F}$. Corollary 3.2.1 asserts that, for all $\mathbf{i} \in V(\mathcal{F})^{n_F}$, $\left(\prod_{h=1}^{n_F} \alpha_{i_h}\right)(F(\mathbf{i}) - G(\sigma(\mathbf{i}))) = 0$. Since each $\alpha_i \neq 0$ and our choice of F and G was arbitrary, this implies

$$F(\mathbf{i}) = G(\sigma(\mathbf{i})) \text{ for every } \mathbf{i} \in V(\mathcal{F})^{n_F} \text{ and every pair } \mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}\,. \tag{3.4.7}$$

Combined with (3.4.4), (3.4.7) implies that σ is a domain-weighted isomorphism of $(\mathcal{F},\alpha)$ and $(\mathcal{G},\beta)$.

It remains to show that $\psi = \sigma \circ \theta$. Again let $\mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}$, with common arity n_F. Fix $c \in [(n-1)k]$. If $n_F = 1$, let $\mathbf{K} \in \mathcal{PLI}^{\text{simp}}[\mathcal{F};(n-1)k]$ be an instance with no unlabeled/free variables and a single constraint (F, v_c), where v_c is the variable labeled c. Then

$$F(\theta(c)) = Z_\alpha^\theta(\mathbf{K}) = Z_\beta^\psi(\mathbf{K}_{\mathcal{F}\to\mathcal{G}}) = G(\psi(c)). \tag{3.4.8}$$

Otherwise, if $n_F \geq 2$, define a third family of #CSP($\mathcal{F}$) instances as follows. Define

$$V_3 := V_2 \setminus \{v_{n_F-1}\} = V_1 \cup \{v_h \mid h \in [n_F - 1]\}.$$

Write $c = a_c + d_c k$ (so that $u_{a_c}^{(d_c)}$ is labeled c). Fix $\rho \in [n_F]$. For $\mathbf{p}^{(0)}, \dots, \mathbf{p}^{(n_F-2)} \in [|V(\mathcal{F})|]^{\mathcal{J}(\mathcal{F})}$, define the following set of constraints on V_2:

$$C^{F,\mathbf{p}^{(0)},\dots,\mathbf{p}^{(n_F-2)}} := \{(F, v_0, \dots v_{\rho-1}, u_{a_c}^{(d_c)}, v_\rho, \dots, v_{n_F-2})\} \cup \bigcup_{h \in [n_F-1]} C^{\mathbf{p}^{(h)}}(v_h). \tag{3.4.9}$$

Define the instance $\mathbf{K}^{F,\mathbf{p}^{(0)},\dots,\mathbf{p}^{(n_F-2)}} := (V_3, C^{F,\mathbf{p}^{(0)},\dots,\mathbf{p}^{(n_F-2)}}) \in \mathcal{PLI}^{\text{simp}}[\mathcal{F}; (n-1)k]$. Now the assumption $Z_\alpha^\theta\big(\mathbf{K}^{F,\mathbf{p}^{(0)},\dots,\mathbf{p}^{(n_F-2)}}\big) = Z_\beta^\psi\big(\mathbf{K}_{\mathcal{F}\to\mathcal{G}}^{F,\mathbf{p}^{(0)},\dots,\mathbf{p}^{(n_F-2)}}\big)$ is equivalent to

$$\begin{aligned} 0 = &\sum_{\mathbf{i} \in V(\mathcal{F})^{n_F-1}} \left(\prod_{h \in [n_F-1]} \alpha_{i_h} \right) F\big(\mathbf{i}^{\rho \hookleftarrow \theta(c)}\big) \prod_{(j,\mathbf{x},r) \in \mathcal{J}(\mathcal{F}), h \in [n_F-1]} F_j\big(\mathbf{x}^{r \hookleftarrow i_h}\big)^{p_{j\,\mathbf{x}\,r}^{(h)}} \\ &+ \sum_{\mathbf{i} \in V(\mathcal{G})^{n_F-1}} \left(- \prod_{h \in [n_F-1]} \beta_{i_h} \right) G\big(\mathbf{i}^{\rho \hookleftarrow \psi(c)}\big) \prod_{(j,\mathbf{x},r) \in \mathcal{J}(\mathcal{F}), h \in [n_F-1]} G_j\big(s(\mathbf{x})^{r \hookleftarrow i_h}\big)^{p_{j\,\mathbf{x}\,r}^{(h)}}. \end{aligned}$$

Applying (3.4.4) and (3.4.3) gives

$$\begin{aligned} 0 = &\sum_{\mathbf{i} \in V(\mathcal{F})^{n_F-1}} \left(\prod_{h \in [n_F-1]} \alpha_{i_h} \right) \Big(F\big(\mathbf{i}^{\rho \hookleftarrow \theta(c)}\big) - G\big(\sigma(\mathbf{i})^{\rho \hookleftarrow \psi(c)}\big) \Big) \\ &\cdot \prod_{(j,\mathbf{x},r) \in \mathcal{J}(\mathcal{F}), h \in [n_F-1]} F_j\big(\mathbf{x}^{r \hookleftarrow i_h}\big)^{p_{j\,\mathbf{x}\,r}^{(h)}}. \end{aligned}$$

As above, $\big(F_j\big(\mathbf{x}^{r \hookleftarrow i_h}\big)\big)_{h \in [n_F-1], (j,\mathbf{x},r) \in \mathcal{J}(\mathcal{F})}$ are distinct for distinct $\mathbf{i} \in V(\mathcal{F})^{n_F-1}$. Hence by a similar application of Corollary 3.2.1 with $m := n_F - 1$, we have $F\big(\mathbf{i}^{\rho \hookleftarrow \theta(c)}\big) = G\big(\sigma(\mathbf{i})^{\rho \hookleftarrow \psi(c)}\big)$ for all $\mathbf{i} \in V(\mathcal{F})^{n_F-1}$. This holds for any pair $F \leftrightsquigarrow G$ (with unary F and G handled by (3.4.8)) and any $\rho \in [n_F]$. Hence, for all $(j, \mathbf{i}, \rho) \in \mathcal{J}(\mathcal{F})$,

$$F_j\big(\mathbf{i}^{\rho \leftarrow \theta(c)}\big) = G_j\big(\sigma(\mathbf{i})^{\rho \leftarrow \psi(c)}\big) = F_j\big(\mathbf{i}^{\rho \leftarrow \sigma^{-1}(\psi(c))}\big),$$

where the second equality is (3.4.7). Since $\mathcal{F}$ is twin-free, we have $\theta(c) = \sigma^{-1}(\psi(c))$, hence $\sigma(\theta(c)) = \psi(c)$. We chose $c \in [(n-1)k]$ arbitrarily, so $\psi = \sigma \circ \theta$. □

Now we remove the requirement that θ be well-balanced, which removes the requirement that k be large. The proof of the next theorem follows and generalizes the proof of [CG21, Theorem 3.1].

Lemma 3.4.2. *Let $(\mathcal{F}, \alpha)$ and $(\mathcal{G}, \beta)$ be finite bijective domain-weighted constraint function sets with $|V(\mathcal{F})| \geq |V(\mathcal{G})|$ and $\mathcal{F}$ twin-free. Let $k \geq 0$ and $\varphi : [k] \to V(\mathcal{F})$ and $\psi : [k] \to V(\mathcal{G})$. If $(\mathcal{F}, \alpha, \varphi)$ and $(\mathcal{G}, \beta, \psi)$ are simple-#*CSP*-indistinguishable and there exist $\ell \geq k$ and an extension $\Theta : [\ell] \to V(\mathcal{F})$ of φ that is an isomorphism pinning of $(\mathcal{F}, \alpha)$, then $|V(\mathcal{F})| = |V(\mathcal{G})|$ and $(\mathcal{F}, \alpha, \varphi) \cong (\mathcal{G}, \beta, \psi)$.*

Proof. First, if $\ell = k$, then φ itself is an isomorphism pinning, and we are done. Otherwise, assume $\ell > k$. Let $E(\varphi) = \{\Phi : [\ell] \to V(\mathcal{F}) \mid \Phi \text{ extends } \varphi\}$ and $E(\psi) = \{\Psi : [\ell] \to V(\mathcal{G}) \mid \Psi \text{ extends } \psi\}$ be the sets of all extensions of φ and ψ to $[\ell]$, respectively. Define

$$A := \{\Phi \in E(\varphi) \mid \exists \sigma \in \mathrm{Aut}(\mathcal{F}, \alpha) \text{ s.t. } \Phi = \sigma \circ \Theta\} \text{ and}$$

$$B := \{\Psi \in E(\psi) \mid \exists \text{ isomorphism } \sigma \text{ between } (\mathcal{F}, \alpha) \text{ and } (\mathcal{G}, \beta) \text{ s.t. } \Psi = \sigma \circ \Theta\}.$$

We will show $B \neq \varnothing$. As Θ is an isomorphism pinning of $(\mathcal{F}, \alpha)$, if $\Phi \in E(\varphi) \setminus A$, then $(\mathcal{F}, \alpha, \Theta)$ and $(\mathcal{F}, \alpha, \Phi)$ are not simple-#CSP-indistinguishable, so there is a $\mathbf{K}^\Phi \in \mathcal{PLI}^{\mathrm{simp}}[\mathcal{F}; \ell]$ such that

$$Z_\alpha^\Theta(\mathbf{K}^\Phi) \neq Z_\alpha^\Phi(\mathbf{K}^\Phi). \tag{3.4.10}$$

Similarly, if $\Psi \in E(\psi) \setminus B$, then the triples $(\mathcal{F}, \alpha, \Theta)$ and $(\mathcal{G}, \beta, \Psi)$ are not simple-#CSP-indistinguishable, so there is a $\mathbf{K}^\Psi \in \mathcal{PLI}^{\mathrm{simp}}[\mathcal{F}; \ell]$ such that

$$Z_\alpha^\Theta(\mathbf{K}^\Psi) \neq Z_\beta^\Psi(\mathbf{K}^\Psi_{\mathcal{F}\to\mathcal{G}}). \tag{3.4.11}$$

Define the set

$$J := (E(\varphi) \setminus A) \sqcup (E(\psi) \setminus B)$$

to index these nonisomorphism witnesses, and for every choice of $\mathbf{q} := (q_\Lambda)_{\Lambda \in J} \in [|E(\varphi) \sqcup E(\psi)|]^J = [|V(\mathcal{F})|^{\ell-k} + |V(\mathcal{G})|^{\ell-k}]^J$, define the ℓ-labeled instance product

$$\mathbf{K}^{\mathbf{q}} := \prod_{\Lambda \in J} (\mathbf{K}^\Lambda)^{q_\Lambda} \in \mathcal{PLI}^{\mathrm{simp}}[\mathcal{F}; \ell].$$

Then, by multiplicativity (3.2.1),

$$Z_\alpha^\Phi(\mathbf{K}^{\mathbf{q}}) = \prod_{\Lambda \in J} \left(Z_\alpha^\Phi(\mathbf{K}^\Lambda)\right)^{q_\Lambda} \tag{3.4.12}$$

for any $\Phi \in E(\varphi)$. For any $\mathbf{K} \in \mathcal{PLI}^{\mathrm{simp}}[\mathcal{F}; \ell]$, define $\pi_k(\mathbf{K}) \in \mathcal{PLI}^{\mathrm{simp}}[\mathcal{F}; k]$ by removing the labels (but not the underlying variables) in $[\ell] \setminus [k]$ from $\mathbf{K}'$. Then, for any $\mathbf{K} \in \mathcal{PLI}^{\mathrm{simp}}[\mathcal{F}; \ell]$,

$$\begin{aligned} &\sum_{\Phi \in E(\varphi)} \left(\prod_{m \in [\ell]\setminus[k]} \alpha_{\Phi(m)}\right) Z_\alpha^\Phi(\mathbf{K}) = Z_\alpha^\varphi(\pi_k(\mathbf{K})) \\ &= Z_\beta^\psi(\pi_k(\mathbf{K})_{\mathcal{F}\to\mathcal{G}}) = \sum_{\Psi \in E(\psi)} \left(\prod_{m \in [\ell]\setminus[k]} \beta_{\Psi(m)}\right) Z_\beta^\Psi(\mathbf{K}_{\mathcal{F}\to\mathcal{G}}), \end{aligned} \tag{3.4.13}$$

where the second equality holds because the triples $(\mathcal{F},\alpha,\varphi)$ and $(\mathcal{G},\beta,\psi)$ are simple-#CSP-indistinguishable. Substituting $\mathbf{K} := \mathbf{K}^{\mathbf{q}}$ in (3.4.13) and applying (3.4.12), we obtain

$$0 = \sum_{\Phi\in E(\varphi)} \left(\prod_{m\in[\ell]\setminus[k]} \alpha_{\Phi(m)}\right) \prod_{\Lambda\in J} \left(Z_\alpha^\Phi(\mathbf{K}^\Lambda)\right)^{q_\Lambda} + \sum_{\Psi\in E(\psi)} \left(-\prod_{m\in[\ell]\setminus[k]} \beta_{\Psi(m)}\right) \prod_{\Lambda\in J} \left(Z_\beta^\Psi(\mathbf{K}^\Lambda_{\mathcal{F}\to\mathcal{G}})\right)^{q_\Lambda}. \tag{3.4.14}$$

Considering (3.4.14) for every choice of $\mathbf{q} \in [|E(\varphi)\sqcup E(\psi)|]^J$, we may apply Lemma 3.2.1 with

$$I := E(\varphi)\sqcup E(\psi),\ a_i := \begin{cases} \prod_{m\in[\ell]\setminus[k]} \alpha_{\Phi(m)} & i=\Phi\in E(\varphi) \\ -\prod_{m\in[\ell]\setminus[k]} \beta_{\Psi(m)} & i=\Psi\in E(\psi) \end{cases}$$
$$b_{i,\Lambda} := \begin{cases} Z_\alpha^\Phi(\mathbf{K}^\Lambda) & i=\Phi\in E(\varphi) \\ Z_\beta^\Psi(\mathbf{K}^\Lambda_{\mathcal{F}\to\mathcal{G}}) & i=\Psi\in E(\psi) \end{cases}.$$

For every $\mathbf{K}^\Lambda$ (indeed, for every $\mathbf{K}\in\mathcal{PLI}^{\text{simp}}[\mathcal{F};\ell]$), any $\Phi\in A$ and $\Psi\in B$ satisfy

$$b_{\Phi,\Lambda} = Z_\alpha^\Phi(\mathbf{K}^\Lambda) = Z_\alpha^\Theta(\mathbf{K}^\Lambda) = Z_\beta^\Psi(\mathbf{K}_{\mathcal{F}\to\mathcal{G}}) = b_{\Psi,\Lambda}. \tag{3.4.15}$$

Thus, when we apply Lemma 3.2.1, there is a single equivalence class containing $A\sqcup B\subset I$. Furthermore, if $\Phi'\in E(\varphi)\setminus A$ and $\Phi\in A$, then, using (3.4.10) and (3.4.15) (with $\Lambda := \Phi'$),

$$b_{\Phi',\Phi'} = Z_\alpha^{\Phi'}(\mathbf{K}^{\Phi'}) \neq Z_\alpha^\Theta(\mathbf{K}^{\Phi'}) = b_{\Phi,\Phi'}.$$

Similarly, if $\Psi'\in E(\psi)\setminus B$ and $\Psi\in B$, then, using (3.4.11) and (3.4.15),

$$b_{\Psi',\Psi'} = Z_\beta^{\Psi'}(\mathbf{K}^{\Psi'}_{\mathcal{F}\to\mathcal{G}}) \neq Z_\alpha^\Theta(\mathbf{K}^{\Psi'}) = b_{\Psi,\Psi'}$$

Therefore $A\sqcup B$ constitutes an *entire* equivalence class. So, by Lemma 3.2.1,

$$\sum_{\Phi\in A}\left(\prod_{m\in[\ell]\setminus[k]} \alpha_{\Phi(m)}\right) - \sum_{\Psi\in B}\left(\prod_{m\in[\ell]\setminus[k]} \beta_{\Psi(m)}\right) = 0. \tag{3.4.16}$$

For any $\Phi\in A$, there is a $\sigma\in\operatorname{Aut}(\mathcal{F},\alpha)$ satisfying $\Phi=\sigma\circ\Theta$. Then $\alpha_{\Phi(m)} = \alpha_{\sigma(\Theta(m))} = \alpha_{\Theta(m)}$, a value independent of the choice of Φ. So (3.4.16) becomes

$$|A|_{\mathbb{K}} \prod_{m\in[\ell]\setminus[k]} \alpha_{\Theta(m)} = \sum_{\Psi\in B}\left(\prod_{m\in[\ell]\setminus[k]} \beta_{\Psi(m)}\right), \tag{3.4.17}$$

where $|A|_{\mathbb{K}}$ is a sum of $|A|$ copies of the multiplicative identity $1_{\mathbb{K}} \in \mathbb{K}$. Each $\alpha_{\Theta(m)} \neq 0$ and, since $\mathbb{K}$ has nonzero characteristic and $A \ni \Theta$ (so $|A| > 1$) , $|A|_{\mathbb{K}} \neq 0$. Hence (3.4.17) is nonzero, so B is nonempty. Any $\Psi \in B$ extends ψ and gives an isomorphism σ between $(\mathcal{F}, \alpha)$ and $(\mathcal{G}, \beta)$ such that $\Psi = \sigma \circ \Theta$. Since Θ itself extends φ, restricting $\Psi = \sigma \circ \Theta$ to $[k]$ gives $\psi = \sigma \circ \varphi$, as desired. □

Now Lemmas 3.4.1 and 3.4.2 (and Lemma 3.2.2) combine to prove the following theorem for finite $\mathcal{F}$ and $\mathcal{G}$. Then we bootstrap the finite case to $\mathcal{F}$ and $\mathcal{G}$ of arbitrary cardinality. The result generalizes [Lov06, Lemma 2.4] from real-weighted graph homomorphism to $\mathbb{K}$-valued #CSP (see Remark 5.1.1 below).

Theorem 3.4.1. *Let $(\mathcal{F}, \alpha)$ and $(\mathcal{G}, \beta)$ be bijective domain-weighted constraint function sets with $|V(\mathcal{F})| \geq |V(\mathcal{G})|$, and $\mathcal{F}$ twin-free. Let $k \geq 0$ and $\varphi : [k] \to V(\mathcal{F})$ and $\psi : [k] \to V(\mathcal{G})$. If $(\mathcal{F}, \alpha, \varphi)$ and $(\mathcal{G}, \beta, \psi)$ are simple-*#CSP*-indistinguishable, then $|V(\mathcal{F})| = |V(\mathcal{G})|$ and $(\mathcal{F}, \alpha, \varphi) \cong (\mathcal{G}, \beta, \psi)$.*

Proof. First suppose $\mathcal{F}$ and $\mathcal{G}$ are finite. If all constraint functions in $\mathcal{F}$ and $\mathcal{G}$ are unary, then apply Lemma 3.2.2. Otherwise, by Lemma 3.4.2, it suffices to find an extension $\Theta : [\ell] \to V(\mathcal{F})$ of φ that is an isomorphism pinning of $(\mathcal{F}, \alpha)$. By choosing sufficiently large ℓ (no larger than $k + 2(n-1)|V(\mathcal{F})|^{2n-1}$), we can always extend φ to a well-balanced Θ, and, by Lemma 3.4.1, such a Θ is an isomorphism pinning of $(\mathcal{F}, \alpha)$. This gives the desired result for finite $\mathcal{F}$ and $\mathcal{G}$.

Next, consider $\mathcal{F}$ and $\mathcal{G}$ with arbitrary cardinality. Define the finite set

$$\Sigma = \{\sigma : V(\mathcal{F}) \to V(\mathcal{G}) \mid \sigma \text{ is a bijection and } \psi = \sigma \circ \varphi \text{ and } \alpha_i = \beta_{\sigma(i)} \text{ for all } i \in V(\mathcal{F})\}.$$

For any finite twin-free $\mathcal{F}' \subset \mathcal{F}$ and corresponding $\mathcal{G}' \subset \mathcal{G}$, $(\mathcal{F}', \alpha, \varphi)$ and $(\mathcal{G}', \beta, \psi)$ inherit the simple-#CSP-indistinguishablity of $(\mathcal{F}, \alpha, \varphi)$ and $(\mathcal{G}, \beta, \psi)$, so, by the finite case above, there is a $\sigma \in \Sigma$ that is an isomorphism of $(\mathcal{F}', \alpha, \varphi)$ and $(\mathcal{G}', \beta, \psi)$. In particular, Σ is nonempty. Suppose towards contradiction that there is no $\sigma \in \Sigma$ that is an isomorphim of the full sets $\mathcal{F}$ and $\mathcal{G}$. Then, for every $\sigma \in \Sigma$, there is a pair $\mathcal{F} \ni F_\sigma \leftrightsquigarrow G_\sigma \in \mathcal{G}$ such that σ is not an isomorphism of F_σ and G_σ. Define the finite, nonempty sets

$$\mathcal{F}' = \{F_\sigma \mid \sigma \in \Sigma\} \text{ and } \mathcal{G}' = \{G_\sigma \mid \sigma \in \Sigma\}.$$

If $\mathcal{F}'$ is not twin-free, then, for each $i, i' \in V(\mathcal{F})$ that are twins for $\mathcal{F}'$, find an $F_{i,i'} \in \mathcal{F}$ such that there exist $\mathbf{x}$ and r for which $F_{i,i'}(\mathbf{x}^{r \hookleftarrow i}) \neq F_{i,i'}(\mathbf{x}^{r \hookleftarrow i'})$ (such an $F_{i,i'}$ must exist because $\mathcal{F}$ is twin-free), and define

$$\mathcal{F}'' = \mathcal{F}' \cup \{F_{i,i'} \mid i, i' \in V(\mathcal{F}) \text{ are twins for } \mathcal{F}'\}.$$

Define $\mathcal{G}''$ to be the subset of $\mathcal{G}$ corresponding to $\mathcal{F}''$. As $\mathcal{F}''$ and $\mathcal{G}''$ are finite and $\mathcal{F}''$ is twin-free, by the finite case above, there is a $\sigma_0 \in \Sigma$ that is an isomorphism of $\mathcal{F}''$ and $\mathcal{G}''$. In particular, since $\mathcal{F}'' \ni F_{\sigma_0} \leftrightsquigarrow G_{\sigma_0} \in \mathcal{G}''$, σ_0 is an isomorphism of F_{σ_0} and G_{σ_0}, a contradiction. □

Next, we introduce domain weights to unweighted constraint function sets to remove the twin-free requirement. The proof of the next corollary generalizes the proof of [CG21, Corollary 6.2].

Corollary 3.4.1. *Let $\mathcal{F}$ and $\mathcal{G}$ be bijective constraint function sets, and let $k \geq 0$, $\varphi : [k] \to V(\mathcal{F})$, and $\psi : [k] \to V(\mathcal{G})$. If $(\mathcal{F}, \varphi)$ and $(\mathcal{G}, \psi)$ are simple-#*CSP*-indistinguishable, then $|V(\mathcal{F})| = |V(\mathcal{G})|$ and there is an isomorphism σ from $\mathcal{F}$ to $\mathcal{G}$ such that $\psi' = \sigma \circ \varphi$, where $\psi'(i)$ is a twin of $\psi(i)$ for every $i \in [k]$.*

Proof. Let $I_0, \dots, I_{s-1}$ be the partition of $V(\mathcal{F})$ into equivalence classes under the twin relation. Define the twin-contracted constraint function set $\widetilde{\mathcal{F}}$ with domain $V(\widetilde{\mathcal{F}}) = [s]$ by replacing each $F \in \mathcal{F}$ with $\widetilde{F}$ defined by $\widetilde{F}(\mathbf{y}) := F(\mathbf{x})$ for arbitrary choices of $x_i \in I_{y_i}$ for $i \in [n_F]$. Introduce domain weights α defined by $\alpha_\ell = |I_\ell|$ for $\ell \in [s]$. Define $\widetilde{\varphi} : [k] \to V(\widetilde{\mathcal{F}})$ by setting $\widetilde{\varphi}(m) = \ell$ if $\varphi(m) \in I_\ell$. Now $\widetilde{\mathcal{F}}$ is twin-free and $(\mathcal{F}, \varphi)$ and $(\widetilde{\mathcal{F}}, \alpha, \widetilde{\varphi})$ are #CSP-indistinguishable. Similarly partition $V(\mathcal{G})$ into equivalence classes $J_0, \dots, J_{s'-1}$ and define a twin-contracted $(\widetilde{\mathcal{G}}, \beta, \widetilde{\psi})$. Then, since $(\mathcal{F}, \varphi)$ and $(\mathcal{G}, \psi)$ are simple-#CSP-indistinguishable by assumption, $(\widetilde{\mathcal{F}}, \alpha, \widetilde{\varphi})$ and $(\widetilde{\mathcal{G}}, \beta, \widetilde{\psi})$ are simple-#CSP-indistinguishable. As $\widetilde{\mathcal{F}}$ is now twin-free, we may apply Theorem 3.4.1 to conclude that $s = s'$ and there is an isomorphism ξ from $(\widetilde{\mathcal{F}}, \alpha)$ to $(\widetilde{\mathcal{G}}, \beta)$ such that $\widetilde{\psi} = \xi \circ \widetilde{\varphi}$.

Define $\sigma : V(\mathcal{F}) \to V(\mathcal{G})$ as follows. For each $\ell \in [s] = V(\widetilde{\mathcal{F}})$, ξ associates $I_\ell \subset V(\mathcal{F})$ with $J_{\xi(\ell)} \subset V(\mathcal{G})$, and $|I_\ell| = \alpha_\ell = \beta_{\xi(\ell)} = |J_{\xi(\ell)}|$. Choose an arbitrary bijection between I_ℓ and $J_{\xi(\ell)}$ and have σ map I_ℓ into $J_{\xi(\ell)}$ according to this bijection. This σ satisfies the desired properties. □

Finally, the $k = 0$ case of Corollary 3.4.1 is equivalent to Theorem 3.2.1, but slightly stronger because it only assumes simple indistinguishablity:

Corollary 3.4.2. *Let $\mathcal{F}$ and $\mathcal{G}$ be bijective constraint function sets. Then $\mathcal{F} \cong \mathcal{G}$ if and only if $\mathcal{F}$ and $\mathcal{G}$ are simple-#*CSP*-indistinguishable.*

3.5 Constructiveness of the interpolation proof

For bijective domain-weighted constraint function sets $(\mathcal{F}, \alpha)$ and $(\mathcal{G}, \beta)$ with $\mathcal{F}$ twin-free and $|V(\mathcal{F})| \geq |V(\mathcal{G})|$, and $\varphi : [k] \to V(\mathcal{F})$, $\psi : [k] \to V(\mathcal{G})$, Theorem 3.4.1 asserts that if $(\mathcal{F}, \alpha, \varphi) \ncong (\mathcal{G}, \beta, \psi)$, then there is some witness instance $\mathbf{K} \in \mathcal{PLI}^{\text{simp}}[\mathcal{F}; k]$ such that $Z_\alpha^\varphi(\mathbf{K}) \neq Z_\alpha^\psi(\mathbf{K}_{\mathcal{F}\to\mathcal{G}})$. If $\mathcal{F}$ and $\mathcal{G}$ are finite, then the proofs of Lemma 3.4.1 and Lemma 3.4.2 construct an explicit finite list of instances in $\mathcal{PLI}^{\text{simp}}[\mathcal{F}; k]$ guaranteed to contain such a witness as follows (if $\mathcal{F}$ and $\mathcal{G}$ are infinite, then we cannot expect to produce such a list in bounded time). The proof of Theorem 3.4.1 first extends φ to an isomorphism pinning Θ with domain $[\ell]$, where $\ell \leq k + 2(n-1)|V(\mathcal{F})|^{2n-1}$ (n is the maximum arity among functions in $\mathcal{F}$). Then the proof of Lemma 3.4.2 shows that

$$Z_\alpha^\varphi(\pi_k(\mathbf{K}^{\mathbf{q}})) = Z_\beta^\psi(\pi_k(\mathbf{K}^{\mathbf{q}})_{\mathcal{F}\to\mathcal{G}}) \text{ for every } \mathbf{q} \in [|V(\mathcal{F})|^{\ell-k} + |V(\mathcal{G})|^{\ell-k}]^J$$

if and only if $(\mathcal{F}, \alpha, \varphi) \cong (\mathcal{G}, \beta, \psi)$. Therefore, if $(\mathcal{F}, \alpha, \varphi) \ncong (\mathcal{G}, \beta, \psi)$, then the finite set

$$\left\{\pi_k(\mathbf{K}^{\mathbf{q}}) \mid \mathbf{q} \in [|V(\mathcal{F})|^{\ell-k} + |V(\mathcal{G})|^{\ell-k}]^J\right\}$$

is guaranteed to contain a nonisomorphism witness. Constructing this set entails constructing the instances $\{\mathbf{K}^\Lambda \mid \Lambda \in J\}$ composing $\mathbf{K}^{\mathbf{q}}$, which are defined existentially using (3.4.10) and (3.4.11). The proof of Lemma 3.4.1 explicitly constructs a finite set to which each $\mathbf{K}^\Lambda$ must belong: let

$$\begin{aligned} W_\Theta = &\left\{ \mathbf{K}^{\mathbf{p}} \mid \mathbf{p} = [2|V(\mathcal{F})|]^{\mathcal{J}(\mathcal{F})} \right\} \\ &\cup \left\{ \mathbf{K}^{F,\mathbf{p}^{(0)},\ldots,\mathbf{p}^{(n_F-1)}} \mid F \in \mathcal{F}, \mathbf{p}^{(0)}, \ldots, \mathbf{p}^{(n_F-1)} \in [|V(\mathcal{F})|]^{\mathcal{J}(\mathcal{F})} \right\} \\ &\cup \left\{ \mathbf{K}^{F,\mathbf{p}^{(0)},\ldots,\mathbf{p}^{(n_F-2)}} \mid F \in \mathcal{F}, \mathbf{p}^{(0)}, \ldots, \mathbf{p}^{(n_F-2)} \in [|V(\mathcal{F})|]^{\mathcal{J}(\mathcal{F})} \right\} \subset \mathcal{PLI}^{\text{simp}}[\mathcal{F}; \ell] \end{aligned}$$

be the (finite) set of all #CSP($\mathcal{F}$) instances whose constraint sets are defined in (3.4.1), (3.4.5), and (3.4.9), the three steps of the proof of Lemma 3.4.1 (with $\theta := \Theta$ in the statement of Lemma 3.4.1). If $\Phi \in E(\varphi) \setminus A$, then $(\mathcal{F}, \alpha, \Theta) \ncong (\mathcal{F}, \alpha, \Phi)$, so, by the proof of Lemma 3.4.1, there is a $\mathbf{K}^\Phi = \mathbf{K}^\Lambda \in W_\Theta$ satisfying (3.4.10). Similarly, if $\Psi \in E(\psi) \setminus B$, then $(\mathcal{F}, \alpha, \Theta) \ncong (\mathcal{G}, \beta, \Psi)$, so there is a $\mathbf{K}^\Psi = \mathbf{K}^\Lambda \in W_\Theta$ satisfying (3.4.11). Thus each $\mathbf{K}^{\mathbf{q}}$ is composed of instances in the explicitly constructed set W_Θ.

Chapter 4

Generators, Invariants, and Duality

In this chapter, based on parts of [CY26] (joint work with Jin-Yi Cai) and [You25b], we introduce the algebraic perspective on gadgets that we will take throughout the rest of this thesis. Since the theorems we wish to use apply to vector spaces, we work with *quantum gadgets*, which are formal linear combinations of gadgets. With the gadget operation $\otimes$, quantum gadget tensors form a subalgebra of $\mathcal{V}$. We consider two related types of quantum gadget subalgebras of $\mathcal{V}$. Both are concretizations of certain types of monoidal (or tensor) categories, and inherit the diagrammatic calculus associated with these categories [Sel10].

The first is a *tensor category with duals*, a concept from category theory with concrete realizations via tensors in the representation theory of classical linear-algebraic groups and compact matrix quantum groups – see e.g. [BS09]. The tensors of quantum *planar* gadgets form a tensor category with duals; Mančinska and Roberson [MR20] introduced tensor categories with duals into the study of counting indistinguishability in their proof of the equivalence between quantum isomorphism and homomorphism indistinguishability over planar graphs. Specifically, they use Woronowicz's *Tannaka-Krein duality* [Wor88] to associate the quantum automorphism group of a graph with a unique tensor category with duals. However, we postpone further discussion of this theorem to Chapter 6.

The second type of subalgebra of $\mathcal{V}$ that we consider is a *wheeled PROP* (short for wheeled product and permutation category). Like tensor categories with duals, abstract wheeled PROPs come from category theory [Mac65; MMS09]. Derksen and Makam [DM23] study concrete sub-wheeled PROPs of $\mathcal{V}$, which correspond to the tensors of quantum bipartite (Holant$(\mathcal{F} \mid \mathcal{F}')$, or more generally Bi-Holant) gadgets. Derksen and Makam prove a duality theorem, analogous to Woronowicz's Tannaka-Krein duality, for certain sub-wheeled PROPs of $\mathcal{V}$. However, we postpone discussion of this theorem to Chapter 10. In this chapter, we present an *orthogonal duality* theorem that is a special case of both Woronowicz's Tannaka-Krein duality [BS09] and of Derksen and Makam's duality [Sch08b]. Orthogonal duality applies to *symmetric* tensor categories with duals, which contain the swap tensor $\times$ and thereby circumvent planarity, and to *pivotal* wheeled PROPs, which contain the pivot tensors $\supset$ and $\subset$ and thereby circumvent bipartiteness. We identify both of these spaces with general quantum Holant gadget tensors. Then the main result of this chapter, Theorem 4.5.2, is an algebraic characterization of quantum Holant gadget tensors. This result generalizes a result of Regts [Reg12] for edge coloring models, or Holant($\mathcal{F}$) for signature sets $\mathcal{F}$ containing only symmetric signatures and at most one signature of each arity.

4.1 Bi-Holant

Holant($\mathcal{F}$) is defined for a set $\mathcal{F} \subset \mathcal{S}$ of signatures, which are shapeless tensors: in a Holant signature grid, we can contract any input of any signature with any other input. More generally, considering sets $\mathcal{F}$ of shaped tensors and demanding that a signature grid respect the shapes, we obtain what we call Bi-Holant.

Definition 4.1.1 (Bi-Holant)**.** For $\mathcal{F} \subset \mathcal{V}(\mathbb{K}^q)$, a Bi-Holant($\mathcal{F}$) signature grid Ω or gadget $\mathbf{K}$ is a Holant($\mathcal{F}$) signature grid respecting the shapes of its signatures – that is, the edge between any adjacent u and v must be a contravariant/left input to F_u and a covariant/right input to F_v, or vice-versa, and all left and right dangling edges must be left and right inputs to their dangling vertices, respectively.

The dangling edge condition ensures that a Bi-Holant gadget with ℓ left and r right dangling edges defines a tensor in ${}_{\ell}\mathcal{V}_r$, and that, like a Holant gadget, a Bi-Holant gadget can replace a vertex in a Bi-Holant grid. If $\mathcal{F}$ and $\mathcal{F}'$ are sets of contravariant and covariant tensors, then Bi-Holant$_{\mathcal{F} \cup \mathcal{F}'}$ is equivalent to Holant$_{\mathcal{F} \mid \mathcal{F}'}$. Therefore, by (2.4.1), Bi-Holant generalizes Holant.

Bi-Holant$_{\mathcal{F}}(\Omega)$ is the value of the contraction of Ω as a tensor network via the standard bilinear form $\langle \cdot, \cdot \rangle : {}_{\ell}\mathcal{V}_r \times {}_{r}\mathcal{V}_{\ell} \to \mathbb{K}$ from Definition 2.2.7 that contracts covariant inputs with contravariant inputs and vice-versa. For example, slicing the n edges of an $\mathcal{F} \,|\, \mathcal{F}'$-grid Ω yields two Bi-Holant gadgets with signatures $F_1 \in {}_n\mathcal{V}_0$ and $F_2 \in {}_0\mathcal{V}_n$ such that $\text{Holant}_{\mathcal{F}\,|\,\mathcal{F}'}(\Omega) = \langle F_1, F_2 \rangle = F_2(F_1)$.

Definition 4.1.2 ($\mathfrak{Q}_{\mathcal{F}}, \langle \mathcal{F} \rangle$)**.** Given $\mathcal{F} \subset \mathcal{V}$, an (ℓ, r) *quantum Bi-Holant*($\mathcal{F}$) *gadget* $\mathbf{K}$ is a formal $\mathbb{K}$-linear combination of (ℓ, r) Bi-Holant($\mathcal{F}$) gadgets. Define $\mathfrak{Q}_{\mathcal{F}}$ and $\langle \mathcal{F} \rangle$ to be the spaces of all quantum Bi-Holant($\mathcal{F}$) gadgets and their tensors, respectively (extending the gadget tensor function M linearly), and ${}_\ell\langle \mathcal{F} \rangle_r := \langle \mathcal{F} \rangle \cap {}_\ell\mathcal{V}_r$.

See Figure 4.1. The terminology "quantum gadget" – $\mathbf{K}$ is a "superposition" of gadgets – comes from quantum labeled graphs [FLS07] (these labeled graphs being a special case of the labeled #CSP instances in Chapter 3), of which quantum gadgets are a broad generalization.

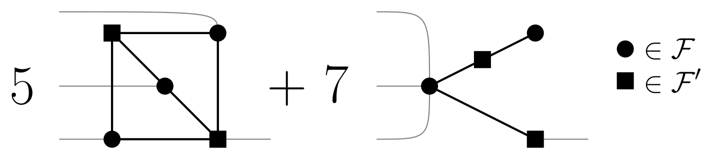

Figure 4.1: A $(3, 1)$-quantum-$\mathcal{F} \,|\, \mathcal{F}'$-gadget. Note that left/right dangling edges are incident to vertices in $\mathcal{F}/\mathcal{F}'$, respectively.

Extend left/right dangling edge contraction and $\top$ linearly to $\mathfrak{Q}_{\mathcal{F}}$, and extend $\circ$ and $\otimes$ bilinearly to $\mathfrak{Q}_{\mathcal{F}}$. Extend $\dagger$ conjugate-linearly to $\mathfrak{Q}_{\mathcal{F}}$ for $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$ – that is, $\left(\sum_i c_i \mathbf{K}_i\right)^\dagger = \sum_i \overline{c_i}\, \mathbf{K}_i^\dagger$. Note that $\mathfrak{Q}_{\mathcal{F}}$ and $\langle \mathcal{F} \rangle$ are closed under quantum gadget construction, as every $\langle \mathcal{F} \rangle$-gadget $\mathbf{K}$ decomposes into an quantum-$\mathcal{F}$-gadget after replacing every $F \in \langle \mathcal{F} \rangle \setminus \mathcal{F}$ in $\mathbf{K}$ with the quantum-$\mathcal{F}$-gadget realizing F and expanding linearly. While ${}_1\langle \mathcal{F} \rangle_1$ always contains $I = (=_2)^{1,1}$ as the signature of a crossing wire, we do not always have the pivots $\subset \in {}_0\langle \mathcal{F} \rangle_2$ or $\supset \in {}_2\langle \mathcal{F} \rangle_0$. Such a co/contravariant $(=_2)$ is quite powerful, as it allows connecting two left or two right dangling edges with each other, circumventing bipartiteness, and allows reshaping tensors – e.g. construct $A^{1,1}$ from $A \in (\mathbb{K}^q)^{\otimes 2}$ by connecting $\subset$ to the second input of A.

Next, we extend Definition 2.4.2 to quantum gadgets and Bi-Holant

Definition 4.1.3 ($\mathbf{K}_{\mathcal{F} \to \mathcal{G}}$, Bi-Holant-indistinguishable)**.** Given a $\mathbf{K} \in \mathfrak{Q}_{\mathcal{F}}$, construct $\mathbf{K}_{\mathcal{F} \to \mathcal{G}} \in \mathfrak{Q}_{\mathcal{G}}$ by replacing every $F \in \mathcal{F}$ in every gadget comprising $\mathbf{K}$ with the corresponding $G \in \mathcal{G}$.

Bijective $\mathcal{F}, \mathcal{G} \subset \mathcal{V}$ are *Bi-Holant-indistinguishable* if $\text{Bi-Holant}_{\mathcal{F}}(\Omega) = \text{Bi-Holant}_{\mathcal{G}}(\Omega_{\mathcal{F} \to \mathcal{G}})$ for every $\mathcal{F}$-grid Ω.

Viewing $\langle\mathcal{F}\rangle$ and $\langle\mathcal{G}\rangle$ as multisets in bijection with $\mathfrak{Q}_\mathcal{F}$ and $\mathfrak{Q}_\mathcal{G}$, the $(\cdot)_{\mathcal{F}\to\mathcal{G}}$ operation induces a (multiset) bijection $\leftrightsquigarrow$ between $\langle\mathcal{F}\rangle$ and $\langle\mathcal{G}\rangle$, where $K \leftrightsquigarrow \widetilde{K}$ if K and $\widetilde{K}$ are the tensors of $\mathbf{K}$ and $\mathbf{K}_{\mathcal{F}\to\mathcal{G}}$. Under this bijection, if $\mathcal{F}$ and $\mathcal{G}$ are (Bi-)Holant-indistinguishable, then so are $\langle\mathcal{F}\rangle$ and $\langle\mathcal{G}\rangle$:

Proposition 4.1.1. *Let $\mathcal{F},\mathcal{G} \subset \mathcal{V}$. Then $\mathcal{F}$ and $\mathcal{G}$ are Bi-Holant-indistinguishable if and only if $\langle\mathcal{F}\rangle$ and $\langle\mathcal{G}\rangle$ are Bi-Holant-indistinguishable.*

Proof. Any $\mathcal{F}$-grid or $\mathcal{G}$-grid is a $\langle\mathcal{F}\rangle$ or $\langle\mathcal{G}\rangle$-grid, respectively, giving the ($\Longleftarrow$) direction. For ($\Longrightarrow$), we can express any $\langle\mathcal{F}\rangle$-grid Ω as a quantum $\mathcal{F}$-grid by, for every vertex v in Ω assigned a signature $F^v \in \langle\mathcal{F}\rangle \setminus \mathcal{F}$, replacing v and its incident edges with the quantum $\mathcal{F}$-gadget $\mathbf{K}^v$ with signature F^v, then linearly expanding to obtain a quantum $\mathcal{F}$-grid Ω' with the same Bi-Holant value as Ω. In $\Omega_{\langle\mathcal{F}\rangle\to\langle\mathcal{G}\rangle}$, v is assigned $G^v \leftrightsquigarrow F^v$, which is the signature of $\mathbf{K}^v_{\mathcal{F}\to\mathcal{G}}$, so similarly replace v with $\mathbf{K}^v_{\mathcal{F}\to\mathcal{G}}$. Expanding $\Omega_{\langle\mathcal{F}\rangle\to\langle\mathcal{G}\rangle}$ to a quantum $\mathcal{G}$-grid, we obtain $\Omega'_{\mathcal{F}\to\mathcal{G}}$. By assumption, every term of Ω' has the same Bi-Holant value as the corresponding term of $\Omega'_{\mathcal{F}\to\mathcal{G}}$, so

$$\text{Bi-Holant}_\Omega(\langle\mathcal{F}\rangle) = \text{Bi-Holant}_{\Omega'}(\mathcal{F}) = \text{Bi-Holant}_{\Omega'_{\mathcal{F}\to\mathcal{G}}}(\mathcal{G}) = \text{Bi-Holant}_{\Omega_{\langle\mathcal{F}\rangle\to\langle\mathcal{G}\rangle}}(\langle\mathcal{G}\rangle). \qquad \square$$

We conclude this section with the following well-known generalization of Theorem 2.4.1.

Proposition 4.1.2. *For $T \in \mathrm{GL}_q$ and $\mathcal{F} \subset \mathcal{V}$ and any quantum Bi-Holant $\mathcal{F}$-gadget $\mathbf{K}$ with tensor K, the tensor of $\mathbf{K}_{\mathcal{F}\to T\mathcal{F}}$ is $T \cdot K$. Therefore*

$$T\langle\mathcal{F}\rangle = \langle T\mathcal{F}\rangle.$$

Proof. The tensor K is constructed by a sequence of left-right contractions between tensors in $\mathcal{F}$, and the tensor of the gadget $\mathbf{K}_{\mathcal{F}\to T\mathcal{F}}$ is constructed by the same sequence of contractions between the corresponding tensors in $T\mathcal{F}$. Since $u(v) = uT^{-1}Tv = (T\cdot u)(T\cdot v)$ for any $v \in \mathbb{K}^q$ and $u \in (\mathbb{K}^q)^*$, these contractions are GL_q equivariant, meaning that if ${}_i\partial_j(F) \in {}_{\ell-1}\mathcal{V}_{r-1}$ is the result of contracting left input i and right input j of $F \in {}_\ell\mathcal{V}_r$, then ${}_i\partial_j(T\cdot F) = T \cdot {}_i\partial_j(F)$ (see e.g. [Sch08b, Equation 11]). Hence, inductively, the tensor of $\mathbf{K}_{\mathcal{F}\to T\mathcal{F}}$ is $T\cdot K$. The extension to quantum gadgets follows by linearity.

Equivalently, in the diagrammatic terms of Figure 2.1 and Figure 2.2, the T transformations cancel on every internal edge of $\mathbf{K}_{\mathcal{F}\to T\mathcal{F}}$, and only survive on the dangling edges. $\square$

Specializing to 0-ary gadgets in $\langle\mathcal{F}\rangle$ – that is, (quantum) Bi-Holant($\mathcal{F}$) grids – gives an extension of Theorem 2.4.1 to Bi-Holant.

Theorem 4.1.1 (The Bi-Holant Theorem)**.** *If $\mathcal{F}, \mathcal{G} \subset \mathcal{V}$ satisfy $\mathcal{F} = T\,\mathcal{G}$ for $T \in \mathrm{GL}_q$, then $\mathcal{F}$ and $\mathcal{G}$ are Bi-Holant-indistinguishable.*

4.2 Wheeled PROPs and TCWDs

Definition 4.2.1 ([DM23, Definition 2.1])**.** A *pre-wheeled PROP* is a bigraded $\mathbb{K}$-vector space $\mathfrak{R} = \bigoplus_{\ell,r\geq 0} {}_{\ell}\mathfrak{R}_r$ together with

(i) a special element $1_{\mathfrak{R}} \in {}_0\mathfrak{R}_0$,

(ii) a special element $I_{\mathfrak{R}} \in {}_1\mathfrak{R}_1$,

(iii) a bilinear map $\otimes : {}_{\ell_1}\mathfrak{R}_{r_1} \times {}_{\ell_2}\mathfrak{R}_{r_2} \to {}_{\ell_1+\ell_2}\mathfrak{R}_{r_1+r_2}$, and

(iv) a linear map ${}_i\partial_j : {}_{\ell}\mathfrak{R}_r \to {}_{\ell-1}\mathfrak{R}_{r-1}$ for every $1 \leq i \leq \ell$ and $1 \leq j \leq r$.

The mixed tensor algebra $\mathcal{V}$ is a pre-wheeled PROP, where $1_{\mathcal{V}} = 1_{\mathbb{K}}$, $I_{\mathcal{V}} = I$ (the identity map), $\otimes$ is the usual tensor product, and ${}_i\partial_j$ contracts the ith contravariant input with the jth covariant input. For any $\mathcal{F}$, the algebra $\mathfrak{Q}_{\mathcal{F}}$ of quantum-$\mathcal{F}$-gadgets (the formal linear combinations of the diagrams themselves) is also a pre-wheeled PROP, where ${}_{\ell}(\mathfrak{Q}_{\mathcal{F}})_r$ is the algebra of (ℓ, r)-quantum-$\mathcal{F}$-gadgets, $1_{\mathfrak{Q}_{\mathcal{F}}}$ is the empty gadget, $I_{\mathfrak{Q}_{\mathcal{F}}}$ is the wire gadget, $\otimes$ is gadget tensor product, and ${}_i\partial_j$ is the operation of connecting the ith left input and jth right input. In fact, $\mathfrak{Q}_{\mathcal{F}}$ is (isomorphic to) the *free wheeled PROP* generated by $\mathcal{F}$ [DM23, Definition 2.16]. A *wheeled PROP* is a pre-wheeled PROP which is the image of a free wheeled PROP under a pre-wheeled PROP homomorphism (a linear map respecting the bigrading and the four elements/operations listed in Definition 4.2.1) [DM23, Definition 2.17]. Therefore $\langle\mathcal{F}\rangle \subset \mathcal{V}$ is a wheeled PROP, as it is the image of the free wheeled PROP $\mathfrak{Q}_{\mathcal{F}}$ under the pre-wheeled PROP homomorphism mapping a quantum-$\mathcal{F}$-gadget to its signature. Specifically, $\langle\mathcal{F}\rangle$ is a sub-wheeled PROP of $\mathcal{V}$ (which is the image of the free wheeled PROP $\mathfrak{Q}_{\mathcal{V}}$ under the same signature-evaluation map), and every sub-wheeled PROP of $\mathcal{V}$ is $\langle\mathcal{F}\rangle$ for some $\mathcal{F} \subset \mathcal{V}$. In this thesis, we only consider sub-wheeled PROPs of $\mathcal{V}$, so we use "wheeled PROP" to refer to a sub-wheeled PROP of $\mathcal{V}$.

Any wheeled PROP $\mathcal{F} \subset \mathcal{V}$ is closed under arbitrary permutations of left inputs and arbitrary permutations of right inputs: to permute e.g. the left inputs of $F \in {}_\ell\mathcal{F}_r$ by $\sigma \in S_\ell$, construct $F \otimes I^{\otimes \ell}$, then contract the right ends of the various copies of I with the left inputs of F according to σ. However, in a general wheeled PROP, reflecting bipartiteness, we cannot permute between left and right inputs. In other words, the allowed permutations are "bipartite, but not planar".

Now we introduce our other main family of tensor subalgebras, which are "planar, but not bipartite".

Definition 4.2.2 (TCWD, $\top$-TCWD)**.** A bigraded $\mathbb{C}$-vector space $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$ is a *tensor category with duals* (TCWD) if it satisfies the following properties:

(i) $\mathcal{F}$ is closed under $\circ : {}_\ell\mathcal{V}_r \times {}_r\mathcal{V}_m \to {}_\ell\mathcal{V}_m$ for every ℓ, r, m.

(ii) $\mathcal{F}$ is closed under $\otimes : {}_\ell\mathcal{V}_r \times {}_{\ell'}\mathcal{V}_{r'} \to {}_{\ell+\ell'}\mathcal{V}_{r+r'}$ for every ℓ, ℓ', r, r'.

(iii) $\mathcal{F}$ is closed under $\dagger : {}_\ell\mathcal{V}_r \to {}_r\mathcal{V}_\ell$ for every ℓ, r.

(iv) $I \in {}_1\mathcal{F}_1$

(v) $\supset \in {}_2\mathcal{F}_0$

$\mathcal{F}$ is a $\top$-TCWD if it satisfies the same properties, but is closed under $\top$ instead of $\dagger$.

Since both wheeled PROPs and TCWDs must contain I (and both bipartite and planar Holant gadgets contain the crossing wire), we implicitly assume all tensor sets contain I.

Definition 4.2.3 ($\langle \mathcal{F} \rangle_{+,\circ,\otimes,\dagger}$)**.** For $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$ and a set $S \subset \{\circ, \otimes, \top, \dagger\}$ of operations, define $\langle \mathcal{F} \rangle_S \subset \mathcal{V}(\mathbb{C}^q)$ to be the multiset generated by $\mathcal{F}$ and I under the operations in S, and define $\langle \mathcal{F} \rangle_{+,S} := \operatorname{span}_{\mathbb{C}}(\langle \mathcal{F} \rangle_S)$.

The presence of $\supset$ and $\subset = \supset^\dagger$ enables viewing a TCWD as a set of signatures, rather than tensors, as it is closed under (cyclic) pivoting (recall from (2.1.2) and Figure 2.2 that pivoting does not change the underlying signature of a tensor):

Lemma 4.2.1. *If F is an n-ary signature such that $F^{\ell,r} \in \mathcal{F}$ for some ℓ, r with $\ell + r = n$, then $F^{\ell',r'} \in \langle \mathcal{F}, \supset, \subset \rangle_{\circ,\otimes}$ for every ℓ', r' with $\ell' + r' = n$.*

Proof. If $r < r'$ and $\ell > \ell'$, we use $r' - r = \ell - \ell'$ horizontally nested copies of $\subset = \supset^{\dagger}$ to pivot the bottom $\ell - \ell'$ left dangling edges to the right (see Figure 4.1).

$$F^{\ell',r'} = (I^{\otimes \ell'} \otimes \subset) \circ (I^{\otimes r-(\ell-\ell'-1)} \otimes \subset \otimes I^{\otimes 1}) \circ \ldots \circ (I^{\otimes \ell-1} \otimes \subset \otimes I^{\otimes \ell-\ell'-1}) \circ (F^{\ell,r} \otimes I^{\otimes \ell-\ell'}).$$

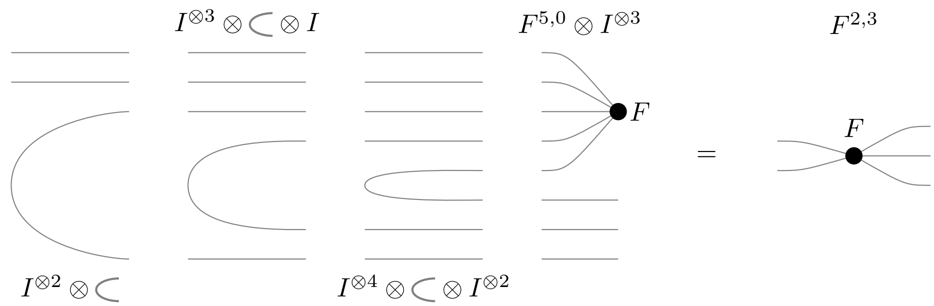

Figure 4.1: Reshaping $F^{5,0}$ to $F^{2,3}$ in a TCWD.

Similarly, if $r > r'$ and $\ell < \ell'$, we apply horizontally reflected reasoning, using $\subset$, to pivot the bottom $r - r'$ output right dangling edges to the left.

$$F^{\ell',r'} = (F^{\ell,r} \otimes I^{\otimes r-r'}) \circ (I^{\otimes r-1} \otimes \supset \otimes I^{\otimes r-r'-1}) \circ \ldots \circ (I^{\otimes r-(r-r'-1)} \otimes \supset \otimes I^{\otimes 1}) \circ (I^{\otimes r'} \otimes \supset). \qquad \square$$

In particular, we may pivot every tensor in a TCWD to its fully contravariant form to view it as a signature. The idea that a TCWD is invariant under changing the dimensions of its matrices due to the presence of $\supset$ is not new, and is an instance of "Frobenius reciprocity" in e.g. [Gro21; Ban05].

The gadget algebra of a complex signature set is not necessarily closed under conjugation, but TCWDs are closed under †, so we must impose the following condition on complex $\mathcal{F}$ for its gadget algebra to be a TCWD (see Proposition 4.4.2).

Definition 4.2.4 (conjugate-closed)**.** A set $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$ is *conjugate-closed* if for every $F \in {}_{\ell}\mathcal{F}_r$, its entrywise conjugate $\overline{\mathcal{F}} \in {}_{\ell}\mathcal{F}_r$, where $\overline{\mathcal{F}}$ is constructed by conjugating the coefficients of F.

Of course, any real-valued $\mathcal{F}$ is conjugate-closed. The signatures in a (conjugate-closed, in the case of $\top$) TWCD are closed under the following operations, which together constitute all dihedral symmetries of a single tensor.

Definition 4.2.5 (Signature $F^{(r)}$, $F^\top$, $F^\dagger$)**.** Given $F \in {}_n\mathcal{S}$, define the m-th rotation $F^{(m)} \in {}_n\mathcal{S}$ by

$$F^{(m)}(x_1, \ldots, x_n) = F(x_{m+1}, \ldots, x_n, x_1, \ldots, x_m).$$

Define the reflection $F^\top \in {}_n\mathcal{S}$ by

$$F^\top(x_1, \ldots, x_n) = F(x_n, \ldots, x_1).$$

For $F \in {}_n\mathcal{S}(\mathbb{C}^q)$, define $F^\dagger := \overline{F}^\top \in {}_n\mathcal{S}(\mathbb{C}^q)$.

The next proposition implies that any TCWD is closed under signature rotations.

Proposition 4.2.1. *For any signature $F \in {}_n\mathcal{S}$ and $m \in [n]$, we have $F^{(m)} \in \langle F, \supset, \subset \rangle_{\circ,\otimes}$.*

Proof. We realize rotations similarly to pivoting:

$$F^{(1)} = (\subset \otimes I^{\otimes n}) \circ (I \otimes F \otimes I) \circ \supset.$$

Iterating this construction r times yields $F^{(m)} \in \langle F, \supset, \subset \rangle_{\circ,\otimes}$. □

Proposition 4.2.2. *Let $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$ be a TCWD. For any signature $F \in {}_n\mathcal{F}_0$,*

- *$F^\dagger \in \mathcal{F}$,*
- *$F^\top \in \mathcal{F} \iff \overline{F} \in \mathcal{F}$.*

Proof. Recall from Section 2.1 that, if $F^{\ell,r}$ is a flattening of signature F, the input order of F is counterclockwise around the diagram representing $F^{\ell,r}$. Applying $\dagger$ or $\top$ to $F^{\ell,r}$ reflects this diagram, which reverses the input order to F from counterclockwise to clockwise. Formally, for $\top$,

$$\begin{aligned}
(F^{\ell,r})^\top &= \sum_{x_1,\ldots,x_\ell} \sum_{y_1,\ldots,y_r} F(x_1, \ldots, x_\ell, y_r, \ldots, y_1) e_{y_1} \otimes \ldots \otimes e_{y_r} \otimes e^*_{x_1} \otimes \ldots \otimes e^*_{x_\ell} \\
&= \sum_{y_1,\ldots,y_r} \sum_{x_1,\ldots,x_\ell} F^\top(y_1, \ldots, y_r, x_\ell, \ldots, x_1) e_{y_1} \otimes \ldots \otimes e_{y_r} \otimes e^*_{x_1} \otimes \ldots \otimes e^*_{x_\ell} \\
&= (F^\top)^{r,\ell}, \qquad (4.2.1)
\end{aligned}$$

so the underlying signature of $(F^{\ell,r})^\top \in {}_r\mathcal{F}_\ell$ is $F^\top$. Similarly,

$$(F^{\ell,r})^\dagger = (F^\dagger)^{r,\ell}, \qquad (4.2.2)$$

so the underlying signature of $(F^{\ell,r})^\dagger$ is $F^\dagger$. Now, since any TCWD is closed under $\dagger$ and $F^{0,n} \in \mathcal{F}$ by Lemma 4.2.1, we recover $F^\dagger = (F^\dagger)^{n,0} = (F^{0,n})^\dagger \in {}_n\mathcal{F}_0$. For the final item, $\mathcal{F}$ does not necessarily contain $F^\top$ and $\overline{F}$, but if $F^\top \in \mathcal{F}$, then $\mathcal{F} \ni (F^\top)^\dagger = \overline{F}$, and if $\overline{F} \in \mathcal{F}$, then $\mathcal{F} \ni \overline{F}^\dagger = F^\top$. □

4.3 Wire gadgets

Recall the assumption that left-right wires are always present as $\mathcal{F}$-gadgets for any $\mathcal{F} \subset \mathcal{V}$. For $\mathcal{F} = \varnothing$, no gadget can contain any vertices (as each vertex must be assigned a signature from $\mathcal{F}$), so every $\varnothing$-gadget is a *braid* composed exclusively of left-right wires. The tensor of any such gadget has the following form.

Definition 4.3.1 (S_σ, S)**.** For permutation $\sigma \in S_n$, define the $2n$-ary *braid* tensor $S_\sigma \in {}_n\langle\varnothing\rangle_n$ by $S_\sigma(\mathbf{x},\mathbf{y}) = 1$ iff $x_i = y_{\sigma(i)}$ for every $i \in [n]$, and $S_\sigma(\mathbf{x}\,\mathbf{y}) = 0$ otherwise.

Define the 4-ary "swap" tensor $S := S_{(01)} =$ . See Figure 4.1

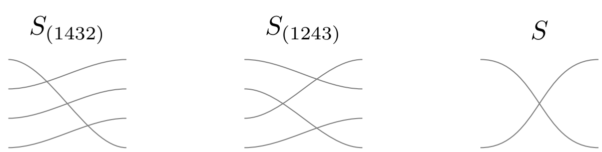

Figure 4.1: Braid gadgets.

We next introduce two standard concepts from invariant theory.

Definition 4.3.2 ($\mathcal{V}^Q$)**.** For a subgroup $Q \subset \mathrm{GL}_q$, define

$$\mathcal{V}^Q := \{F \in \mathcal{V} : T \cdot F = F \text{ for every } T \in Q\} \subset \mathcal{V}$$

to be the set of Q-invariant tensors.

Definition 4.3.3 ($\mathrm{Stab}(\mathcal{F})$)**.** Define the stabilizer group of $\mathcal{F} \subset \mathcal{V}$ by

$$\mathrm{Stab}(\mathcal{F}) := \{T \in \mathrm{GL}_q : T \cdot F = F \text{ for every } F \in \mathcal{F}\} \subset \mathrm{GL}_q\,.$$

In any wheeled PROP (equivalently, $\langle\mathcal{F}\rangle$ for any $\mathcal{F}$), we can construct all braid tensors by permuting the left (or right) inputs of $I^{\otimes n}$. In this sense, the braid tensors are the minimal sub-wheeled PROP of $\mathcal{V}$. Therefore, they should have the maximal stabilizer group, as the following classical result of invariant theory, originally due to Weyl (see also [GW09, Theorem 5.3.1]) shows. See Figure 4.2.

Theorem 4.3.1 (Tensor First Fundamental Theorem for GL_q [Wey66])**.** $\mathcal{V}^{\mathrm{GL}_q} = \langle \varnothing \rangle$.

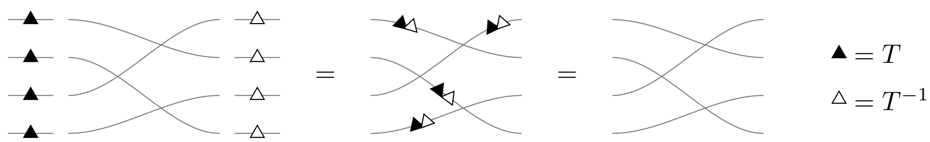

Figure 4.2: The $\supset$ inclusion of Theorem 4.3.1: $T \cdot S_{(1243)} = S_{(1243)}$ for every $T \in \mathrm{GL}_q$.

TCWDs, which don't admit arbitrary contraction, do not necessarily contain braid tensors. The next result shows that, in a TCWD, containing all braid tensors is equivalent to containing $\times$.

Proposition 4.3.1. *For any n and $\sigma \in S_n$, we have $S_\sigma \in \langle I, \times \rangle_{\circ,\otimes}$.*

Proof. Decompose σ into adjacent transpositions as $\sigma = (a_1\ a_1+1)(a_2\ a_2+1)\dots(a_s\ a_s+1)$. Then, since $\times$ swaps the position of adjacent dangling edges, we have (see Figure 4.3)

$$S_\sigma = \bigcirc_{i=1}^{s}(I^{\otimes a_i - 1} \otimes \times \otimes I^{\otimes n - a_i - 1}). \qquad \square$$

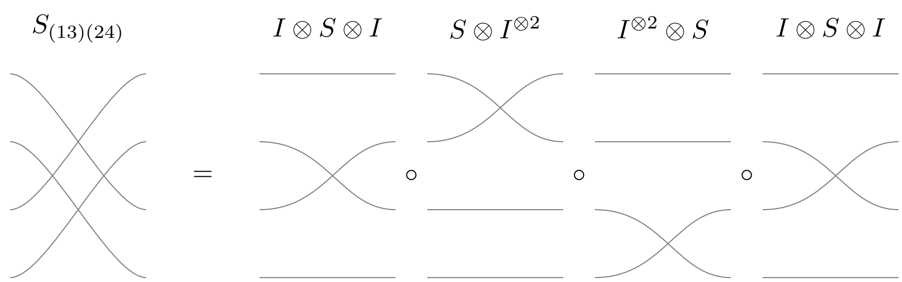

Figure 4.3: Proposition 4.3.1, with $\sigma = (1\ 3)(2\ 4) = (2\ 3)(1\ 2)(3\ 4)(2\ 3)$.

4.4 Liberated wheeled PROPs and TCWDs

Informally, contractions in a wheeled PROP may be nonplanar but are subject to the bipartite left/right restriction, and the presence of the pivots in a TCWD removes the bipartite restriction, but contractions are restricted to be planar. In this section, we "liberate" wheeled PROPs and TCWDs from these restrictions by adding to them the pivots and the swap, respectively.

Definition 4.4.1. A sub-wheeled PROP $\mathcal{F}$ of $\mathcal{V}$ is *pivotal* if it contains $\supset$ and $\subset$.

Contracting a right input with one input of $\supset$ moves this input to the left, and contracting a left input with one input of $\subset$ moves this input to the right. Combining this with closure under all left- and right-preserving permutations, as discussed in Section 4.2, shows that a pivotal wheeled PROP is closed under all input permutations.

Definition 4.4.2. A TCWD or $\top$-TCWD $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$ is *symmetric* if it contains $\times$.

We will see that symmetric TCWDs also admit arbitrary input permutations. In fact, once we liberate (conjugate-closed) wheeled PROPs and TCWDs, they become identical: both become the space of all quantum-$\mathcal{F}$-gadgets in the context of unconstrained Holant, with shapeless tensors (signatures). See Figure 4.3.

Proposition 4.4.1. *Let $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$. Then*

- *$\mathcal{F}$ is a pivotal wheeled PROP if and only if $\mathcal{F}$ is a symmetric $\top$-TCWD.*
- *$\mathcal{F}$ is a conjugate-closed pivotal wheeled PROP if and only if $\mathcal{F}$ is a symmetric TCWD.*

Proof. By definition, wheeled PROPs over $\mathbb{C}$, TCWD, and $\top$-TCWDs are bigraded $\mathbb{C}$-vector spaces, both contain I and the scalar 1, and both are closed under $\otimes$. We now show the remaining implications simultaneously for both claims.

($\Longrightarrow$): Let $\mathcal{F}$ be a pivotal wheeled PROP. Then $\mathcal{F}$ satisfies item (v) of Definition 4.2.2, and satisfies item (i) because we construct $F \circ F'$ for $F \in {}_\ell\mathcal{V}_r$ and $F' \in {}_r\mathcal{V}_m$ by contracting the r left inputs of F with the corresponding right inputs of F'. Next, $\mathcal{F}$ satisfies item (iii) because we construct $F^\top$ or $F^\dagger$ by connecting $\supset$ to every right input and $\subset$ to every left input of F or $\overline{F}$, if $\mathcal{F}$ is conjugate closed, respectively. Finally, $S \in \mathcal{F}$ because, as mentioned earlier, any wheeled PROP contains all braid tensors (formally, start with $I^{\otimes 4}$, then connect the first wire to the fourth and the second to the third).

($\Longleftarrow$): Let $\mathcal{F}$ be a symmetric TCWD or $\top$-TCWD. First, $\mathcal{F}$ contains $\supset$ and $\subset = \supset^\dagger = \supset^\top$. Next, we show $\mathcal{F}$ is closed under arbitrary left-right contractions. Let $F \in {}_\ell\mathcal{F}_r$ with $\ell, r > 0$ and $i \leq \ell$ and $j \leq r$. We will construct ${}_i\partial_j(F) \in \mathcal{F}$. By Lemma 4.2.1, $F^{\ell+r,0} \in \mathcal{F}$. To contract arbitrary inputs x and y of $F^{\ell+r,0}$, define $\sigma \in S_n$ to move inputs x and y to positions 1 and 2, and shift all other inputs down while maintaining their order. By Proposition 4.3.1, $S_\sigma \in \mathcal{F}$. Then connect inputs x and y using $\subset$ (see Figure 4.1):

$$(\subset \otimes I^{\otimes n-2}) \circ S_\sigma \circ F^{\ell+r,0} \in \mathcal{F} \tag{4.4.1}$$

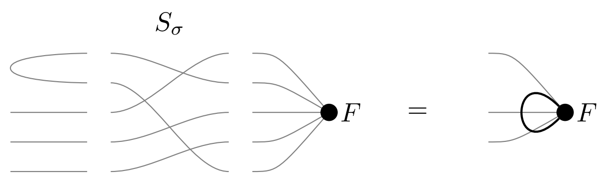

Figure 4.1: Contracting the 2nd and 5th inputs of $F^{\ell+r,0}$ as in (4.4.1).

Now contract inputs $x := i$ and $y := \ell + j$ in $F^{\ell+r,0}$, then use Lemma 4.2.1 again to pivot the uncontracted original right dangling edges back to the right. The result is ${}_i\partial_j(F) \in \mathcal{F}$. Therefore $\mathcal{F}$ is a pivotal wheeled PROP. Now, by the ($\Longrightarrow$) direction of the first claim, if $\mathcal{F}$ is a TCWD and $F \in \mathcal{F}$, then $\overline{F} = (F^\dagger)^\top \in \mathcal{F}$. Hence $\mathcal{F}$ is conjugate-closed. □

Definition 4.4.3 ($[\mathcal{F}]$)**.** Given a signature set $\mathcal{F} \subset \mathcal{S}(\mathbb{K}^q)$, let $[\mathcal{F}] := \langle \mathcal{F}, \supset, \subset \rangle \subset \mathcal{V}(\mathbb{K}^q)$ be the pivotal wheeled PROP generated by $\mathcal{F}$. Equivalently, $[\mathcal{F}]$ is the algebra of quantum Holant $\mathcal{F}$-gadget tensors.

Note that $[\mathcal{F}]$ depends only on the signatures of $\mathcal{F}$, not the specific shapes of the tensors, because $\supset$ and $\subset$ enable arbitrary reshaping in $\langle \mathcal{F}, \supset, \subset \rangle$. By the same logic, we can and will view $[\mathcal{F}] \subset \mathcal{S}$. With this perspective, define ${}_n\mathcal{F} \subset \mathcal{S}$ to be the space of n-ary signatures in $[\mathcal{F}]$. We obtain from Proposition 4.1.2 the following well-known generalization of Corollary 2.4.1.

Corollary 4.4.1. *For $H \in O_q(\mathbb{K})$, signature set $\mathcal{F} \subset \mathcal{S}(\mathbb{K}^q)$ and any quantum Holant $\mathcal{F}$-gadget* $\mathbf{K}$ *with tensor K, the tensor of* $\mathbf{K}_{\mathcal{F} \to H\mathcal{F}}$ *is $H \cdot K$. Therefore*

$$H[\mathcal{F}] = [H\mathcal{F}].$$

Proof. By definition, $[\mathcal{F}] = \langle \mathcal{F}, \supset, \subset \rangle$, so $\mathbf{K}$ is a quantum Bi-Holant $(\mathcal{F} \cup \{\supset, \subset\})$-gadget. Similarly, $\mathbf{K}_{\mathcal{F} \to H\mathcal{F}}$ is the corresponding quantum Bi-Holant $(H\mathcal{F} \cup \{\supset, \subset\})$-gadget. Since $H(\mathcal{F} \cup \{\supset, \subset\}) = H\mathcal{F} \cup \{\supset, \subset\}$ by Proposition 2.4.1, the result follows from Proposition 4.1.2 with $\mathcal{F} := \mathcal{F} \cup \{\supset, \subset\}$. □

In particular, for any signature $F \in \mathcal{S}$ and $\ell + r = \text{arity}(F)$, $F^{\ell,r}$ is the tensor of the gadget consisting of a single vertex assigned F with ℓ left and r right dangling edges, so $F^{\ell,r} \in [\mathcal{F}]$. This yields another well-known and useful fact (see Figure 4.2): the action of an orthogonal transformation is independent of a tensor's shape (in fact, since every invertible transformation preserves I but only orthogonal transformations preserve $\supset$ and $\subset$, only orthogonal transformations have this property):

$$(H^{\otimes n}F)^{\ell,r} = (H\cdot F)^{\ell,r} = H^{\otimes \ell}F^{\ell,r}(H^{\top})^{\otimes r} = H\cdot F^{\ell,r} \text{ for } H\in O_q(\mathbb{K}) \text{ and } F\in\mathcal{S}(\mathbb{K}). \tag{4.4.2}$$

In fact, it is well-known, and follows from direct calculation and the same diagrammatic calculus in Figure 4.2, that the second equality in (4.4.2) holds for any matrix $T\in\mathrm{GL}_q$ and $F\in {}_{\ell+r}\mathcal{S}(\mathbb{K}^q)$:

$$(T\cdot F)^{\ell,r} = T^{\otimes \ell}F^{\ell,r}(T^{\top})^{\otimes r}. \tag{4.4.3}$$

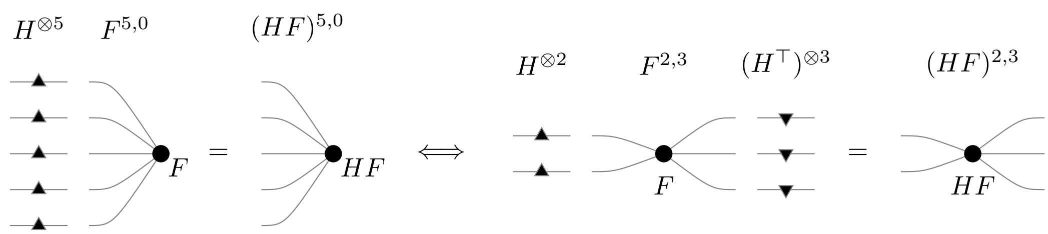

Figure 4.2: Illustrating $H^{\otimes n}F^{n,0} = (HF)^{n,0}$, or equivalently $H^{\otimes m}F^{m,d}(H^{\top})^{\otimes d} = (HF)^{m,d}$. Pivoting the left dangling edges to the right flips (transposes) the H on those edges.

We now formally prove the first equivalence in Figure 4.3 (which in turn implies the other two equivalences): quantum Holant($\mathcal{F}$) gadgets are exactly the pivotal wheeled PROP and (for conjugate-closed $\mathcal{F}$) the symmetric TCWD generated by $\mathcal{F}$.

Proposition 4.4.2. *Let $\mathcal{F}\subset\mathcal{S}(\mathbb{C}^q)$. Then*

- $[\mathcal{F}] = \langle \mathcal{F}, \supset, \times \rangle_{+,\circ,\otimes,\top}$.
- $[\mathcal{F}] = \langle \mathcal{F}, \supset, \times \rangle_{+,\circ,\otimes,\dagger}$ *if $\mathcal{F}$ is conjugate-closed.*

Proof. We prove the second claim; the first follows similarly from the first point of Proposition 4.4.1. If $K\in\langle\mathcal{F}\rangle$ be the signature of $\mathbf{K} = \sum_i \alpha_i\,\mathbf{K}_i$, then $\overline{K}$ is the signature of $\sum_i \overline{\alpha_i}(\mathbf{K}_i)_{\mathcal{F}\to\overline{\mathcal{F}}}$. Thus

$$\text{If } \mathcal{F}\subset\mathcal{V}(\mathbb{C}^q) \text{ is conjugate-closed, then } \langle\mathcal{F}\rangle \text{ is conjugate-closed.} \tag{4.4.4}$$

Therefore, since $(\mathcal{F},\supset,\subset)$ is conjugate-closed, $[\mathcal{F}] = \langle\mathcal{F},\supset,\subset\rangle$ is a conjugate-closed pivotal wheeled PROP. Now Proposition 4.4.1 asserts that $[\mathcal{F}]$ is a symmetric TCWD. Every symmetric TCWD containing $\mathcal{F}$ contains $\langle \mathcal{F}, \supset, \times \rangle_{+,\circ,\otimes,\dagger}$, so $[\mathcal{F}] \supset \langle \mathcal{F}, \supset, \times \rangle_{+,\circ,\otimes,\dagger}$. Conversely, $\langle \mathcal{F}, \supset, \times \rangle_{+,\circ,\otimes,\dagger}$ is a pivotal wheeled PROP by Proposition 4.4.1, and every pivotal wheeled PROP containing $\mathcal{F}$ contains $[\mathcal{F}] = \langle \mathcal{F}, \supset, \subset \rangle$, so $\langle \mathcal{F}, \supset, \times \rangle_{+,\circ,\otimes,\dagger} \subset [\mathcal{F}]$. □

Recall that $\langle \mathcal{F} \rangle$, the not-necessarily-pivotal wheeled PROP generated by $\mathcal{F}$, is the algebra of all quantum Bi-Holant $\mathcal{F}$ gadget tensors. One now asks: what is the not-necessarily-symmetric TCWD generated by $\mathcal{F}$? We will see in Theorem 6.3.1 that the answer is the algebra of all *planar* quantum Holant gadget tensors.

Remark 4.4.1. Figure 4.3 shows the quantum gadget algebras for various counting problems for conjugate-closed $\mathcal{F}$ over $\mathbb{C}$. In the gadgets 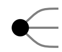 and 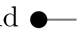, the vertices are assigned $=_3$ and $=_1$, respectively. The problem $\#\text{CSP}^{(3)}(\mathcal{F})$ is the restriction of $\#\text{CSP}(\mathcal{F})$ to instances in which each variable appears at most three times across all constraints, and the special case $\text{hom}^{(3)}(\cdot, X)$ is the problem of counting homomorphisms from graphs of degree at most three to X (Chapter 10). $\text{Pl}^\dagger$-Holant (Chapter 6) and $\text{Pl}^\dagger$-#CSP (Chapter 7) are the respective restrictions of Holant and #CSP to planar signature grids and constraint-variable incident graphs in which clockwise-oriented signatures/constraints are conjugated. Pl-hom$(\cdot, X)$, a special case of $\text{Pl}^\dagger$-#CSP (and, since A_X is real-valued and symmetric, of the other variants of planar #CSP in Chapter 7), is the problem of counting homomorphisms from planar graphs to X.

If problem A is an ancestor of problem B, then problem B is a special case of problem A with certain additional signatures assumed to be present; by considering problem A parameterized by the union of $\mathcal{F}$ and these additional signatures, we recover problem B parameterized by $\mathcal{F}$. Hence an indistinguishability characterization for problem A implies one for problem B, so our most general results are for Bi-Holant$(\mathcal{F})$ (Chapter 10) and $\text{Pl}^\dagger$-Holant$(\mathcal{F})$ (Chapter 9). The indistinguishability characterizations of hom$(\cdot, X)$ and Pl-hom$(\cdot, X)$ are known to be isomorphism [Lov67] and quantum isomorphism [MR20]; the characterizations for all other problems in Figure 4.3 are novel results of this thesis and the papers on which it is based, as shown in Table 4.1.

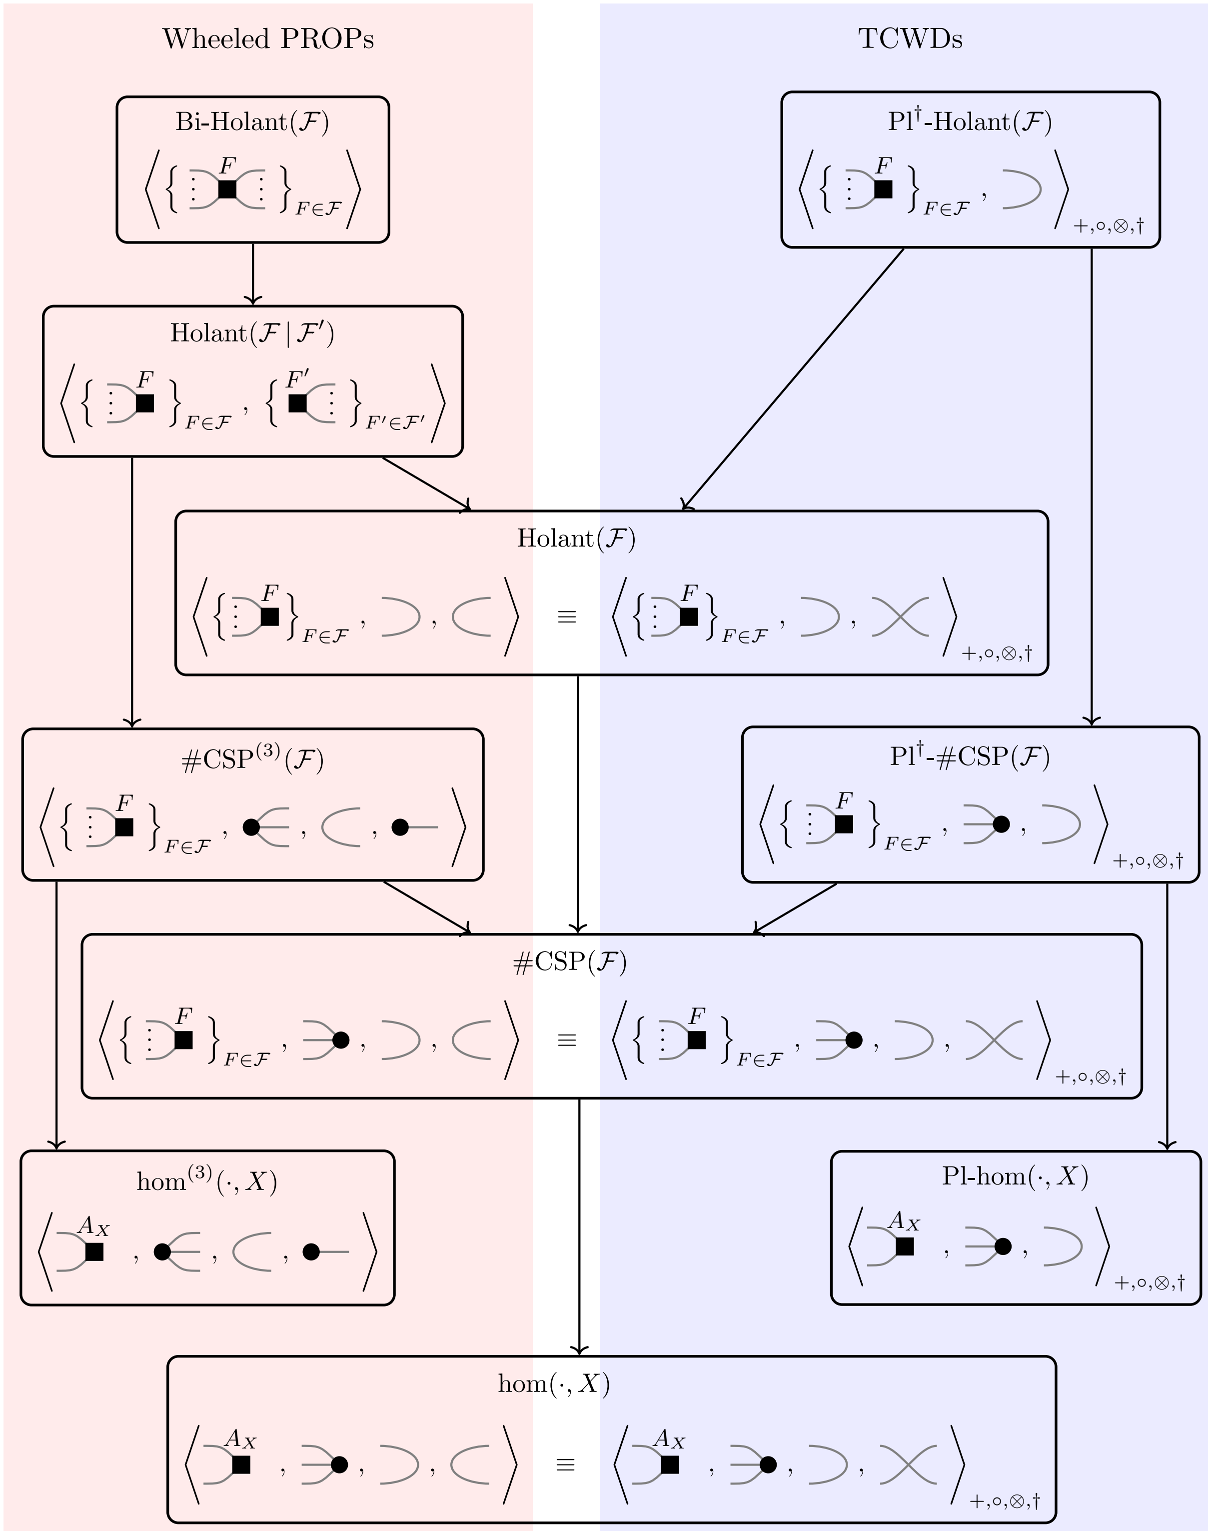

Figure 4.3: See Remark 4.4.1

Problem	Characterization	Location
#CSP($\mathcal{F}$)	Isomorphism	Theorem 3.2.1
$\mathrm{Pl}^{\dagger}$-#CSP($\mathcal{F}$)	Quantum isomorphism	Theorem 7.1.1
Holant($\mathcal{F}$)	Orthogonal transformation	Theorem 8.1.1
$\mathrm{Pl}^{\dagger}$-Holant($\mathcal{F}$)	Quantum orthogonal transformation	Theorem 9.0.1
Bi-Holant($\mathcal{F}$)	GL_q orbit closure intersection	Theorem 10.2.4
Holant($\mathcal{F} \mid \mathcal{F}'$)	Invertible transformation (if quantum-nonvanishing)	Theorem 10.3.2
$\#\mathrm{CSP}^{(d)}(\mathcal{F})$	GL_q orbit closure intersection with $\mathcal{EQ}_{\leq d}$	Corollary 10.5.5
$\hom^{(d)}(\cdot, X)$	GL_q orbit closure intersection with $\mathcal{EQ}_{\leq d}$	Corollary 10.5.6

Table 4.1: The various indistinguishability characterizations proved throughout this thesis, in chronological order. Theorem 10.3.2 is not a full characterization for Holant($\mathcal{F} \mid \mathcal{F}'$); the full characterization is a special case of that for Bi-Holant($\mathcal{F}$).

4.5 Orthogonal duality

In this section, we introduce two equivalent formulations – one in terms of symmetric TCWDs and another in terms of pivotal wheeled PROPs – of a *duality* theorem for real-valued Holant that underlies the main results of Chapter 5 and Chapter 8. We will later use two extensions of this theorem – one for general TCWDs and another for general wheeled PROPs – to prove the main results of Chapter 7 and Chapter 9 (on planar #CSP and Holant) and of Chapter 10 (on bipartite Holant), respectively.

The symmetric TCWD formulation of gadget duality is a classical version of Woronowicz's *Tannaka-Krein duality* [Wor88], as expressed by Banica and Speicher [BS09]. It states that every symmetric tensor category with duals is exactly the set of tensors invariant under some group of real orthogonal transformations.

Theorem 4.5.1 ([BS09, Theorem 1.3])**.** *The mapping $Q \mapsto \mathcal{V}(\mathbb{C}^q)^Q$ induces an inclusion-reversing bijection between subgroups $Q \subset O_q$ and symmetric tensor categories with duals.*

In particular, $\mathcal{V}(\mathbb{C}^q)^Q$ is a symmetric TCWD for any $Q \subset O_q$; this also follows from Corollary 4.4.1 and the fact that $[\mathcal{F}]$ is a symmetric TCWD (for conjugate-closed $\mathcal{F}$; otherwise also use that, since $H \in O_q$ is real-valued, $(H \cdot F)^\dagger = H \cdot F^\dagger$). The literature on TCWDs typically uses the *intertwiner space*

$$C_Q := \{F \in {}_\ell\mathcal{V}(\mathbb{C}^q)_r : \ell, r \geq 0, H^{\otimes \ell} F = F H^{\otimes r} \text{ for every } H \in Q\} \tag{4.5.1}$$

Observe that C_Q is exactly $\mathcal{V}(\mathbb{C}^q)^Q$; however, we will see the utility of the intertwiner formulation in Chapter 6.

Definition 4.5.1 (Stab_O)**.** For $\mathcal{F} \subset \mathcal{S}(\mathbb{C}^q)$, define the orthogonal stabilizer group

$$\mathrm{Stab}_O(\mathcal{F}) := \{H \in O_q : H \cdot F = F \text{ for every } F \in \mathcal{F}\}.$$

By (4.4.2), Stab_O is well-defined on sets of tensors, as it is equal to the orthogonal stabilizer group of their underlying signatures.

Corollary 4.5.1. *For any $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$, we have $\mathcal{V}(\mathbb{C}^q)^{\mathrm{Stab}_O(\mathcal{F})} = \langle \mathcal{F}, \supset, \times \rangle_{+,\circ,\otimes,\dagger}$.*

Proof. By Theorem 4.5.1, there is a $Q \subset O_q$ such that $\langle \mathcal{F}, \supset, \times \rangle_{+,\circ,\otimes,\dagger} = \mathcal{V}(\mathbb{C}^q)^Q$. Since $\mathcal{F} \subset \langle \mathcal{F}, \supset, \times \rangle_{+,\circ,\otimes,\dagger}$, we have $Q \subset \mathrm{Stab}_O(\mathcal{F})$. Suppose toward contradiction that $Q \subsetneq \mathrm{Stab}_O(\mathcal{F})$. Then, applying Theorem 4.5.1 again, $\mathcal{V}(\mathbb{C}^q)^{\mathrm{Stab}_O(\mathcal{F})}$ is a symmetric TCWD strictly smaller than $\mathcal{V}(\mathbb{C}^q)^Q = \langle \mathcal{F}, \supset, \times \rangle_{+,\circ,\otimes,\dagger}$. But $\mathcal{V}(\mathbb{C}^q)^{\mathrm{Stab}_O(\mathcal{F})}$ contains $\mathcal{F}$ and $\langle \mathcal{F}, \supset, \times \rangle_{+,\circ,\otimes,\dagger}$ is the smallest symmetric TCWD containing $\mathcal{F}$, a contradiction. □

The results of the last two sections combine into an algebraic characterization of quantum Holant $\mathcal{F}$-gadget signatures that generalizes a result of Regts [Reg12, Theorem 3] from edge coloring models. In Chapter 7 and Chapter 10, we will see analogous characterizations of Pl-Holant $\mathcal{F}$-gadgets and certain Bi-Holant $\mathcal{F}$-gadgets, respectively.

Theorem 4.5.2. *For any conjugate-closed $\mathcal{F} \subset \mathcal{S}(\mathbb{C}^q)$, we have $[\mathcal{F}] = \mathcal{V}(\mathbb{C}^q)^{\mathrm{Stab}_O(\mathcal{F})}$.*

Proof. Proposition 4.4.2 and Corollary 4.5.1 give $[\mathcal{F}] = \langle \mathcal{F}, \supset, \times \rangle_{+,\circ,\otimes,\dagger} = \mathcal{V}(\mathbb{C}^q)^{\mathrm{Stab}_O(\mathcal{F})}$. □

From the wheeled PROP perspective, Schrijver [Sch08b] studied tensors invariant under the action of the unitary group $U_q \subset \mathrm{GL}_q(\mathbb{C})$ on $\mathcal{V}(\mathbb{C}^q)$ and proved the following.

Theorem 4.5.3 ([Sch08b, Theorem 1 and Equation 13])**.** *Let $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$. Then there is a subgroup $Q \subset U_q$ with $\mathcal{F} = \mathcal{V}(\mathbb{C}^q)^Q$ if and only if $\mathcal{F}$ is a wheeled PROP closed under $\dagger$.*

We now give a corollary generalizing [Sch08b, Corollary 1e] from real to conjugate-closed complex sets via a similar derivation from Theorem 4.5.3.

Corollary 4.5.2. *Let $\mathcal{F} \subset \mathcal{S}(\mathbb{C}^q)$. Then there is a subgroup $Q \subset O_q$ with $\mathcal{F} = \mathcal{S}(\mathbb{C}^q)^Q$ if and only if $\mathcal{F}$ is a conjugate-closed pivotal wheeled PROP.*

Proof. Every pivotal wheeled PROP is closed under $\top$ (connect $\supset$ to every right input and $\subset$ to every left input), so every conjugate-closed pivotal wheeled PROP $\mathcal{F}$ is closed under $\dagger$. Therefore $\mathcal{F} = \mathcal{V}(\mathbb{C}^q)^Q$ for some $Q \subset U_q$ by Theorem 4.5.3. But $Q \subset O_q(\mathbb{C})$ by Proposition 2.4.1, and any matrix $H \in U_q \cap O_q(\mathbb{C})$ must be in O_q, because $H^\top = H^{-1} = H^\dagger$. □

By Proposition 4.4.1, Corollary 4.5.2 is equivalent to Theorem 4.5.1. The statement of [Sch08b, Corollary 1e] concerns (real) subalgebras $\mathcal{F}$ of $\mathcal{S}$ (signature sets) closed under arbitrary contraction. However, such shapeless $\mathcal{F}$ are just the underlying signatures of the pivotal wheeled PROP $[\mathcal{F}]$, and (4.4.2) shows that invariance of a tensor under $Q \subset O_q$ is independent of that tensor's shape (cf. [Sch08b, Equation 38]). Thus [Sch08b, Corollary 1e] is the real-valued case of Corollary 4.5.2.

Chapter 5

#CSP and Isomorphism: Intertwiners

In this chapter, based on the first half of [You25a], we use the orthogonal duality results in Section 4.5 to give an alternate proof of the #CSP indistinguishability Theorem 3.2.1 for $\mathbb{K} = \mathbb{C}$. This alternate *intertwiner* proof (recall Equation 4.5.1) is a 'classical' version of Mančinska and Roberson's [MR20] proof of their quantum/planar homomorphism indistinguishability theorem, generalized to #CSP. Mančinska and Roberson's proof uses *quantum permutation groups*, in particular the quantum automorphism group $\text{Qut}(X)$ of a graph X, analogous to the classical automorphism group $\text{Aut}(X)$. A key component of their proof is Woronowicz's *Tannaka-Krein duality* [Wor88], which implies that $\text{Qut}(X)$, a highly abstract object, is uniquely determined by its much more concrete *intertwiner space*, the tensors 'invariant' under $\text{Qut}(\mathcal{F})$. They then characterize the intertwiner space of $\text{Qut}(\mathcal{F})$ as exactly the span of the matrices of planar bi-labeled graphs, which in our terminology are quantum planar $\text{Holant}(A_X \cup \mathcal{EQ})$ gadget tensors. The orthogonal duality in Section 4.5 gives us the classical version of this characterization, for $\text{Aut}(F)$. Specifically, Theorem 4.5.2 characterizes quantum $\text{Holant}(\mathcal{F})$ gadget tensors as exactly the tensors invariant under $\text{Stab}_O(\mathcal{F})$, so, using the facts that $\#\text{CSP}(\mathcal{F}) \equiv \text{Holant}(\mathcal{F} \cup \mathcal{EQ})$ and that an orthogonal matrix stabilizing $\mathcal{EQ}$ must be a permutation matrix, we conclude that the quantum $\text{Holant}(\mathcal{F} \cup \mathcal{EQ})$ gadget tensors, corresponding to the signatures of k-labeled $\#\text{CSP}(\mathcal{F})$ instances, are exactly the tensors invariant under $\text{Aut}(\mathcal{F})$. Schrijver [Sch09] also proves a version of this characterization for polynomials invariant under $\text{Aut}(F)$ for a symmetric complex matrix (binary signature) F.

While the main result of this chapter is only a special case of Theorem 3.2.1, proved in Chapter 3, this chapter's intertwiner proof has a short, clean, combinatorial presentation that avoids the detailed notation of interpolation and the technical graph algebra calculations of [FLS07; Lov06; LS09; Sch09; GRV16]. Along the way, we prove results about the expressivity of labeled #CSP instance signatures that we will use in Section 7.4. In Chapter 7, we will present a quantum version of the intertwiner proof in this chapter, using many of the same ideas and constructions, giving a generalization to #CSP of Mančinska and Roberson's original quantum result. Finally, in Chapter 8, we further demonstrate the power of orthogonal duality for proving indistinguishability theorems when we apply Theorem 4.5.2 to prove the converse of the orthogonal Holant theorem for real-valued sets, without assuming the presence of $\mathcal{EQ}$.

5.1 From orthogonal transformation to isomorphism

Recall from Chapter 3 that, in a k-labeled #CSP instance, no variable has more than one label. In this chapter and the next, we relax this restriction: a #CSP instance is *k-multilabeled* if the labels $[k]$ are placed on $\leq k$ not necessarily distinct variables in $\mathbf{K}$. Define the signature of a multilabeled #CSP instance as in Definition 3.2.2; note that $Z^{\psi}(\mathbf{K}) = 0$ unless $\psi(\ell) = \psi(\ell')$ for every pair of labels ℓ, ℓ' on the same vertex. Multilabeled instances are more natural objects to identify with Holant gadgets, which do not forbid multiple dangling edges incident to the same vertex.

Proposition 5.1.1. *Let $\mathcal{F} \subset \mathcal{S}(\mathbb{K}^q)$. Then $[\mathcal{F}, =_3]$ is exactly the space of all quantum multilabeled #CSP($\mathcal{F}$) instance signatures.*

Proof. Connecting an input of $=_a$ with an input of $=_b$ yields a gadget with signature $=_{a+b-2}$. To construct any $=_n$ with $n \geq 3$, chain together $n-2$ copies of $=_3$. Construct $=_1$ by connecting two inputs of the same $=_3$, and $=_2$ by connecting $=_1$ and $=_3$ (or just as the signature of the ever-present wire). Thus

$$\mathcal{EQ} \subset [=_3]_{\text{Pl}}, \tag{5.1.1}$$

so $[\mathcal{F}, =_3] = [\mathcal{F}, \mathcal{EQ}]$. As discussed in Section 2.3, #CSP($\mathcal{F}$) instances I correspond uniquely to Holant($\mathcal{F} \mid \mathcal{EQ}$) signature grids $\Omega(I)$. Also, recall from (2.3.2) that the signatures in $[\mathcal{F}, \mathcal{EQ}]$ are exactly the signatures underlying the tensors in $\langle \mathcal{F} \mid \mathcal{EQ} \rangle$ because we may insert dummy $=_2$ vertices between adjacent $\mathcal{F}$ vertices and merge adjacent $\mathcal{EQ}$ vertices. By inserting more dummy $=_2$ vertices on any dangling edges incident to an $\mathcal{F}$ vertex, we may assume all dangling edges are incident to $\mathcal{EQ}$ vertices. Given a k-multilabeled #CSP($\mathcal{F}$) instance $\mathbf{K}$ with underlying unlabeled #CSP($\mathcal{F}$) instance I, construct a k-ary $(\mathcal{F}\,|\,\mathcal{EQ})$-gadget $\mathbf{K}'$ by starting with the Holant$(\mathcal{F}\,|\,\mathcal{EQ})$ grid $\Omega(I)$, then attach k dangling edges to $\Omega(I)$ such that the ith dangling edge is incident to the vertex whose corresponding variable in $\mathbf{K}$ is labeled with $i \in [k]$ (since we only care about the underlying signature, we do not specify whether these dangling edges face left or right). Comparing Definition 2.2.3 and Definition 3.2.2, we find

$$K(\mathbf{x}) = Z^{[n]\mapsto\mathbf{x}}(\mathbf{K}) = K'(\mathbf{x}),$$

where K and K' are the signatures of $\mathbf{K}$ and $\mathbf{K}'$, so $K = K'$. Conversely, given a k-ary Holant($\mathcal{F}, =_3$)-gadget $\mathbf{K}$, first construct the Holant$(\mathcal{F} \mid \mathcal{EQ})$-gadget $\mathbf{K}'$ with the same signature, then from $\mathbf{K}'$ construct a k-multilabeled #CSP($\mathcal{F}$) instance $\mathbf{K}''$ with label $i \in [n]$ on the variable whose corresponding vertex in $\mathbf{K}'$ is incident to the ith dangling edge. Similarly, the signatures of $\mathbf{K}$ and $\mathbf{K}''$ are equal. □

The next result is well-known, and is equivalent to the fact that a finite dimensional commutative algebra with a (multiplicative) unit has a basis of idempotents, as used in a similar context in [Lov06]. Like the original proof of Theorem 3.2.1 in Chapter 3, it exploits the entrywise product to perform Vandermonde interpolation.

Proposition 5.1.2 ([GRS25, Lemma 2.3])**.** *Let I be a finite index set and $V \subset \mathbb{K}^I$ be vector space closed under entrywise product $\bullet$ and containing the all-ones vector $\mathbb{1}_I$. Let $I = \bigsqcup_{\ell\in[s]} I_\ell$ be the partition of I into equivalence classes defined by relation $\sim$, where $i \sim i'$ iff $v(i) = v(i')$ for every $v \in V$. Define the indicator vectors $\mathbb{1}[I_\ell]$ by, for $i \in I$,*

$$\mathbb{1}[I_\ell](\mathbf{x}) = \begin{cases} 1 & i \in I_\ell \\ 0 & \textit{otherwise} \end{cases}.$$

Then $\{\mathbb{1}[I_\ell]\}_{\ell\in[s]}$ form a basis for V.

Proof. By definition of $\mathbb{1}[I_\ell]$, it suffices to show that each $\mathbb{1}[I_\ell] \in V$. There is a finite set $\{b^j\}_{j\in J}$ of vectors in V such that $i \sim i'$ iff $b^j(i) = b^j(i')$ for all $j \in J$. Define a matrix M with columns indexed by tuples $(p_j)_{j\in J}$ with each $0 \le p_j < |I|$ such that the column indexed by $(p_j)_{j\in J}$ is the vector $\bullet_{j\in J}(b^j)^{\bullet p_j} \in V$ (where $(b^j)^{\bullet 0} = \mathbb{1}_d$). Any two rows of M indexed by i, i' such that $i \sim i'$ are identical, so M contains s distinct rows. In any linear combination of the rows of M, we combine the coefficients of the rows indexed by each I_ℓ (as these rows are identical), so Lemma 3.2.1 concludes that the s distinct rows of M are linearly independent. The column span of M is contained in $\text{span}(\{\mathbb{1}[I_\ell]\}_{\ell\in[s]})$, so, since both of these spaces have dimension s, the reverse inclusion also holds. Thus $\text{span}(\{\mathbb{1}[I_\ell]\}_{\ell\in[s]}) \subset V$. □

Proposition 5.1.3. *An orthogonal matrix $H \in O_q(\mathbb{K})$ is a permutation matrix if and only if $H \cdot (=_3) = (=_3)$.*

Proof. By (5.1.1) and Corollary 4.4.1, any $H \in O_q(\mathbb{K})$ such that $H \cdot (=_3) = (=_3)$ satisfies $H \cdot \mathcal{EQ} = \mathcal{EQ}$. Xia [Xia10] showed using a standard Vandermonde interpolation argument that $H \cdot \mathcal{EQ} = \mathcal{EQ}$ if and only if H is a permutation matrix. □

We now obtain the #CSP corollary of Theorem 4.5.2: while quantum Holant $\mathcal{F}$-gadget signatures are exactly those invariant under $\text{Stab}_O(\mathcal{F})$, quantum multilabeled #CSP($\mathcal{F}$) instance signatures (equivalently, quantum Holant $[\mathcal{F}, =_3]$-gadget signatures) are exactly those invariant under $\text{Aut}(\mathcal{F}) \subset \text{Stab}_O(\mathcal{F})$. In the #CSP setting, we exploit the entrywise product, via Proposition 5.1.2, to remove the conjugate-closed assumption from Theorem 4.5.2. The result is a generalization of [MR20, Theorem 8.5] from graph homomorphism to #CSP, and a generalized variation of [Lov06, Lemma 2.5] and [CG21, Theorem 9.3] (see Remark 5.1.1).

Theorem 5.1.1. *For any $\mathcal{F} \subset \mathcal{S}(\mathbb{C}^q)$, we have $[\mathcal{F}, =_3] = \mathcal{V}(\mathbb{C}^q)^{\text{Aut}(\mathcal{F})}$.*

Proof. By Proposition 5.1.1 and the fact (3.2.1) that the n-labeled instance product induces entrywise product on the signatures, every ${}_n[\mathcal{F}, =_3]$ is closed under entrywise product (explicitly, given $F, G \in {}_n\mathcal{F}$, construct an n-ary gadget with signature $F \bullet G$ from F, G, and n copies of $(=_3)$ as follows: for the ith copy of $(=_3)$, connect one input to to ith input of F, another input to the ith input of G, and leave the third input dangling). Furthermore, the n-ary all-ones signature $\mathbb{1}_{[q]^n} = (=_1)^{\otimes n} \in {}_n[\mathcal{F}, =_3]$ for every n. Therefore Proposition 5.1.2 applies to the vector space ${}_n[\mathcal{F}, =_3] \subset \mathbb{K}^{[q]^n}$: every ${}_n[\mathcal{F}, =_3]$ has a basis of 0-1-valued signatures. Let $\mathcal{O}$ be the union of all such basis sets for every $n \geq 1$, so $[\mathcal{F}, =_3] = [\mathcal{O}, =_3]$, as the quantum gadgets form a graded vector space. Since $(\mathcal{O}, =_3)$ is conjugate-closed, Theorem 4.5.2 gives

$$[\mathcal{F}, =_3] = [\mathcal{O}, =_3] = (\mathcal{V}(\mathbb{C}^q))^{\text{Stab}_O(\mathcal{O}, =_3)}. \tag{5.1.2}$$

It also follows from $[\mathcal{F}, =_3] = [\mathcal{O}, =_3]$ and Corollary 4.4.1 that $\text{Stab}_O(\mathcal{O}, =_3) = \text{Stab}_O(\mathcal{F}, =_3)$, so, by (5.1.2) and Proposition 5.1.3,

$$[\mathcal{F}, =_3] = (\mathcal{V}(\mathbb{C}^q))^{\text{Stab}_O(\mathcal{O},=_3)} = (\mathcal{V}(\mathbb{C}^q))^{\text{Stab}_O(\mathcal{F},=_3)} = (\mathcal{V}(\mathbb{C}^q))^{\text{Aut}(\mathcal{F})}. \qquad \square$$

Definition 5.1.1 (n-orbitals)**.** For a permutation group $Q \subset S_q$, define the *n-orbitals* of Q as the equivalence classes of the relation $\sim_n$ on $[q]^n$ defined by $\mathbf{x} \sim_n \mathbf{y}$ if and only if there is a $\sigma \in Q$ such that $\sigma(\mathbf{x}) := (\sigma(x_1), \dots, \sigma(x_n)) = \mathbf{y}$.

If $\mathbf{x}$ and $\mathbf{y}$ are in the same n-orbital of $\text{Aut}(\mathcal{F})$, then every $F \in \mathcal{V}(\mathbb{C}^q)^{\text{Aut}(\mathcal{F})}$ satisfies $F(\mathbf{x}) = F(\mathbf{y})$. Conversely, if $\mathbf{x}$ and $\mathbf{y}$ are not in the same n-orbital of $\text{Aut}(\mathcal{F})$, then the indicator signature $\mathbb{1}[\text{Aut}(\mathcal{F})\,\mathbf{x}] \in \mathcal{V}(\mathbb{C}^q)^{\text{Aut}(\mathcal{F})}$ of the n-orbital of $\mathbf{x}$ satisfies $\mathbb{1}[\text{Aut}(\mathcal{F})\,\mathbf{x}](\mathbf{x}) = 1 \neq 0 = \mathbb{1}[\text{Aut}(\mathcal{F})\,\mathbf{x}](\mathbf{y})$. Hence

$$\mathbf{x}, \mathbf{y} \in [q]^n \text{ are in the same } n\text{-orbital of } \text{Aut}(\mathcal{F}) \iff F(\mathbf{x}) = F(\mathbf{y}) \text{ for every } F \in \mathcal{V}(\mathbb{C}^q)^{\text{Aut}(\mathcal{F})}. \tag{5.1.3}$$

We now obtain a version of Theorem 3.4.1 for $\mathcal{F} = \mathcal{G}$, and without domain weights or a twin-freeness assumption.

Corollary 5.1.1. *The following are equivalent for $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$ and $\mathbf{x}, \mathbf{y} \in [q]^n$.*

(1) $\mathbf{x}$ and $\mathbf{y}$ are in the same n-orbital of $\text{Aut}(\mathcal{F})$,

(2) $Z^{[n]\mapsto\mathbf{x}}(\mathbf{K}) = Z^{[n]\mapsto\mathbf{y}}(\mathbf{K})$ for every n-multilabeled #CSP($\mathcal{F}$) instance $\mathbf{K}$.

(3) $F(\mathbf{x}) = F(\mathbf{y})$ for every $F \in [\mathcal{F}, =_3]$.

Proof. The equivalence of (1) and (3) follows from Theorem 5.1.1 and (5.1.3). The equivalence of (2) and (3) follows from Proposition 5.1.1 and the fact that all instances/gadgets agree on $\mathbf{x}$ and $\mathbf{y}$ if and only if all quantum instances/gadgets agree on $\mathbf{x}$ and $\mathbf{y}$. □

In particular, with $I := [q]^n$ in Proposition 5.1.2, the 0-1-valued basis signatures in $\mathcal{O}$ in the proof of Theorem 5.1.1 are exactly the indicator signatures of the n-orbitals of $\text{Aut}(\mathcal{F})$ for every n.

Remark 5.1.1. This chapter's Theorem 5.1.1 differs from the graph homomorphism results [Lov06, Lemma 2.5] and [CG21, Theorem 9.3], and similarly Corollary 5.1.1 differs from the graph homomorphism result [Lov06, Lemma 2.4] (a special case of Theorem 3.4.1) in that the results of this chapter do not have any twin-freeness assumption. We avoid assuming twin-freeness in this chapter by using *multilabeled* #CSP($\mathcal{F}$) instances (($\mathcal{F}, =_3$)-gadgets) instead of labeled instances, which allow at most one label per vertex (equivalently, in the ($\mathcal{F}, =_3$)-gadget world, at most one dangling edge incident to each $\mathcal{EQ}$ vertex after combining adjacent $\mathcal{EQ}$ vertices to make a ($\mathcal{F} \mid \mathcal{EQ}$)-gadget). Cai and Govorov [CG21, Remark 5] state that the analogue of Theorem 5.1.1 does not hold for (singly) k-labeled graphs in the context of $\hom(\cdot, X) \equiv \#\mathrm{CSP}(A_X)$ if X is not twin-free. Indeed, consider the problem $\#\mathrm{CSP}\left(\left[\begin{smallmatrix}1&1\\1&1\end{smallmatrix}\right]\right)$. In any 2-labeled instance $\mathbf{K}$ with labels on distinct variables, no information is transmitted between the two labeled variables: every variable assignment yields the same value because every constraint takes value 1 regardless of the assignment. Thus $\mathbf{K}$ has matrix $\left[\begin{smallmatrix}a&a\\a&a\end{smallmatrix}\right]$ for some a. In particular, it is impossible to obtain I, which is nevertheless invariant under any automorphism, as the matrix of such a 2-labeled instance. However, I is trivially the matrix of the 2-multilabeled instance consisting of a single variable with two labels (it is important to note that Lovász and Szegedy [LS09, Theorem 1.4] show that any weighted graph X has a *contractor* – a 2-labeled instance with matrix I – but only after introducing vertex weights and moving to a smaller domain size to handle twins).

5.2 The direct sum

We now introduce a construction that will play a key role in the proofs of all remaining indistinguishability results in this thesis.

Definition 5.2.1 ($\oplus$)**.** For a bijective pair $\mathcal{F} \subset \mathcal{V}(\mathbb{K}^{V(\mathcal{F})})$ and $\mathcal{G} \subset \mathcal{V}(\mathbb{K}^{V(\mathcal{G})})$, define a set $\mathcal{F} \oplus \mathcal{G} = \{F \oplus G \mid \mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}\}$ bijective with $\mathcal{F}$ and $\mathcal{G}$, where, for $F \in {}_\ell\mathcal{V}(\mathbb{K}^{V(\mathcal{F})})_r$ and $G \in {}_\ell\mathcal{V}(\mathbb{K}^{V(\mathcal{G})})_r$, the *direct sum* $F \oplus G \in {}_\ell\mathcal{V}(\mathbb{K}^{V(\mathcal{F}) \sqcup V(\mathcal{G})})_r$ of F and G has underlying signature defined by

$$(F \oplus G)(\mathbf{x}) = \begin{cases} F(\mathbf{x}) & \mathbf{x} \in V(F)^{\ell+r} \\ G(\mathbf{x}) & \mathbf{x} \in V(G)^{\ell+r} \\ 0 & \text{otherwise.} \end{cases}$$

for $\mathbf{x} \in (V(F) \sqcup V(G))^{\ell+r}$.

For $\ell = r = 1$, $A_X \oplus A_Y$ is the adjacency matrix of the disjoint union of graphs X and Y. Generalizing the fact that homomorphism from a connected graph K to $X \oplus Y$ for connected graphs X and Y must map K entirely into X or Y, providing any input from $V(\mathcal{F})$ or $V(\mathcal{G})$ to a connected $\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle$-gadget forces all edges in the gadget take values in $V(\mathcal{F})$ or $V(\mathcal{G})$, respectively (all other edge assignments give 0).

Now we introduce *subdomain restriction*, which will play a significant role in Chapters 8 and 10.

Definition 5.2.2 ($F|_X$,$\mathcal{F}|_X$,$K|_\mathcal{F}$)**.** For $F \in {}_\ell\mathcal{V}(\mathbb{K}^q)_r$ and $X \subset [q]$, define $F|_X \in {}_\ell\mathcal{V}(\mathbb{K}^X)_r$ to be the subtensor of F induced by X. Define $\mathcal{F}|_X := \{F|_X \mid F \in \mathcal{F}\}$, a (multiset) bijective with $\mathcal{F}$.

For $K \in \langle\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle\rangle$ or $[[\mathcal{F}] \oplus [\mathcal{G}]]$, use $K|_\mathcal{F}$ as shorthand for $K|_{V(\mathcal{F})}$.

The following useful statement, which says roughly that an $\mathcal{F} \oplus \mathcal{G}$ gadget behaves like an $\mathcal{F}$ gadget on inputs from $V(\mathcal{F})$ and like the corresponding $\mathcal{G}$ gadget on inputs from $V(\mathcal{G})$, also holds for $[\cdot]$ and its planar variants in place of $\langle\cdot\rangle$.

Proposition 5.2.1. *If $K \in \langle\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle\rangle$, then $\langle\mathcal{F}\rangle \ni K|_\mathcal{F} \leftrightsquigarrow K|_\mathcal{G} \in \langle\mathcal{G}\rangle$.*

Proof. By definition, K is the signature of some quantum $\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle$-gadget $\mathbf{K}$. Any connected component of a term of $\mathbf{K}$ without a dangling edge contributes only a constant factor; by absorbing this factor into the term's coefficient, we may assume no term of $\mathbf{K}$ has such a component. To construct $K|_\mathcal{F}$, restrict all inputs to $\mathbf{K}$ to $V(\mathcal{F})$. As discussed above, this restricts all edges of all gadgets composing $\mathbf{K}$ to $V(\mathcal{F})$. Thus $K|_\mathcal{F}$ is the signature of $\mathbf{K}_{\langle\mathcal{F}\rangle\oplus\langle\mathcal{G}\rangle\to\langle\mathcal{F}\rangle}$. Similarly, $\mathbf{K}_{\langle\mathcal{F}\rangle\oplus\langle\mathcal{G}\rangle\to\langle\mathcal{G}\rangle}$ has signature $K|_\mathcal{G}$, and the result follows. □

Definition 5.2.3 (Connected)**.** Say domain elements $y, z \in [q]$ are *adjacent* in $\mathcal{F} \subset \mathcal{S}(\mathbb{K}^q)$ if there exists an $F \in \mathcal{F}$ and a tuple $\mathbf{x} \in [q]^{\text{arity}(F)}$ containing y and z such that $F(\mathbf{x}) \neq 0$. Then say y and z are *connected* in $\mathcal{F}$ if y and z are in the same relation of the reflexive and transitive closure $\sim_c$ of the adjacency relation.

Say $\mathcal{F}$ is *connected* if $\sim_c$ has a single equivalence class (*connected component*), and is *disconnected* otherwise.

If $\mathcal{F} = \{A_X\}$ for graph X, then these definitions specialize to the familiar concepts for graphs. The following proposition also follows the same intuition as the graph special case.

Proposition 5.2.2. *If $I, J \subset [q]$ are distinct connected components of $\mathcal{F}$ and $\sigma \in \mathrm{Aut}(\mathcal{F})$ satisfies $\sigma(i) = j$ for some $i \in I$ and $j \in J$, then σ is an isomorphism between $\mathcal{F}|_I$ and $\mathcal{F}|_J$.*

Proof. It suffices to show that $\sigma(I) = J$. If any index tuple $\mathbf{x} \in [q]^n$ contains an element of I (resp. J) and an element of $[q]\setminus I$ (resp. $[q]\setminus J$), then $F(\mathbf{x}) = 0$ for every n-ary $F \in \mathcal{F}$. Let $i \neq i' \in I$. There is a sequence of tuples $\mathbf{x}^1, \ldots, \mathbf{x}^p \in \bigcup_n I^n$ such that $\mathbf{x}^1$ contains i, $\mathbf{x}^p$ contains i', and every consecutive pair $\mathbf{x}^\ell, \mathbf{x}^{\ell+1}$ share some element of I. Thus $\sigma(i) \in J$ inductively gives $\sigma(\mathbf{x}^\ell) \in \bigcup_n J^n$ for every ℓ, so $\sigma(i') \in J$. Hence $\sigma(I) \subset J$. Applying similar reasoning to $\sigma^{-1} \in \mathrm{Aut}(\mathcal{F})$ gives $\sigma^{-1}(J) \subset I$, so $J \subset \sigma(I)$. □

In particular, if $\mathcal{F}$ and $\mathcal{G}$ are connected, then $V(\mathcal{F})$ and $V(\mathcal{G})$ are the two connected components of $V(\mathcal{F} \oplus \mathcal{G}) = V(\mathcal{F}) \sqcup V(\mathcal{G})$, so if there is a $\sigma \in \mathrm{Aut}(\mathcal{F} \oplus \mathcal{G})$ mapping some $x \in V(\mathcal{F})$ to some $y \in V(\mathcal{G})$, then $\mathcal{F} \cong \mathcal{G}$.

We now have all the tools needed to prove Theorem 3.2.1 for $\mathbb{K} = \mathbb{C}$.

Theorem 5.2.1. *Let $\mathcal{F}$, $\mathcal{G}$ be $\mathbb{C}$-valued signature sets. Then $\mathcal{F} \cong \mathcal{G}$ if and only if $\mathcal{F}$ and $\mathcal{G}$ are* #CSP*-indistinguishable.*

Proof. We only need the ($\Longleftarrow$) direction. Suppose $\mathcal{F}$ and $\mathcal{G}$ are #CSP-indistinguishable. Let 0_F and 0_G be new domain elements. For each $F \in \mathcal{F}$, $G \in \mathcal{G}$ of arity $n \geq 2$, define signatures F' and G' on $V(\mathcal{F}') := V(\mathcal{F}) \cup \{0_F\}$ and $V(\mathcal{G}') := V(\mathcal{G}) \cup \{0_G\}$ by, for $\mathbf{x} \in V(\mathcal{F}')^n$ and $\mathbf{x} \in V(\mathcal{G}')^n$, respectively,

$$F'(\mathbf{x}) = \begin{cases} F(\mathbf{x}) & \mathbf{x} \in V(\mathcal{F})^n \\ 1 & \text{otherwise} \end{cases}, \qquad G'(\mathbf{x}) = \begin{cases} G(\mathbf{x}) & \mathbf{x} \in V(\mathcal{G})^n \\ 1 & \text{otherwise} \end{cases}$$

In other words, if any entry of $\mathbf{x}$ is 0_F, then $F'(\mathbf{x}) = 1$, and similarly for G'. For unary $F \in \mathcal{F}$, $G \in \mathcal{G}$, define *binary* F' and G' by, for $x, y \in V(\mathcal{F}')$ and $x, y \in V(\mathcal{G}')$, respectively,

$$F'(x,y) = \begin{cases} F(x) & x = y \in V(\mathcal{F}) \\ 1 & x = 0_F \text{ or } y = 0_F \\ 0 & \text{otherwise} \end{cases}, \qquad G'(x,y) = \begin{cases} G(x) & x = y \in V(\mathcal{G}) \\ 1 & x = 0_G \text{ or } y = 0_G \\ 0 & \text{otherwise.} \end{cases}$$

The arity increase is necessary because connectivity only makes sense for signatures with arity at least two. Now let $\mathcal{F}' = \{F' \mid F \in \mathcal{F}\}$ and $\mathcal{G}' = \{G' \mid G \in \mathcal{G}\}$.

Let $\mathbf{K} = (V, C)$ be a 1-labeled #CSP($\mathcal{F}' \oplus \mathcal{G}'$) instance, with labeled variable $v_0 \in V$. We will show that

$$Z^{0 \mapsto 0_F}(\mathbf{K}) = Z^{0 \mapsto 0_G}(\mathbf{K}). \tag{5.2.1}$$

If $\mathbf{K}$ is not connected (i.e. the underlying graph of the Holant($\mathcal{F}, \mathcal{EQ}$) signature grid corresponding to $\mathbf{K}$ is not connected), then the components of $\mathbf{K}$ that do not contain v_0 contribute the same value to the partition regardless of the assignment to v_0. Hence, to establish (5.2.1), we may assume $\mathbf{K}$ is connected. Therefore, as in Proposition 5.2.1, if $\phi : V \to V(\mathcal{F}' \oplus \mathcal{G}')$ satisfies $\phi(v_0) = 0_F$, then, since each $F' \oplus G'$ evaluates to 0 unless all its inputs are in $V(\mathcal{F}')$ or all its inputs are in $V(\mathcal{G}')$, we must have $\phi(V) \subset V(\mathcal{F}')$ for any ϕ contributing a nonzero value to $Z^{0 \mapsto 0_F}(\mathbf{K})$.

Any ϕ with $\phi(v_0) = 0_F$ maps some $S \subset V$ to 0_F, with $v_0 \in S$. For a fixed $S \subset V$, the remaining variables $V \setminus S$ take all values in $V(\mathcal{F}') \setminus \{0_F\} = V(\mathcal{F})$ as ϕ ranges over $\{\phi \mid \phi^{-1}(0_F) = S\}$. Additionally, any constraint containing a variable in S always evaluates to 1, regardless of the assignments to the other variables. Construct a #CSP($\mathcal{F}$) instance $\mathbf{K}^{\mathcal{F}}_{V \setminus S}$ from $\mathbf{K}$ as follows. First eliminate all variables in S and all constraints containing any variable in S. Then, for each constraint applying $F' \oplus G'$, if F and G have arity > 1, replace $F' \oplus G'$ with F. If F and G are unary, then merge the two variables to which the binary $F' \oplus G'$ is applied and replace the constraint with a constraint applying F to the merged variable. Assuming all inputs to $F' \oplus G'$ are in $V(\mathcal{F})$, this variable merging procedure does not change the value of the partition function, since by construction $F' \oplus G'$ acts as the function $(x, y) \mapsto \delta_{xy} F_x$. Now the total contribution to $Z^{0 \mapsto 0_F}(\mathbf{K})$ of the assignments ϕ satisfying $\phi^{-1}(0_F) = S$ is $Z(\mathbf{K}^{\mathcal{F}}_{V \setminus S})$. Thus

$$Z^{0 \mapsto 0_F}(\mathbf{K}) = \sum_{S \subset V, S \ni v_0} Z(\mathbf{K}^{\mathcal{F}}_{V \setminus S}).$$

A similar expression holds for $Z^{0 \mapsto 0_G}(\mathbf{K})$, with the #CSP($\mathcal{G}$) instance $\mathbf{K}^{\mathcal{G}}_{V \setminus S} = (\mathbf{K}^{\mathcal{F}}_{V \setminus S})_{\mathcal{F} \to \mathcal{G}}$ in place of $\mathbf{K}^{\mathcal{F}}_{V \setminus S}$. So, by assumption,

$$Z^{0 \mapsto 0_F}(\mathbf{K}) = \sum_{S \subset V, S \ni v_0} Z(\mathbf{K}^{\mathcal{F}}_{V \setminus S}) = \sum_{S \subset V, S \ni v_0} Z(\mathbf{K}^{\mathcal{G}}_{V \setminus S}) = Z^{0 \mapsto 0_G}(\mathbf{K}),$$

proving (5.2.1). Now by Corollary 5.1.1 with $n = 1$, there is a $\sigma \in \text{Aut}(\mathcal{F}' \oplus \mathcal{G}')$ satisfying $\sigma(0_F) = 0_G$. By construction, 0_F is adjacent to every $x \in V(\mathcal{F})$, and 0_G is adjacent to every $y \in V(\mathcal{G})$, so $\mathcal{F}'$ and $\mathcal{G}'$ are connected. Therefore, Proposition 5.2.2 asserts that $\sigma|_{V(\mathcal{F}')}$ is an isomorphism between

$\mathcal{F}'$ and $\mathcal{G}'$. But $\sigma(0_F) = 0_G$, so $\sigma\mid_{V(\mathcal{F})}$ is an isomorphism between $\mathcal{F}$ and $\mathcal{G}$ (if F and G are unary, then $\sigma\mid_{V(\mathcal{F})}$ is really an isomorphism between the functions $(x,y) \mapsto \delta_{xy}F_y$ and $(x,y) \mapsto \delta_{xy}G_y$, but this implies an isomorphism between F and G, since unary functions are isomorphic if and only if they have the same multiset of entries). □

The proof of Theorem 5.2.1 is a generalization of Lovász's proof of [Lov06, Corollary 2.6], which is essentially Theorem 5.2.1 restricted to real-weighted graph homomorphism. Both proofs use the idea of adding a universal vertex to connect the graph/signature, since for weighted such objects we cannot take the complement to assume connectedness.

Remark 5.2.1. By Corollary 4.4.1 and Proposition 5.1.3, $\mathcal{F} \cong \mathcal{G} \iff [\mathcal{F}, =_3] \cong [\mathcal{G}, =_3]$. Furthermore, since every Holant($[\mathcal{F}, =_3], =_3$) grid is, by gadget substitution, a quantum Holant($\mathcal{F}, =_3$) grid, $\mathcal{F}$ and $\mathcal{G}$ are #CSP-indistinguishable if and only if $[\mathcal{F}, =_3]$ and $[\mathcal{G}, =_3]$ are #CSP-indistinguishable. Therefore, to prove Theorem 5.2.1, we may replace $\mathcal{F}$ and $\mathcal{G}$ with $[\mathcal{F}, =_3]$ and $[\mathcal{G}, =_3]$. Thereby we may assume that $\mathcal{F}$ and $\mathcal{G}$ are connected without adding a universal domain element, as $[\mathcal{F}, =_3]$ and $[\mathcal{G}, =_3]$ contain the connected all-ones signatures $(=_1)^{\otimes n}$ of every arity. Then one could alternatively apply the techniques of [MR20] to prove Theorem 5.2.1. We will take this approach for planar #CSP in Chapter 7.

Chapter 6

Planarity, Quantum Matrices, and Quantum Groups

This chapter is mostly based on [CY24] (joint work with Jin-Yi Cai), with improvements and new content in Section 6.2 Section 6.3.

Quantum isomorphism of (undirected, unweighted) graphs, introduced in [Ats+19], is a relaxation of classical isomorphism. Graphs X and Y are *quantum isomorphic* if there is a perfect winning strategy in a two-player *graph isomorphism game* in which the players share and can perform measurements on an entangled quantum state. This condition is equivalent to the existence of a *quantum permutation matrix* U – a relaxation of a permutation matrix whose entries do not necessarily commute – satisfying $UA_X = A_Y U$, or equivalently $U^{\otimes 2}(A_X)^{2,0} = (A_Y)^{2,0}$ [LMR17]. Analogously to classical isomorphism, we say n-ary signatures F and G are *quantum isomorphic* if there is a quantum permutation matrix U satisfying $U^{\otimes n}F = G$.

Mančinska and Roberson [MR20] proved that two graphs are quantum isomorphic if and only if they are homomorphism-indistinguishable over all planar graphs. In Chapter 7, we generalize this result to #CSP and sets of constraint functions. Then, in Chapter 9, we generalize it further to show that $\mathcal{F}$ and $\mathcal{G}$ are planar-Holant-indistinguishable if and only if there is a quantum orthogonal matrix transforming $\mathcal{F}$ to $\mathcal{G}$. In this chapter, we lay the groundwork for these quantum/planar indistinguishability results.

First, in Section 6.2, we introduce two new variants of the standard definition of planar Holant, motivated by the need for the space of quantum planar Holant gadgets to be a TCWD. Next, in Section 6.3, we prove that every planar gadget is generated by simple building block gadgets under the TCWD operations $\otimes, \circ, \dagger$. This planar gadget decomposition generalizes Mančinska and Roberson's similar decomposition of planar bi-labeled graphs (Holant($A_X \cup \mathcal{EQ}$) gadgets for counting homomorphisms to X), but the possible asymmetry of our signatures presents new subtleties beyond the graph homomorphism special case. As a consequence of the planar gadget decomposition, we obtain the quantum orthogonal Holant theorem, a generalization of the orthogonal Holant theorem (Corollary 2.4.1) to quantum orthogonal transformations. Such a quantum holographic transformation does not work on general signature grids, since viewing U itself as a signature in a signature grid is not in general well-defined, as U's entries do not commute and the definition of the Holant value does not specify an order to multiply the signature evaluations. However, the planarity of the signature grid and the resulting gadget decomposition and matrix product expression for the Holant value implicitly provide a multiplication order. The success of the quantum holographic transformation for asymmetric signatures is also dependent on the fact that the holographic transformation action of a quantum permutation matrix is invariant under signature rotations and (conjugate) reflections. The asymmetry and rotation and reflection issues are only relevant in the context of planar signature grids, since in nonplanar grids, one can simply cross and twist the incident edges to achieve the desired input order. Hence this is another interesting connection between quantum permutation matrices and the structural properties of planar graphs.

Finally, in Section 6.4 we develop a quantum/planar version of the orthogonal duality of Section 4.5. This is based on the theory of the *quantum orthogonal group* [Wan95], a compact matrix quantum group [Wor87] whose behavior mimics that of the real orthogonal group O_q. Specifically, we use Woronowicz's Tannaka-Krein duality [Wor88] to prove a quantum/planar analogue (Theorem 6.4.2) of Theorem 4.5.2: quantum planar Holant($\mathcal{F}$) gadget tensors are exactly the tensors 'invariant' under the quantum orthogonal stabilizer group of $\mathcal{F}$.

6.1 Preliminaries

6.1.1 The graph isomorphism game

We begin this chapter with a short exposition of the *graph isomorphism game*, introduced in [Ats+19]. The graph isomorphism game is a two-player *nonlocal game* played between a verifier and two cooperating players, Alice and Bob. In a general nonlocal game, the verifier gives Alice and Bob inputs x_A and x_B from finite sets X_A and X_B, respectively. After receiving their inputs, Alice and Bob respond with outputs y_A and y_B from finite sets Y_A and Y_B, respectively. The verifier then determines from x_A, x_B, y_A, y_B whether Alice and Bob have won or lost. We are concerned with whether Alice and Bob have a *perfect strategy* - a strategy that wins the game with probability 1. Alice and Bob can agree on a predetermined strategy, but cannot communicate after receiving their inputs. However, if allowed a *quantum* strategy, they can make measurements on a shared quantum state after receiving their inputs. In particular, we consider *quantum commuting strategies*. In this framework, Alice and Bob share a state in a possibly infinite-dimensional Hilbert space H and Alice's measurement operators – positive bounded linear operators on H – must commute with Bob's measurement operators (see [Ats+19] for background on quantum measurements).

The graph isomorphism game concerns two unweighted graphs F and G. We take $X_A = X_B = Y_A = Y_B = V(F) \cup V(G)$. Let $f_A = \{x_A, y_A\} \cap V(F)$ and $g_A = \{x_A, y_A\} \cap V(G)$, and define f_B and g_B similarly for Bob. f_A and g_A are well-defined (assuming the players win) given condition (i) below. The players win if and only if the following three conditions are satisfied:

(i) $x_A \in V(F) \iff y_A \in V(G)$ and $x_B \in V(F) \iff y_B \in V(G)$;

(ii) $f_A = f_B \iff g_A = g_B$;

(iii) $(f_A, f_B) \in E(F) \iff (g_A, g_B) \in E(G)$.

The three conditions ensure that the maps between each player's input and output vertices define isomorphisms between F and G, so that Alice and Bob have a perfect classical strategy if and only if $F \cong G$ [Ats+19]. Say F and G are *quantum isomorphic*, denoted $F \cong_{qc} G$, if Alice and Bob have a perfect quantum commuting strategy for the (F, G)-isomorphism game. Examples of non-isomorphic but quantum isomorphic graphs were exhibited in [Ats+19]; more such examples were later found by Chan and Martin [CM24b].

Mančinska and Roberson [MR20] showed that $F \cong_{qc} G$ if and only if F and G are Pl-hom-indistinguishable. In Section 7.3, we generalize the graph isomorphism game to $\mathbb{C}$-valued matrices F and G, and show that $F \cong_{qc} G$ if and only if F and G are Pl-#CSP-indistinguishable. Ideally, we would also show equivalence between Pl-#CSP-indistinguishability of n-ary signatures F and G and the existence of a perfect strategy for some n-player tensor isomorphism game. However, there appear to be fundamental obstacles to proving such a result (see Remark 7.3.1). Instead, we prove equivalence between Pl-#CSP-indistinguishability of any signature sets $\mathcal{F}$ and $\mathcal{G}$ and a different definition of quantum isomorphism – introduced in the next section – that, for binary signatures, agrees with the nonlocal game definition.

6.1.2 Quantum matrices and quantum tensor isomorphism

The first step to connecting quantum isomorphism and homomorphisms from planar graphs is an alternative characterization of quantum isomorphism in terms of a *quantum permutation matrix*. Such a matrix has nonscalar entries from a unital C^*-algebra. We will not formally define a C^*-algebra here; it suffices for our purposes to say that it is an algebra over $\mathbb{C}$ with a multiplicative unit $\mathbf{1}$ an operator $*$ satisfying $(a^*)^* = a$, $(a+b)^* = a^* + b^*$, $(ab)^* = b^*a^*$, and $(ca)^* = \overline{c}a^*$ for every a, b in the C^*-algebra and $c \in \mathbb{C}$. The algebra of bounded linear operators on some Hilbert space, with adjoint $*$ (in the finite-dimensional case, complex matrices with conjugate transpose), is a prototypical example of a C^*-algebra; for example, the entries of the quantum permutation matrix defining quantum isomorphism are Alice and Bob's measurement operators from a quantum strategy for the graph isomorphism game [Ats+19]. Critically, a C^*-algebra is not necessarily commutative.

Definition 6.1.1 (Quantum matrix of self-adjoints)**.** A *quantum matrix of self-adjoints* $U = (u_{ij})_{i,j\in[q]}$ is a matrix with entries from a unital C^*-algebra whose entries are self-adjoint: $u_{ij}^* = u_{ij}$.

Define the Kronecker product $\otimes$ for quantum matrices U, V as for scalar-valued matrices: $(U \otimes V)_{x_1x_2,y_1y_2} = u_{x_1y_1}v_{x_2y_2}$. Since the entries of U and V do not necessarily commute, $u_{x_1y_1}v_{x_2y_2}$ may be distinct from $v_{x_2y_2}u_{x_1y_1}$. Nevertheless, many natural properties of scalar-valued matrices carry over to quantum matrices. It follows from direct calculation that the familiar mixed-product property still holds for quantum matrices U, V and scalar-valued matrices M and N:

$$(U \otimes V)(M \otimes N) = (UM) \otimes (VN), \qquad (M \otimes N)(U \otimes V) = (MU) \otimes (NV). \tag{6.1.1}$$

Similarly, since the entries of a quantum matrix commute with scalars, we have $(UM)^\top = M^\top U^\top$. Furthermore, $(U^{\otimes n})^\top$ and $(U^\top)^{\otimes n}$ both have $\mathbf{x}, \mathbf{y}$-entry equal to $u_{y_1 x_1} \dots u_{y_n x_n}$, so

$$(U^{\otimes n})^\top = (U^\top)^{\otimes n}.$$

Finally, associativity of the product of quantum matrices follows from associativity of the underlying C^*-algebra. However, noncommutativity invalidates the mixed-product property for two pairs U, V and U', V' of quantum matrices:

$$\begin{aligned}
((U \otimes V)(U' \otimes V'))_{x_1x_2,y_1y_2} &= \sum_{z_1z_2} u_{x_1z_1} v_{x_2z_2} u'_{z_1y_1} v'_{z_2y_2} \\
&\neq \sum_{z_1z_2} u_{x_1z_1} u'_{x_2z_2} v_{z_1y_1} v'_{z_2y_2} \\
&= (UU')_{x_1y_1} (VV')_{x_2y_2} = ((UU') \otimes (VV'))_{x_1x_2,y_1y_2}
\end{aligned}$$

so in general $(U \otimes V)(U' \otimes V') \neq (UU') \otimes (VV')$.

Quantum matrices act as transformations on scalar-valued signatures via the same formula as scalar-valued matrices: for $q \times q$ quantum matrix U and n-ary $F \in \mathcal{S}(\mathbb{C}^q)$, define $U \cdot F := U^{\otimes n} F$. Here we view F as a length-q^n column vector, so $U \cdot F$ is a column vector whose entries are not necessarily scalars. However, we usually consider the case where $U \cdot F = G \otimes \mathbf{1}$ for scalar-valued $G \in \mathcal{S}(\mathbb{C}^q)$. In this case, we for brevity write $U \cdot F = G$, as $G \otimes \mathbf{1}$ behaves identically to G (in particular, its entries commute with all entries of U).

Definition 6.1.2 (Quantum orthogonal matrix)**.** A *quantum orthogonal matrix* U is a quantum matrix of self-adjoints satisfying the following equivalent conditions:

(1) $U^{\otimes 2} \supset = \supset$ and $\subset U^{\otimes 2} = \subset$.

(2) $UU^\top = I \otimes \mathbf{1} = U^\top U$.

If U is defined over the C^*-algebra $\mathbb{C}$, then $U \in O_q(\mathbb{C})$ is an ordinary orthogonal matrix. The equivalence between the two conditions follows the same direct calculation as Proposition 2.4.1

(taking some care now to account for noncommutativity). A quantum orthogonal matrix U has an inverse $U^\top$. However, since in general $U^{\otimes n}(U^\top)^{\otimes n} \neq (UU^\top)^{\otimes n}$ by the failure of the mixed-product property, $(U^\top)^{\otimes n} = (U^{\otimes n})^\top$ is not the inverse of $U^{\otimes n}$. Instead, for a quantum matrix V, define $V^\dagger$ to be the matrix whose (x,y)-entry is v_{yx}^*. Since the entries of quantum orthogonal U are self-adjoint, we have $U^\dagger = U^\top$. However, the entries of $U^{\otimes n}$ may not be self-adjoint due to the order-reversing property of $*$ (technically, this means $U^{\otimes n}$ is not a quantum orthogonal matrix). As a result, in general

$$(U^{\otimes n})^\dagger \neq (U^{\otimes n})^\top = (U^\top)^{\otimes n} = (U^\dagger)^{\otimes n}$$

and it turns out that $(U^{\otimes n})^\dagger$ is the inverse of $U^{\otimes n}$:

$$\begin{aligned}
(U^{\otimes n}(U^{\otimes n})^\dagger)_{\mathbf{x},\mathbf{y}} &= \sum_{\mathbf{z}} u_{x_1z_1}\dots u_{x_nz_n}(u_{y_1z_1}\dots u_{y_nz_n})^* \\
&= \sum_{\mathbf{z}} u_{x_1z_1}\dots u_{x_nz_n}u_{y_nz_n}\dots u_{y_1z_1} \\
&= \sum_{z_1,\dots,z_{n-1}} u_{x_1z_1}\dots u_{x_{n-1}z_{n-1}}\sum_{z_n}(u_{x_nz_n}u_{y_nz_n})u_{y_{n-1}z_{n-1}}\dots u_{y_1z_1} \\
&= \sum_{z_1,\dots,z_{n-1}} u_{x_1z_1}\dots u_{x_{n-1}z_{n-1}}(UU^\top)_{x_ny_n}u_{y_{n-1}z_{n-1}}\dots u_{y_1z_1} \\
&= \delta_{x_n,y_n}\mathbf{1}\sum_{z_1,\dots,z_{n-2}} u_{x_1z_1}\dots u_{x_{n-2}z_{n-2}}\sum_{z_{n-1}}(u_{x_{n-1}z_{n-1}}u_{y_{n-1}z_{n-1}})u_{y_{n-2}z_{n-2}}\dots u_{y_1z_1} \\
&= \dots = \delta_{\mathbf{x},\mathbf{y}}\mathbf{1},
\end{aligned}$$

and similarly for $(U^{\otimes n})^\dagger U^{\otimes n}$, so

$$U^{\otimes n}(U^{\otimes n})^\dagger = I^{\otimes n} \otimes \mathbf{1} = (U^{\otimes n})^\dagger U^{\otimes n}. \tag{6.1.2}$$

With an inverse in hand, we extend the action of a quantum orthogonal matrix from signatures to tensors, paralleling the action of GL:

Definition 6.1.3 ($U \cdot F$)**.** For $F \in {}_\ell\mathcal{V}(\mathbb{C}^q)_r$, viewed as a $q^\ell \times q^r$ complex matrix, define $U \cdot F := U^{\otimes \ell}F(U^{\otimes r})^\dagger$.

Again, $U \cdot F$ may not have scalar entries, but if $U \cdot F = G \otimes \mathbf{1}$ for $G \in {}_\ell\mathcal{V}(\mathbb{C}^q)_r$, then U transforms F to G. In this case, it is useful to take the intertwiner view

$$U \cdot F = G \iff U^{\otimes \ell}F = GU^{\otimes r}. \tag{6.1.3}$$

Recall from Proposition 5.1.3 that a classical permutation matrix is an orthogonal matrix stabilizing $=_3$. We use this perspective to define a *quantum* permutation matrix.

Definition 6.1.4 (Quantum permutation matrix)**.** A quantum permutation matrix U is a quantum orthogonal matrix satisfying the following two equivalent conditions:

(1) $U \cdot (=_3) = (=_3)$

(2) For every i, j, k, $u_{ij}u_{kj} = \delta_{ik}u_{ij}$ and $u_{ij}u_{ik} = \delta_{jk}u_{ij}$.

Item (2), which states that each row and column of U consist of mutually orthogonal projectors, is directly equivalent to $U^{\otimes 2}(=_3)^{2,1} = (=_3)^{2,1}U$, which, by (6.3.3) below, is, for quantum orthogonal U, equivalent to $U \cdot (=_3) = (=_3)$. Recall from Proposition 5.1.3 that a scalar-valued orthogonal matrix preserves $=_3$ if and only if it is a permutation matrix. Thus a quantum permutation matrix over the C^*-algebra $\mathbb{C}$ is an ordinary permutation matrix (alternatively, a scalar-valued matrix satisfying item (2) has 0-1 entries and at most one 1 entry per row and column; any orthogonal such matrix is a permutation matrix).

The following concept was introduced in [Ats+19], and expressed in this form in [LMR17]. It is motivated by the fact that graphs X and Y with adjacency matrices A_X and A_Y are isomorphic iff there is a permutation matrix P such that $PA_X = A_Y P$.

Definition 6.1.5 (Graph $\cong_{qc}$)**.** Graphs X and Y with adjacency matrices A_X and A_Y are *quantum isomorphic* ($X \cong_{qc} Y$) if there is a quantum permutation matrix U satisfying $UA_X = A_Y U$.

By (6.3.3) below, we may, as in the classical case, rewrite $UA_X = A_Y U$ as $U \cdot (A_X)^{2,0} = (A_Y)^{2,0}$. This perspective leads to a natural generalization of quantum isomorphism to higher-arity signatures.

Definition 6.1.6 ($\cong_{qc}$)**.** Signatures $F, G \in {}_n\mathcal{S}(\mathbb{C}^q)$ are *quantum isomorphic* ($F \cong_{qc} G$) if there is a $q \times q$ quantum permutation matrix matrix U satisfying $U \cdot F = G$.

Bijective signature sets $\mathcal{F}$ and $\mathcal{G}$ with common domain $[q]$ are quantum isomorphic ($\mathcal{F} \cong_{qc} \mathcal{G}$) if there is a $q \times q$ quantum permutation matrix U satisfying $U \cdot F = G$ for every $\mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}$.

As with classical isomorphism, we require every corresponding pair of functions in $\mathcal{F}$ and $\mathcal{G}$ to be quantum isomorphic via the same U.

6.2 Three types of planar Holant

Quantum orthogonal matrices are intimately related to planar Holant problems. In this subsection, we introduce three variants of planar Holant that interact differently with quantum orthogonal matrices. Recall that, in a Holant signature grid or gadget, each vertex is assigned a signature F, along with an ordering of the edges incident to v to determine which edge corresponds to each input of F. The most natural such ordering is counterclockwise around the vertex, in agreement with the counterclockwise order of the dangling edges around a gadget when defining its signature. In this way, if K is the signature of an $\mathcal{F}$-gadget $\mathbf{K}$, when we construct from any $\mathcal{F} \cup \{K\}$-grid Ω an $\mathcal{F}$-grid Ω' with the same Holant value by replacing every vertex assigned K with $\mathbf{K}$, we can orient $\mathbf{K}$ so that its dangling edges align with the edges incident to v in Ω, without any local edge crossings. Thereby a plane embedding of Ω becomes a plane embedding of Ω'. Without any planarity restriction, however, no ordering is *a proiri* preferred, because, by admitting local edge crossings, we can identify the edges incident to v with the dangling edges of $\mathbf{K}$ in any order. The standard definition of planar Holant [Cai+15; CLX17; CFS; CF22], which we call Pl-Holant, specifies the counterclockwise orientation.

Definition 6.2.1 (Pl-Holant)**.** A Pl-Holant($\mathcal{F}$) gadget is a planar (in the sense of Definition 2.2.8) $\mathcal{F}$-gadget admitting a planar embedding such that, for every vertex v, the order of the edges incident to v proceeds counterclockwise around v. Pl-Holant denotes the restriction of Holant to Pl-Holant signature grids.

Globally specifying counterclockwise orientations is equivalent to globally specifying clockwise orientations (in fact, this is the definition of planar Holant in [CC17a]). However, since Holant concerns local, rather than global, constraints, it seems more fitting to allow any locally planar orientation – that is, allow both counterclockwise and clockwise orientations in the same signature grid.

Definition 6.2.2 (Pl$^\top$-Holant)**.** A Pl$^\top$-Holant($\mathcal{F}$) gadget is a planar $\mathcal{F}$-gadget admitting a planar embedding such that, for every vertex v, the order of the edges incident to v proceeds counterclockwise or clockwise around v. Pl$^\top$-Holant denotes the restriction of Holant to Pl$^\top$-Holant signature grids.

We still define the signature K of a $\mathrm{Pl}^\top$-Holant($\mathcal{F}$) gadget $\mathbf{K}$ using the counterclockwise dangling edge order (independent of the orientations of the vertices in the gadget). However, we can still ensure gadget substitutions preserve planarity. Let v be a clockwise-oriented vertex assigned K in a $\mathrm{Pl}^\top$-Holant $\mathcal{F}\cup\{K\}$-grid Ω. Reflecting $\mathbf{K}$ along any axis preserves its planarity while reversing its dangling edge order from counterclockwise to clockwise (and also flipping the orientations of every vertex between counterclockwise and clockwise). So we construct a $\mathrm{Pl}^\top$-Holant $\mathcal{F}$-grid Ω' with the same Holant value as Ω by replacing v with the reflected $\mathbf{K}$. Equivalently, the signature of $\mathbf{K}^\top$ is $K^\top$ by Proposition 2.2.1 and (4.2.1), and a clockwise-oriented vertex assigned K is identical to a counterclockwise-oriented vertex assigned $K^\top$. In other words,

$$\mathrm{Pl}^\top\text{-Holant}(\mathcal{F}) \equiv \text{Pl-Holant}(\mathcal{F}\cup\mathcal{F}^\top), \tag{6.2.1}$$

where $\mathcal{F}^\top = \{F^\top \mid F \in \mathcal{F}\}$.

To motivate our final variant of planar Holant, recall from Proposition 4.2.2 that any TCWD is closed under $\dagger$, but only conjugate-closed TCWDs are closed under $\top$. We will see that quantum orthogonal matrices are intimately related to TCWDs, so the most natural definition of planar Holant from the quantum perspective uses $\dagger$ instead of $\top$.

Definition 6.2.3 ($\mathrm{Pl}^\dagger$-Holant)**.** A $\mathrm{Pl}^\dagger$-Holant($\mathcal{F}$) gadget is an planar $\mathcal{F}\cup\{\overline{F} \mid F \in \mathcal{F}\}$ gadget admitting a planar embedding such that, for every vertex v, either v is assigned a signature from F and the order of the edges incident to v proceeds counterclockwise, or v is assigned a signature from $\overline{F}$ and the order of the edges incident to v proceeds clockwise. $\mathrm{Pl}^\dagger$-Holant denotes the restriction of Holant to $\mathrm{Pl}^\dagger$-Holant signature grids.

Vertices assigned signatures in $\mathcal{F}\cap\{\overline{F} \mid F \in \mathcal{F}\}$ can have arbitrary orientation. A clockwise-oriented vertex assigned $\overline{F}$ is identical to a counterclockwise-oriented vertex assigned $F^\dagger$, so

$$\mathrm{Pl}^\dagger\text{-Holant}(\mathcal{F}) \equiv \text{Pl-Holant}(\mathcal{F}\cup\mathcal{F}^\dagger), \tag{6.2.2}$$

where $\mathcal{F}^\dagger = \{F^\dagger \mid F \in \mathcal{F}\}$. Planar gadget substitution still works as desired, similarly to $\mathrm{Pl}^\top$-Holant: first, if $\mathbf{K}$ is a $\mathrm{Pl}^\dagger$-Holant($\mathcal{F}$)-gadget, then $\mathbf{K}^\dagger$ is a well-defined $\mathrm{Pl}^\dagger$-Holant($\mathcal{F}$)-gadget (with signature $K^\dagger$ by Proposition 2.2.1) because every counterclockwise-oriented $F \in \mathcal{F}$ assigned to a vertex in $\mathbf{K}$ is, by definition of the $\dagger$ operation on gadgets, flipped and conjugated to a clockwise-oriented $\overline{F}$ in $\mathbf{K}^\dagger$ and vice-versa. Second, if $\mathbf{K}$ is a $\mathrm{Pl}^\dagger$-Holant$(\mathcal{F})$-gadget with signature K, then, in any $\mathrm{Pl}^\dagger$-Holant$(\mathcal{F} \cup K)$-grid, we can substitute $\mathbf{K}$ for any counterclockwise-oriented vertex assigned K, and substitute $\mathbf{K}^\dagger$, with signature $K^\dagger$, for any clockwise-oriented vertex assigned $\overline{K}$ (equivalently, counterclockwise- oriented vertex assigned $K^\dagger$).

Definition 6.2.4 ($[\mathcal{F}]_{\mathrm{Pl}}, [\mathcal{F}]_{\mathrm{Pl}^\top}, [\mathcal{F}]_{\mathrm{Pl}^\dagger}$). Define $[\mathcal{F}]_{\mathrm{Pl}}, [\mathcal{F}]_{\mathrm{Pl}^\top}, [\mathcal{F}]_{\mathrm{Pl}^\dagger} \subset \mathcal{V}(\mathbb{C}^q)$ to be the algebras of tensors of quantum Pl-Holant$(\mathcal{F})$, $\mathrm{Pl}^\top$-Holant$(\mathcal{F})$, and $\mathrm{Pl}^\dagger$-Holant$(\mathcal{F})$ gadgets, respectively.

For a quantum Pl-Holant$(\mathcal{F})$ or $\mathrm{Pl}^\top$-Holant$(\mathcal{F})$ gadget $\mathbf{K}$, define $\mathbf{K}_{\mathcal{F}\to\mathcal{G}}$ by replacing every vertex assigned $F \in \mathcal{F}$ with a vertex with the same orientation assigned the corresponding $G \in \mathcal{G}$. For a quantum $\mathrm{Pl}^\dagger$-Holant$(\mathcal{F})$-gadget $\mathbf{K}$, define $\mathbf{K}_{\mathcal{F}\to\mathcal{G}}$ by replacing every counterclockwise-oriented vertex assigned $F \in \mathcal{F}$ with a counterclockwise-oriented vertex assigned $G \in \mathcal{G}$ for $G \leftrightsquigarrow F$, and replacing every clockwise-oriented vertex assigned $\overline{F}$ for $F \in \mathcal{F}$ with a clockwise-oriented vertex assigned $\overline{G}$ for $\mathcal{G} \ni G \leftrightsquigarrow F$. Now, as for general gadgets, $[\mathcal{F}]_{\mathrm{Pl}}$ and $[\mathcal{G}]_{\mathrm{Pl}}$ are bijective via the map $\leftrightsquigarrow$ mapping the tensor of a Pl-Holant$(\mathcal{F})$ gadget $\mathbf{K}$ to the tensor of $\mathbf{K}_{\mathcal{F}\to\mathcal{G}}$ and similarly for $\mathrm{Pl}^\top$-Holant and $\mathrm{Pl}^\dagger$-Holant (then extend $\leftrightsquigarrow$ linearly to quantum gadgets). Also say bijective $\mathcal{F}$ and $\mathcal{G}$ are Pl-Holant-indistinguishable if $\text{Pl-Holant}_{\mathcal{F}}(\Omega) = \text{Pl-Holant}_{\mathcal{G}}(\Omega_{\mathcal{F}\to\mathcal{G}})$, and similarly for $\mathrm{Pl}^\top$-Holant and $\mathrm{Pl}^\dagger$-Holant.

Remark 6.2.1. If $\mathcal{F}$ is conjugate-closed, then $\mathcal{F} \cup \mathcal{F}^\dagger = \mathcal{F} \cup \mathcal{F}^\top$, so $\mathrm{Pl}^\top$-Holant$(\mathcal{F}) \equiv \mathrm{Pl}^\dagger$-Holant$(\mathcal{F})$. However, there is a subtle difference between $\mathrm{Pl}^\top$-Holant and $\mathrm{Pl}^\dagger$-Holant gadgets for bijective pairs of conjugate-closed sets $\mathcal{F}$ and $\mathcal{G}$. We can view a $\mathrm{Pl}^\dagger$-Holant$(\mathcal{F})$ gadget $\mathbf{K}$ as a $\mathrm{Pl}^\top$-Holant gadget for conjugate-closed $\mathcal{F}$, as every clockwise-oriented vertex is still assigned a signature in $\mathcal{F}$. Viewing $\mathbf{K}$ as a $\mathrm{Pl}^\dagger$-Holant$(\mathcal{F})$ gadget, if a clockwise-oriented vertex is assigned F', then $F' = \overline{F}$ for some $F \in \mathcal{F}$, and in $\mathbf{K}_{\mathcal{F}\to\mathcal{G}}$ we replace F' with $\overline{G}$ for $\mathcal{G} \ni G \leftrightsquigarrow F$. But, viewing $\mathbf{K}$ as a $\mathrm{Pl}^\top$-Holant$(\mathcal{F})$ gadget, we replace this $F' \in \mathcal{F}$ with the G' in $\mathcal{G}$ corresponding to F', and we do not necessarily have $G' = \overline{G}$ unless $\overline{F} \leftrightsquigarrow \overline{G}$. This difference prevents us from treating the conjugate-closed $\mathrm{Pl}^\top$-Holant setting as a special case of the $\mathrm{Pl}^\dagger$-Holant setting in the proofs below.

Finally, we need one more technical proposition.

Proposition 6.2.1. *Let $\mathcal{F}, \mathcal{G}$ be bijective. If $\mathbf{K}$ and $\mathbf{J}$ are quantum Pl-Holant$(\mathcal{F})$, $Pl^\top$-Holant$(\mathcal{F})$, or $Pl^\dagger$-Holant$(\mathcal{F})$ gadgets, then*

- $(\mathbf{K} \circ \mathbf{J})_{\mathcal{F}\to\mathcal{G}} = \mathbf{K}_{\mathcal{F}\to\mathcal{G}} \circ \mathbf{J}_{\mathcal{F}\to\mathcal{G}}$ *and*
- $(\mathbf{K} \otimes \mathbf{J})_{\mathcal{F}\to\mathcal{G}} = \mathbf{K}_{\mathcal{F}\to\mathcal{G}} \otimes \mathbf{J}_{\mathcal{F}\to\mathcal{G}}$.

If $\mathbf{K}$ *is a quantum Pl*†*-Holant*$(\mathcal{F})$ *gadget, or if* $\mathbf{K}$ *is a quantum Pl*$^\top$*-Holant*$(\mathcal{F})$ *gadget and* $\mathcal{F}$ *and* $\mathcal{G}$ *are conjugate-closed and satisfy* $F \leftrightsquigarrow G$ *iff* $\overline{F} \leftrightsquigarrow \overline{G}$*, then*

- $(\mathbf{K}^\dagger)_{\mathcal{F}\to\mathcal{G}} = (\mathbf{K}_{\mathcal{F}\to\mathcal{G}})^\dagger$.

Proof. The first two points hold because $\circ$ and $\otimes$ do not change signatures or orientations of individual vertices. For the third point, suppose $\mathbf{K}$ is a quantum Pl†-Holant$(\mathcal{F})$ gadget. If v is oriented counterclockwise and assigned $F \in \mathcal{F}$ in (some gadget term of) $\mathbf{K}$, then, with $\mathcal{G} \ni G \leftrightsquigarrow F$, v is oriented clockwise and assigned $\overline{G}$ in both $(\mathbf{K}^\dagger)_{\mathcal{F}\to\mathcal{G}}$ and $(\mathbf{K}_{\mathcal{F}\to\mathcal{G}})^\dagger$. If v is oriented clockwise and assigned $\overline{F}$ for some $F \in \mathcal{F}$, then, with $\mathcal{G} \ni G \leftrightsquigarrow F$, v is oriented counterclockwise and assigned G in both $(\mathbf{K}^\dagger)_{\mathcal{F}\to\mathcal{G}}$ and $(\mathbf{K}_{\mathcal{F}\to\mathcal{G}})^\dagger$. Suppose $\mathbf{K}$ is a quantum Pl$^\top$-Holant$(\mathcal{F})$ gadget. If v is assigned $F \in \mathcal{F}$ in $\mathbf{K}$, then, with $\mathcal{G} \ni G \leftrightsquigarrow F$, v is assigned $\overline{G} \in \mathcal{G}$ in both $(\mathbf{K}^\dagger)_{\mathcal{F}\to\mathcal{G}}$ and $(\mathbf{K}_{\mathcal{F}\to\mathcal{G}})^\dagger$ (here we apply the extra assumption $F \leftrightsquigarrow G \iff \overline{F} \leftrightsquigarrow \overline{G}$). □

6.3 The planar gadget decomposition

In this section, we prove the planar version of Proposition 4.4.2. There, the space $[\mathcal{F}]$ of quantum Holant $\mathcal{F}$-gadgets is the symmetric TCWD generated by $\mathcal{F}$. Here, the space of quantum Pl†-Holant $\mathcal{F}$-gadgets is the (not-necessarily-symmetric) TCWD generated by $\mathcal{F}$. Intuitively, the difference between planar and non-planar gadgets is . The proof follows the special case for planar bi-labeled graphs (Pl-Holant$(A_X, \mathcal{EQ})$ gadgets) in [MR20], but with some new subtleties because our signatures $\mathcal{F}$ can be, unlike A_X and $\mathcal{EQ}$, asymmetric.

We begin with the following result, a restatement of [MR20, Corollary 6.5]. For completeness, we sketch the proof.

Lemma 6.3.1. *Let* $\mathbf{K}$ *be a plane gadget with at least one dangling edge and no wires. There is some* v *on the outer face of* $\mathbf{K}$ *such that the dangling ends of all dangling edges incident to* v *occur consecutively in the counterclockwise cyclic ordering of dangling edges around* $\mathbf{K}$*.*

Proof. By planarity, any vertex of $\mathbf{K}$ with an incident dangling edge must lie on the outer face. Label the $d > 0$ dangling edges of $\mathbf{K}$ in counterclockwise cyclic order by $[d]$. Let $u \in V(\mathbf{K})$ be a vertex with at least one incident dangling edge. If every dangling edge of $\mathbf{K}$ is incident to u, or if u has exactly one incident dangling edge, then we are done. Otherwise, WLOG, we may assume there is an $r < d$ such that dangling edges 1 and r are incident to u, and dangling edges $2, \ldots, r-1$ are not incident to u. If u' is a vertex incident to some dangling edge with index in $\{2, \ldots, r-1\}$, then every dangling edge incident to u' must have index in $\{2, \ldots, r-1\}$, as any dangling edge incident to u' with index outside of this range would cross dangling edge 1 or r, violating the planarity of $\mathbf{K}$. Of all such u', let v be the vertex with minimal distance $h - \ell$ between its lowest-indexed and highest-indexed dangling edges $\ell, h \in \{2, \ldots, r-1\}$. Any vertex w with an incident dangling edge between ℓ and h must either have all incident dangling edges between ℓ and h, violating the minimality of v, or have a dangling edge crossing a dangling edge of v, violating planarity. Thus the dangling edges incident to v occur consecutively as $\ell, \ldots, h$. □

The possible asymmetry of signatures in our setting necessitates the following refinement of Lemma 6.3.1.

Corollary 6.3.1. *Let* $\mathbf{K}$ *be a plane gadget with no wires and no self-loops. There is a vertex* v *on the outer face of* $\mathbf{K}$ *satisfying one of the following conditions.*

(i) v *is incident to no dangling edges of* $\mathbf{K}$ *and it is possible to add a a dummy left dangling edge to* v *without violating planarity.*

(ii) v *is incident to no dangling edges of* $\mathbf{K}$ *and it is possible to add a a dummy right dangling edge to* v *without violating planarity.*

(iii) v *is incident to* $\ell > 0$ *consecutive left dangling edges and no right dangling edges of* $\mathbf{K}$*, and these edges occur consecutively in the counterclockwise cyclic order of edges incident to* v*.*

(iv) v *is incident to* $r > 0$ *consecutive right dangling edges and no left dangling edges of* $\mathbf{K}$*, and these edges occur consecutively in the counterclockwise cyclic order of edges incident to* v*.*

(v) v is incident to exactly the top ℓ left dangling edges and top r right dangling edges of $\mathbf{K}$*, and these dangling edges occur consecutively in the counterclockwise cyclic order of edges incident to v.*

(vi) v is incident to exactly the bottom ℓ left dangling edges and bottom r right dangling edges of $\mathbf{K}$*, and these dangling edges occur consecutively in the counterclockwise cyclic order of edges incident to v.*

Proof. Let v be the vertex guaranteed by Lemma 6.3.1. First, suppose the dangling edges incident to v occur consecutively in the counterclockwise cyclic order of edges incident to v. If v is incident to only left or only right dangling edges of $\mathbf{K}$, then we are in case (iii) or (iv), respectively. If v is incident to at least one left dangling edge and at least one right dangling edge of $\mathbf{K}$, then v must be incident to the top left and right dangling edges or the bottom left and right dangling edges (or both). If v is incident to all right dangling edges of $\mathbf{K}$ and to the first and last, but not all, left dangling edges of $\mathbf{K}$ (this is possible under Lemma 6.3.1), then apply Lemma 6.3.1 to the gadget formed from $\mathbf{K}$ by removing all dangling edges incident to v to obtain a different v' with at least one left dangling edge and no right dangling edges, reducing to item (iii). Handle the horizontally reflected case similarly. Otherwise, the dangling edges incident to v wrap around either the top or bottom of $\mathbf{K}$, but not both. These are cases (v) and (vi).

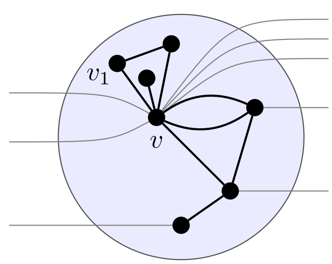

Figure 6.1: An example where the dangling edges incident to v do not occur consecutively around v.

Otherwise, the dangling edges incident to v do not occur consecutively around v. See Figure 6.1. Then there is a consecutive sequence $e, e_1, \ldots, e_p, e'$ of edges incident to v where e and e' are consecutive dangling edges of $\mathbf{K}$ and $e_1, \ldots, e_p$ for $p > 0$ are edges between v and, by the no-loops assumption, other vertices $v_1, \ldots, v_p$ of $\mathbf{K}$. None of $v_1, \ldots, v_p$ can have any dangling edges, as, by planarity, any such dangling edge would lie between e and e' in the order of dangling edges of $\mathbf{K}$.

Recall that v is on the outer face of $\mathbf{K}$ and, by planarity, e must lie entirely within the outer face of $\mathbf{K}$. If we place a dummy node u on e, then u and v are on the outer face of $\mathbf{K}$, so u, v, v_1 is a sequence of consecutive nodes on the outer face. Thus we may route a new dangling edge incident to v_1 infinitesimally close to the path between v_1 and u, then parallel to e, without introducing any crossings. Depending on whether e was a left or right dangling edge of $\mathbf{K}$, v_1 witnesses that we are in case (i) or (ii). □

The single-vertex gadgets we extract from a larger planar gadget during the decomposition will have the following form.

Definition 6.3.1 ($\mathbf{V}(F, \ell, r, m)$)**.** For $\ell, r \geq 0$, $m \in [\ell{+}r]$, and $(\ell{+}r)$-ary signature F, the Pl-Holant gadget $\mathbf{V}(F, \ell, r, m)$ consists of a single vertex assigned F incident to ℓ left dangling edges $e_1, \ldots, e_\ell$ and r right dangling edges $e_{\ell+r}, \ldots, e_{\ell+1}$ (both in top-down order, so the counterclockwise cyclic order around v is $e_1, \ldots, e_{\ell+r}$), where v is oriented counterclockwise with first input along edge e_{m+1}.

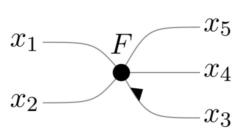

$$= F(x_3, x_4, x_5, x_1, x_2) = F^{(2)}(x_1, x_2, x_3, x_4, x_5) =$$

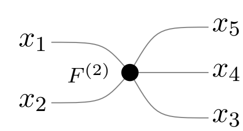

Figure 6.2: Computing the tensor of $\mathbf{V}(F, 2, 3, 2)$ (left) as the flattening $(F^{(2)})^{2,3}$ of a rotation of F (right). The black arrow indicates the first input and orientation of F. The diagram on the right is a diagrammatic tensor representation (as in Figure 2.1), rather than a Pl-Holant gadget, but can be viewed as a Pl-Holant gadget in which the central vertex has the standard counterclockwise orientation starting with the dangling edge with input x_1.

The evaluation of $\mathbf{V}(F, \ell, r, m)$ on inputs x_i along dangling edge e_i for every i is (see Figure 6.2)

$$F(x_{m+1}, \ldots, x_n, x_1, \ldots, x_m) = F^{(m)}(x_1, \ldots, x_n),$$

so the signature of $\mathbf{V}(F, \ell, r, m)$ is $F^{(m)}$ (we are already considering the inputs to $\mathbf{V}(F, \ell, r, m)$ is counterclockwise cyclic order, so no reversal of right inputs is required) and the tensor of $\mathbf{V}(F, \ell, r, m)$ is $(F^{(m)})^{\ell,r}$.

Following the planar bi-labeled graph decomposition of Mančinska and Roberson [MR20], we extract the vertices of a gadget one at a time, choosing at each step the vertex v given by Corollary 6.3.1. The possible asymmetry of our sigatures introduces two new difficulties. First, the extracted vertex may be rotated – that is, be the center of a gadget $\mathbf{V}(F, \ell, r, m)$ for $m > 0$ – in which case we need to apply Proposition 4.2.1. Second, and more subtly (in fact, this point was missed in [CY24]), there is a difference between the cases where the dangling edges incident to v occur consecutively around v, as in Figure 6.4 below, or when they are interrupted by edges between v and other vertices, as Figure 6.1. The decomposition in [MR20], which only deals with symmetric signatures, does not distinguish between these two cases because we can always reorder the edges incident to a vertex assigned a symmetric signature. In our asymmetric setting, however, we cannot directly extract a vertex v of the latter kind without moving some dangling edges past non-dangling edges in the cyclic order of edges incident to v.

Theorem 6.3.1. *The following hold for* $\mathcal{F} \subset \mathcal{S}(\mathbb{C}^q)$*:*

- $[\mathcal{F}]_{Pl} = \langle \mathcal{F}, \supset, \subset \rangle_{+,\circ,\otimes} \subset \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\dagger}$.
- $[\mathcal{F}]_{Pl^\top} = \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\top}$ *and, if* $\mathcal{F}$ *is conjugate-closed, then* $[\mathcal{F}]_{Pl^\top} = \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\dagger}$.
- $[\mathcal{F}]_{Pl^\dagger} = \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\dagger}$.

Proof. The inclusion $\langle \mathcal{F}, \supset, \subset \rangle_{+,\circ,\otimes} \subset \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\dagger}$ follows from the fact that $\subset = \supset^\dagger$. Also, if $\mathcal{F}$ is conjugate-closed, then $\mathcal{F} \cup \mathcal{F}^\top = \mathcal{F} \cup \mathcal{F}^\dagger$, so $[\mathcal{F}]_{\mathrm{Pl}^\top} = [\mathcal{F}]_{\mathrm{Pl}^\dagger}$ by (6.2.1) and (6.2.2). Therefore it suffices to prove the three equalities at the start of the three points.

We first show the three claimed $\supset$ inclusions. Every n-ary $F \in \mathcal{F}$ is the signature of the simple planar gadget consisting of a single counterclockwise-oriented vertex assigned F with n left dangling edges. The pivots $\supset$ and $\subset$ are the signatures of ever-present wire gadgets. Therefore $\mathcal{F} \cup \{\supset, \subset\} \subset [\mathcal{F}]_{\mathrm{Pl}} \cap [\mathcal{F}]_{\mathrm{Pl}^\top} \cap [\mathcal{F}]_{\mathrm{Pl}^\dagger}$. The definition of quantum gadgets yields closure under $\mathbb{C}$-linear combinations. Inductively, if K, J are the tensors of gadgets with planar underlying multigraphs, then $K \circ J$ and $K \otimes J$ are the tensors of gadgets with planar underlying multigraphs by Proposition 2.2.1 and Proposition 2.2.2 (see also [MR20, Lemmas 5.12-5.14], which show that the corresponding operations on bilabeled graphs also preserve planarity). Since $\circ$ and $\otimes$ do not affect vertex orientation, we conclude $[\mathcal{F}]_{\mathrm{Pl}} \supset \langle \mathcal{F}, \supset, \subset \rangle_{+,\circ,\otimes}$. Recall from Section 6.2 that the operations $\top$ and $\dagger$ well-defined on $\mathrm{Pl}^\top$-Holant and $\mathrm{Pl}^\dagger$-Holant gadgets, respectively (but not on Pl-Holant gadgets, because they reverse vertex orientation). By Proposition 2.2.1 and Proposition 2.2.2, if K is the tensor of a $\mathrm{Pl}^\top$-Holant($\mathcal{F}$) or $\mathrm{Pl}^\dagger$-Holant($\mathcal{F}$) gadget, then $K^\top$ or $K^\dagger$ is the signature of a

$\mathrm{Pl}^\top$-Holant($\mathcal{F}$) or $\mathrm{Pl}^\dagger$-Holant($\mathcal{F}$) gadget, respectively. Thus $[\mathcal{F}]_{\mathrm{Pl}^\top} \supset \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\top}$ and $[\mathcal{F}]_{\mathrm{Pl}^\dagger} \supset \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\dagger}$.

We next show the three $\subset$ inclusions. Assume that the first holds: $[\mathcal{F}]_{\mathrm{Pl}} \subset \langle \mathcal{F}, \supset, \subset \rangle_{+,\circ,\otimes}$ for any $\mathcal{F}$. Then, by (6.2.1),

$$[\mathcal{F}]_{\mathrm{Pl}^\top} = [\mathcal{F} \cup \mathcal{F}^\top]_{\mathrm{Pl}} \subset \langle \mathcal{F} \cup \mathcal{F}^\top, \supset, \subset \rangle_{+,\circ,\otimes} \subset \langle \mathcal{F} \cup \mathcal{F}^\top, \supset \rangle_{+,\circ,\otimes,\top}.$$

For signature $F \in \mathcal{F}$, (4.2.1) and Lemma 4.2.1 give $F^\top = (F^\top)^{n,0} = (F^{0,n})^\top \in \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\top}$. Therefore

$$[\mathcal{F}]_{\mathrm{Pl}^\top} \subset \langle \mathcal{F} \cup \mathcal{F}^\top, \supset \rangle_{+,\circ,\otimes,\top} \subset \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\top},$$

giving the second claim of the theorem. Similarly, (6.2.2), (4.2.2) and Lemma 4.2.1 give

$$[\mathcal{F}]_{\mathrm{Pl}^\dagger} = [\mathcal{F} \cup \mathcal{F}^\dagger]_{\mathrm{Pl}} \subset \langle \mathcal{F} \cup \mathcal{F}^\dagger, \supset, \subset \rangle_{+,\circ,\otimes} \subset \langle \mathcal{F} \cup \mathcal{F}^\dagger, \supset \rangle_{+,\circ,\otimes,\dagger} \subset \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\dagger},$$

giving the third claim. So it suffices to prove $[\mathcal{F}]_{\mathrm{Pl}} \subset \langle \mathcal{F}, \supset, \subset \rangle_{+,\circ,\otimes}$. For this, it suffices to show that, for any Pl-Holant gadget $\mathbf{K}$ with tensor K, we have $K \in \langle \mathcal{F}, \supset, \subset \rangle_{\circ,\otimes}$. To do this, we decompose $\mathbf{K}$ into copies of I, $\supset$, $\subset$, and simple gadgets consisting of a single vertex assigned a (possibly rotated) signature in $\mathcal{F}$ and its dangling edges. Since the signature of $=_2$ is $\supset$, we may assume WLOG that $\mathcal{F}$ contains $=_2$. Now, if the underlying multigraph of $\mathbf{K}$ contains any loops, subdivide this loop and assign the new vertex $=_2$. This does not change the tensor K of $\mathbf{K}$, and the theorem statement concerns only gadget tensors, not the gadgets themselves. Consequently, we may assume every edge in the underlying graph of $\mathbf{K}$ is between two distinct vertices. Similarly, we may, by placing a vertex assigned $=_2$ on any wire, assume $\mathbf{K}$ contains no wires.

Suppose that $\mathbf{K}$ is disconnected. Any connected components of $\mathbf{K}$ without a dangling edge contribute only an overall scalar factor to the tensor of $\mathbf{K}$, so, for the purpose of showing $[\mathcal{F}]_{\mathrm{Pl}} \subset \langle \mathcal{F}, \supset, \subset \rangle_{+,\circ,\otimes}$, can be disregarded. Applying Corollary 6.3.1 to the gadget constructed from $\mathbf{K}$ by contracting every connected component of $\mathbf{K}$ into a vertex (keeping all dangling edges intact), we conclude that there is some connected component $\mathbf{J}$ of $\mathbf{K}$ whose incident dangling edges occur consecutively. Let $\mathbf{J}'$ be the union of the other connected components of $\mathbf{K}$, and let $e_1, \ldots, e_\ell$ and $e'_1, \ldots, e'_r$ be the left and right dangling edges of $\mathbf{K}$, in top-down order. If $\mathbf{J}$ is incident to only $e_a, \ldots, e_b$ for $1 \le a \le b \le \ell$, then

$$K = (I^{\otimes a-1} \otimes J \otimes I^{\ell-b}) \circ J',$$

where J and J' are the tensors of $\mathbf{J}$ and $\mathbf{J}'$, respectively. The case where $\mathbf{J}$ is incident to only right dangling edges is similar. If $\mathbf{J}$ is incident to only $e_1, \ldots, e_a$ and $e'_1, \ldots, e'_b$ for $a \leq \ell$ and $b \leq r$, then

$$K = J \otimes J'.$$

Up to exchanging the roles of $\mathbf{J}$ and $\mathbf{J}'$, we have exhausted all cases. Hence, in what follows, we may assume $\mathbf{K}$ is connected.

We proceed by induction on the number of vertices in $\mathbf{K}$. Since $\mathbf{K}$ has no wires, it must have at least one vertex unless it is empty. If $\mathbf{K}$ has one vertex, then $\mathbf{K} = \mathbf{V}(F, \ell, r, m)$ for some $F \in \mathcal{F}$ and ℓ, r, m, so $K = (F^{(m)})^{\ell, r}$. By Proposition 4.2.1 and Lemma 4.2.1, $K \in \langle \mathcal{F}, \supset, \subset \rangle_{\circ, \otimes}$.

If $\mathbf{K}$ has multiple vertices, then let v, assigned n-ary signature $F \in \mathcal{F}$, satisfy the conditions of Corollary 6.3.1. Broadly, we will extract v from $\mathbf{K}$ by a procedure depending on the case in Corollary 6.3.1, then show that the remaining gadget is still planar.

Case 1. Suppose that v satisfies items (i) or (iii) of Corollary 6.3.1, so v is incident to $\ell \geq 0$ left dangling edges $e_1, \ldots, e_\ell$ in counterclockwise cyclic order around v (if v satisfies item (i), then $\ell = 0$ and this list is empty).

The extraction procedure, shown in Figure 6.3, is as follows: First, break the $r > 0$ edges between v and the other vertices of $\mathbf{K}$. Orient the halves of these r edges still attached to v to become new right dangling edges $f_1, \ldots, f_r$ incident to v in top-down (or clockwise relative to v) order, and orient the respective other halves to become left dangling edges $f'_1, \ldots, f'_r$ of $\mathbf{K}'$ in top-down (or counterclockwise relative to $\mathbf{K}'$) order, where $\mathbf{K}'$ is defined to be the gadget such that reconnecting f'_i with f_i for $i = 1, \ldots, r$ reconstructs $\mathbf{K}$. To construct the top-down order of left dangling edges of $\mathbf{K}'$, replace $e_1, \ldots, e_\ell$ with $f'_1, \ldots, f'_r$ in the top-down order of left dangling edges of $\mathbf{K}$, as $e_1, \ldots, e_\ell$ remain left dangling edges incident to v and are not present in $\mathbf{K}'$. Since $e_1, \ldots, e_\ell$ occur consecutively around v and constitute all dangling edges incident to v, $f_1, \ldots, f_r$ also occur consecutively around v, as every edge incident to v was either a left dangling edge of $\mathbf{K}$ or connected v to another vertex of $\mathbf{K}$. Therefore the counterclockwise cyclic order of dangling edges incident to v, starting from the top left dangling edge, is now $e_1, \ldots, e_\ell, f_r, \ldots, f_1$, and we can orient $e_1, \ldots, e_\ell$ to the left and $f_r, \ldots, f_1$ to the right without crossing any edges or changing the order of edges incident to v.

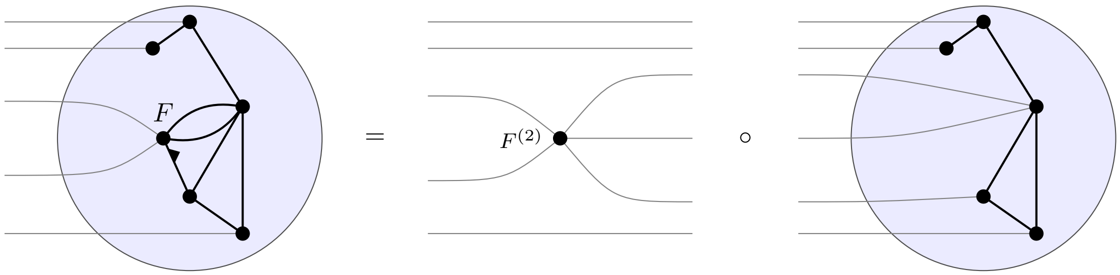

Figure 6.3: Extracting v as in (6.3.1), with v as in Figure 6.2: $K = \big(I^{\otimes 2} \otimes (F^{(2)})^{2,3} \otimes I\big) \circ K'$.

After this extraction, v and its dangling edges form a copy of $\mathbf{V}(F, \ell, r, m)$ for some m depending on the rotation of v in $\mathbf{K}$. The top-down order of right dangling edges of $\mathbf{V}(F, \ell, r, m)$ is $f_1, \ldots, f_r$, and the top-down order of left dangling edges of $\mathbf{K}'$ consists of $\ell' \geq 0$ unaffected left dangling edges from $\mathbf{K}'$, then $f'_1, \ldots, f'_r, e_{\ell+1}$, then $\ell'' \geq 0$ more left dangling edges. Composing $\mathbf{V}(F, \ell, r, m)$, along with ℓ' wires above and ℓ'' wires below, with $\mathbf{K}'$ reconnects f_i with f'_i for $i = 1, \ldots, r$, reforming $\mathbf{K}$. Let K' be the tensor of $\mathbf{K}'$, and recall that the tensor of $\mathbf{V}(F, \ell, r, m)$ is $(F^{(m)})^{\ell,r}$. Then

$$K = (I^{\otimes \ell'} \otimes (F^{(m)})^{\ell,r} \otimes I^{\otimes \ell''}) \circ K'. \tag{6.3.1}$$

To apply induction, we must also show that $\mathbf{K}'$ is planar. We obtain the underlying multigraph with dangling edges of $\mathbf{K}'$ by adding $f_1, \ldots, f_r$, in top-down order, to the underlying multigraph with dangling edges of $\mathbf{K} \setminus \{v\}$. The latter is planar by planarity of $\mathbf{K}$, so it suffices to add $f'_1, \ldots, f'_r$ to the plane embedding of $\mathbf{K} \setminus \{v\}$ while preserving planarity. Recall that $f'_1, \ldots, f'_r$ come from the r internal edges of $\mathbf{K}$ incident to v, in clockwise order around v. Route $f'_1, \ldots, f'_r$ along the original paths of these edges to the location of v in the plane embedding of $\mathbf{K}$. Then move them infinitesimally apart at this point and route them along the path of any of $e_1, \ldots, e_\ell$ in the plane embedding of $\mathbf{K}$ or, if $\ell = 0$, along the hypothetical path of a left dangling edge incident to v from the assumption of this case. The ends of $f'_1, \ldots, f'_r$ end up in top-down order without crossing along this new route because the clockwise order of $f'_1, \ldots, f'_r$ relative to v is a counterclockwise (or top-down on the left) order relative to $\mathbf{K}'$. Furthermore, $f'_1, \ldots, f'_r$ replace $e_1, \ldots, e_\ell$ in the overall top-down order of dangling edges of $\mathbf{K}$, as specified above. Any other edge crossing $f'_1, \ldots, f'_r$ along these new paths would have crossed some (possibly dangling) edge incident to v in the plane embedding of $\mathbf{K}$, a contradiction.

Hence $\mathbf{K}'$ is planar and contains one fewer vertex than $\mathbf{K}$, so, by induction, $K' \in \langle \mathcal{F}, \supset, \subset \rangle_{\circ,\otimes}$. Furthermore, $(F^{(m)})^{\ell,r} \in \langle \mathcal{F}, \supset, \subset \rangle_{\circ,\otimes}$ by Proposition 4.2.1 and Lemma 4.2.1. Thus, (6.3.1) implies that $K \in \langle \mathcal{F}, \supset, \subset \rangle_{\circ,\otimes}$.

Case 2. Suppose that v satisfies items (ii) or (iv) of Corollary 6.3.1. Applying similar reasoning to Case 1 with the roles of right and left dangling edges exchanged, and considering the horizontal reflection of $\mathbf{K}$ for the planarity argument, gives

$$K = K' \circ (I^{\otimes r'} \otimes (F^{(m)})^{\ell,r} \otimes I^{\otimes r''}) \tag{6.3.2}$$

for some m, ℓ, r', r''. As in Case 1, it follows that $K \in \langle \mathcal{F}, \supset, \subset \rangle_{\circ,\otimes}$.

Case 3. Suppose that v satisfies item (v) of Corollary 6.3.1, so v is incident to the top ℓ left dangling edges $e_1, \ldots, e_\ell$ and top r right dangling edges $f_1, \ldots, f_r$ of $\mathbf{K}$, both in top-down order, or in counterclockwise order $f_r, \ldots, f_1, e_1, \ldots, e_\ell$ around v. See Figure 6.4 Suppose $\mathbf{K}$ has $\ell' \geq 0$ left dangling edges below e_ℓ and $r' \geq 0$ right dangling edges below f_r. Since $f_r, \ldots, f_1, e_1, \ldots, e_\ell$ are consecutive edges in the outer face incident to v, we may freely pivot $f_r, \ldots, f_1$ to the left above e_1, producing a planar gadget $\mathbf{K}_p$ with $r + \ell + \ell'$ left dangling edges and r' right dangling edges, and whose first $r + \ell$ left dangling edges in the top-down order are $f_r, \ldots, f_1, e_1, \ldots, e_\ell$. Let G be the signature of $\mathbf{K}$, so $K = G^{\ell+\ell',r+r'}$. Since cyclic input/dangling edge pivoting does not affect the signature of a gadget, the tensor of $\mathbf{K}_p$ is $G^{r+\ell+\ell',r'}$. Now $\mathbf{K}_p$ falls into Case 1, so $G^{r+\ell+\ell',r'} \in \langle \mathcal{F}, \supset, \subset \rangle_{\circ,\otimes}$. Another application of Lemma 4.2.1 gives $K = G^{\ell+\ell',r+r'} \in \langle \mathcal{F}, \supset, \subset \rangle_{\circ,\otimes}$.

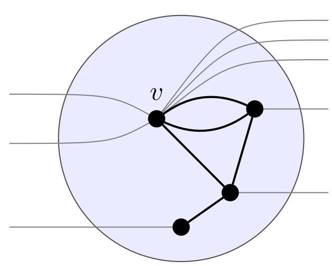

Figure 6.4: Case 3 with $\ell = 2, r = 3, \ell' = 1, r' = 2$.

Case 4. Suppose v satisfies item (vi) of Corollary 6.3.1. Reasoning similar to Case 3 again gives $K \in \langle \mathcal{F}, \supset, \subset \rangle_{\circ,\otimes}$.

We have shown in every case of Corollary 6.3.1 that $K \in \langle \mathcal{F}, \supset, \subset \rangle_{\circ,\otimes}$, so we are done. $\square$

We now obtain the planar version of Corollary 4.4.1.

Corollary 6.3.2. *Let U be a quantum orthogonal matrix and $\mathcal{F} \subset \mathcal{S}(\mathbb{C}^q)$ such that $U\,\mathcal{F}$ is scalar-valued. If $\mathbf{K}$ is a quantum Pl-Holant($\mathcal{F}$) gadget or a quantum Pl†-Holant($\mathcal{F}$) gadget with tensor K, then the tensor of $\mathbf{K}_{\mathcal{F}\to U\,\mathcal{F}}$ is $U \cdot K$. If $\mathbf{K}$ is a quantum Pl$^\top$-Holant($\mathcal{F}$) gadget with tensor K and furthermore $U \cdot \overline{F} = \overline{U \cdot F}$ for every $F \in \mathcal{F}$, then the tensor of $\mathbf{K}_{\mathcal{F}\to U\,\mathcal{F}}$ is $U \cdot K$. Therefore*

- $U[\mathcal{F}]_{Pl} = [U\,\mathcal{F}]_{Pl}$,
- $U[\mathcal{F}]_{Pl^\top} = [U\,\mathcal{F}]_{Pl^\top}$ *if* $U \cdot \overline{F} = \overline{U \cdot F}$ *for every* $F \in \mathcal{F}$,
- $U[\mathcal{F}]_{Pl^\dagger} = [U\,\mathcal{F}]_{Pl^\dagger}$.

Proof. The idea is that the action of U is linear and respects I, the pivots, and the operations $\circ, \otimes, \dagger$ that, by Theorem 6.3.1, generate all quantum gadgets. This is essentially the proof of [MR20, Lemma 7.4]; for completeness, we present it here. Write $U\,\mathcal{F} = \mathcal{G}$. By definition, $U \cdot I = UU^\dagger = I$ and $U \cdot \supset = \supset$ and $U \cdot \subset = \subset$. To avoid noncommutativity-induced complications, we avoid products involving more than one quantum matrix using (6.1.3): $U \cdot F = G$ iff $U^{\otimes \ell} F = G U^{\otimes r}$. First, linearity follows from the linearity of matrix multiplication: if $U^{\otimes \ell} F_i = G_i U^{\otimes r}$ for $\{F_i\}_i, \{G_i\}_i \subset {}_\ell\mathcal{V}(\mathbb{C}^q)_r$, then, for $\{c_i\}_i \subset \mathbb{C}$,

$$U^{\otimes \ell}\left(\sum_i c_i F_i\right) = \sum_i c_i U^{\otimes \ell} F_i = \sum_i c_i G_i U^{\otimes \ell} = \left(\sum_i c_i G_i\right) U^{\otimes r}$$

so $U \cdot (\sum_i c_i F_i) = \sum_i c_i (U \cdot F_i)$. For $F \in {}_\ell\mathcal{V}(\mathbb{C}^q)_r$ and $F' \in {}_r\mathcal{V}(\mathbb{C}^q)_d$,

$$U \cdot (F \circ F') = U^{\otimes \ell} F (U^{\otimes r})^\dagger U^{\otimes r} F' (U^{\otimes d})^\dagger = (U \cdot F) \circ (U \cdot F').$$

If $U^{\otimes \ell} F = G U^{\otimes r}$ and $U^{\otimes \ell'} F' = G' U^{\otimes r'}$ then, by (6.1.1),

$$U^{\otimes \ell + \ell'}(F \otimes F') = (U^{\otimes \ell} F) \otimes (U^{\otimes \ell'} F') = (G U^{\otimes r}) \otimes (G' U^{\otimes r'}) = (G \otimes G') U^{\otimes r + r'}$$

so $U \cdot (F \otimes F') = (U \cdot F) \otimes (U \cdot F')$.

We show the first claim: $U[\mathcal{F}]_{\mathrm{Pl}} = [U\,\mathcal{F}]_{\mathrm{Pl}}$. Let $\mathbf{K}$ be a quantum Pl-Holant($\mathcal{F}$) gadget with tensor K. By Theorem 6.3.1, $K \in [\mathcal{F}]_{\mathrm{Pl}} = \langle \mathcal{F}, \supset, \subset \rangle_{+,\circ,\otimes}$ is a linear combination of tensors constructed via (6.3.1) and (6.3.2) from $\mathcal{F}, I, \supset$, and $\subset$ by a sequence of $\circ$ and $\otimes$ operations. Applying the same sequence of operations with every F in (6.3.1) or (6.3.2) replaced by $U \cdot F$, we obtain the tensor K' of $\mathbf{K}_{\mathcal{F}\to U\,\mathcal{F}}$. Since the action of U commutes with all of these operations, induction on the steps of the sequence yields $K' = U \cdot K$. Thus $U[\mathcal{F}]_{\mathrm{Pl}} = [U\,\mathcal{F}]_{\mathrm{Pl}}$.

The other two claims follow from the first after we also show that the action of U respects the † operation on signatures. First, since $F^{\ell,r} \in [F]_{\mathrm{Pl}}$ (as the tensor of $\mathbf{V}(F,\ell,r,0)$), the first claim $U[\mathcal{F}]_{\mathrm{Pl}} = [U\,\mathcal{F}]_{\mathrm{Pl}}$ gives a quantum version of (4.4.2), showing that the action of a quantum orthogonal matrix U, like the action of a classical orthogonal matrix, on a tensor only depends on the tensor's underlying signature:

$$(U^{\otimes n}F)^{\ell,r} = (U\cdot F)^{\ell,r} = U\cdot F^{\ell,r} = U^{\otimes\ell}F^{\ell,r}(U^{\otimes r})^{\dagger} \tag{6.3.3}$$

(the idea that the action of a quantum orthogonal matrix on a tensor is independent of the shape of the tensor is not new, and is an instance of a wider phenomenon called "Frobenius reciprocity" in [Gro21; Ban05]). Applying (6.3.3) with $(\ell,r)=(0,n)$ and (4.2.2) give, for signature F,

$$U\cdot F^{\dagger} = U^{\otimes n}(F^{\dagger})^{n,0} = U^{\otimes n}(F^{0,n})^{\dagger} = (F^{0,n}(U^{\otimes n})^{\dagger})^{\dagger} = (U\cdot F^{0,n})^{\dagger} = ((U\cdot F)^{0,n})^{\dagger} = (U\cdot F)^{\dagger}. \tag{6.3.4}$$

Now, applying (6.2.2),

$$U[\mathcal{F}]_{\mathrm{Pl}^{\dagger}} = U[\mathcal{F}\cup\mathcal{F}^{\dagger}]_{\mathrm{Pl}} = [U(\mathcal{F}\cup\mathcal{F}^{\dagger})]_{\mathrm{Pl}} = [U\,\mathcal{F}\cup(U\,\mathcal{F})^{\dagger}]_{\mathrm{Pl}} = [U\,\mathcal{F}]_{\mathrm{Pl}^{\dagger}},$$

giving the third claim. For the second claim, the extra assumption gives

$$U\cdot F^{\top} = U\cdot(\overline{F})^{\dagger} = (U\cdot\overline{F})^{\dagger} = (\overline{U\cdot F})^{\dagger} = (U\cdot F)^{\top},$$

so, similarly, by (6.2.1),

$$U[\mathcal{F}]_{\mathrm{Pl}^{\top}} = U[\mathcal{F}\cup\mathcal{F}^{\top}]_{\mathrm{Pl}} = [U(\mathcal{F}\cup\mathcal{F}^{\top})]_{\mathrm{Pl}} = [U\,\mathcal{F}\cup(U\,\mathcal{F})^{\top}]_{\mathrm{Pl}} = [U\,\mathcal{F}]_{\mathrm{Pl}^{\top}}. \qquad \square$$

The extra condition in the second point of Corollary 6.3.2 is, by (6.3.4), equivalent to $U\,\mathcal{F}^{\top} = (U\,\mathcal{F})^{\top}$, but we state the condition in terms of conjugation instead to emphasize that it always holds for real-valued $\mathcal{F}$.

For an easy but useful first application of Corollary 6.3.2, recall that any quantum permutation matrix U by definition satisfies $U\cdot(=_3) = (=_3)$, and that $\mathcal{EQ} \subset [=_3]_{\mathrm{Pl}}$ by (5.1.1). It follows that $U\,\mathcal{EQ} = \mathcal{EQ}$ for any quantum permutation matrix U. More generally, the tensor of any planar $\mathcal{EQ}$ gadget is invariant under U. These planar $\mathcal{EQ}$ gadgets are exactly the *non-crossing partitions* from the theory of easy quantum groups [BS09]. In particular, $U(=_1)^{1,0} = (=_1)^{1,0}$ and $(=_1)^{0,1}U = (=_1)^{0,1}$, which are equivalent to the useful identities

$$\sum_x U_{xy} = \mathbf{1} \text{ for every } y \text{ and } \sum_y U_{xy} = \mathbf{1} \text{ for every } x. \tag{6.3.5}$$

In fact, the typical definition of a quantum permutation matrix is by (6.3.5) and the second condition in Definition 6.1.4.

The special case of Corollary 6.3.2 for signature grids is a quantum version of Corollary 2.4.1.

Corollary 6.3.3 (The quantum orthogonal Holant theorem)**.** *Let $\mathcal{F}$ and $\mathcal{G}$ be bijective complex-valued signature sets such that $U\mathcal{F} = \mathcal{G}$ for some quantum orthogonal matrix U. Then*

- *$\mathcal{F}$ and $\mathcal{G}$ are Pl-Holant-indistinguishable.*
- *$\mathcal{F}$ and $\mathcal{G}$ are $Pl^\top$-Holant-indistinguishable if $U \cdot \overline{F} = \overline{G}$ for every $\mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}$.*
- *$\mathcal{F}$ and $\mathcal{G}$ are $Pl^\dagger$-Holant-indistinguishable.*

6.3.1 What is a quantum holographic transformation?

If $\mathcal{F} = \{A_X, =_3\}$ and $\mathcal{G} = \{A_Y, =_3\}$ for quantum isomorphic graphs X and Y, then there is a quantum permutation matrix U satisfying $U\mathcal{F} = \mathcal{G}$. Suppose counterfactually that the quantum orthogonal Holant theorem applied to all Holant signature grids, not merely planar ones. Then $\mathcal{F}$ and $\mathcal{G}$ would be Holant-indistinguishable, so, by Proposition 5.1.1, X and Y would be homomorphism-indistinguishable. Then, by Theorem 2.5.1, X and Y are isomorphic. However, there are known examples of graphs that are quantum isomorphic but not isomorphic [Ats+19]. So something must go wrong when applying a quantum holographic transformation in a nonplanar signature grid.

Recall that the classical orthogonal Holant theorem Corollary 2.4.1, which applies to nonplanar signature grids, follows from Valiant's bipartite Holant theorem (Theorem 2.4.1). The proof of the bipartite Holant theorem goes roughly as follows, following the intuition in Figure 2.2 (see [CC17a]): Since $TT^{-1} = I$, we can subdivide each edge of Ω into three edges by introducing two arity-2 vertices, assigned T and T^{-1}, without changing the Holant value. Since Ω is bipartite, partitioned according to whether a vertex is assigned a signature from $\mathcal{F}$ or $\mathcal{F}'$, we can subdivide every edge in such a way that every u assigned $F \in \mathcal{F}$ or v assigned $F' \in \mathcal{F}'$ is now adjacent to only vertices assigned T or T^{-1}, respectively. Then we can associate u and its surrounding T vertices into a single 'gadget' with signature $T \cdot F$, and similarly associate v and its surrounding

T^{-1} vertices into $T \cdot F'$, giving the result. To obtain the classical orthogonal Holant theorem, we set $F' = {=_2}$ and used the fact that $T \cdot (=_2) = (=_2)$ if and only if T is orthogonal.

Two things go wrong with this argument if $T = U$ is a quantum orthogonal matrix with noncommutative entries. First, since in general

$$(U^{\otimes 2})^{-1} = (U^{\otimes 2})^\dagger \neq (U^\dagger)^{\otimes 2} = (U^{-1})^{\otimes 2},$$

the disparate copies of U^{-1} introduced on each edge do not associate into the desired $(U^{\otimes 2})^\dagger$, but into $(U^\dagger)^{\otimes 2}$ (recall that, by definition, $U \cdot \subset = \subset \iff \subset \circ (U^{\otimes 2})^\dagger = \bigcirc$). If we had set $\mathcal{F} = \{=_2\}$, instead, we would similarly see U fail to transform $\mathcal{F}'$ into $\mathcal{G}'$.

Second, even though the original and transformed signatures and the ultimate Holant values are scalars (technically scalar multiples of $\mathbb{1}$, which commute), the intermediate steps of the above proof sketch have a hidden dependence on the Holant value of a signature grid containing vertices assigned 'signature' U. Without further specification, the Holant value for such a signature grid is not even well-defined: the product $\prod_{v \in V} F_v(\sigma(\delta(v)))$ in (2.2.1) does not specify an ordering of V, but when F_v can take noncommutative values, different orderings give different products. When Ω_K is planar, however, the decomposition procedure in the proof of Theorem 6.3.1 produces a sequence of gadgets whose signature matrices multiply to the Holant value, in a sense defining an ordering of Ω_K's vertices.

Both of these obstacles arise specifically due to noncommutativity U's entries, rather than the fact that these entries are operators instead of scalars. If the entries of U are commuting operators, then, as we will see in Corollary 6.4.2 in the next section, $U\mathcal{F} = \mathcal{G}$ does imply that $\mathcal{F}$ and $\mathcal{G}$ are indistinguishable on all, not necessarily planar, Holant grids. In this sense, the above intuition for the proof of the Holant theorem does hold for operator-valued T, as long as the operators commute.

6.4 Quantum orthogonal and permutation groups

In this section, we introduce some definitions from quantum group theory and prove a few lemmas that will play crucial roles in our planar/quantum indistinguishability theorems. We do not provide the full mathematical background on these definitions, and instead refer the reader to [BS09; LMR17; MR20]

The following object was introduced by Wang [Wan95]. It is a *compact matrix quantum groups* (CMQGs), as originally defined by Woronowicz [Wor87].

Definition 6.4.1 (Quantum orthogonal group, O_q^+)**.** The *quantum orthogonal group* $\mathcal{Q}$ of order q is defined by the universal (unital) C^*-algebra $C(\mathcal{Q})$ generated by the entries of a $q \times q$ matrix U, called the *fundamental representation* of $\mathcal{Q}$, subject to the algebraic relations making U a quantum orthogonal matrix (that is, $u_{ij}^* = u_{ij}$ and $\sum_k u_{ik}u_{jk} = \delta_{ij}\,\mathbb{1} = \sum_k u_{ki}u_{kj}$ for every i, j).

Roughly, the universal C^*-algebra generated by the entries $(u_{ij})_{ij}$ of U subject to certain algebraic relations is analogous to the generator-relation presentation of a group, or equivalently the quotient of the free group on the generators by the relations. That is, every relation in $C(\mathcal{Q})$ can be derived from the defining relations and the C^*-algebra axioms. In this sense, imposing additional relations yields a 'subgroup' in the following sense: a *quantum subgroup* $\mathcal{Q} \subset O_q^+$ is defined by the universal C^*-algebra $C(\mathcal{Q})$ generated by the entries of a $q \times q$ matrix U – the fundamental representation of $\mathcal{Q}$ – subject to algebraic relations including those making U a quantum orthogonal matrix.

The following important class of quantum subgroups of O_q^+ was introduced by Wang [Wan98].

Definition 6.4.2 (Quantum permutation group, S_q^+)**.** The *quantum symmetric group* S_q^+ is the quantum subgroup of O_q^+ whose fundamental representation U is subject to exactly the relations making U a quantum permutation matrix.

A *quantum permutation group* $\mathcal{Q}$ of order q is a quantum subgroup of S_q^+ (i.e. the fundamental representation of $\mathcal{Q}$ is subject at least to the relations making it a quantum permutation matrix).

Next, we give a first nontrivial example of a quantum permutation group.

Definition 6.4.3 (Qut(X) [Ban05])**.** For an undirected, unweighted graph X, the *quantum automorphism group* Qut(X) of X is a quantum permutation group whose fundamental representation U is subject to the additional relation $UA_X = A_XU$, where A_X is the adjacency matrix of X.

The notation $C(\mathcal{Q})$ is motivated by the notation $C(\mathrm{Aut}(X))$ for the commutative algebra of continuous complex linear functionals on $\mathrm{Aut}(X)$ (this algebra is isomorphic to the algebra $\mathbb{C}^{\mathrm{Aut}(X)}$ under entrywise multiplication). Indeed, if we also impose the relations that the entries of U commute, then the resulting quantum permutation group is isomorphic to $C(\mathrm{Aut}(X))$ (see e.g. [BBC07, Theorem 1.4]). This isomorphism maps u_{ij} to the characteristic function of the automorphisms sending vertex i to j. Without the commutativity condition, $\mathrm{Qut}(X)$ doesn't actually exist as a group, but we still think of $C(\mathrm{Qut}(X))$ as the algebra of continuous complex linear functionals on it. Hence the absence of commutativity is what makes things "quantum."

We generalize $\mathrm{Qut}(X)$ to the quantum permutation group $\mathrm{Qut}(\mathcal{F})$ for any set $\mathcal{F}$ of arbitrary-arity tensors over $\mathbb{C}$. By (6.3.3), $UA_X = A_XU$ is equivalent to $U^{\otimes 2}A_X^{2,0} = A_X^{2,0}$. This, along with the fact that a classical permutation matrix P defines an automorphism of arity-n tensor F if and only if $P \cdot F = F$, motivates the following definition.

Definition 6.4.4 ($\mathrm{Qut}(\mathcal{F})$)**.** For a set $\mathcal{F}$ of signatures on domain $[q]$, define the *quantum automorphism group* $\mathrm{Qut}(\mathcal{F})$ of $\mathcal{F}$ as the quantum permutation group whose fundamental representation is subject to the additional relations $U\mathcal{F} = \mathcal{F}$.

For graph X, $\mathrm{Qut}(\{A_X\})$ coincides with $\mathrm{Qut}(X)$. In the next subsection, we need a generalization, analogous to Stab_O, of $\mathrm{Qut}(\mathcal{F})$ for the quantum orthogonal group:

Definition 6.4.5 (Stab_{O^+})**.** For a set $\mathcal{F}$ of signatures on domain $[q]$, define the *quantum orthogonal stabilizer group* $\mathrm{Stab}_{O^+}(\mathcal{F}) \subset O_q^+$ of $\mathcal{F}$ to be the quantum subgroup of O_q^+ whose fundamental representation is subject to the additional relations $U\mathcal{F} = \mathcal{F}$.

By (6.3.3), Stab_{O^+} is, like Stab_O, well-defined on sets of tensors, as it is equal to the quantum orthogonal stabilizer group of their underlying signatures.

6.4.1 Quantum orthogonal duality

While quantum groups are highly abstract, they admit a duality theory for O_q^+ analogous to the duality for O_q in Section 4.5. In particular, a quantum subgroup of O_q^+ is completely determined by its very concrete *intertwiner space*:

Definition 6.4.6 ($C_{\mathcal{Q}}$)**.** For an order-q quantum permutation group $\mathcal{Q}$ with fundamental representation U, let

$$C_{\mathcal{Q}}(\ell, r) = \{F \in {}_\ell\mathcal{V}(\mathbb{C}^q)_r : \ell, r \geq 0, U^{\otimes \ell}F = FU^{\otimes r}\} = \{F \in {}_\ell\mathcal{V}(\mathbb{C}^q)_r : \ell, r \geq 0, U \cdot F = F\}$$

be the *intertwiner space* of $\mathcal{Q}$ (where the second equality applies (6.1.3)).

In other words, $C_{\mathcal{Q}}$ is the quantum analogue of $\mathcal{V}(\mathbb{C}^q)^Q = C_Q$ (recall (4.5.1)) for a classical subgroup $Q \subset O_q$: it is the set of tensors invariant under the action of $\mathcal{Q}$, specified by the 'universal' fundamental representation. By definition, $\mathcal{F} \subset C_{\text{Stab}_{O^+}(\mathcal{F})}$.

It is well-known (see e.g. [BS09]) and follows from the proof of Corollary 6.3.2 that, for any $\mathcal{Q} \subset O_q^+$, $C_{\mathcal{Q}}$ is a TCWD (alternatively, this follows from the conclusion of Corollary 6.3.2 and the fact that $[\mathcal{F}]_{\text{Pl}^\dagger}$ is a TCWD). The following important result, referred to as *Tannaka-Krein duality*, was proved in [Wor88], and expressed in the following form in [BS09; MR20] and elsewhere. It is the quantum/asymmetric analogue of Theorem 4.5.1. It says that a quantum subgroup of O_q^+ is determined by its intertwiner space and that every TCWD defines a quantum subgroup of O_q^+.

Theorem 6.4.1. *The mapping $\mathcal{Q} \mapsto C_{\mathcal{Q}}$ induces a bijection between quantum subgroups $\mathcal{Q} \subset O_q^+$ and tensor categories with duals.*

If $C_{\mathcal{Q}} \subset C_{\mathcal{Q}'}$, then $\mathcal{Q}'$ is a quantum subgroup of $\mathcal{Q}$, because the fundamental representation of $\mathcal{Q}'$ is subject to all of the relations to which the fundamental representation of $\mathcal{Q}$ is subject. Therefore, as in the classical version Theorem 4.5.1, the bijection in Theorem 6.4.1 is inclusion-reversing.

Now we obtain the quantum/asymmetric version of Corollary 4.5.1 via a similar proof.

Corollary 6.4.1. *For any $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$, we have $C_{\text{Stab}_{O^+}(\mathcal{F})} = \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\dagger}$.*

Proof. By Theorem 6.4.1, there is a quantum subgroup $\mathcal{Q} \subset O_q^+$ such that $\langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\dagger} = C_{\mathcal{Q}}$. Suppose toward contradiction that $\mathcal{Q} \neq \text{Stab}_{O^+}(\mathcal{F})$, so, applying Theorem 4.5.1 again, $C_{\text{Stab}_{O^+}(\mathcal{F})}$ is a TCWD but $C_{\text{Stab}_{O^+}(\mathcal{F})} \neq C_{\mathcal{Q}} = \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\dagger}$. Since $\mathcal{F} \subset C_{\text{Stab}_{O^+}(\mathcal{F})}$ and $\langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\dagger}$ is the smallest TCWD containing $\mathcal{F}$, we must have $\langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\dagger} \subset C_{\text{Stab}_{O^+}(\mathcal{F})}$. Hence there exists some $A \in C_{\text{Stab}_{O^+}(\mathcal{F})} \setminus \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\dagger}$. But then the fundamental representation U of $\text{Stab}_{O^+}(\mathcal{F})$ is subject to the relation $AU = UA$, which does not follow from $U\mathcal{F} = \mathcal{F}$ or U being a quantum orthogonal matrix, contradicting the definition of $\text{Stab}_{O^+}(\mathcal{F})$. ☐

Now the characterization of quantum planar gadget tensors in Theorem 6.3.1 gives the key result for proving our quantum/planar indistinguishability theorems, a quantum/asymmetric version of Theorem 4.5.2. However, we must now drop Pl-Holant from our list of three planar Holants, because the quantum gadget algebra must be a TCWD to apply Corollary 6.4.1, and thus must be closed under some form of transpose.

Theorem 6.4.2. *The following hold for any $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$:*

- $C_{\mathrm{Stab}_{O^+}(\mathcal{F})} = [\mathcal{F}]_{Pl^\top}$ *if $\mathcal{F}$ is conjugate-closed,*
- $C_{\mathrm{Stab}_{O^+}(\mathcal{F})} = [\mathcal{F}]_{Pl^\dagger}$.

Proof. This follows from Theorem 6.3.1 and Corollary 6.4.1. □

For complex-valued $\mathcal{F}$, Proposition 4.4.2 and Theorem 6.3.1 give $[\mathcal{F}] = \langle \mathcal{F}, \supset, \times \rangle_{+,\circ,\otimes,\top}$ and $[\mathcal{F}]_{\mathrm{Pl}^\top} = \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\top}$ and the only difference between general and quantum $\mathrm{Pl}^\top$-Holant gadgets is $\times$. Given the results of Corollary 4.5.1 and Corollary 6.4.1 – that $C_{\mathrm{Stab}_{O^+}(\mathcal{F})} = \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\dagger}$ and $C_{\mathrm{Stab}_{O}(\mathcal{F})} = \langle \mathcal{F}, \supset, \times \rangle_{+,\circ,\otimes,\dagger}$ – one should also expect $\times$, or really the lack thereof, to play a role in the 'quantumness' of the quantum orthogonal group. Indeed, a direct calculation shows that, for any quantum subgroup $\mathcal{Q} \subset O_q^+$ with fundamental representation U,

$$\times \in C_{\mathcal{Q}} \iff U^{\otimes 2} \times = \times\, U^{\otimes 2} \iff \text{the entries of } U \text{ commute.} \tag{6.4.1}$$

One symptom of the loss of quantumness due to the presence of $\times$ is the result promised in Section 6.3.1: a quantum orthogonal matrix with commuting entries can perform a holographic transformation on any, not necessarily planar, signature grid. In fact, we obtain a stronger version: a generalization of (the $\mathbb{C}$-valued case of) Corollary 4.4.1.

Corollary 6.4.2. *Let $\mathcal{F}, \mathcal{G} \subset \mathcal{V}(\mathbb{C}^q)$. If U is a quantum orthogonal matrix with commuting entries such that $U\,\mathcal{F} = \mathcal{G}$, then $U[\mathcal{F}] = [\mathcal{G}]$.*

Proof. By (6.4.1), $U(\mathcal{F}, \times) = (\mathcal{G}, \times)$. Furthermore, commutativity implies that $(U^{\otimes n})^\dagger = (U^{\otimes n})^\top$, so $(U \cdot F)^\top = U \cdot F^\top$, or equivalently $\overline{U \cdot F} = U \cdot \overline{F}$. Now, by Corollary 6.3.2,

$$U[\mathcal{F}] = U[\mathcal{F}, \times]_{\mathrm{Pl}^\top} = [\mathcal{G}, \times]_{\mathrm{Pl}^\top} = [\mathcal{G}].$$

□

Chapter 7

Planar #CSP and Quantum Isomorphism

Most of this chapter is based on [CY24] (joint work with Jin-Yi Cai), and Section 7.4 is based on [CMY26] (joint work with Jin-Yi Cai and Ashwin Maran).

The main goal of this chapter is to extend to #CSP the result of Mančinska and Roberson [MR20] that two graphs are quantum isomorphic if and only if they are homomorphism-indistinguishable over planar graphs. Recall the construction, shown in Figure 2.1, of the Holant($\mathcal{F} \cup \mathcal{EQ}$) grid $\Omega(I_K)$ modeling the #CSP(A_X) instance I_K whose partition function counts the number of homomorphisms from graph K to X. Since the underlying graph of $\Omega(I_K)$ (the constraint-variable incidence graph of I_K) is a subdivision of K, $\Omega(I_K)$ is planar if and only if K is planar. Therefore, when we define *planar* #CSP to be the restriction of #CSP to instances whose constraint-variable incidence graphs are planar, Mančinska and Roberson's theorem is a special case of the following main theorem of this chapter.

Theorem 7.0.1 (Theorem 7.1.1, simplified)**.** *Bijective* $\mathcal{F}, \mathcal{G} \subset \mathcal{S}(\mathbb{R}^q)$ *are quantum isomorphic if and only if they are planar-*#CSP*-indistinguishable.*

The ($\Longrightarrow$) direction follows from the quantum orthogonal Holant theorem (Corollary 6.3.3) from Chapter 6. For the converse direction, we develop theory parallel to Mančinska and Roberson's original proof, and parallel to the intertwiner proof for #CSP indistinguishability in Chapter 5. First, we apply the quantum orthogonal duality Theorem 6.4.2 from Chapter 6 to prove a quantum version of Theorem 5.1.1: quantum planar #CSP($\mathcal{F}$) gadget tensors capture all tensors invariant under the action of Qut($\mathcal{F}$). Then we study the *orbits* of Qut($\mathcal{F}$) [LMR17], which behave similarly to the classical orbits of Aut($\mathcal{F}$). The final step, again following the final step of the classical proof, uses the fact that, for connected $\mathcal{F}$ and $\mathcal{G}$, if a vertex of $\mathcal{F}$ and a vertex of $\mathcal{G}$ are in the same orbit of Qut($\mathcal{F} \oplus \mathcal{G}$), then $\mathcal{F}$ and $\mathcal{G}$ are quantum isomorphic.

In Section 7.3, we extend the connection between quantum isomorphism and nonlocal games. We define a graph isomorphism game for complex-weighted directed graphs and prove that, as in the unweighted case, complex-weighted graphs F and G admit the same number of homomorphisms from all planar graphs if and only if there is a perfect quantum commuting strategy for the (F, G)-isomorphism game. Finally, in Section 7.4 we present an application of the theory of quantum planar #CSP($\mathcal{F}$) gadget expressivity and the orbits of Qut($\mathcal{F}$) to the computational complexity of counting homomorphisms from planar graphs. The complexity dichotomy of [CMY26] relies on finding a binary planar #CSP(A_X) gadget whose matrix has distinct diagonal entries. We show using the undecidability of quantum isomorphism that determining the existence of a such a gadget separating two given diagonal entries is in general undecidable.

7.1 From quantum orthogonal transformation to quantum isomorphism

First, we must define planar #CSP. Recall from Section 2.3 that every #CSP($\mathcal{F}$) instance corresponds I to a unique Holant($\mathcal{F} \mid \mathcal{EQ}$) grid $\Omega(I)$, which we can view as a Holant($\mathcal{F} \cup \mathcal{EQ}$) grid.

Definition 7.1.1. For signature set $\mathcal{F} \subset \mathcal{S}(\mathbb{C}^q)$, a Pl-#CSP($\mathcal{F}$), $\mathrm{Pl}^{\top}$-#CSP($\mathcal{F}$), or $\mathrm{Pl}^{\dagger}$-#CSP($\mathcal{F}$) instance is a #CSP($\mathcal{F}$) instance I for which $\Omega(I)$ is a Pl-Holant($\mathcal{F} \cup \mathcal{EQ}$), $\mathrm{Pl}^{\top}$-Holant($\mathcal{F} \cup \mathcal{EQ}$), or $\mathrm{Pl}^{\dagger}$-Holant($\mathcal{F} \cup \mathcal{EQ}$) grid, respectively.

Say $\mathcal{F}$ and $\mathcal{G}$ are Pl-#CSP, $\mathrm{Pl}^{\top}$-#CSP, or $\mathrm{Pl}^{\dagger}$-#CSP-indistinguishable if $\mathcal{F} \cup \mathcal{EQ}$ and $\mathcal{G} \cup \mathcal{EQ}$ are Pl-Holant, $\mathrm{Pl}^{\top}$-Holant, or $\mathrm{Pl}^{\dagger}$-Holant-indistinguishable, respectively.

Recall from the proof of Proposition 5.1.1 that every k-multilabeled #CSP($\mathcal{F}$) instance corresponds to a Holant($\mathcal{F} \cup \mathcal{EQ}$) gadget with the same signature and vice-versa. The 'vice-versa' requires transforming a Holant($\mathcal{F} \cup \mathcal{EQ}$) gadget into a Holant($\mathcal{F} \mid \mathcal{EQ}$) gadget using (2.3.2). This transformation contracts and subdivides edges, both of which do not affect planarity (and the orientations of the merged $\mathcal{EQ}$ vertices are irrelevant because $\mathcal{EQ}$ signatures and symmetric). Therefore Pl-Holant($\mathcal{F} \mid \mathcal{EQ}$), $\text{Pl}^\top$-Holant($\mathcal{F} \mid \mathcal{EQ}$), and $\text{Pl}^\dagger$-Holant($\mathcal{F} \mid \mathcal{EQ}$) are equivalent to Pl-Holant($\mathcal{F} \cup \mathcal{EQ}$), $\text{Pl}^\top$-Holant($\mathcal{F} \cup \mathcal{EQ}$), and $\text{Pl}^\dagger$-Holant($\mathcal{F} \cup \mathcal{EQ}$), respectively. So define a k-multilabeled Pl-#CSP($\mathcal{F}$), $\text{Pl}^\top$-#CSP($\mathcal{F}$), or $\text{Pl}^\dagger$-#CSP($\mathcal{F}$) instance as a k-multilabeled #CSP($\mathcal{F}$) instance for which the corresponding $\mathcal{F} \cup \mathcal{EQ}$ grid is a Pl-Holant, $\text{Pl}^\top$-Holant, or $\text{Pl}^\dagger$-Holant grid, respectively. Furthermore, since every equality signature is realizable as a planar $=_3$-gadget by (5.1.1), we have $[\mathcal{F}, \mathcal{EQ}]_{\text{Pl}} = [\mathcal{F}, =_3]_{\text{Pl}}$ and similarly for the other planar variants (again, the orientations of signatures in $\mathcal{EQ}$ are irrelevant). Hence we have arranged the definitions to obtain a planar version of Proposition 5.1.1:

Proposition 7.1.1. *Let $\mathcal{F} \subset \mathcal{S}(\mathbb{C}^q)$. Then $[\mathcal{F}, =_3]_{Pl}$, $[\mathcal{F}, =_3]_{Pl^\top}$, and $[\mathcal{F}, =_3]_{Pl^\dagger}$ are exactly the spaces of all quantum multilabeled Pl-#CSP($\mathcal{F}$), $Pl^\top$-#CSP($\mathcal{F}$), and $Pl^\dagger$-#CSP($\mathcal{F}$) signatures, respectively.*

We now formally state the main theorem of this chapter.

Theorem 7.1.1. *The following hold for bijective complex-valued signature sets $\mathcal{F}$ and $\mathcal{G}$.*

- *If $\mathcal{F}$ and $\mathcal{G}$ are conjugate-closed and satisfy $F \leftrightsquigarrow G$ if and only if $\overline{F} \leftrightsquigarrow \overline{G}$, then $\mathcal{F}$ and $\mathcal{G}$ are $Pl^\top$-#CSP-indistinguishable if and only if $\mathcal{F} \cong_{qc} \mathcal{G}$.*
- *If $\mathcal{F}$ and $\mathcal{G}$ consist of signatures with arity at most 2, then $\mathcal{F}$ and $\mathcal{G}$ are $Pl^\top$-#CSP-indistinguishable if and only if $\mathcal{F} \cong_{qc} \mathcal{G}$.*
- *$\mathcal{F}$ and $\mathcal{G}$ are $Pl^\dagger$-#CSP-indistinguishable if and only if $\mathcal{F} \cong_{qc} \mathcal{G}$.*

We next show the #CSP corollary of Theorem 6.4.2, or the quantum/planar version of Theorem 5.1.1: quantum planar k-multilabeled #CSP($\mathcal{F}$) instance tensors are exactly the tensors invariant under Qut($\mathcal{F}$). Theorem 5.1.1 removed the conjugate-closed requirement using the fact that $[\mathcal{F}, =_3]$ is closed under entrywise product. This construction is in general nonplanar. However, it is planar for signatures of arity one or two.

Proposition 7.1.2. *For any $\mathcal{F} \in \mathcal{S}(\mathbb{K}^n)$, the spaces $_{\le 2}[\mathcal{F}, =_3]_{Pl}$, $_{\le 2}[\mathcal{F}, =_3]_{Pl^\top}$ and $_{\le 2}[\mathcal{F}, =_3]_{Pl^\dagger}$ are closed under entrywise product.*

Proof. We work with Pl-Holant; the proofs for the other variants are identical. Given $F, G \in {}_2[\mathcal{F}, =_3]_{\mathrm{Pl}}$, let **K** and **J** be the gadgets with tensors $F^{1,1}$ and $G^{1,1}$, respectively. Then the gadget consisting of **K** and **J** running in parallel between two vertices assigned $=_3$ (one on the left and one on the right) is planar and has tensor $F^{1,1} \bullet G^{1,1}$, so signature $F \bullet G$. Alternatively,

$$F^{1,1} \bullet G^{1,1} = (=_3)^{1,2} \circ (F^{1,1} \otimes G^{1,1}) \circ (=_3)^{2,1}.$$

Similarly, if $F, G \in {}_1[\mathcal{F}, =_3]_{\mathrm{Pl}}$, then construct $F \bullet G$ by connecting F and G to the same copy of $=_3$:

$$F^{1,0} \bullet G^{1,0} = (=_3)^{1,2} \circ (F^{1,0} \otimes G^{1,0}). \qquad \square$$

Theorem 7.1.2. *The following hold for any $\mathcal{F} \subset \mathcal{S}(\mathbb{C}^q)$.*

- $C_{\mathrm{Qut}(\mathcal{F})} = [\mathcal{F}, =_3]_{Pl^\top}$ *if $\mathcal{F}$ is conjugate-closed or consists of signatures of arity at most 2.*
- $C_{\mathrm{Qut}(\mathcal{F})} = [\mathcal{F}, =_3]_{Pl^\dagger}$.

Proof. The second claim follows from Theorem 6.4.2 and the fact that $C_{\mathrm{Stab}_{O^+}(\mathcal{F}, =_3)} = C_{\mathrm{Qut}(\mathcal{F})}$ (recall that a quantum permutation matrix is a quantum orthogonal matrix that intertwines $=_3$). The first claim for conjugate-closed $\mathcal{F}$ follows similarly. Suppose $\mathcal{F}$ consists of signatures of arity at most 2. By Proposition 7.1.2 and the fact that $(=_1)^{\otimes n} \in {}_n[\mathcal{F}, =_3]_{\mathrm{Pl}^\top}$ for every n, Proposition 5.1.2 applies to the vector space ${}_{\leq n}[\mathcal{F}, =_3]_{\mathrm{Pl}^\top}$, so ${}_1[\mathcal{F}, =_3]_{\mathrm{Pl}^\top}$ and ${}_2[\mathcal{F}, =_3]_{\mathrm{Pl}^\top}$ have bases of 0-1-valued signatures, which we denote ${}_1\mathcal{O}$ and ${}_2\mathcal{O}$, respectively. By assumption,

$$\mathcal{F} \subset {}_{\leq 2}[\mathcal{F}, =_3]_{\mathrm{Pl}^\top} \subset [{}_1\mathcal{O}, {}_2\mathcal{O}]_{\mathrm{Pl}^\top},$$

so $[\mathcal{F}, =_3]_{\mathrm{Pl}^\top} = [{}_1\mathcal{O}, {}_2\mathcal{O}, =_3]_{\mathrm{Pl}^\top}$. Hence $[\mathcal{F}, =_3]_{\mathrm{Pl}^\top}$ has real-valued generators, so is conjugate-closed. Unary signatures F satisfy $F^\top = F$, binary signatures F satisfy $F^\top = F^{(1)} \in [\mathcal{F}]_{\mathrm{Pl}}$, and $(=_3)^\top = (=_3)$ by symmetry. Hence $[\mathcal{F}, =_3]_{\mathrm{Pl}^\top} = [\mathcal{F}, =_3]_{\mathrm{Pl}}$ and $[\mathcal{F}, =_3]_{\mathrm{Pl}^\dagger} = [\mathcal{F}, =_3, \{\overline{F} \mid F \in \mathcal{F}\}]_{\mathrm{Pl}}$. But, since $[\mathcal{F}, =_3]_{\mathrm{Pl}^\top}$ is conjugate-closed,

$$[\mathcal{F}, =_3]_{\mathrm{Pl}^\top} = [\mathcal{F}, =_3]_{\mathrm{Pl}} = [\mathcal{F}, =_3, \{\overline{F} \mid F \in \mathcal{F}\}]_{\mathrm{Pl}} = [\mathcal{F}, =_3]_{\mathrm{Pl}^\dagger} = C_{\mathrm{Qut}(\mathcal{F})}. \qquad \square$$

The n-orbitals of the classical permutation group $\mathrm{Aut}(\mathcal{F})$ played a key role in Chapter 5. Lupini, Mančinska, and Roberson [LMR17] show that, for $n \leq 2$, the quantum permutation group $\mathrm{Qut}(\mathcal{F})$ admits n-orbitals with similar properties.

Definition 7.1.2 ([LMR17, Lemma 3.2, Lemma 3.4, Definition 3.5])**.** For an order-q quantum permutation group $\mathcal{Q}$ with fundamental representation U, define relations $\sim_1$ on X and $\sim_2$ on $[q] \times [q]$ by

$$x \sim_1 y \iff u_{xy} \neq 0, \qquad (x_1, x_2) \sim_2 (y_1, y_2) \iff u_{x_1 y_1} u_{x_2 y_2} \neq 0.$$

Then $\sim_1$ and $\sim_2$ are equivalence relations, and their equivalence classes are called the 1-orbitals (or *orbits*) and 2-orbitals (or just *orbitals*), respectively, of $\mathcal{Q}$.

The next two lemmas provide key connections between the intertwiners and orbits and orbitals of $\mathrm{Qut}(\mathcal{F})$, nontrivial quantum versions of (5.1.3).

Lemma 7.1.1 ([LMR17, Lemma 3.7])**.** *Let $\mathcal{Q}$ be an order-q quantum permutation group and $v \in {}_1\mathcal{V}(\mathbb{C}^q)_0$. Then $v \in C_{\mathcal{Q}}$ if and only if v is constant on the orbits of $\mathcal{Q}$ (that is, $x \sim_1 y$ implies $v_x = v_y$).*

Lemma 7.1.2 ([LMR17, Theorem 3.10])**.** *Let $\mathcal{Q}$ be an order-q quantum permutation group and $M \in {}_1\mathcal{V}(\mathbb{C}^q)_1$. Then $M \in C_{\mathcal{Q}}$ if and only if M is constant on the orbitals of $\mathcal{Q}$ (that is, $(i, i') \sim_2 (j, j')$ implies $M_{ii'} = M_{jj'}$.*

We exploit this nice behavior of quantum n-orbitals for $n \leq 2$ to obtain a quantum/planar version of Corollary 5.1.1

Corollary 7.1.1. *For $n \leq 2$, the following are equivalent for $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$ and $\mathbf{x}, \mathbf{y} \in [q]^n$.*

(1) $\mathbf{x}$ and $\mathbf{y}$ are in the same n-orbital of $\mathrm{Qut}(\mathcal{F})$,

(2) $Z^{[n]\mapsto\mathbf{x}}(\mathbf{K}) = Z^{[n]\mapsto\mathbf{y}}(\mathbf{K})$ for every n-multilabeled Pl†-#CSP($\mathcal{F}$) instance $\mathbf{K}$.

(3) $F(\mathbf{x}) = F(\mathbf{y})$ for every $F \in [\mathcal{F}, =_3]_{Pl^\dagger}$.

If $\mathcal{F}$ is conjugate-closed or consists of signatures of arity at most 2, then the following conditions are equivalent to the first three:

(4) $Z^{[n]\mapsto\mathbf{x}}(\mathbf{K}) = Z^{[n]\mapsto\mathbf{y}}(\mathbf{K})$ for every n-multilabeled Pl$^\top$-#CSP($\mathcal{F}$) instance $\mathbf{K}$.

(5) $F(\mathbf{x}) = F(\mathbf{y})$ for every $F \in [\mathcal{F}, =_3]_{Pl^\top}$.

Proof. The equivalences of (2) and (3) and of (4) and (5) follow from Proposition 7.1.1 and the fact that all instances/gadgets agree on $\mathbf{x}$ and $\mathbf{y}$ if and only if all quantum instances/gadgets agree on $\mathbf{x}$ and $\mathbf{y}$. The equivalence of (3) and (5) follows from Theorem 7.1.2. Finally, we show the equivalence of (1) and (3). Theorem 7.1.2 gives $C_{\text{Qut}(\mathcal{F})} = [\mathcal{F}, =_3]_{\text{Pl}^\dagger}$. If $\mathbf{x}, \mathbf{y}$ are in different n-orbitals of $\text{Qut}(\mathcal{F})$, then the indicator signature $\mathbb{1}[\text{Qut}(\mathcal{F})\,\mathbf{x}]$ of the orbit of $\mathbf{x}$ satisfies $\mathbb{1}[\text{Qut}(\mathcal{F})\,\mathbf{x}](\mathbf{x}) = 1 \neq 0 = \mathbb{1}[\text{Qut}(\mathcal{F})\,\mathbf{x}](\mathbf{y})$ and is in $C_{\text{Qut}(\mathcal{F})} = [\mathcal{F}, =_3]_{\text{Pl}^\dagger}$ by Lemma 7.1.1 or Lemma 7.1.2. Conversely, if F is constant on the n-orbitals of $\text{Qut}(\mathcal{F})$, then, by Lemma 7.1.1 or Lemma 7.1.2, $F(\mathbf{x}) = F(\mathbf{y})$ for every $F \in C_{\text{Qut}(\mathcal{F})} = [\mathcal{F}, =_3]_{\text{Pl}^\dagger}$. □

Again, with $I := [q]^n$ in Proposition 5.1.2, the sets ${}_1\mathcal{O}$ and ${}_2\mathcal{O}$ of 0-1-valued basis signatures in the proof of Theorem 7.1.2 are exactly the sets of indicator signatures of the orbits and orbitals of $\text{Qut}(\mathcal{F})$, respectively. As discussed in [LMR17, Section 5], one can give an analogous definition of n-orbitals of a quantum permutation group for $n > 2$, but the proof of transitivity of this relation breaks down due to noncommutativity. From the perspective of this section, another reason for the lack of well-defined n-orbitals for $n > 2$ is the nonplanarity of the entrywise product construction beyond arity two, which prevents applying Proposition 5.1.2 to realize the indicator signatures of these higher-order orbitals.

7.2 The quantum isomorphism converse

Now we prove a nontrivial quantum version of Proposition 5.2.2. The result is a generalization from graphs to general signature sets of [LMR17, Theorem 4.5], with a similar proof. Let $J \in {}_1\mathcal{V}_1$ be the all-ones matrix. If $\mathcal{F}$ and $\mathcal{G}$ contain J, then they are connected in the sense of Definition 5.2.3.

Lemma 7.2.1. *Let $\mathcal{F}$ and $\mathcal{G}$ be $\mathbb{C}$-valued signature sets with domains $V(\mathcal{F})$ and $V(\mathcal{G})$, respectively (not necessarily of the same size). Assume that $\mathcal{F}$ and $\mathcal{G}$ contain no unary signatures and $\mathcal{F} \ni J \leftrightsquigarrow J \in \mathcal{G}$. If there are some $\hat{x} \in V(\mathcal{F})$, $\hat{y} \in V(\mathcal{G})$ in the same orbit of $\text{Qut}(\mathcal{F} \oplus \mathcal{G})$, then $|V(\mathcal{F})| = |V(\mathcal{G})|$ and $\mathcal{F} \cong_{qc} \mathcal{G}$.*

Proof. Let U be the fundamental representation of $\text{Qut}(\mathcal{F} \oplus \mathcal{G})$, so U is indexed by $V(\mathcal{F}) \sqcup V(\mathcal{G})$. By assumption, $J \oplus J \in \mathcal{F} \oplus \mathcal{G}$, so, by Lemma 7.1.2, $J \oplus J$ is constant on the orbitals of $\text{Qut}(\mathcal{F} \oplus \mathcal{G})$. For every $x, x', x'' \in V(\mathcal{F})$ and $y \in V(\mathcal{G})$, we have $(J \oplus J)_{x'x} = 1 \neq 0 = (J \oplus J)_{x''y}$, so (x', x) and

(x'', y) cannot be in the same orbital of $\mathrm{Qut}(\mathcal{F} \oplus \mathcal{G})$. Similarly, (x, x') and (x'', y) cannot be in the same orbital of $\mathrm{Qut}(\mathcal{F} \oplus \mathcal{G})$, so

$$u_{xy} u_{x'x''} = 0 \text{ and } u_{xx''} u_{x'y} = 0. \tag{7.2.1}$$

For each $x \in V(\mathcal{F})$, let $p_x = \sum_{y \in V(\mathcal{G})} u_{xy}$. Recall from (6.3.5) that the rows and columns of U sum to $\mathbf{1}$. Thus, for $x, x' \in V(\mathcal{F})$,

$$p_x - p_x p_{x'} = p_x(\mathbf{1} - p_{x'}) = \left(\sum_{y \in V(\mathcal{G})} u_{xy} \right) \left(\sum_{x'' \in V(\mathcal{F})} u_{x'x''} \right) = 0$$

by the first equation in (7.2.1), so $p_x = p_x p_{x'}$. Similarly, applying the second equation in (7.2.1),

$$p_{x'} - p_x p_{x'} = (\mathbf{1} - p_x) p_{x'} = \left(\sum_{x'' \in V(\mathcal{F})} u_{xx''} \right) \left(\sum_{y \in V(\mathcal{G})} u_{x'y} \right) = 0,$$

so $p_{x'} = p_x p_{x'}$. Therefore $p_x = p_{x'}$. If we define $p_y = \sum_{x \in V(\mathcal{F})} u_{xy}$, for $y \in V(\mathcal{G})$, then a symmetric calculation shows $p_y = p_{y'}$ for any $y, y' \in V(\mathcal{G})$. Finally, for $x \in V(\mathcal{F})$ and $y \in V(\mathcal{G})$,

$$|V(\mathcal{F})| p_x = \sum_{x \in V(\mathcal{F})} p_x = \sum_{x \in V(\mathcal{F}), y \in V(\mathcal{G})} u_{xy} = \sum_{y \in V(\mathcal{G})} p_y = |V(\mathcal{G})| p_y. \tag{7.2.2}$$

Recall from Definition 6.1.4 of a quantum permutation matrix that the entries of a row or column of U are mutually orthogonal idempotents. For any $x \in V(\mathcal{F})$ and $y \in V(\mathcal{G})$, applying this fact and (7.2.2) gives

$$\begin{aligned} |V(\mathcal{F})| u_{xy} &= |V(\mathcal{F})| \sum_{y' \in V(\mathcal{G})} u_{xy'} u_{xy} = (|V(\mathcal{F})| p_x) u_{xy} \\ &= (|V(\mathcal{G})| p_y) u_{xy} = |V(\mathcal{G})| \sum_{x' \in V(\mathcal{F})} u_{x'y} u_{xy} = |V(\mathcal{G})| u_{xy}. \end{aligned} \tag{7.2.3}$$

Specializing to $u_{\hat{x}\hat{y}} \neq 0$, we get $|V(\mathcal{F})| = |V(\mathcal{G})|$. Thus $p_x = p_y$, so p_z does not depend on the choice of $z \in V(\mathcal{F}) \sqcup V(\mathcal{F})$. Call this common element p. If $p_{\hat{x}} = 0$, then multiplying $p_{\hat{x}}$ by $u_{\hat{x}\hat{y}}$ would produce $u_{\hat{x}\hat{y}} = 0$, a contradiction. Hence $p_{\hat{x}} \neq 0$, so $p \neq 0$.

Consider the C^*-subalgebra C' of $C(\mathrm{Qut}(\mathcal{F} \oplus \mathcal{G}))$ generated by u_{xy} for $x \in V(\mathcal{F}), y \in V(\mathcal{G})$. The first two equalities in (7.2.3) show $p u_{xy} = u_{xy}$, and a similar calculation shows $u_{xy} p = u_{xy}$. Thus p is the identity in C'. Consider the square matrix $U' = (u_{xy})$ for $x \in V(\mathcal{F}), y \in V(\mathcal{G})$ – the top right quadrant of U. The entries of U', being entries of the quantum permutation matrix

U, are still mutually orthogonal projectors by condition (2) of Definition 6.1.4. Furthermore, the rows and columns of U' sum to p, the identity in C', so U' is a quantum orthogonal matrix over C'. Thus U' is a quantum permutation matrix over C'.

Any corresponding $F \in \mathcal{F}$ and $G \in \mathcal{G}$ satisfy $U \cdot (F \oplus G) = F \oplus G$. Suppose F and G have arity n. By assumption, $n \geq 2$. By (6.3.3), $U^{\otimes n-1}(F \oplus G)^{n-1,1} = (F \oplus G)^{n-1,1}U$. For $x_1, \dots, x_{n-1} \in V(\mathcal{F})$, and $y \in V(\mathcal{G})$, using the property that $F \oplus G$ is zero on mixed inputs from $V(\mathcal{F})$ and $V(\mathcal{G})$,

$$\begin{aligned}
((U')^{\otimes n-1} G^{n-1,1})_{x_1 \dots x_{n-1}, y} &= \sum_{z_1,\dots,z_{n-1} \in V(\mathcal{G})} u_{x_1 z_1} \dots u_{x_{n-1} z_{n-1}} G_{z_1 \dots z_{n-1} y} \\
&= \sum_{z_1,\dots,z_{n-1} \in V(\mathcal{F}) \sqcup V(\mathcal{G})} u_{x_1 z_1} \dots u_{x_{n-1} z_{n-1}} (F \oplus G)_{z_1 \dots z_{n-1} y} \\
&= (U^{\otimes n-1}(F \oplus G)^{n-1,1})_{x_1 \dots x_{n-1}, y} \\
&= ((F \oplus G)^{n-1,1} U)_{x_1 \dots x_{n-1}, y} \\
&= \sum_{z \in V(\mathcal{F}) \sqcup V(\mathcal{G})} (F \oplus G)_{x_1 \dots x_{n-1} z} u_{zy} \\
&= \sum_{z \in V(\mathcal{F})} F_{x_1 \dots x_{n-1} z} u_{zy} \\
&= (F^{n-1,1} U')_{x_1 \dots x_{n-1}, y}.
\end{aligned}$$

Thus $(U')^{\otimes n-1} G^{n-1,1} = F^{n-1,1} U'$, so by (6.3.3), $U' \cdot G = F$. $\square$

The extra conditions that $\mathcal{F}$ and $\mathcal{G}$ contain J but no unary signatures are not a significant obstacle because, as in Remark 5.2.1, we can WLOG assume that $\mathcal{F}$ and $\mathcal{G}$ contain the tensors of all $\text{Pl}^\dagger$-Holant($\mathcal{F}$) and $\text{Pl}^\dagger$-Holant($\mathcal{G}$) gadgets. This approach, suggested by David Roberson, is more efficient than the original proof in [CY24].

Corollary 7.2.1. *Let $\mathcal{F}$ and $\mathcal{G}$ be bijective complex-valued signature sets of domains $V(\mathcal{F})$ and $V(\mathcal{G})$, respectively. If there exist $x \in V(\mathcal{F})$ and $y \in V(\mathcal{G})$ such that $F(x) = G(y)$ for every $[\mathcal{F}, =_3]_{Pl^\dagger} \ni F \leftrightsquigarrow G \in [\mathcal{G}, =_3]_{Pl^\dagger}$, then $\mathcal{F} \cong_{qc} \mathcal{G}$.*

If $\mathcal{F}$ and $\mathcal{G}$ are conjugate-closed or consist of signatures of arity at most 2, then the same conclusion holds with $Pl^\top$-Holant in place of $Pl^\dagger$-Holant.

Proof. Since Corollary 7.1.1 applies in both cases, the proofs for $\text{Pl}^\dagger$-Holant and $\text{Pl}^\top$-Holant will be identical. We prove the statement for $\text{Pl}^\dagger$-Holant.

Construct $\mathcal{F}'$ by adding $=_1$ to $\mathcal{F}$, then replacing every unary $F \in \mathcal{F}$ (including $F = (=_1)$) with the binary $F \otimes (=_1) \in [\mathcal{F}, =_3]_{\mathrm{Pl}^\dagger}$, and construct $\mathcal{G}'$ similarly. Since every $\mathrm{Pl}^\dagger$-Holant$(\mathcal{F}', =_3)$ gadget $\mathbf{K}$ is a $\mathrm{Pl}^\dagger$-Holant$(\mathcal{F}, =_3)$ gadget, and $\mathbf{K}_{(\mathcal{F}',=_3)\to(\mathcal{G}',=_3)}$ is the corresponding $\mathrm{Pl}^\dagger$-Holant$(\mathcal{G}, =_3$) gadget, the assumption still holds if we replace $\mathcal{F}$ and $\mathcal{G}$ by $\mathcal{F}'$ and $\mathcal{G}'$. Furthermore, by (6.3.5), every quantum permutation matrix U satisfies $U(=_1) = (=_1)$, so

$$U \cdot (F \otimes =_1) = U \cdot (G \otimes =_1) \iff (U \cdot F) \otimes (=_1) = (U \cdot G) \otimes (=_1) \iff U \cdot F = U \cdot G.$$

Thus $\mathcal{F}' \cong_{qc} \mathcal{G}' \iff \mathcal{F} \cong_{qc} \mathcal{G}$. Therefore, by replacing $\mathcal{F}$ and $\mathcal{G}$ by $\mathcal{F}'$ and $\mathcal{G}'$, we may assume $\mathcal{F} \ni J \leftrightsquigarrow J \in \mathcal{G}$ and that $\mathcal{F}$ and $\mathcal{G}$ contain no unary signatures, as required for Lemma 7.2.1 (the former holds because $J = ((=_1)^{\otimes 2})^{1,1} \in \mathcal{F}' \cap \mathcal{G}'$).

Let $\mathbf{K}$ be a $\mathrm{Pl}^\dagger$-Holant$(\mathcal{F} \oplus \mathcal{G}, =_3)$ gadget with signature K. Since $=_3$ is already a direct sum of $=_3$ on smaller domains, $[\mathcal{F} \oplus \mathcal{G}, =_3]_{\mathrm{Pl}^\dagger} = [(\mathcal{F}, =_3) \oplus (\mathcal{G}, =_3)]_{\mathrm{Pl}^\dagger}$. Thus, by Proposition 5.2.1 and the assumption on $\mathcal{F}$ and $\mathcal{G}$, we have $K(x) = K(y)$ for every $K \in [\mathcal{F} \oplus \mathcal{G}, =_3]_{\mathrm{Pl}^\dagger}$. Now Corollary 7.1.1 asserts that x and y are in the same orbit of $\mathrm{Qut}(\mathcal{F} \oplus \mathcal{G})$, so, by Lemma 7.2.1, $\mathcal{F} \cong_{qc} \mathcal{G}$. □

Finally, we prove the 'hard' direction of Theorem 7.1.1. To recover a quantum isomorphism from $\mathcal{F}$ and $\mathcal{G}$ from Corollary 7.2.1, we must find some indistinguishable $x \in V(\mathcal{F})$ and $y \in V(\mathcal{G})$. Since there is no natural choice of such an x and y, we add new 'universal' domain elements to play the part of x and y, as in the proof of the classical case Theorem 5.2.1.

Lemma 7.2.2. *The following hold for bijective complex-valued signature sets $\mathcal{F}$ and $\mathcal{G}$.*

- *If $\mathcal{F}$ and $\mathcal{G}$ are $Pl^\top$-#CSP-indistinguishable and either conjugate-closed or consist of signatures of arity at most 2, then $\mathcal{F} \cong_{qc} \mathcal{G}$.*
- *If $\mathcal{F}$ and $\mathcal{G}$ are $Pl^\dagger$-#CSP-indistinguishable, then $\mathcal{F} \cong_{qc} \mathcal{G}$.*

Proof. Since Corollary 7.2.1 applies in both cases, the proofs of both claims will be identical; we prove the latter. Let 0_F and 0_G be new domain elements. For each $F \in \mathcal{F}$, $G \in \mathcal{G}$ with arity $n \geq 2$, define signatures F' and G' on $V(\mathcal{F}) \sqcup \{0_F\}$ and $V(\mathcal{G}) \sqcup \{0_G\}$, by letting, for $\mathbf{x} \in V(\mathcal{F}) \sqcup \{0_F\}$ or

$\mathbf{x} \in V(\mathcal{G}) \sqcup \{0_G\}$, respectively,

$$F'(\mathbf{x}) = \begin{cases} F(\mathbf{x}) & \mathbf{x} \in V(\mathcal{F})^n \\ \gamma & \mathbf{x} = (0_F, \ldots, 0_F, c),\ c \neq 0_F \\ 0 & \text{otherwise} \end{cases}, \quad G'(\mathbf{x}) = \begin{cases} G(\mathbf{x}) & \mathbf{x} \in V(\mathcal{G})^n \\ \gamma & \mathbf{x} = (0_G, \ldots, 0_G, c),\ c \neq 0_G \\ 0 & \text{otherwise} \end{cases},$$

for some fixed $\gamma \in \mathbb{R} \setminus \{0\}$. For each $F \in \mathcal{F}$, $G \in \mathcal{G}$ with arity 1, define F' and G' on $V(\mathcal{F}) \sqcup \{0_F\}$ and $V(\mathcal{G}) \sqcup \{0_G\}$, respectively, by

$$F'_x = \begin{cases} F_x & x \in V(\mathcal{F}) \\ 0 & x = 0_F \end{cases}, \qquad G'_x = \begin{cases} G_x & x \in V(\mathcal{G}) \\ 0 & x = 0_G \end{cases}.$$

Let $\mathcal{F}' = \{F' \mid F \in \mathcal{F}\}$ and $\mathcal{G}' = \{G' \mid G \in \mathcal{G}\}$, with $V(\mathcal{F}') = V(\mathcal{F}) \sqcup \{0_F\}$ and $V(\mathcal{G}') = V(\mathcal{G}) \sqcup \{0_G\}$.

We now show that

$$Z^{0 \mapsto 0_F}(I) = Z^{0 \mapsto 0_G}(I) \tag{7.2.4}$$

for every 1-labeled $\text{Pl}^\dagger$-#CSP$(\mathcal{F}' \oplus \mathcal{G}')$ instance $I = (V, C)$. Let v_0 be the labeled variable in V. If I is not connected, then the components of I that do not contain v_0 contribute the same value to the partition function regardless of the value taken by v_0. Hence, to establish (7.2.4) we may assume I is connected. As such, when v_0 takes value $0_F \in V(\mathcal{F}')$, all variables of I must take values in $V(\mathcal{F}')$ (otherwise the assignment contributes 0 to the partition function, as every $F' \oplus G'$ takes value 0 unless all its inputs are in $V(\mathcal{F}')$ or all its inputs are in $V(\mathcal{G}')$). We stratify the sum over assignments $\sigma : V \to V(\mathcal{F}')$ on the left side of (7.2.4) by the set of variables $S = \sigma^{-1}(0_F) \subseteq V$ assigned 0_F. For all σ, $v_0 \in \sigma^{-1}(0_F)$. Let $\chi_S = 1$ if no unary constraints $F' \oplus G'$ are applied to any variable in S, and if any variable in S appears in a constraint of arity $n \geq 2$, then the first $n-1$ arguments of the constraint are in S and the nth argument is not. By the definition of F', $\chi_S = 1$ captures the property that an assignment σ with $\sigma^{-1}(0_F) = S$ does not contribute 0 to the partition function. We define $\chi_S = 0$ otherwise. Suppose $\chi_S = 1$. The remaining 'free' variables $v \in V \setminus S$ satisfy $\sigma(v) \in V(\mathcal{F})$ and either appear in no constraints with any variables in S, or appear as the last argument of a constraint with all other variables in S, in which case the constraint takes value γ.

Let $I^{\mathcal{F}}_{V \setminus S}$ be the #CSP$(\mathcal{F})$ instance formed by eliminating all variables in S and replacing every function $F' \oplus G'$ with F. In other words, $I^{\mathcal{F}}_{V \setminus S}$ is the #CSP$(\mathcal{F})$ instance corresponding to the subgraph of I's signature grid induced by the vertices corresponding to variables in $V \setminus S$ and the constraints applied to only variables in $V \setminus S$, with the above function substitutions. Let d_S be the number of constraints containing any variable in S (each of which, as $\chi_S = 1$, is of the form $F'(v_1, \ldots, v_{n-1}, u)$ for some $v_1, \ldots, v_{n-1} \in S$ and $u \in V \setminus S$).

Critically, χ_S is determined by S, and does not depend on individual assignments σ which has S as the preimage set of 0_F. Thus for a fixed S with $\chi_S = 1$, the remaining variables $V \setminus S$ take all values in $V(\mathcal{F}') \setminus \{0_F\} = V(\mathcal{F})$ as σ ranges over $\{\sigma \mid \sigma^{-1}(0_F) = S\}$. Hence the sum over all assignments σ with $\sigma^{-1}(0_F) = S$ of the product of the constraints containing no variables in S is $Z(I^{\mathcal{F}}_{V\setminus S})$. The d_S constraints containing variables in S each contribute γ regardless of the value of the free variable in the nth argument. So

$$Z^{0 \mapsto 0_F}(I) = \sum_{S \subseteq V, S \ni v_0} \chi_S \gamma^{d_S} Z(I^{\mathcal{F}}_{V\setminus S}).$$

The above reasoning depends only on the relationship between 0_F and $V(\mathcal{F})$, which by construction is exactly the same as the relationship between 0_G and $V(\mathcal{G})$. So we get a similar expression for $Z^{0 \mapsto 0_G}(I)$. Since $I^{\mathcal{F}}_{V\setminus S}$ and $I^{\mathcal{G}}_{V\setminus S}$ were created from I by deleting vertices, they are planar $\text{Pl}^\dagger$-#CSP($\mathcal{F}$) and $\text{Pl}^\dagger$-#CSP($\mathcal{G}$) instances, respectively, and by construction $I^{\mathcal{G}}_{V\setminus S} = (I^{\mathcal{F}}_{V\setminus S})_{\mathcal{F}\to\mathcal{G}}$. Thus by assumption we have

$$Z^{0 \mapsto 0_F}(I) = \sum_{S \subseteq V, S \ni v_0} \chi_S \gamma^{d_S} Z(I^{\mathcal{F}}_{V\setminus S}) = \sum_{S \subseteq V, S \ni v_0} \chi_S \gamma^{d_S} Z(I^{\mathcal{G}}_{V\setminus S}) = Z^{0 \mapsto 0_G}(I),$$

proving (7.2.4). By Proposition 7.1.1, this implies that $K(0_F) = K(0_G)$ for every $K \in [\mathcal{F}' \oplus \mathcal{G}', =_3]_{\text{Pl}^\dagger}$. Now, as in the proof of Corollary 7.2.1, Proposition 5.2.1 implies that $F(0_F) = G(0_G)$ for every $[\mathcal{F}, =_3]_{\text{Pl}^\dagger} \ni F \leftrightsquigarrow G \in [\mathcal{G}, =_3]_{\text{Pl}^\dagger}$. Then Corollary 7.2.1 asserts that $\mathcal{F}' \cong_{qc} \mathcal{G}'$.

The final step is to extract a quantum isomorphism between $\mathcal{F}$ and $\mathcal{G}$ from the quantum permutation matrix $U = (u_{uv})_{u \in V(\mathcal{G}'), v \in V(\mathcal{F}')}$ satisfying $U\mathcal{F}' = \mathcal{G}'$ Define the matrix $\widehat{U} = (u_{uv})_{u \in V(\mathcal{G}), v \in V(\mathcal{F})}$ (in other words, we eliminate row 0_G and column 0_F from U). We must show that $\widehat{U}$ is a quantum permutation matrix and that $\widehat{U}\mathcal{F} = \mathcal{G}$. Let $\mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}$

have arity $n \geq 2$. For $\mathbf{x} \in V(\mathcal{G})^n$,

$$\begin{aligned} G(\mathbf{x}) = G'(\mathbf{x}) &= \sum_{\mathbf{y} \in V(\mathcal{F}')^n} u_{x_1 y_1} \dots u_{x_n y_n} F(\mathbf{y}) \\ &= \sum_{\mathbf{y} \in V(\mathcal{F})^n} u_{x_1 y_1} \dots u_{x_n y_n} F(\mathbf{y}) + \gamma \cdot u_{x_1 0_F} \dots u_{x_{n-1} 0_F} \sum_{y_n \in V(\mathcal{F})} u_{x_n y_n} \\ &= (\widehat{U} \cdot F)(\mathbf{x}) + \gamma \cdot u_{x_1 0_F} \dots u_{x_{n-1} 0_F} (\mathbf{1} - u_{x_n 0_F}), \end{aligned} \tag{7.2.5}$$

with the last equality using (6.3.5). If $\mathbf{x}$ does not satisfy $x_1 = x_2 = \ldots = x_{n-1} \neq x_n$, then the second term of the RHS of (7.2.5) is 0, giving the desired $G(\mathbf{x}) = (\widehat{U} \cdot F)(\mathbf{x})$. If $x_1 = x_2 = \ldots = x_{n-1} \neq x_n$, then (7.2.5) becomes

$$G(\mathbf{x}) = (\widehat{U} \cdot F)(\mathbf{x}) + \gamma \cdot u_{x_1 0_F}. \tag{7.2.6}$$

Since each u_{xy} is a (self-adjoint) projector, we have $\|u_{x_1 0_F}\| \in \{0, 1\}$ and $\|(\widehat{U} \cdot F)(\mathbf{x})\| \leq \sum_{\mathbf{y} \in V(\mathcal{F})^n} |F(\mathbf{y})|$. So if we choose

$$\gamma > \sum_{\mathbf{y} \in V(\mathcal{F})^n} |F(\mathbf{y})| + \max_{\mathbf{x} \in V(\mathcal{G})^n} |G(\mathbf{x})|,$$

then (7.2.6) is impossible unless $u_{x_1 0_F} = 0$, again giving $G(\mathbf{x}) = (\widehat{U} \cdot F)_{\mathbf{x}}$. By applying the above to all $\mathbf{x} \in V(\mathcal{G})^n$, we obtain $G = \widehat{U} \cdot F$, as well as $u_{x_1 0_F} = 0$ for all $x_1 \in V(\mathcal{G})$. The latter implies $u_{0_G 0_F} = \mathbf{1}$ by (6.3.5), so $u_{0_G y_1} = 0$ for all $y_1 \in V(\mathcal{F})$ as well. Thus $\widehat{U}$ is orthogonal (as U was orthogonal), so $\widehat{U}$ is a quantum permutation matrix.

If $n = 1$, for any $x \in V(\mathcal{G})$,

$$G(x) = G'(x) = \sum_{y \in V(\mathcal{F}')} u_{xy} F'(y) = \sum_{y \in V(\mathcal{F})} u_{xy} F(y),$$

so $\widehat{U} F = G$. Therefore $\mathcal{F} \cong_{qc} \mathcal{G}$. □

Now we obtain the main theorem of this chapter, Theorem 7.1.1

Proof of Theorem 7.1.1. The indistinguishability to quantum isomorphism directions follow from Lemma 7.2.2. The converse directions for the first and third claims follow from Proposition 7.1.1 (with $k = 0$) and the quantum orthogonal Holant theorem (Corollary 6.3.3). To apply the quantum orthogonal Holant theorem to the second claim, we must also observe that, by (6.3.4),

$$U \cdot \overline{F} = \overline{G} \iff U \cdot F^\top = G^\top \iff U \cdot F = G$$

for F and G of arity at most 2, because unary F satisfies $F^\top = F$ and binary F satisfies $F^\top = F^{(1)} \in [F]_{\text{Pl}}$. □

7.3 The complex-weighted graph isomorphism game

Throughout this chapter, we have defined quantum isomorphism by the existence of a quantum permutation matrix transforming $\mathcal{F}$ to $\mathcal{G}$. However, recall that quantum isomorphism of graphs was originally defined by the two-player quantum isomorphism game. Ideally, we would define an n-player tensor isomorphism game and connect this game to the existence of a quantum permutation matrix transforming one n-ary tensor to another, showing an equivalence between a perfect winning strategy in the game and planar #CSP indistinguishability. Unfortunately, we will not achieve this result for $n > 2$ (see Remark 7.3.1). However, for $n = 2$, we will extend the results of [Ats+19] and [LMR17] from graphs to complex-weighted matrices, proving equivalence between the quantum permutation matrix definition of quantum isomorphism for these matrices and the existence of a perfect winning strategy in a $\mathbb{C}$-*weighted graph isomorphism game*.

Let $F \in \mathbb{C}^{V(F)\times V(F)}$ and $G \in \mathbb{C}^{V(G)\times V(G)}$ be complex matrices, thought of as the adjacency matrices of $\mathbb{C}$-weighted graphs. We define the nonlocal (F, G)-*isomorphism game* as follows (recall the notation for nonlocal games in Section 6.1.1). Let $X_A = X_B = Y_A = Y_B = V(F) \cup V(G)$. Let $f_A = \{x_A, y_A\} \cap V(F)$ and $g_A = \{x_A, y_A\} \cap V(G)$, and define f_B and g_B similarly for Bob. f_A and g_A are well-defined (assuming the players win) given condition (i) below. The players win if and only if the following three conditions are satisfied:

(i) $x_A \in V(F) \iff y_A \in V(G)$ and $x_B \in V(F) \iff y_B \in V(G)$;

(ii) $f_A = f_B \iff g_A = g_B$;

(iii) $F_{f_A f_B} = G_{g_A g_B}$.

Say $F \cong_{qcg} G$ if Alice and Bob have a perfect quantum commuting strategy for the (F, G)-isomorphism game. We will show $F \cong_{qcg} G \iff F \cong_{qc} G$.

If F and G are symmetric and 0-1 valued, then the (F, G)-isomorphism game is the graph isomorphism game defined [Ats+19] and presented in Section 6.1.1 for the graphs whose adjacency matrices are F and G. Many of the results in [Ats+19] extend to the $\mathbb{C}$-weighted graph isomorphism game. First, since Alice and Bob must define a common bijection $V(F) \to V(G)$ by conditions (i) and (ii), $F \cong G$ if and only if Alice and Bob have a perfect *classical* strategy for the (F, G)-isomorphism game. Second, by (i) and (ii) and since any quantum strategy is *non-signalling*, $F \cong_{qcg} G \implies |V(F)| = |V(G)|$ [Ats+19, Lemma 4.1]. Third, since the $\mathbb{C}$-weighted graph isomorphism game is, like the graph game, a *synchronous game* (the players share the same input set, the same output set, and the players lose if Alice's input and output are disequal but Bob's input and output are equal, or vice-versa), we have the following extension of [Ats+19, Theorem 5.14]:

Lemma 7.3.1. *Let $F \in \mathbb{C}^{V(F)\times V(F)}$ and $G \in \mathbb{C}^{V(G)\times V(G)}$. Then $F \cong_{qcg} G$ if and only if there is a unital C^* algebra $\mathcal{A}$ which admits a faithful tracial state, and projectors $u_{fg} \in \mathcal{A}$ for $f \in V(F)$, $g \in V(G)$ such that*

1. *$U = (u_{fg})_{f\in V(F), g\in V(G)}$ is a quantum permutation matrix.*
2. *$u_{f_1g_1}u_{f_2g_2} = 0$ if $F_{f_1f_2} \neq G_{g_1g_2}$*

A *state* is a linear functional $s : \mathcal{A} \to \mathbb{C}$ satisfying $s(\mathbf{1}) = 1$ and $s(a^*a) \geq 0$ for all $a \in \mathcal{A}$, and a state is *tracial* if $s(ab) = s(ba)$ for all $a, b \in \mathcal{A}$, and is *faithful* if $s(a^*a) = 0 \iff a = 0$.

$F \cong_{qcg} G$ and $F \cong_{qc} G$ both imply $|V(F)| = |V(G)|$, so let $V(F) = V(G) = [q]$. For graph isomorphism, it is known [LMR17, Theorem 2.5] (see also [Ats+19, Lemma 5.8]) that condition 2 is equivalent to $FU = UG$. We have the following extension for $\mathbb{C}$-weighted graphs F and G:

Proposition 7.3.1. *Let $F \in \mathbb{C}^{q\times q}$, $G \in \mathbb{C}^{q\times q}$ and let $U = (u_{fg})_{f,g\in[q]}$ be a quantum permutation matrix. Then $FU = UG$ if and only if $[F_{f_1f_2} \neq G_{g_1g_2} \implies u_{f_1g_1}u_{f_2g_2} = 0]$.*

Proof. The proof is a simple modification of the proof of [LMR17, Lemma 3.9], which is the case $F = G$. Suppose $F_{f_1f_2} \neq G_{g_1g_2} \implies u_{f_1g_1}u_{f_2g_2} = 0$. Then for any $g_1, g_2 \in [q]$,

$$\left(U^\top \cdot F\right)_{g_1,g_2} = \sum_{f_1,f_2} u_{f_1g_1}u_{f_2g_2}F_{f_1,f_2} = \sum_{f_1,f_2\colon\, F_{f_1,f_2}=G_{g_1,g_2}} u_{f_1g_1}u_{f_2g_2}G_{g_1,g_2} = G_{g_1,g_2}.$$

Now by (6.3.3), $FU = UG$. Conversely, suppose $U^\top \cdot F = G$, so $\sum_{f_1,f_2} u_{f_1g_1}u_{f_2g_2}F_{f_1,f_2} = G_{g_1,g_2}$ for every g_1, g_2. Multiplying on the left by $u_{f'g_1}$ and on the right by $u_{f''g_2}$ for any f' and f'' yields $F_{f'f''}u_{f'g_1}u_{f''g_2} = G_{g_1g_2}u_{f'g_1}u_{f''g_2}$. Hence, if $u_{f'g_1}u_{f''g_2} \neq 0$, then $F_{f'f''} = G_{g_1g_2}$. □

Next, we extend [LMR17, Theorem 4.4] from graphs to $\mathbb{C}$-weighted graphs.

Lemma 7.3.2. *Let $F \in \mathbb{C}^{V(F)\times V(F)}$, $G \in \mathbb{C}^{V(G)\times V(G)}$. Then $F \cong_{qcg} G \iff F \cong_{qc} G$.*

Proof. By Proposition 7.3.1 and Lemma 7.3.1 (applied to $U^\top$, also a quantum permutation matrix), $F \cong_{qcg} G \implies F \cong_{qc} G$. However, our definition of $\cong_{qc}$ did not assume that the C^*-algebra over which U is defined admits a faithful tracial state. To show $F \cong_{qc} G \implies F \cong_{qcg} G$, it suffices to construct a faithful tracial state on the C^*-algebra $\mathcal{A}$ over which the quantum permutation matrix witnessing $F \cong_{qc} G$ is defined. Aside from a few small differences, the construction follows the proof of [LMR17, Theorem 4.4] (this theorem restricted to unweighted graphs), which doubles as the proof of [LMR17, Theorem 4.5], the graph case of our Lemma 7.2.1. The proof in [LMR17] assumes F and G are connected by taking the complement, which is not possible for $\mathbb{C}$-weighted graphs. Instead, as in the proof of Corollary 7.2.1, we may assume that $\mathcal{F} = \{F\}$ and $\mathcal{G} = \{G\}$ also contain corresponding copies of J without affecting whether $\mathcal{F} \cong_{qc} \mathcal{G}$. Now the proof of [LMR17, Theorem 4.4] goes through on these connected signature sets, where we use Proposition 7.3.1 in place of the corresponding statement for unweighted graphs. $\square$

Say complex-weighted graphs $F \in \mathbb{C}^{V(F)\times V(F)}$ and $G \in \mathbb{C}^{V(G)\times V(G)}$ are planar homomorphism indistinguishable if F and G are Pl-#CSP-indistinguishable (recall the discussion after (2.3.1)). Since $F^\top = F^{(1)} \in [\mathcal{F}]_{\mathrm{Pl}}$ and similarly for G, this in turn holds if and only if F and G are $\mathrm{Pl}^\top$-#CSP-indistinguishable. Now combining Lemma 7.3.2 and Theorem 7.1.1 gives the following.

Theorem 7.3.1. *Let $F \in \mathbb{C}^{V(F)\times V(F)}$ and $G \in \mathbb{C}^{V(G)\times V(G)}$. There is a perfect quantum commuting strategy for the (F, G)-isomorphism game if and only if F and G are planar homomorphism indistinguishable.*

Remark 7.3.1. It is natural to attempt a generalization of the results of this section from binary tensors (matrices) to tensors of arity $n > 2$ – a "tensor isomorphism game" with n players. However, any such extension must overcome the following observation to the author by David Roberson, that the natural analogue of Proposition 7.3.1 does not in general hold for F and G of higher arity. Recall the quantum symmetric group $S_q^+ = \mathrm{Qut}(\varnothing) = \mathrm{Stab}_{O^+}(=_3)$ from Definition 6.4.2. By Corollary 6.4.1, $C_{S_q^+} = \langle =_3, \supset \rangle_{+,\circ,\otimes,\dagger}$ (the span of the "non-crossing partitions"). Put $F = G :=$ $(I \otimes (=_1)^{1,0} \otimes I) \circ \supset$, with $F(a, b, c) = \delta_{ac}$. Since $F \in C_{S_q^+}$, we have $U \cdot F = F$ for every quantum permutation matrix U. For $q \geq 3$, we have $F(0,1,0) = 1 \neq 0 = F(0,1,2)$, so an $n = 3$ version of Proposition 7.3.1 would then imply $u_{11}u_{22}u_{13} = 0$, where U is the fundamental representation of S_q^+. However, this is false for $q \geq 4$, since in this case S_q^+ has *free orbitals*: any relation $u_{x_1y_1} \dots u_{x_ky_k} = 0$ is a consequence of some pair of adjacent factors being distinct entries of the same row or column of U [McC23, Theorem 2.3].

7.4 Vertex separation by edge gadgets

In this section, we present an application of the theory developed in this chapter (and the corresponding classical theory in Chapter 5) to the complexity of counting planar graph homomorphisms. This section is based on [CMY26], which is joint work with Jin-Yi Cai and Ashwin Maran.

Let $\text{Sym}_q(\mathbb{R}_{\geq 0}) \subset \mathbb{R}^{q \times q}$ denote the set of symmetric matrices with nonnegative real entries. For $M \in \text{Sym}_q(\mathbb{R}_{\geq 0})$, let Pl-GH$(M)$ $\equiv$ Pl-#CSP(M) denote the problem of counting (real-weighted) homomorphisms from planar graphs K to M (since M is symmetric and real-valued, Pl-#CSP(M), Pl$^\top$-#CSP(M), and Pl†-#CSP(M) are equivalent). Classifying the complexity of Pl-GH(M) is a difficult problem, and has only been fully achieved for graphs M on $q = 3$ or 4 vertices [CM23; CM24a]. These classifications are trichotomies: depending on M, Pl-GH(M) is either #P-hard on planar input graphs, tractable (in FP) on all input graphs, or #P-hard on general input graphs but tractable on planar input graphs by a reduction to the FKT algorithm for counting perfect matchings in planar graphs. It is conjectured that the same trichotomy holds for M on any number of vertices.

Contrary to the planar case, a complete dichotomy is known for the problem of counting homomorphisms from general graphs to M with any number of vertices, even for M with complex entries [CCL13]. In this section, we present one potential underlying explanation for the difficulty of the planar case relative to the general case. Gadget construction is the main reduction technique in the study of the complexity of both planar and general homomorphisms. While the expressive power of (quantum) #CSP(M) $\equiv$ Holant$(M, \mathcal{EQ})$ gadgets (equivalently, k-labeled #CSP(M) instances) is governed by Aut(M), which has a straightforward and decidable theory, the expressive power of (quantum) Pl-GH(M) $\equiv$ Pl-#CSP(M) $\equiv$ Pl-Holant$(M, \mathcal{EQ})$ gadgets is governed by Qut(M), whose theory is significantly more mysterious, and, as we will see, even undecidable.

If K is the signature of a k-labeled Pl-GH(M) instance $\mathbf{K}$, then Pl-#CSP(K) reduces to Pl-GH(M) because in any Pl-#CSP(K) instance I, we can replace every constraint applying K with $\mathbf{K}$ to obtain a Pl-#CSP(M) $\equiv$ Pl-GH(M) instance with the same value. Since we wish to reduce between Pl-GH problems, we mainly consider 2-labeled Pl-GH(M) instances whose matrix K is symmetric, and hence, since M has nonnegative entries, satisfies $K^{1,1} \in \mathrm{Sym}_q(\mathbb{R}_{\geq 0})$.

Definition 7.4.1 ($\mathfrak{B}(M)$, Pl-$\mathfrak{B}(M)$, $\mathfrak{E}(M)$,Pl-$\mathfrak{E}(M)$)**.** For $M \in \mathrm{Sym}_q(\mathbb{R}_{\geq 0})$, define the sets $\mathfrak{B}(M), \text{Pl-}\mathfrak{B}(M) \subset (\mathbb{R}_{\geq 0})^{q \times q}$ of matrices (i.e. flattened signatures) of arity-two Holant($M, \mathcal{EQ}$) and Pl-Holant($M, \mathcal{EQ}$) gadgets, called *binary* and *planar binary M* gadgets, respectively.

An *edge gadget* is a 2-labeled #CSP(M) instance, and a *planar edge gadget* is a 2-labeled Pl-#CSP(M) instance. Define the set

$$\mathfrak{E}(M) = \{K^{1,1} \mid \text{edge gadget } \mathbf{K}\} \cap \mathrm{Sym}_q(\mathbb{R}_{\geq 0})$$

of edge gadget matrices in the context of M, and define the set

$$\text{Pl-}\mathfrak{E}(M) = \{K^{1,1} \mid \text{planar edge gadget}\} \cap \mathrm{Sym}_q(\mathbb{R}_{\geq 0})$$

of planar edge gadget matrices.

Recall from Proposition 5.1.1 and Proposition 7.1.1 that Holant($M, \mathcal{EQ}$) and Pl-Holant($M, \mathcal{EQ}$) gadgets correspond to multilabeled #CSP(M) and Pl-#CSP(M) instances, so edge gadgets are special not only in that their matrices are symmetric, but that their labels are on distinct vertices. The terminology "edge gadget" comes from the Holant($M \mid \mathcal{EQ}$) representation of Pl-GH(M). Recall from Figure 2.1 that the underlying graph of the Holant($M \mid \mathcal{EQ}$) grid whose Holant value is the number of homomorphisms $G \to M$ is a subdivision of G by vertices assigned M. In the theory of homomorphism counting, we typically ignore these vertices and just work with G. With this perspective, replacing a vertex v assigned M with an edge gadget $\mathbf{K}$ is equivalent to replacing the edge of G on which v sits with the underlying graph of $\mathbf{K}$ (also without the M vertices), then counting homomorphisms from the resulting graph to M. We specify that the labeled vertices of an edge gadget are distinct so that this substitution replaces the edge of G without contracting it.

Theorem 5.1.1 and Theorem 7.1.2 characterize the matrices of quantum binary and quantum planar binary M gadgets as those matrices invariant under Aut(M) and Qut(M), respectively.

In general, if K is the signature of a quantum Holant($\mathcal{F}$) gadget $\mathbf{K}$, then Holant(K) does not necessarily reduce *in polynomial time* to Holant($\mathcal{F}$), because, after replacing every vertex in a Holant(K) grid Ω with $\mathbf{K}$, we obtain not a single Holant($\mathcal{F}$) grid but a Holant($[\mathcal{F}]$) grid that expands linearly into a potentially exponential number of Holant($\mathcal{F}$) grids. Since the invariant-theoretic methods we employ only apply to vector spaces (in our context, quantum gadget spaces), this fact generally precludes applying the theory of Holant indistinguishability to the study of Holant complexity. However, in this section we are concerned with the following question: given $x, y \in [q]$, is there a $K \in \mathfrak{E}(M)$ or Pl-$\mathfrak{E}(M)$ such that $K_{xx} \neq K_{yy}$? If K is *quantum* edge gadget matrix such that $K_{xx} \neq K_{yy}$, then one of the terms of K, a true edge gadget matrix, must also have disequal (x, x) and (y, y) entries. Thus our question has the same answer whether we study edge gadgets or quantum edge gadgets, so we can apply our indistinguishability theory.

7.4.1 Diagonal distinctness from pairwise distinguishability

Call a $q \times q$ matrix M *diagonal distinct* if its diagonal entries M_{xx} for $x \in [q]$ are pairwise distinct. The following theorem, due to Cai and Maran in [CMY26], is a dichotomy theorem for Pl-GH(M) for diagonal distinct M on any number of vertices.

Theorem 7.4.1 ([CMY26, Theorem 33]). *Let $M \in Sym_q(\mathbb{R}_{\geq 0})$ be diagonal distinct. Then Pl-GH(M) is polynomial time tractable if (after a renaming of $[q]$), $M = M_1 \oplus \cdots \oplus M_r$ for rank 0 or rank 1 matrices M_i. In all other cases, Pl-GH(M) is #P-hard.*

In particular, if $K \in$ Pl-$\mathfrak{E}(M)$ is diagonal distinct and hard, then, since Pl-GH(K) reduces to Pl-GH(M), we can obtain hardness for Pl-GH(M) from Theorem 7.4.1. In this subsection, we develop some tools to weaken the assumption of Theorem 7.4.1 from simultaneous diagonal distinctness to the existence, for each x, y, of some planar edge gadget whose (x, x) and (y, y) diagonal entries are distinct (cf. the question mentioned immediately above this subsection).

Lemma 7.4.1. *Let $M \in Sym_q(\mathbb{R}_{\geq 0})$. Assume that for all $x < y \in [q]$, there is a $K^{xy} \in$ Pl-$\mathfrak{E}(M)$ such that $(K^{xy})_{xx} \neq (K^{xy})_{yy}$ and $(K^{xy})_{ii} > 0$ for every $i \in [q]$. Then there is a diagonal distinct $K \in$ Pl-$\mathfrak{E}(M)$.*

Proof. Let $\mathbf{K}^{xy}$ be the planar edge gadget with matrix K^{xy}. Given $\mathbf{z} = (z_{xy})_{x<y\in[q]} \in (\mathbb{Z}_{>0})^{\frac{q(q-1)}{2}}$, construct the edge gadget $\mathbf{K}^{\mathbf{z}}$ as the 2-labeled instance product of z_{xy} copies of each $\mathbf{K}^{xy}$. Since the copies $\mathbf{K}^{xy}$ run in parallel between the two labeled vertices, $\mathbf{K}^{\mathbf{z}}$ is planar (also recall Proposition 7.1.2), so $(K^{\mathbf{z}}) \in \text{Pl-}\mathfrak{C}(M)$. By (3.2.1), the matrix of $\mathbf{K}^{\mathbf{z}}$ is

$$K^{\mathbf{z}} = \bigodot_{x<y\in[q]} (K^{xy})^{\bullet z_{xy}}.$$

Every K^{xy} has positive diagonal entries, so $(K^{\mathbf{z}})_{ii} > 0$ for all $i \in [q]$.

We will now show that there exists some $\mathbf{z}^*$ such that $K^{\mathbf{z}^*}$ is diagonal distinct. Pairwise distinctness of $(K^{\mathbf{z}})_{ii}$ for $i \in [q]$ is equivalent to pairwise distinctness of the values $\log((K^{\mathbf{z}})_{ii}) = \sum_{x<y} z_{xy} \cdot \log((K^{xy})_{ii})$. This is equivalent to a choice of $\mathbf{z}^*$ on which the polynomial

$$\begin{aligned}\zeta(\mathbf{z}) &= \prod_{i<j\in[q]} \left(\sum_{x<y\in[q]} z_{xy} \cdot \log((K^{xy})_{ii}) - \sum_{x<y\in[q]} z_{xy} \cdot \log(K^{xy}_{jj}) \right) \\ &= \prod_{i<j\in[q]} \sum_{x<y\in[q]} z_{xy} \cdot \big(\log((K^{xy})_{ii}) - \log(K^{xy}_{jj}) \big)\end{aligned}$$

is nonzero. For any $i < j \in [q]$, by our assumption about K^{ij}, the coefficient of z_{ij} in

$$\sum_{x<y\in[q]} z_{xy} \cdot \big(\log((K^{xy})_{ii}) - \log(K^{xy}_{jj}) \big)$$

is: $\log((K^{ij})_{ii}) - \log((K^{ij})_{jj}) \neq 0$. Thus each factor of $\zeta(\mathbf{z})$ is nonzero, and so $\zeta(\mathbf{z})$ is not the zero polynomial. Since $(\mathbb{Z}_{>0})$ is infinite, and ζ is non-zero, it cannot vanish on all of $(\mathbb{Z}_{>0})^{\frac{q(q-1)}{2}}$. Hence there exists some $\mathbf{z}^* \in (\mathbb{Z}_{>0})^{\frac{q(q-1)}{2}}$, such that $K^{\mathbf{z}^*} \in \text{Pl-}\mathfrak{C}(M)$ is diagonal distinct. □

We need the $\mathbf{K}$ given by Lemma 7.4.1 to have strictly positive diagonal entries because in the key step where we construct $K^{\mathbf{z}}$, if any of the gadgets K^{xy} has a zero diagonal entry: $(K^{xy})_{ii} = 0$ for some $i \in [q]$, then $(K^{\mathbf{z}})_{ii} = 0$. For $M \in \text{Sym}_q(\mathbb{R}_{\geq 0})$, let $\Gamma_M = ([q], E(\Gamma_M))$ be the *underlying graph* of M, namely $(i, j) \in E(\Gamma_M)$ if and only if $M_{ij} > 0$. Say that M is *connected* if Γ_M is connected, and is *bipartite* if Γ_M is bipartite. We show that, for connected M, we can remove the strict positivity assumption in Lemma 7.4.1. We split the proof into two lemmas, depending on whether or not M is bipartite.

Lemma 7.4.2. *Let $M \in Sym_q(\mathbb{R}_{\geq 0})$ (for $q \geq 2$) be connected and bipartite. Let $K \in Pl\text{-}\mathfrak{C}(M)$ satisfy $K_{xx} \neq K_{yy}$ for some $x, y \in [q]$. Then $K_{ii} \neq 0$ for all $i \in [q]$.*

Proof. Let $\mathbf{K}$ be the edge gadget with matrix K. Identify $\mathbf{K}$ with the graph G such that the $\#\mathrm{CSP}(M)$ value of the unlabeled version of $\mathbf{K}$ equals the number of real-weighted homomorphisms from G to M (recall Figure 2.1; the vertices of G are the variables of $\mathbf{K}$, and two vertices are adjacent if the corresponding variables appear in the same M constraint). It is well-known that, if there is at least one homomorphism from graph G to bipartite graph Γ, then G must be bipartite. Since $(K)_{xx} \neq (K)_{yy}$, one of $(K)_{xx}$ or $(K)_{yy}$ is nonzero, so $\mathbf{K}$ must be bipartite. Let ℓ_1 and ℓ_2 be the two labeled variables – thought of as vertices – of $\mathbf{K}$. First, assume that ℓ_1 and ℓ_2 belong to the same connected component of $\mathbf{K}$, but to different partitions. Since one of $(K)_{xx}$ or $(K)_{yy}$ is nonzero, this implies that the odd-length path from ℓ_1 to ℓ_2 in $\mathbf{K}$ had to be mapped to an even length path between x and x, or y and y in Γ_M. This contradiction therefore implies that if ℓ_1 and ℓ_2 belong to the same connected component of $\mathbf{K}$, then ℓ_1 and ℓ_2 must be in the same partition of $\mathbf{K}$.

Since Γ_M is connected, every $i \in [q]$ is adjacent to some other $j \in [q]$ in Γ_M. If ℓ_1 and ℓ_2 belong to different connected components of $\mathbf{K}$, we can trivially define a bipartition of $\mathbf{K}$ such that ℓ_1 and ℓ_2 belong to the same partition. We have also seen that if ℓ_1 and ℓ_2 belong to the same connected component, then they must belong to the same partition. So, in any case, there exists a bipartition of $\mathbf{K}$ such that ℓ_1, and ℓ_2 belong to the same partition. Now, we can define a map $\mathbf{K}$ to Γ_M that sends the partition containing ℓ_1 and ℓ_2 to i and the other partition to j. This map is a homomorphism, so $(K)_{ii} \geq (M_{ij})^{|E(\mathbf{K})|} > 0$, as $M_{ij} > 0$. □

For the non-bipartite case, we apply the "connector" idea of Lovász and Szegedy [LS09, Theorem 1.4] and Dvořák [Dvo10, Lemma 11]: by replacing every edge of an edge gadget by a linear combination of sufficiently long paths, we may assume the underlying graph admits homomorphisms to M.

Lemma 7.4.3. *Let $M \in Sym_q(\mathbb{R}_{\geq 0})$ (for $q \geq 2$) be connected but not bipartite. Let $x, y \in [q]$. If there is a $K \in Pl\text{-}\mathfrak{E}(M)$ such that $K_{xx} \neq K_{yy}$, then there exists $J \in Pl\text{-}\mathfrak{E}(M)$ such that $J_{xx} \neq J_{yy}$, and $J_{ii} \neq 0$ for all $i \in [q]$.*

Proof. By nonnegativity, $(M^k)_{ii} \neq 0$ if and only if there is a length-k walk in Γ_M from vertex i back to itself. Since Γ_M is connected and not bipartite, there exists a smallest $m \geq 2$ such that, for every $k \geq m$, every vertex i has a length-k walk to itself, so $(M^k)_{ii} > 0$ for every $i \in [q]$. We first show that

$$M \in \operatorname{span}_{\mathbb{R}}(M^m, \dots M^{m+q}). \tag{7.4.1}$$

Write $M = H \operatorname{diag}(\lambda_1, \dots, \lambda_q) H^\top$ for orthogonal H, so $M^k = H \operatorname{diag}(\lambda_1^k, \dots, \lambda_q^k) H^\top$, and hence (7.4.1) is equivalent to $\operatorname{diag}(\lambda_1, \dots, \lambda_q) \in \operatorname{span}_{\mathbb{R}}(\operatorname{diag}(\lambda_1^m, \dots, \lambda_q^m), \dots \operatorname{diag}(\lambda_1^{m+q}, \dots, \lambda_q^{m+q}))$. Without loss of generality, let $\lambda_1, \dots, \lambda_s$, with $s \le q$, be the unique non-zero eigenvalues of M. Then (7.4.1) is equivalent to

$$\operatorname{diag}(\lambda_1, \dots, \lambda_s) \in \operatorname{span}_{\mathbb{R}}(\operatorname{diag}(\lambda_1^m, \dots, \lambda_s^m), \dots \operatorname{diag}(\lambda_1^{m+q}, \dots, \lambda_s^{m+q})). \tag{7.4.2}$$

Upon multiplying its ith column by the nonzero scalar λ_i^{-m} for every $i \in [s]$, the matrix

$$L = \begin{bmatrix} \lambda_1^m & \lambda_2^m & \dots & \lambda_s^m \\ \lambda_1^{m+1} & \lambda_2^{m+1} & \dots & \lambda_s^{m+1} \\ \vdots & \vdots & \dots & \vdots \\ \lambda_1^{m+q} & \lambda_2^{m+q} & \dots & \lambda_s^{m+q} \end{bmatrix}$$

becomes a Vandermonde matrix with q rows and $s \le q$ columns. Since $\lambda_1, \dots, \lambda_s$ are distinct, L has rank s, so the rows of L span $\mathbb{R}^s$. This proves (7.4.2), hence also (7.4.1).

Now write $M = \sum_{i=0}^q c_i M^{m+i}$ for $c_0, \dots, c_q \in \mathbb{R}_{\neq 0}$. Let $\mathbf{K}$ be the planar edge gadget with matrix K. Then, by Proposition 7.1.1, $\mathbf{K}$ expands linearly into a quantum binary Pl-Holant$(M^m, \dots, M^{m+q}, \mathcal{EQ})$ gadget $\mathbf{K} = \sum_j d_j \mathbf{K}^j$ with the same matrix K. Since $K_{xx} \neq K_{yy}$, there is some j such that $(K^j)_{xx} \neq (K^j)_{yy}$, where K^j is the matrix of $\mathbf{K}^j$. Furthermore, since every diagonal entry of each M^{m+i} is nonzero, and every entry of each M^{m+i} is nonnegative, we have $(K^j)_{ii} \neq 0$, as the assignment of i to every edge of K^j yields a nonzero value. By replacing every vertex of $\mathbf{K}^j$ assigned M^{m+i} with a path of $m+i$ vertices assigned M (a gadget whose signature is $M \circ \dots \circ M = M^{m+i}$), we may assume $\mathbf{K}^j$ is a Pl-Holant$(M, \mathcal{EQ})$ gadget, so $K^j \in$ Pl-$\mathfrak{B}(M)$. This $\mathbf{K}^j$ does still have labels on distinct variables, but K^j may not be symmetric, so construct a binary gadget $\mathbf{J}$ by starting with two copies of $\mathbf{K}^j$, then merging the variable labeled 0 in the first copy with the variable labeled 1 in the second copy and vice-versa. Then $J_{ij} = (K^j)_{ij}(K^j)_{ji} = J_{ji}$ for all $i, j \in [q]$, so $J \in$ Pl-$\mathfrak{E}(M)$. Furthermore, by nonnegativity, $J_{xx} = ((K^j)_{xx})^2 \neq ((K^j)_{yy})^2 = J_{yy}$, and, for every $i \in [q]$, $J_{ii} = ((K^j)_{ii})^2 \neq 0$. $\square$

Using Lemma 7.4.2 and Lemma 7.4.3 to meet the conditions of Lemma 7.4.1 for $M \in \operatorname{Sym}_q(\mathbb{R}_{\ge 0})$ with a connected Γ_M, we can now construct a diagonal distinct $K \in$ Pl-$\mathfrak{E}(M)$.

Theorem 7.4.2. *Let $M \in Sym_q(\mathbb{R}_{\geq 0})$, such that the underlying graph Γ_M is connected. Assume that for all $x < y \in [q]$, there is a $K^{xy} \in Pl\text{-}\mathfrak{E}(M)$ such that $(K^{xy})_{xx} \neq (K^{xy})_{yy}$. Then there is a diagonal distinct $K \in Pl\text{-}\mathfrak{E}(M)$.*

Using Theorem 7.4.2, Cai and Maran prove a stronger version of Theorem 7.4.1 that does not assume that M separates all diagonal entries simultaneously.

Theorem 7.4.3 ([CMY26, Theorem 34])**.** *Let $M \in Sym_q(\mathbb{R}_{\geq 0})$. Assume that for all $x < y \in [q]$, there exists some $K^{xy} \in Pl\text{-}\mathfrak{E}(M)$ such that $(K^{xy})_{xx} \neq (K^{xy})_{yy}$. Then, Pl-GH($M$) is polynomial-time tractable if (after a renaming of $[q]$), $M = M_1 \oplus \cdots \oplus M_r$, for rank 0, or rank 1, or bipartite rank 2 matrices M_i. In all other cases, Pl-GH(M) is #P-hard.*

It is natural to ask whether Lemmas 7.4.2 and 7.4.3 hold for all $M \in \mathrm{Sym}_q(\mathbb{R}_{\geq 0})$, even those for which Γ_M is disconnected. The proofs still go through when all components of Γ_M are bipartite or all components of Γ_M are not bipartite. However, we show using techniques of Dvořák [Dvo10] that when Γ_M contains a mix of bipartite and non-bipartite components, the existence of a planar edge gadget matrix K^{xy} separating x and y does not guarantee the existence of the desired J^{xy} with nonzero diagonal entries. Given a graph H with adjacency matrix M, let $H \times K_2$ be the bipartite double cover of H, with adjacency matrix $M \otimes \left[\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right]$. Equivalently, $H \times K_2$ has vertex set $V(H) \times \{0, 1\}$ and $((h, b), (h', b')) \in E(H \times K_2)$ if and only if $(h, h') \in E(H)$ and $b \neq b'$ ($H \times K_2$ is bipartite between $\{(h, 0) \mid h \in V(H)\}$ and $\{(h, 1) \mid h \in V(H)\}$).

Proposition 7.4.1. *For any connected non-bipartite graph H, let $H_1 := H \oplus H$ and $H_2 := H \times K_2$, and let $M \in Sym_{V(H_1) \sqcup V(H_2)}(\mathbb{R}_{\geq 0})$ be the adjacency matrix of $H_1 \oplus H_2$. Then for every $x \in V(H_1)$, there is a $y \in V(H_2)$ such that*

- *there exists a $K \in Pl\text{-}\mathfrak{E}(M)$ such that $K_{xx} \neq K_{yy}$, but*
- *every $J \in \mathfrak{B}(M)$ such that $J_{xx} \neq J_{yy}$ also satisfies $J_{yy} = 0$.*

Proof. For every vertex x of $H_1 = H \oplus H$, there is an odd r such that there is a length-r closed walk from x to itself in H. So, if $\mathbf{K}$ is the edge gadget consisting of two copies of the cycle C_r, each with one labeled vertex, then $K_{xx} \neq 0$. Additionally, $K_{yy} = 0$ for every $y \in V(H_2)$, as $H_2 = H \times K_2$ is bipartite. This proves the first claim.

For the second claim, we may assume that the gadget $\mathbf{J}$ with signature $J \in \mathfrak{B}(M)$ contains no connected components without a labeled vertex, as such components simply multiply J_{xx} and J_{yy} by the same constant. Again, identify $\mathbf{J}$ with its underlying graph of variable vertices. Since H_2 is not bipartite, any non-bipartite $\mathbf{J}$ will satisfy $J_{yy} = 0$ for every $y \in V(H_2)$. So assume $\mathbf{J}$ is bipartite.

Dvořák [Dvo10, Theorem 13] shows that there is an isomorphism $f : H_1 \times K_2 \to H_2 \times K_2$ preserving the second coordinate – that is, for every $(h, b) \in V(H_1 \times K_2)$, we have $f((h, b)) = ((h', b))$ for some $h' \in V(H_2)$ – and, for any bipartite $\mathbf{J}$, uses f to construct a family of bijections between the sets $\mathrm{Hom}(\mathbf{J}, H_1)$ and $\mathrm{Hom}(\mathbf{J}, H_2)$ of homomorphisms from $\mathbf{J}$ to H_1 and H_2, respectively (ignoring for now the vertex labels of $\mathbf{J}$). Each $c : V(\mathbf{J}) \to \{0, 1\}$ defining a valid bipartition of $\mathbf{J}$ yields a bijection as follows. For $\beta \in \{1, 2\}$, define $\pi_1^{H_\beta \times K_2} : V(H_\beta \times K_2) \to V(H_\beta)$ by $\pi_1(h, b) = h$. Then map $g_1 \in \mathrm{Hom}(\mathbf{J}, H_1)$ to $g_2 \in \mathrm{Hom}(\mathbf{J}, H_2)$ defined by

$$g_2(\ell) = \pi_1^{H_2 \times K_2}(f(g_1(\ell), c(\ell))) \tag{7.4.3}$$

for $\ell \in V(\mathbf{J})$. The inverse of this map is given by

$$g_1(\ell) = \pi_1^{H_1 \times K_2}(f^{-1}(g_2(\ell), c(\ell))) \tag{7.4.4}$$

for $\ell \in V(\mathbf{J})$. Let $\ell_1, \ell_2 \in V(\mathbf{J})$ be the labeled vertices. If there is an odd-length path in $\mathbf{J}$ between ℓ_1 and ℓ_2, then again $J_{yy} = 0$ for every $y \in V(K_2)$ and we are done. Assume otherwise, so we can choose c such that $c(\ell_1) = c(\ell_2) = 0$. Let $y := \pi_1^{H_2 \times K_2}(f(x, 0)) \in V(H_2)$. If $g_1(\ell_1) = g_1(\ell_2) = x$, then, by (7.4.3), $g_2(\ell_1) = g_2(\ell_2) = y$. Conversely, if $g_2(\ell_1) = g_2(\ell_2) = y$, then, by (7.4.4),

$$g_1(\ell_1) = g_1(\ell_2) = \pi_1^{H_1 \times K_2}(f^{-1}(y, 0)) = \pi_1^{H_1 \times K_2}(x, 0) = x.$$

Therefore the bijection between $\mathrm{Hom}(\mathbf{J}, H_1)$ and $\mathrm{Hom}(\mathbf{J}, H_2)$ restricts to a bijection between those homomorphisms in $\mathrm{Hom}(\mathbf{J}, H_1)$ sending ℓ_1 and ℓ_2 to x and those homomorphisms in $\mathrm{Hom}(\mathbf{J}, H_2)$ sending ℓ_1 and ℓ_2 to y. Since $\mathbf{J}$ contains no connected components without a labeled vertex, any homomorphism from $\mathbf{J}$ to $H_1 \oplus H_2$ sending ℓ_1 and ℓ_2 to x (resp. y) is a homomorphism from $\mathbf{J}$ to H_1 (resp. H_2). Hence $J_{xx} = J_{yy}$. □

7.4.2 Undecidability of planar vertex separation

Theorem 7.4.2 reduces the problem of finding a diagonal distinct planar edge gadget to the following problem: given $x, y \in [q]$, determine whether there is a $K \in \text{Pl-}\mathfrak{E}(M)$ such that $K_{xx} \neq K_{yy}$. In this subsection, we show that the latter problem is undecidable (it is an open problem whether the former problem is undecidable) via a reduction from quantum isomorphism, which is known to be undecidable [Ats+19; Slo19].

First, we show that, for $M \in \text{Sym}_q(\mathbb{R}_{\geq 0})$, planar edge gadgets capture the full vertex-separation power of binary M gadgets.

Proposition 7.4.2. *Let $M \in Sym_q(\mathbb{R}_{\geq 0})$. Then $K_{xx} = K_{yy}$ for every $K \in Pl\text{-}\mathfrak{E}(M)$ if and only if x and y are in the same orbit of* $\text{Qut}(M)$.

Proof. By the $n = 2$ case of Corollary 7.1.1, (x, x) and (y, y) are in the same orbital of $\text{Qut}(M)$ if and only if $K'_{xx} = K'_{yy}$ for every $K' \in \text{Pl-}\mathfrak{B}(M)$. For diagonal pairs (x, x) and (y, y), the notion of orbit and orbital are the same: $(x, x) \sim_2 (y, y)$ if and only if $x \sim_1 y$, as $u_{xy}u_{xy} = u_{xy}$.

Thus it suffices to show that, if $K'_{xx} \neq K'_{yy}$ for some $K' \in \text{Pl-}\mathfrak{B}(M)$, then there is a $K \in \text{Pl-}\mathfrak{E}(M)$ such that $K_{xx} \neq K_{yy}$. This K' can fail to be in $\text{Pl-}\mathfrak{E}(M)$ in two ways: its underlying binary gadget $\mathbf{K}'$ may have both labels on the same vertex, or K' may not be symmetric. In the first case, we have $K' = \text{diag}(\mathbf{x})$ for some $\mathbf{x} \in (\mathbb{R}_{\geq 0})^q$. Construct an edge gadget $\mathbf{K}$ as the disjoint union of two copies of $\mathbf{K}'$, but with the originally doubly-labeled vertex of $\mathbf{K}'$ now labeled 0 in the first copy and 1 in the second copy. Now $K := \mathbf{x}\mathbf{x}^\top \in \text{Pl-}\mathfrak{E}(M)$ and, by nonnegativity, $K_{xx} = (K'_{xx})^2 \neq (K'_{yy})^2 = K_{yy}$.

In the second case, K' is asymmetric, so the two labels of $\mathbf{K}'$ are on distinct vertices. Construct $\mathbf{K}$ by again creating two copies $\mathbf{K}'$, then merging the variable labeled 0 in the first copy with the variable labeled 1 in the second copy and vice-versa. As in the proof of Lemma 7.4.3, $K \in \text{Pl-}\mathfrak{E}(M)$ and $K_{xx} = (K'_{xx})^2 \neq (K'_{yy})^2 = K_{yy}$. □

Applying the same argument with the classical/nonplanar Corollary 5.1.1 in place of Corollary 7.1.1, we also obtain a classical/nonplanar version of Proposition 7.4.2.

Proposition 7.4.3. *Let $M \in Sym_q(\mathbb{R}_{\geq 0})$. Then $K_{ii} = K_{jj}$ for every $K \in \mathfrak{E}(M)$ if and only if i and j are in the same orbit of* $\text{Aut}(M)$.

Lemma 7.2.1 is a generalization of [LMR17, Theorem 4.5], which states that connected graphs H and H' with adjacency matrices M and M' are quantum isomorphic if and only if some $v \in V(H)$ and $v' \in V(H')$ are in the same orbit of $\mathrm{Qut}(M \oplus M')$. Furthermore, either a graph or its complement is connected, a connected graph cannot be quantum isomorphic to a disconnected graph [LMR17], and two graphs are quantum isomorphic if and only if their complements are [Ats+19]. The matrix $M \oplus M'$ (and the adjacency matrix of $\overline{H} \oplus \overline{H'}$) are in $\mathrm{Sym}_{|V(H)|+|V(H')|}(\mathbb{R}_{\geq 0})$, so we have the following.

Theorem 7.4.4 ([LMR17])**.** *Quantum isomorphism of graphs reduces to the following problem: given $q \geq 1$, $M \in Sym_q(\mathbb{R}_{\geq 0})$ and $i, j \in [q]$, determine whether i and j are in the same orbit of* $\mathrm{Qut}(M)$.

Let $\mathrm{Sym}_q^{pd}(\mathbb{R}_{>0}) \subset \mathrm{Sym}_q(\mathbb{R}_{\geq 0})$ be the set of positive definite matrices with positive entries (such matrices play a role in the dichotomy theorems of [CMY26]). We now obtain undecidability of the vertex separation problem even for $M \in \mathrm{Sym}_q^{pd}(\mathbb{R}_{>0})$.

Theorem 7.4.5. *The following problem is undecidable: given $M \in Sym_q^{pd}(\mathbb{R}_{>0})$ and $x, y \in [q]$, determine whether there is a $K \in$ Pl-$\mathfrak{E}(M)$ such that $K_{xx} \neq K_{yy}$.*

Proof. We reduce the problem in Theorem 7.4.4 to this problem, and the result follows from the undecidability of quantum isomorphism [Ats+19; Slo19]. Let $M \in \mathrm{Sym}_q(\mathbb{R}_{\geq 0})$ and $i, j \in [q]$. Let λ be the smallest eigenvalue of $M + J$, and let $M' := M + J + (|\lambda| + 1)I$, so $M' \in \mathrm{Sym}_q^{pd}(\mathbb{R}_{>0})$. We have $I, J = ((=_1)^{\otimes 2})^{1,1} \in [=_3]_{\mathrm{Pl}} = C_{\mathrm{Qut}(\varnothing)}$ by Theorem 7.1.2, so I and J are invariant under the action of any quantum permutation matrix. Thus $\mathrm{Qut}(M) = \mathrm{Qut}(M')$, and in particular $\mathrm{Qut}(M)$ and $\mathrm{Qut}(M')$ have the same orbitals, so x and y are in the same orbit of $\mathrm{Qut}(M)$ if and only if x and y are in the same orbit of $\mathrm{Qut}(M')$. By Proposition 7.4.2, this in turn is equivalent to $K_{xx} = K_{yy}$ for every symmetric $K \in \text{Pl-}\mathfrak{E}(M')$. □

We also construct explicit M for which there exist a pair $x, y \in [q]$ separable *only* by a nonplanar edge gadget.

Proposition 7.4.4. *For infinitely many $q \geq 1$, there exist $M \in Sym_q^{pd}(\mathbb{R}_{>0})$ and $x, y \in [q]$ such that $K_{xx} = K_{yy}$ for every $K \in$ Pl-$\mathfrak{E}(M)$, but there is a $K' \in \mathfrak{E}(M)$ such that $K'_{xx} \neq K'_{yy}$.*

Proof. Chan and Martin [CM24b, Theorem 4.3] show that any two Hadamard graphs H and H' with the same number of vertices are quantum isomorphic. They also note that, if there exists a Hadamard matrix of order n, then there exist non-isomorphic Hadamard graphs on $8n$ vertices. Since there exist Hadamard matrices of infinitely many distinct orders (e.g. Walsh matrices), there exist infinitely many pairs of graphs that are quantum isomorphic but not isomorphic (we could also use the pairs of graphs with this property constructed from perfect quantum strategies for linear constraint system games in [Ats+19]). Let H, H' be such a pair, with adjacency matrices M, M'. By the discussion before Theorem 7.4.4, we may assume H and H' are connected by taking complements if needed, so there exist $x \in V(H)$ and $y \in V(H')$ in the same orbit of $\mathrm{Qut}(M \oplus M')$ but different orbits of $\mathrm{Aut}(M \oplus M')$. As in the proof of Theorem 7.4.5, let λ be the smallest eigenvalue of $(M \oplus M') + J$. Then $M'' := (M \oplus M') + J + (|\lambda| + 1)I \in \mathrm{Sym}_q^{pd}(\mathbb{R}_{>0})$ and $\mathrm{Qut}(M'')$ and $\mathrm{Qut}(M \oplus M')$ have the same orbits. By the complete symmetry among the $[q]$ vertices of the graphs with adjacency matrices I or J, $\mathrm{Aut}(M'')$ and $\mathrm{Aut}(M \oplus M')$ also have the same orbits. Hence, by Proposition 7.4.3 and Proposition 7.4.2, M'' satisfies the desired property. □

Despite the conclusion of Theorem 7.4.2, the undecidability of separability of a single given x, y by planar edge gadgets does not necessarily imply that the question of existence of a diagonal distinct edge gadget is undecidable. However, we can prove an analogue of Proposition 7.4.4 for the latter question. Say that $\mathrm{Aut}(M)$ and $\mathrm{Qut}(M)$ are *trivial* if all of their orbits have size one. Proposition 7.4.2, Proposition 7.4.3, and Theorem 7.4.2 give a characterization of both conditions for connected M.

Corollary 7.4.1. *The following hold for connected* $M \in Sym_q(\mathbb{R}_{\geq 0})$*.*

- $\mathrm{Aut}(M)$ *is trivial if and only if there exists a diagonal distinct* $\mathbf{K}(M) \in \mathfrak{E}(M)$*.*
- $\mathrm{Qut}(M)$ *is trivial if and only if there exists a diagonal distinct* $\mathbf{K}(M) \in Pl\text{-}\mathfrak{E}(M)$*.*

Corollary 7.4.2. *For infinitely many* $q \geq 1$*, there exists an* $M \in Sym_q^{pd}(\mathbb{R}_{>0})$ *such that there is a diagonal distinct matrix in* $\mathfrak{E}(M)$*, but there is no diagonal distinct matrix in* $Pl\text{-}\mathfrak{E}(M)$*.*

Proof. Using the theory of perfect quantum strategies for linear constraint system games, van Dobben de Bruyn, Roberson, and Schmidt [DRS25, Theorem A] construct an infinite family of graphs X with trivial $\mathrm{Aut}(X)$ but nontrivial $\mathrm{Qut}(X)$. As in the proof of Proposition 7.4.4, we can replace X with a connected $M \in \mathrm{Sym}_q^{pd}(\mathbb{R}_{>0})$ by adding appropriate multiples of I and J. For such M, by Corollary 7.4.1, there is a diagonal distinct edge gadget, but no diagonal distinct planar edge gadget. □

The family of matrices M in Corollary 7.4.2 could present a formidable barrier to proving a dichotomy for Pl-GH on all q because, despite M having no nontrivial automorphisms, we cannot apply Theorem 7.4.1 to obtain hardness from some matrix in Pl-$\mathfrak{E}(M)$.

Chapter 8

Holant and Orthogonal Transformation

This chapter is based on [You25b].

8.1 Introduction

Recall the Holant theorem (Theorem 2.4.1) and its corollary, the orthogonal Holant theorem (Corollary 2.4.1). Xia [Xia10] conjectured that the converse of the Holant theorem holds as long as one of $\mathcal{F}$ or $\mathcal{F}'$ contains a signature with arity greater than one – that is, if $\mathcal{F} \mid \mathcal{F}'$ and $\mathcal{G} \mid \mathcal{G}'$ are Holant-indistinguishable, then there is an $A \in \mathrm{GL}_q$ such that $\mathcal{G} = A\,\mathcal{F}$ and $\mathcal{G}' = \mathcal{F}'\,A^{-1}$. However, as discussed in Chapter 10, Cai, Guo, and Williams [CGW16, Section 4.3] showed that this conjecture is false. Xia also conjectured the converse of the orthogonal Holant theorem and proved that it holds for specific $\mathcal{F}$ and $\mathcal{G}$ consisting of real-valued symmetric signatures with small domain and/or arity. In this chapter, we verify this conjecture for any two sets of real-valued signatures.

Theorem 8.1.1. *Let $\mathcal{F}$, $\mathcal{G}$ be sets of real-valued signatures. Then $\mathcal{F}$ and $\mathcal{G}$ are Holant-indistinguishable if and only if there is a real orthogonal matrix H such that $H\,\mathcal{F} = \mathcal{G}$.*

The orthogonal Holant theorem holds for complex-valued signatures. However, Theorem 8.1.1 does not hold for general sets $\mathcal{F}$ and $\mathcal{G}$ of complex-valued signatures, even when we allow H to be complex. For example, Draisma and Regts [DR13a] consider the *vanishing* unary signature

$F \in {}_1\mathcal{S}(\mathbb{C}^2)$ defined by $F(0) = 1$ and $F(1) = i$. Any F-grid Ω with at least one vertex satisfies $\text{Holant}_\Omega(F) = 0$, as Ω is a disjoint union of K_2 complete graphs, with each component having value $F^\top F = [1, i]^\top [1, i] = 0$. Thus F is Holant-indistinguishable from 0, but there is no orthogonal matrix H, real or complex, satisfying $HF = H[1, i]^\top = [0, 0]^\top$. However, this counterexample relies on the fact that $\{F\}$ is not conjugate-closed, and we will extend Theorem 8.1.1 to conjugate-closed sets in Chapter 10.

The proof of Theorem 8.1.1 uses the orthogonal duality developed in Section 4.5, specifically Theorem 4.5.2. Existing work on the special case of edge-coloring models [Sze07; Sch08a; Dra+12; Reg12; DR13a; Reg13a] also apply the invariant theory of O_q. Schrijver [Sch08a] showed that $\mathcal{F}$ and $\mathcal{G}$ define indistinguishable real edge coloring models if and only if $\mathcal{F}$ and $\mathcal{G}$ are equivalent under a real orthogonal transformation. This is a special case of Theorem 8.1.1. Schrijver's proof exploits the specific nature of edge coloring models – that $\mathcal{F}$ and $\mathcal{G}$ consist of symmetric signatures and exactly one signature per arity – to transform input graphs into polynomials expressible in variables $y_1, \ldots, y_q$ (for $\mathcal{F}$ and $\mathcal{G}$ on domain $[q]$), where a monomial with variable multiset $\{y_{i_1}, \ldots, y_{i_n}\}$ corresponds to the $\{i_1, \ldots, i_n\}$-entry of the unique n-ary signature in $\mathcal{F}$. Another form of this result (in which the orthogonal matrix is allowed to be complex) follows from Regts' proof of [Reg13a, Lemma 5], which similarly encodes $\mathcal{F}$ and $\mathcal{G}$ as polynomials.

Odeco signature sets A real-valued symmetric signature (tensor) is *orthogonally decomposable*, or *odeco* [Rob16], if it is orthogonally transformable to a signature in $\mathcal{GEQ}$, the set of *generalized equality* signatures, which take nonzero values only when all of their inputs are equal. Hence odeco tensors generalize diagonalizable matrices. Call a set $\mathcal{F}$ of signatures odeco if the signatures are simultaneously odeco (there is a single orthogonal transformation mapping $\mathcal{F}$ into $\mathcal{GEQ}$). In counting complexity, if $\mathcal{F}$ is odeco, then $\text{Holant}(\mathcal{F})$ is polynomial-time tractable, as $\mathcal{F}$ maps into $\mathcal{GEQ}$, a trivially tractable set, under an orthogonal holographic transformation. Indeed, the tractability of *Fibonacci* signature sets [CLX08] can, with one exception, be explained by such sets being simultaneously odeco (see e.g. [CC17a, Section 2.2]). Fibonacci sets constitute almost all nontrivial tractable cases of Holant^* problems (an important variant of Holant in which all unary signatures are assumed present) for symmetric signatures on the Boolean domain [CLX11]. Simultaneously odeco sets provide a natural starting point for extending Fibonacci signatures to higher domains

[Liu24; CI25], where no full complexity dichotomy for Holant* is known.

Boralevi, Draisma, Horobeţ, and Robeva [Bor+17], resolving a conjecture of Robeva [Rob16], showed using techniques from algebraic geometry that a single tensor F is odeco if and only if the signature of a certain F-gadget is symmetric. Using Theorem 8.1.1, we in Theorem 8.5.1 extend this characterization to sets of simultaneously odeco signatures: $\mathcal{F}$ is odeco if and only if every connected $\mathcal{F}$-gadget has a symmetric signature. The latter condition is equivalent to the symmetry of the signatures of a set of small gadgets constructed from every pair of signatures in $\mathcal{F}$. Therefore, if $\mathcal{F}$ is finite, our characterization yields a simple $O(q^{2n-2}|\mathcal{F}|^2)$-time algorithm (for $\mathcal{F}$ on domain $[q]$ with maximum arity n) for deciding whether $\mathcal{F}$ is odeco. Our characterization also deepens the connection between Fibonacci and odeco signatures, as the original proof of tractability of any Fibonacci signature set $\mathcal{F}$ [CLX08] relied on the fact that every connected $\mathcal{F}$-gadget has a signature which is itself Fibonacci (in particular, is symmetric). One can view the (iii) $\implies$ (ii) result in Theorem 8.5.1 as a general-domain version of this proof.

Overview. The proof of Theorem 8.1.1, like the intertwiner proof for #CSP indistinguishability in Chapter 5, applies the orthogonal invariant-gadget duality of Theorem 4.5.2. However, the rest of the #CSP proof uses the orbits of $\text{Aut}(\mathcal{F})$, which do not have analogues for $\text{Stab}_O(\mathcal{F})$. Instead, we apply a novel method: induction on the domain size q. We show in Lemma 8.3.2 that, if $\mathcal{F}$ and $\mathcal{G}$ contain a binary signature with nontrivial (i.e., not a multiple of I) diagonal flattened form, then we can separate $[q]$ into smaller subdomains and inductively transform the restrictions of $\mathcal{F}$ and $\mathcal{G}$ to each subdomain, producing a full transformation from $\mathcal{F}$ to $\mathcal{G}$. Using Theorem 4.5.2, we show that there is some nonzero matrix D intertwining $\mathcal{F}$ and $\mathcal{G}$. Exploiting the power of diagonalization afforded by orthogonal transformations, we may assume this D is diagonal. Either $D = \pm I$, giving a trivial orthogonal transformation between $\mathcal{F}$ and $\mathcal{G}$, or D is not a multiple of I, in which case we use the fact that D intertwines $\mathcal{F}$ and $\mathcal{G}$ to add D to both $\mathcal{F}$ and $\mathcal{G}$ while preserving their Holant-indistinguishability, then apply induction. In Section 8.4, we show that Theorem 8.1.1 encompasses a wide range of existing counting indistinguishability theorems, and yields some novel variations of these theorems. In Section 8.5, we prove our combinatorial characterization of odeco signature sets.

8.2 Subtensors

8.2.1 Block tensor actions

In this section, we define block tensors and prove several technical results stating that the action of a block matrix H on a block tensor K follows block matrix multiplication rules as one would expect.

Definition 8.2.1. Let I be an index/domain set, and $X \sqcup Y = I$ be a nontrivial partition of I.

1. For a matrix $H \in \mathbb{K}^{I\times I}$ and $R, C \in \{X, Y\}$, let $H|_{R,C} \in \mathbb{K}^{R\times C}$ be the submatrix of H with rows indexed by R and columns indexed by C. Up to row and column reordering, H is the block matrix
$$H = \begin{bmatrix} H|_{X,X} & H|_{X,Y} \\ H|_{Y,X} & H|_{Y,Y} \end{bmatrix}.$$

2. Let $F^{m,d} \in \mathbb{K}^{I^m \times I^d}$ be a flattened tensor, and let $\mathbf{R} \in \{X, Y\}^m$ and $\mathbf{C} \in \{X, Y\}^d$. Define $F^{m,d}|_{\mathbf{R},\mathbf{C}} \in \mathbb{K}^{\mathbf{R}\times\mathbf{C}}$ as the submatrix of $F^{m,d}$ with rows indexed by $\mathbf{R}$ and columns indexed by $\mathbf{C}$.

We will frequently apply Definition 8.2.1 to the partition $V(\mathcal{F}) \sqcup V(\mathcal{G})$ of the domain of $\mathcal{F} \oplus \mathcal{G}$. It follows directly from the definitions that, for $\mathcal{F} \ni F \rightsquigarrow G \in \mathcal{G}$ and any $m + d = \text{arity}(F)$,

$$(F \oplus G)^{m,d}|_{\mathbf{R},\mathbf{C}} = \begin{cases} F^{m,d} & \mathbf{R} = V(\mathcal{F})^m \wedge \mathbf{C} = V(\mathcal{F})^d \\ G^{m,d} & \mathbf{R} = V(\mathcal{G})^m \wedge \mathbf{C} = V(\mathcal{G})^d \\ 0 & \text{otherwise.} \end{cases} \tag{8.2.1}$$

Proposition 8.2.1. *Let $I = X \sqcup Y$. For $H \in \mathbb{K}^{I\times I}$, $K^{m,d} \in \mathbb{K}^{I^m\times I^d}$ and any $\mathbf{R} \in \{X, Y\}^m$ $\mathbf{C} \in \{X, Y\}^d$,*
$$(H^{\otimes m} K^{m,d})|_{\mathbf{R},\mathbf{C}} = \sum_{\mathbf{J}\in\{X,Y\}^m} \left(\bigotimes_{i=1}^{m} H|_{R_i,J_i} \right) K^{m,d}|_{\mathbf{J},\mathbf{C}}$$
(with $H^{\otimes m} K^{m,d} \in \mathbb{K}^{I^m\times I^d}$ indexed as in part 2 of Definition 8.2.1) and similarly
$$(K^{m,d} H^{\otimes d})|_{\mathbf{R},\mathbf{C}} = \sum_{\mathbf{J}\in\{X,Y\}^d} K^{m,d}|_{\mathbf{R},\mathbf{J}} \left(\bigotimes_{i=1}^{d} H|_{J_i,C_i} \right).$$

That is, with $H^{\otimes m}|_{\mathbf{R},\mathbf{J}} = \bigotimes_{i=1}^{m} H|_{R_i,J_i}$, we can compute $H^{\otimes m}K^{m,d}$ as a block matrix product with

$$H^{\otimes m} = \begin{bmatrix} H^{\otimes m}|_{X^m,X^m} & H^{\otimes m}|_{X^m,X^{m-1}Y} & H^{\otimes m}|_{X^m,X^{m-2}YX} & \cdots & H^{\otimes m}|_{X^m,Y^m} \\ H^{\otimes m}|_{X^{m-1}Y,X^m} & H^{\otimes m}|_{X^{m-1}Y,X^{m-1}Y} & H^{\otimes m}|_{X^{m-1}Y,X^{m-2}YX} & \cdots & H^{\otimes m}|_{X^{m-1}Y,Y^m} \\ \vdots & \vdots & \vdots & & \vdots \\ H^{\otimes m}|_{Y^{m-1}X,X^m} & H^{\otimes m}|_{Y^{m-1}X,X^{m-1}Y} & H^{\otimes m}|_{Y^{m-1}X,X^{m-2}YX} & \cdots & H^{\otimes m}|_{Y^{m-1}X,Y^m} \\ H^{\otimes m}|_{Y^m,X^m} & H^{\otimes m}|_{Y^m,X^{m-1}Y} & H^{\otimes m}|_{Y^m,X^{m-2}YX} & \cdots & H^{\otimes m}|_{Y^m,Y^m} \end{bmatrix}$$

and

$$K^{m,d} = \begin{bmatrix} K^{m,d}|_{X^m,X^d} & K^{m,d}|_{X^m,X^{d-1}Y} & \cdots & K^{m,d}|_{X^m,Y^{d-1}X} & K^{m,d}|_{X^m,Y^d} \\ K^{m,d}|_{X^{m-1}Y,X^d} & K^{m,d}|_{X^{m-1}Y,X^{d-1}Y} & \cdots & K^{m,d}|_{X^{m-1}Y,Y^{d-1}X} & K^{m,d}|_{X^{m-1}Y,Y^d} \\ K^{m,d}|_{X^{m-2}YX,X^d} & K^{m,d}|_{X^{m-2}YX,X^{d-1}Y} & \cdots & K^{m,d}|_{X^{m-2}YX,Y^{d-1}X} & K^{m,d}|_{X^{m-2}YX,Y^d} \\ \vdots & \vdots & & \vdots & \vdots \\ K^{m,d}|_{Y^m,X^d} & K^{m,d}|_{Y^m,X^{d-1}Y} & \cdots & K^{m,d}|_{Y^m,Y^{d-1}X} & K^{m,d}|_{Y^m,Y^d}. \end{bmatrix}.$$

Proof. We prove the first statement. The second is proved similarly. Let $\mathbf{r} \in \mathbf{R}$ and $\mathbf{c} \in \mathbf{C}$. Then

$$\begin{aligned} ((H^{\otimes m}K^{m,d})|_{\mathbf{R},\mathbf{C}})_{\mathbf{r},\mathbf{c}} = (H^{\otimes m}K^{m,d})_{\mathbf{r},\mathbf{c}} &= \sum_{\mathbf{j}\in I^m} \left(\prod_{i=1}^{m} H_{r_i,j_i}\right) K^{m,d}_{\mathbf{j},\mathbf{c}} \\ &= \sum_{\mathbf{J}\in\{X,Y\}^m} \sum_{\mathbf{j}\in\mathbf{J}} \left(\prod_{i=1}^{m} H_{r_i,j_i}\right) K^{m,d}_{\mathbf{j},\mathbf{c}} \\ &= \sum_{\mathbf{J}\in\{X,Y\}^m} \sum_{\mathbf{j}\in\mathbf{J}} \left(\bigotimes_{i=1}^{m} H|_{R_i,J_i}\right)_{\mathbf{r},\mathbf{j}} (K^{m,d}|_{\mathbf{J},\mathbf{C}})_{\mathbf{j},\mathbf{c}} \\ &= \sum_{\mathbf{J}\in\{X,Y\}^m} \left(\left(\bigotimes_{i=1}^{m} H|_{R_i,J_i}\right) K^{m,d}|_{\mathbf{J},\mathbf{C}}\right)_{\mathbf{r},\mathbf{c}} \\ &= \left(\sum_{\mathbf{J}\in\{X,Y\}^m} \left(\bigotimes_{i=1}^{m} H|_{R_i,J_i}\right) K^{m,d}|_{\mathbf{J},\mathbf{C}}\right)_{\mathbf{r},\mathbf{c}}. \end{aligned}$$

□

We will only need Proposition 8.2.1 as written, for partitions of I into two blocks, but it is not hard to see that it extends naturally to partitions of I into more than two blocks.

We will apply Proposition 8.2.1 for block-diagonal H.

Corollary 8.2.1. *If*

$$H = H_X \oplus H_Y = \begin{bmatrix} H_X & 0 \\ 0 & H_Y \end{bmatrix}$$

is block-diagonal, and $K^{n,0} \in \mathcal{S}(\mathbb{K}^{X \sqcup Y})$*, then the block form of* $H^{\otimes n} K^{n,0}$ *is*

$$\begin{bmatrix} H_X^{\otimes n} & 0 & \dots & 0 & 0 \\ 0 & * & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & * & 0 \\ 0 & 0 & \dots & 0 & H_Y^{\otimes n} \end{bmatrix} \begin{bmatrix} (K|_X)^{n,0} \\ * \\ \vdots \\ * \\ (K|_Y)^{n,0} \end{bmatrix} = \begin{bmatrix} H_X^{\otimes n}(K|_X)^{n,0} \\ * \\ \vdots \\ * \\ H_Y^{\otimes n}(K|_Y)^{n,0} \end{bmatrix}.$$

Proof. Here, $H|_{R_i,J_i} = H_{R_i}$ if $R_i = J_i$ and $H|_{R_i,J_i} = 0$ if $R_i \neq J_i$, so H takes the claimed block form in Proposition 8.2.1. □

If $X = V(F)$, $Y = V(G)$, and $K = F \oplus G \in \mathcal{S}(\mathbb{K}^{X \sqcup Y})$, then by (8.2.1), all blocks of K are 0 except $K|_X = F$ and $K|_Y = G$, so, by Corollary 8.2.1,

$$(H_{V(F)} \oplus H_{V(G)}) \cdot (F \oplus G) = (H_{V(F)} \cdot F) \oplus (H_{V(G)} \cdot G). \qquad (8.2.2)$$

8.2.2 Restricting to a subdomain

In this subsection, we prove several more results on subdomain restriction (recall Definition 5.2.2). Since $[\mathcal{F}] := \langle \mathcal{F}, \supset, \subset \rangle$ by definition and the restrictions of $\supset$ and $\subset$ to a subdomain are just $\supset$ and $\subset$ on the subdomain, the results in this subsection concerning $\langle \cdot \rangle$ also apply to $[\cdot]$.

Definition 8.2.2 ($\langle \mathcal{F} \rangle_X$)**.** For $\mathcal{F} \subset \mathcal{V}(\mathbb{K}^q)$, define $\langle \mathcal{F} \rangle_X := \{F|_X : F \in \langle \mathcal{F} \rangle\} \subset \mathcal{V}(\mathbb{K}^X)$, a set on domain X bijective with $\langle \mathcal{F} \rangle$.

Definition 8.2.3 ($(\cdot)^{\uparrow Z}$)**.** Let $X \subset Z$ and $F \in \mathcal{V}(\mathbb{K}^X)$. Define $F^{\uparrow Z} \in \mathcal{V}(\mathbb{K}^Z)$ by

$$F^{\uparrow Z}(\mathbf{x}) = \begin{cases} F(\mathbf{x}) & \mathbf{x} \in X^n \\ 0 & \text{otherwise} \end{cases} \quad \text{for } \mathbf{x} \in Z^n.$$

That is, $F^{\uparrow Z}$ expands the domain of F to all of Z by padding with zeros. We frequently write simply $F^{\uparrow}$ when the ambient domain Z is clear from context.

The next two results show the utility of realizing the *subdomain restrictor* $I_X^{\uparrow} = \left[\begin{smallmatrix} I_X & 0 \\ 0 & 0 \end{smallmatrix}\right] \in {}_1\mathcal{V}_1$, which acts like an edge (I) on inputs from X and zeroes out $[q] \backslash X$. Note that, while $\langle \mathcal{F} \rangle_X \subset \langle \langle \mathcal{F} \rangle_X \rangle$, we may not have $\langle \mathcal{F} \rangle_X \supset \langle \langle \mathcal{F} \rangle_X \rangle$. For example, if $\langle \mathcal{F} \rangle$ contains the $(X, [q] \setminus X)$-block matrix $\left[\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right]$, then $A \in \langle \mathcal{F} \rangle_X$, so $A^2 \in \langle \langle \mathcal{F} \rangle_X \rangle$, but we may not be able to obtain A^2 as the X-block of a matrix in $\langle \mathcal{F} \rangle$. However, if $I_X^{\uparrow} \in \langle \mathcal{F} \rangle$, then the reverse inclusion does hold:

Proposition 8.2.2. *Suppose* $\langle\mathcal{F}\rangle \ni I_X^\uparrow \leftrightsquigarrow I_X^\uparrow \in \langle\mathcal{G}\rangle$. *Then, for any* $\langle\langle\mathcal{F}\rangle_X\rangle \ni F \leftrightsquigarrow G \in \langle\langle\mathcal{G}\rangle_X\rangle$, *we have* $\langle\mathcal{F}\rangle \ni F^\uparrow \leftrightsquigarrow G^\uparrow \in \langle\mathcal{G}\rangle$. *Therefore* $\langle\mathcal{F}\rangle_X = \langle\langle\mathcal{F}\rangle_X\rangle$.

Proof. By definition, F is the tensor of a quantum-$\langle\mathcal{F}\rangle_X$-gadget $\mathbf{K}$ and G is the tensor of $\mathbf{K}_{\langle\mathcal{F}\rangle_X\to\langle\mathcal{G}\rangle_X}$. Construct a quantum-$\langle\mathcal{F}\rangle$-gadget $\mathbf{K}^\uparrow$ as follows. Start with $\mathbf{K}_{\langle\mathcal{F}\rangle_X\to\langle\mathcal{F}\rangle}$, constructed by replacing each $S|_X \in \langle\mathcal{F}\rangle_X$ in $\mathbf{K}$ with the corresponding $S \in \langle\mathcal{F}\rangle$. Then replace each dangling and internal edge – which when viewed alone is a wire gadget with tensor I – with $I_X^\uparrow \in {}_1\langle\mathcal{F}\rangle_1$. This has the effect of forcing all edges in $\mathbf{K}^\uparrow$, including dangling edges, to take values in X, so the tensor of $\mathbf{K}^\uparrow$ is $F^\uparrow$. Similarly, the tensor of $(\mathbf{K}_{\langle\mathcal{F}\rangle_X\to\langle\mathcal{G}\rangle_X})^\uparrow$ is $G^\uparrow$, and, since $(\mathbf{K}_{\langle\mathcal{F}\rangle_X\to\langle\mathcal{G}\rangle_X})^\uparrow = (\mathbf{K}^\uparrow)_{\langle\mathcal{F}\rangle\to\langle\mathcal{G}\rangle}$, we have $F^\uparrow \leftrightsquigarrow G^\uparrow$. See Figure 8.1.

The second claim follows from the first and the fact that $F^\uparrow|_X = F$. □

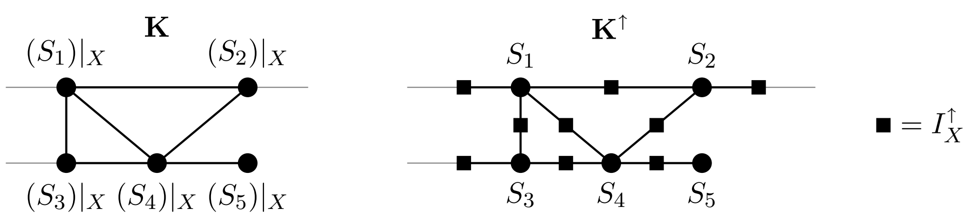

Figure 8.1: The construction in Propositions 8.2.2 and 8.2.3.

Proposition 8.2.3. *If* $\mathcal{F}$ *and* $\mathcal{G}$ *are Bi-Holant-indistinguishable and* $\langle\mathcal{F}\rangle \ni I_X^\uparrow \leftrightsquigarrow I_X^\uparrow \in \langle\mathcal{G}\rangle$, *then* $\langle\mathcal{F}\rangle_X$ *and* $\langle\mathcal{G}\rangle_X$ *are Bi-Holant-indistinguishable.*

Proof. Let Ω be an $\langle\mathcal{F}\rangle_X$-grid. Viewing Ω as a $(0,0)$-$\langle\mathcal{F}\rangle_X$-gadget, construct the quantum $\mathcal{F}$-grid $\Omega^\uparrow$ as in the proof of Proposition 8.2.2. Applying similar reasoning, we obtain

$$\text{Bi-Holant}_{\langle\mathcal{F}\rangle_X}(\Omega) = \text{Bi-Holant}_{\langle\mathcal{F}\rangle}\left(\Omega^\uparrow\right) = \text{Bi-Holant}_{\langle\mathcal{G}\rangle}\left(\Omega^\uparrow_{\langle\mathcal{F}\rangle\to\langle\mathcal{G}\rangle}\right) = \text{Bi-Holant}_{\langle\mathcal{G}\rangle_X}(\Omega_{\langle\mathcal{F}\rangle_X\to\langle\mathcal{G}\rangle_X}).$$

□

8.3 Domain induction

Our proof of this chapter's indistinguishability theorem Theorem 8.1.1, like the proofs of the isomorphism and quantum isomorphism indistinguishability theorems in chapters 5 and 7, uses the direct sum $\mathcal{F}\oplus\mathcal{G}$. However, the orthogonal group $\text{Stab}_O(\mathcal{F}\oplus\mathcal{G})$, unlike $\text{Aut}(\mathcal{F}\oplus\mathcal{G})$ and $\text{Qut}(\mathcal{F}\oplus\mathcal{G})$, does not admit any sort of orbits on $V(\mathcal{F}) \sqcup V(\mathcal{G})$. Thus we cannot follow the path from indistinguishability to orbits to transformation followed by Corollary 5.1.1 and Proposition 5.2.2 in Chapter 5 and by Corollary 7.1.1 and Lemma 7.2.1 in Chapter 7. Instead, the following lemma, the only nonconstructive step in the proof of Theorem 8.1.1, is a 'continuous' version of the first step on this path, from indistinguishability to finding a vertex of $\mathcal{F}$ and a vertex of $\mathcal{G}$ in the same orbit. We show, using the same orthogonal duality Theorem 4.5.2 used in the indistinguishability to orbit step in Chapter 5, that $\mathrm{Stab}_O(\mathcal{F} \oplus \mathcal{G})$ must contain a matrix that nontrivially intertwines $\mathcal{F}$ and $\mathcal{G}$ (recall that Holant-indistinguishable $\mathcal{F}$ and $\mathcal{G}$ must have the same domain size q, the Holant value of the vertexless loop):

Lemma 8.3.1. *If $\mathcal{F}, \mathcal{G} \subset \mathcal{S}(\mathbb{R}^q)$ are Holant-indistinguishable, then there is an $H \in \mathrm{Stab}_O(\mathcal{F} \oplus \mathcal{G})$ with $H|_{\mathcal{F},\mathcal{G}} \neq 0$ or $H|_{\mathcal{G},\mathcal{F}} \neq 0$.*

Proof. Suppose towards contradiction that every $H \in \mathrm{Stab}_O(\mathcal{F} \oplus \mathcal{G})$ satisfies $H|_{V(\mathcal{F}),V(\mathcal{G})} = H|_{V(\mathcal{G}),V(\mathcal{F})} = 0$ (i.e. is block diagonal). Then the block diagonal matrix $A = \left[\begin{smallmatrix} I & 0 \\ 0 & 2I \end{smallmatrix}\right] \in {}_1\mathcal{V}(\mathbb{R}^{2q})_1$ satisfies $HA = AH$, hence $HAH^T = A$, for every $H \in \mathrm{Stab}_O(\mathcal{F} \oplus \mathcal{G})$, so $A \in \mathcal{V}(\mathbb{R}^{2q})^{\mathrm{Stab}_O(\mathcal{F} \oplus \mathcal{G})}$. Thus, by Theorem 4.5.2, $A \in [\mathcal{F} \oplus \mathcal{G}]$ is realizable as the tensor of a quantum gadget. But now, by Proposition 5.2.1,

$$[\mathcal{F}] \ni I = A|_{\mathcal{F}} \rightsquigarrow A|_{\mathcal{G}} = 2I \in [\mathcal{G}]. \tag{8.3.1}$$

By Proposition 4.1.1, $[\mathcal{F}]$ and $[\mathcal{G}]$ should be Holant-indistinguishable, but the corresponding $[\mathcal{F}]$ and $[\mathcal{G}]$-grids consisting of a single looped vertex assigned I or $2I$ have Holant values $\mathrm{tr}(I) = q \neq 2q = \mathrm{tr}(2I)$, a contradiction. $\square$

The following definition and its applications below borrow a simple but powerful idea of Shao and Cai [SC20, Section 8.2]: isolating all vertices of an $\mathcal{F} \cup \{F\}$-grid Ω assigned F, the rest of Ω is an $\mathcal{F}$-gadget.

Definition 8.3.1 (Subgadget, $\overline{\mathbf{K}}$)**.** Let $\mathbf{J}$ be a gadget. A *subgadget* $\mathbf{K} \subset \mathbf{J}$ induced by a subset $U \subset V(\mathbf{J})$ of vertices of $\mathbf{J}$ is a gadget composed of the vertices in U and *all* of their incident edges: internal edges of $\mathbf{J}$ incident to exactly one vertex in U become new dangling edges of $\mathbf{K}$. For any $\mathbf{K} \subset \mathbf{J}$, there is a unique (up to left/right dangling edge pivoting) $\overline{\mathbf{K}} \subset \mathbf{J}$, induced by $V(\mathbf{J}) \setminus U$, called the *complement* of $\mathbf{K}$, such that, upon reconnecting the new dangling edges of $\mathbf{K}$ and $\overline{\mathbf{K}}$, we recover $\mathbf{J}$.

We often take $\mathbf{J}$ to be a signature grid (0-ary gadget) Ω, in which case $\Omega = \langle \mathbf{K}, \overline{\mathbf{K}} \rangle$.

Say $\mathcal{F}$ is *quantum-gadget-closed* if $\mathcal{F} = [\mathcal{F}]$ (as sets, rather than multisets). To replace second step – from orbits to transformation – on the path of the (quantum) isomorphism indistinguishability proof, we turn to a novel approach: induction on the domain size of $\mathcal{F}$ and $\mathcal{G}$.

Lemma 8.3.2. *Let $\mathcal{F}$ and $\mathcal{G}$ be signature sets on domain $[q]$, and suppose Theorem 8.1.1 holds for all $\mathcal{F}'$, $\mathcal{G}'$ on domain smaller than q. If $\mathcal{F}$ and $\mathcal{G}$ are Holant-indistinguishable and contain corresponding copies of a diagonal matrix (binary signature) $D \notin \operatorname{span}(I)$, then there is a real orthogonal transformation between $\mathcal{F}$ and $\mathcal{G}$.*

Proof. By Proposition 4.1.1, we may replace $\mathcal{F}$ and $\mathcal{G}$ with $[\mathcal{F}]$ and $[\mathcal{G}]$ to assume $\mathcal{F}$ and $\mathcal{G}$ are quantum-gadget-closed. Since $D \notin \operatorname{span}(I)$, there exist $x, y \in [q]$ such that $D_{x,x} \neq D_{y,y}$, so the sets

$$X = \{z \in [q] : D_{z,z} = D_{x,x}\} \text{ and } Y = [q] \setminus X$$

are a nontrivial partition of $[q]$. Since $\mathcal{F}$ is quantum-gadget-closed, it contains $I \in \mathcal{W}$. Consider the subalgebra $\langle D, I \rangle_{+,\circ} \subset [\mathcal{F}] = \mathcal{F}$. Since D and I are diagonal, composition $\circ$ is equivalent to entrywise multiplication in $\langle D, I \rangle_{+,\circ}$. Therefore, by Proposition 5.1.2 (identifying the matrix diagonals with $\mathbb{R}^q$), we have $I_X^{\uparrow}, I_Y^{\uparrow} \in \langle D, I \rangle_{+,\circ} \subset \mathcal{F}$. Applying the same interpolation in $\mathcal{G}$, we obtain corresponding copies of $I_X^{\uparrow}, I_Y^{\uparrow} \in \mathcal{G}$. Now Proposition 8.2.3 asserts that $\mathcal{F}\,|_X$ and $\mathcal{G}\,|_X$ are Holant-indistinguishable. Furthermore, $|X| < q$, so, by assumption, there is an orthogonal matrix $H_X \in \mathbb{R}^{X \times X}$ satisfying

$$H_X \, \mathcal{F}\,|_X = \mathcal{G}\,|_X. \tag{8.3.2}$$

For $b \in [q]$, let $\Delta_b \in \mathcal{S}(\mathbb{R}^q)$ be the unary *pinning signature* defined by

$$\Delta_b(x) = \begin{cases} 1 & b = x \\ 0 & b \neq x \end{cases}$$

for $x \in [q]$. Call signatures in $\{\Delta_b \mid b \in X\}$ *X-pins*. Let I_Y be the identity operator on $\mathbb{R}^{Y \times Y}$ and define

$$\mathcal{F}' = \mathcal{F} \sqcup \{\Delta_b \mid b \in X\} \text{ and } \mathcal{G}' = \big((H_X^{-1} \oplus I_Y)\,\mathcal{G}\,\big) \sqcup \{\Delta_b \mid b \in X\}. \tag{8.3.3}$$

Claim 8.3.1. $\mathcal{F}'$ and $\mathcal{G}'$ are Holant-indistinguishable.

To see this, let Ω be a connected $\mathcal{F}'$-grid. If Ω contains no X-pins then it is an $\mathcal{F}$-grid, and $\Omega_{\mathcal{F}'\to\mathcal{G}'}$ is the corresponding $(H_X^{-1} \oplus I_Y)\,\mathcal{G}$-grid, so by assumption and the orthogonal Holant theorem, $\text{Holant}_\Omega = \text{Holant}_{\Omega_{\mathcal{F}'\to\mathcal{G}'}}$. If Ω contains two adjacent X-pins, then, since Ω is connected, its underlying multigraph consists only of these two vertices. X-pins in $\mathcal{F}'$ correspond to identical X-pins in $\mathcal{G}'$, so again $\text{Holant}_\Omega = \text{Holant}_{\Omega_{\mathcal{F}'\to\mathcal{G}'}}$. Otherwise, Ω contains p pairwise non-adjacent vertices assigned X-pins. Let $\mathbf{K} \subset \Omega$ be a subgadget induced by these p vertices, so $\mathbf{K}$'s signature is $\bigotimes_{i=1}^p \Delta_{b_i}$ for some $b_1, \dots, b_p \in X$. Since $\mathbf{K}$ includes *all* the vertices in Ω assigned X-pins, its complement $\overline{\mathbf{K}}$ is an $\mathcal{F}$-gadget. See Figure 8.1. Hence the signature F of $\overline{\mathbf{K}}$ is in $\mathcal{F}$, as $\mathcal{F}$ is

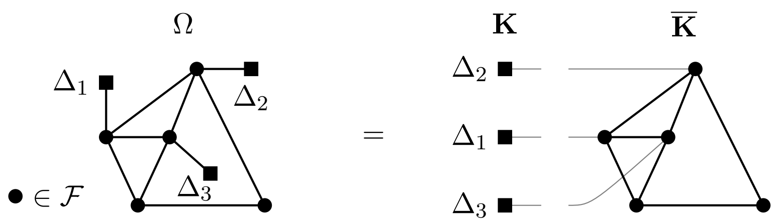

Figure 8.1: An $\mathcal{F}'$-grid Ω and its pin-induced subgadget $\mathbf{K}$ and complement $\overline{\mathbf{K}}$

quantum-gadget-closed. From $\Omega = \langle \mathbf{K}, \overline{\mathbf{K}} \rangle$, we obtain

$$\text{Holant}_\Omega = \left\langle \bigotimes_{i=1}^{p} \Delta_{b_i}, F \right\rangle = F_{\mathbf{b}},$$

where $\mathbf{b} = b_1 b_2 \dots b_p$. A similar calculation, using the fact that $\mathcal{G}$ is also quantum-gadget-closed, shows $\text{Holant}_{\Omega_{\mathcal{F}'\to\mathcal{G}'}} = ((H_X^{-1} \oplus I_Y)G)_{\mathbf{b}}$ for $\mathcal{G} \ni G \leftrightsquigarrow F$. Then, since $\mathbf{b} \in X^p$, we have, using (8.3.2),

$$\text{Holant}_\Omega = F_{\mathbf{b}} = (F|_X)_{\mathbf{b}} = (H_X^{-1} G|_X)_{\mathbf{b}} = (((H_X^{-1} \oplus I_Y)G)|_X)_{\mathbf{b}} = ((H_X^{-1} \oplus I_Y)G)_{\mathbf{b}} = \text{Holant}_{\Omega_{\mathcal{F}'\to\mathcal{G}'}}$$

(where the fourth equality uses Corollary 8.2.1 with $K := G$ and $H := H_X^{-1} \oplus I_Y$). So $\mathcal{F}'$ and $\mathcal{G}'$ are Holant-indistinguishable, completing the proof of Claim 8.3.1.

By Proposition 4.1.1, $[\mathcal{F}']$ and $[\mathcal{G}']$ are also Holant-indistinguishable. Since $\mathcal{F} \subset [\mathcal{F}']$ and $\mathcal{G} \subset [\mathcal{G}']$, it suffices to find an orthogonal transformation between $[\mathcal{F}']$ and $[\mathcal{G}']$. Note that $[\mathcal{F}']$ and $[\mathcal{G}']$ still contain $I_Y^{\updownarrow}$ (as $I_Y^{\updownarrow}$, being 0 on X, is unaffected by the transform $H_X^{-1} \oplus I_Y$), so we may apply Proposition 8.2.3 to conclude $\big([\mathcal{F}']\big)\big|_Y$ and $\big([\mathcal{G}']\big)\big|_Y$ are Holant-indistinguishable. Again, since $|Y| < q$, there is an orthogonal $H_Y \in \mathbb{R}^{Y \times Y}$ such that $H_Y\big([\mathcal{F}']\big)\big|_Y = \big([\mathcal{G}']\big)\big|_Y$. Define

$$\mathcal{F}'' = [\mathcal{F}'] \sqcup \{\Delta_b \mid b \in Y\} \text{ and } \mathcal{G}'' = \big((I_X \oplus H_Y^{-1})[\mathcal{G}']\big) \sqcup \{\Delta_b \mid b \in Y\}. \tag{8.3.4}$$

Since $(I_X \oplus H_Y^{-1}) \in O_q$, it suffices to find an orthogonal transformation between $\mathcal{F}''$ and $\mathcal{G}''$. As $[\mathcal{F}']$ and $[\mathcal{G}']$ are, like $\mathcal{F}$ and $\mathcal{G}$, quantum-gadget-closed and Holant-indistinguishable, we repeat the proof of Claim 8.3.1, with (8.3.4) in place of (8.3.3), to show that $\mathcal{F}''$ and $\mathcal{G}''$ are Holant-indistinguishable. Observe that, by definition, $\mathcal{F}''$ and $\mathcal{G}''$ contain corresponding copies of all Y-pins $\{\Delta_b \mid b \in Y\}$. Furthermore, $\mathcal{F}'' \supset [\mathcal{F}'] \supset \mathcal{F}' \supset \{\Delta_b \mid b \in X\}$ and

$$\mathcal{G}'' \supset (I_X \oplus H_Y^{-1})[\mathcal{G}'] \supset (I_X \oplus H_Y^{-1})\mathcal{G}' \supset (I_X \oplus H_Y^{-1})\{\Delta_b \mid b \in X\} = \{\Delta_b \mid b \in X\},$$

where the final equality holds because the X-pins are zero on Y, so are unaffected by the transform $(I_X \oplus H_Y^{-1})$. Thus $\mathcal{F}''$ and $\mathcal{G}''$ contain corresponding copies of all pins Δ_b for $b \in [q]$. We claim that this implies that $\mathcal{F}'' = \mathcal{G}''$. To see this, consider any $\mathcal{F}'' \ni F \leftrightsquigarrow G \in \mathcal{G}''$ of common arity n. For any $\mathbf{x} \in [q]^n$, since $\mathcal{F}''$ and $\mathcal{G}''$ are Holant-indistinguishable, we have

$$F_{\mathbf{x}} = \left\langle F, \bigotimes_{i=1}^{n} \Delta_{x_i} \right\rangle = \left\langle G, \bigotimes_{i=1}^{n} \Delta_{x_i} \right\rangle = G_{\mathbf{x}}.$$

Thus $F = G$ for every $\mathcal{F}'' \ni F \leftrightsquigarrow G \in \mathcal{G}''$, so $\mathcal{F}'' = \mathcal{G}''$. □

The final step is to realize the diagonal matrix D in the statement of Lemma 8.3.2, and apply induction.

Proof of Theorem 8.1.1. The ($\Longleftarrow$) direction is the orthogonal Holant theorem. We show the converse. Let $\mathcal{F}, \mathcal{G}$ be Holant-indistinguishable. We proceed by induction on the domain size q. If $q = 1$ then the only orthogonal matrices in $\mathbb{R}^{q \times q}$ are $\pm I$ and every signature F, regardless of arity, has a single entry, which we denote by $F_0 \in \mathbb{R}$. Let $F \leftrightsquigarrow G$ have even arity n. Contracting $\frac{n}{2}$ pairs of inputs of both F and G, we obtain corresponding signature grids with values F_0 and G_0, respectively, so $F_0 = G_0$, hence $F = G$. For $F \leftrightsquigarrow G$ of odd arity n,

$$F_0^2 = \|F\|^2 = \langle F, F \rangle = \langle G, G \rangle = \|G\|^2 = G_0^2,$$

so $F_0 = \pm G_0$. Let $F \leftrightsquigarrow G$ have odd arity n and $F' \leftrightsquigarrow G'$ have odd arity n', all nonzero. Then $F^{\otimes n'} \otimes (F')^{\otimes n}$ has even arity $2nn'$; contracting its nn' pairs of inputs gives a $\mathcal{F}$-grid with value $F_0^{n'}(F_0')^n$. The corresponding $\mathcal{G}$-grid has value $G_0^{n'}(G_0')^n$. Let $G_0 = (-1)^a F_0$ and $G_0' = (-1)^{a'} F_0'$ for $a, a' \in \{0, 1\}$. Then

$$F_0^{n'}(F_0')^n = G_0^{n'}(G_0')^n = (-1)^{an'+a'n} F_0^{n'}(F_0')^n = (-1)^{a+a'} F_0^{n'}(F_0')^n$$

(where in the final equality we used that n, n' are odd), so $a = a'$. Thus there is a common $a \in \{0,1\}$ such that $((-1)^a)^n F = G$ for every n-ary $F \rightsquigarrow G$, so $(-I)^a \mathcal{F} = \mathcal{G}$.

Now assume $q > 1$. By Proposition 4.1.1, we may assume $\mathcal{F}$ and $\mathcal{G}$ are quantum-gadget-closed. By Lemma 8.3.1, there is an $H \in \mathrm{Stab}_O(\mathcal{F} \oplus \mathcal{G})$ with either $H|_{\mathcal{F},\mathcal{G}} \neq 0$ or $H|_{\mathcal{G},\mathcal{F}} \neq 0$. Assume WLOG that $H|_{\mathcal{G},\mathcal{F}} \neq 0$. Let $H|_{\mathcal{G},\mathcal{F}} = U^\top D V$ be the singular value decomposition of $H|_{\mathcal{G},\mathcal{F}}$, with U, V orthogonal and $D \neq 0$ diagonal, all real. Replace $\mathcal{F}$ with $V\mathcal{F}$ and $\mathcal{G}$ with $U\mathcal{G}$ (this does not affect the existence of an orthogonal transformation between $\mathcal{F}$ and $\mathcal{G}$ and, by the orthogonal Holant theorem, preserves Holant-indistinguishability of $\mathcal{F}$ and $\mathcal{G}$). This has the effect of replacing $\mathcal{F} \oplus \mathcal{G}$ with $(V\mathcal{F}) \oplus (U\mathcal{G}) = (V \oplus U)(\mathcal{F} \oplus \mathcal{G})$ (by (8.2.2)), which replaces $\mathrm{Stab}_O(\mathcal{F} \oplus \mathcal{G})$ with $(V \oplus U) \circ \mathrm{Stab}_O(\mathcal{F} \oplus \mathcal{G}) \circ (V \oplus U)^{-1} = (V \oplus U) \circ \mathrm{Stab}_O(\mathcal{F} \oplus \mathcal{G}) \circ (V \oplus U)^\top$. In particular, H is replaced with

$$\begin{bmatrix} V & 0 \\ 0 & U \end{bmatrix} \begin{bmatrix} H|_{\mathcal{F},\mathcal{F}} & H|_{\mathcal{F},\mathcal{G}} \\ U^\top D V & H|_{\mathcal{G},\mathcal{G}} \end{bmatrix} \begin{bmatrix} V^\top & 0 \\ 0 & U^\top \end{bmatrix} = \begin{bmatrix} V H|_{\mathcal{F},\mathcal{F}} V^\top & V H|_{\mathcal{F},\mathcal{G}} U^\top \\ U(U^\top D V) V^\top & U H|_{\mathcal{G},\mathcal{G}} U^\top \end{bmatrix} = \begin{bmatrix} * & * \\ D & * \end{bmatrix}.$$

To summarize, after transforming $\mathcal{F}$ by V and $\mathcal{G}$ by U, we have $H = \begin{bmatrix} * & * \\ D & * \end{bmatrix} \in \mathrm{Stab}_O(\mathcal{F} \oplus \mathcal{G})$ for nonzero diagonal D. We consider two cases for D: either $D \in \mathrm{span}(I)$ or $D \notin \mathrm{span}(I)$. First, suppose $D \in \mathrm{span}(I)$, so $D = cI$ for $c \neq 0$. Let $F \rightsquigarrow G$ be nonzero n-ary signatures with $n \geq 2$ (note that $F = 0 \iff G = 0$ because $\|F\|^2 = \langle F, F \rangle = \langle G, G \rangle = \|G\|^2$). By (4.4.2), we have

$$H^{\otimes n-1} (F \oplus G)^{n-1,1} = (F \oplus G)^{n-1,1} H. \tag{8.3.5}$$

Now, by Proposition 8.2.1 with $K := F \oplus G$ and (8.2.1), we can write (8.3.5) as the block matrix equation

$$\begin{bmatrix} * & * & \dots & * \\ \vdots & \vdots & \iddots & \vdots \\ * & * & \dots & * \\ D^{\otimes n-1} & * & \dots & * \end{bmatrix} \begin{bmatrix} F^{n-1,1} & 0 \\ 0 & 0 \\ \vdots & \vdots \\ 0 & 0 \\ 0 & G^{n-1,1} \end{bmatrix} = \begin{bmatrix} F^{n-1,1} & 0 \\ 0 & 0 \\ \vdots & \vdots \\ 0 & 0 \\ 0 & G^{n-1,1} \end{bmatrix} \begin{bmatrix} * & * \\ D & * \end{bmatrix}. \tag{8.3.6}$$

The bottom-left block of (8.3.6) is $D^{\otimes n-1} F^{n-1,1} = G^{n-1,1} D$; using $D = cI$, this is equivalent to

$$c^{n-2} F = G. \tag{8.3.7}$$

Then

$$\|F\|^2 = \langle F, F \rangle = \langle G, G \rangle = c^{2(n-2)} \|F\|^2. \tag{8.3.8}$$

As $\mathcal{F}$ and $\mathcal{G}$ are quantum-gadget-closed, there are some $\mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}$ with arity $n \geq 3$, so (8.3.8) gives $c = \pm 1$. Now applying (8.3.7) to any n-ary pair $F \leftrightsquigarrow G$ with $n \geq 2$ gives $c^n F = c^{n-2} F = G$, so $(cI) \cdot F = G$, with $cI \in O_q$. For unary ($n = 1$) F and G, since $\mathcal{F}$ and $\mathcal{G}$ are quantum-gadget-closed, they contain the ternary signatures $F^{\otimes 3}$ and $G^{\otimes 3}$, respectively. So, by the previous reasoning, $(cI)F^{\otimes 3} = G^{\otimes 3}$, or equivalently $(cF)^{\otimes 3} = G^{\otimes 3}$, which implies $cF = G$, as F and G are real-valued. Combining the non-unary and unary cases gives $(cI)\,\mathcal{F} = \mathcal{G}$.

Otherwise, $D \notin \text{span}(I)$. We will show $\mathcal{F} \cup \{D\}$ and $\mathcal{G} \cup \{D\}$ are Holant-indistinguishable, then apply Lemma 8.3.2. The proof is illustrated in Figure 8.2. Consider a $\mathcal{F} \cup \{D\}$-grid Ω with at

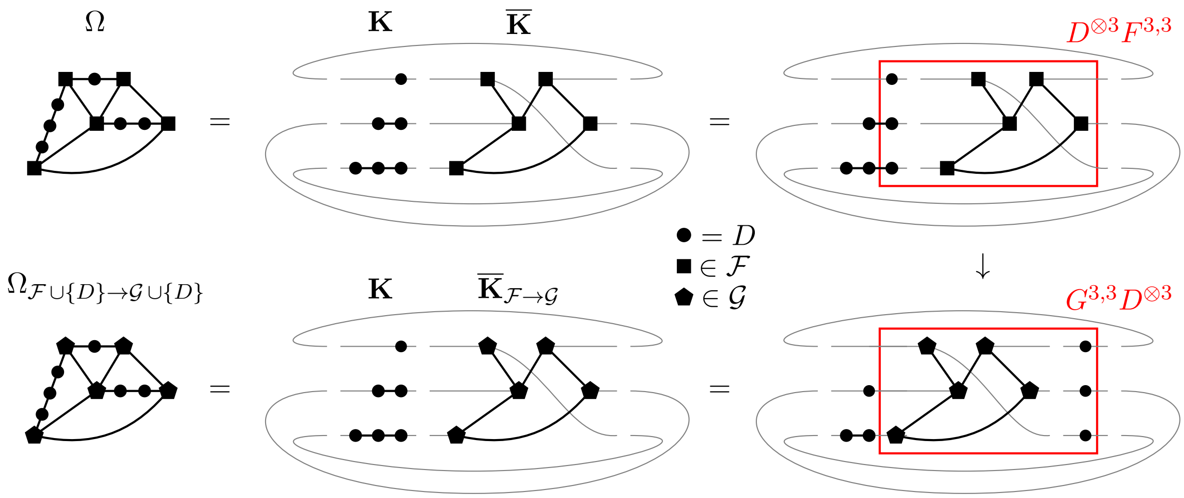

Figure 8.2: The Holant-value-preserving transformation from Ω to $\Omega_{\mathcal{F}\cup\{D\}\to\mathcal{G}\cup\{D\}}$ in the $D \notin \text{span}(I)$ case. The transition from the bottom right grid to the bottom center grid wraps the three D vertices on the right around to the left along their respective wires.

least one vertex assigned D (if Ω has no such vertex then we are done, as $\mathcal{F}$ and $\mathcal{G}$ are Holant-indistinguishable). Let $\mathbf{K} \subset \Omega$ be the subgadget induced by all vertices assigned D. Any connected component of $\mathbf{K}$ is either a cycle or a binary path gadget with signature D^m for some m. The multiplicative factors from corresponding D-cycles in Ω and $\Omega_{\mathcal{F}\cup\{D\}\to\mathcal{G}\cup\{D\}}$ cancel, so, disregarding its cycle components, $\mathbf{K}$ consists of p disconnected path gadgets for some p. By rearranging the dangling edges of $\mathbf{K}$ and $\overline{\mathbf{K}}$, we may assume $\mathbf{K} \in \mathfrak{G}_{\{D\}}(p,p)$ with $M(\mathbf{K}) = \bigotimes_{i=1}^{p} D^{m_i}$ for $m_1, \ldots, m_p \geq 1$, and furthermore that $\overline{\mathbf{K}} \in \mathfrak{G}_{\mathcal{F}}(p,p)$, and that connecting the ith left input and ith right input of $\mathbf{K} \circ \overline{\mathbf{K}}$, for $i \in [p]$, reconstructs Ω (see Figure 8.2). Since $\overline{\mathbf{K}}$ is an $\mathcal{F}$-gadget and $\mathcal{F}$

is quantum-gadget-closed, $\overline{\mathbf{K}}$ has signature F for some $F \in \mathcal{F}$. Then

$$\operatorname{Holant}_\Omega = \operatorname{tr}\left(M(\mathbf{K})M(\overline{\mathbf{K}})\right) = \operatorname{tr}\left(\left(\bigotimes_{i=1}^{p} D^{m_i}\right) F^{p,p}\right). \tag{8.3.9}$$

With $G \rightsquigarrow F$, we similarly have

$$\operatorname{Holant}_{\Omega_{\mathcal{F}\cup\{D\}\to\mathcal{G}\cup\{D\}}} = \operatorname{tr}\left(\left(\bigotimes_{i=1}^{p} D^{m_i}\right) G^{p,p}\right). \tag{8.3.10}$$

As in (8.3.5), (4.4.2) gives

$$H^{\otimes p}(F \oplus G)^{p,p} = (F \oplus G)^{p,p} H^{\otimes p},$$

which has block form

$$\begin{bmatrix} * & * & \dots & * \\ \vdots & \vdots & \iddots & \vdots \\ * & * & \dots & * \\ D^{\otimes p} & * & \dots & * \end{bmatrix} \begin{bmatrix} F^{p,p} & 0 & \dots & 0 & 0 \\ 0 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & \dots & 0 & G^{p,p} \end{bmatrix} = \begin{bmatrix} F^{p,p} & 0 & \dots & 0 & 0 \\ 0 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & \dots & 0 & G^{p,p} \end{bmatrix} \begin{bmatrix} * & * & \dots & * \\ \vdots & \vdots & \iddots & \vdots \\ * & * & \dots & * \\ D^{\otimes p} & * & \dots & * \end{bmatrix}.$$

The bottom left block of this equation is

$$D^{\otimes p} F^{p,p} = G^{p,p} D^{\otimes p}. \tag{8.3.11}$$

Now (8.3.9), (8.3.11), and (8.3.10) give

$$\begin{aligned} \operatorname{Holant}_\Omega &= \operatorname{tr}\left(\left(\bigotimes_{i=1}^{p} D^{m_i}\right) F^{p,p}\right) \\ &= \operatorname{tr}\left(\left(\bigotimes_{i=1}^{p} D^{m_i-1}\right) D^{\otimes p} F^{p,p}\right) \\ &= \operatorname{tr}\left(\left(\bigotimes_{i=1}^{p} D^{m_i-1}\right) G^{p,p} D^{\otimes p}\right) \\ &= \operatorname{tr}\left(D^{\otimes p}\left(\bigotimes_{i=1}^{p} D^{m_i-1}\right) G^{p,p}\right) \\ &= \operatorname{tr}\left(\left(\bigotimes_{i=1}^{p} D^{m_i}\right) G^{p,p}\right) \\ &= \operatorname{Holant}_{\Omega_{\mathcal{F}\cup\{D\}\to\mathcal{G}\cup\{D\}}}. \end{aligned}$$

Thus $\mathcal{F}\cup\{D\}$ and $\mathcal{G}\cup\{D\}$ are Holant-indistinguishable. This fact and the induction hypothesis satisfy the hypotheses of Lemma 8.3.2, which gives a real orthogonal transformation between $\mathcal{F}\cup\{D\}$ and $\mathcal{G}\cup\{D\}$, hence between $\mathcal{F}$ and $\mathcal{G}$. □

8.4 Indistinguishability consequences of the orthogonal converse

In this section, we exploit the expressiveness of the Holant framework to show that Theorem 8.1.1 encompasses a variety of existing results, and derive a few novel consequences. First, recall from (2.3.2) and Proposition 5.1.3 that $\#\mathrm{CSP}(\mathcal{F}) \equiv \mathrm{Holant}(\mathcal{F} \cup \mathcal{EQ})$, and that an orthogonal matrix H satisfies $H\,\mathcal{EQ} = \mathcal{EQ}$ if and only if H is a permutation matrix. Therefore the real-valued case of the #CSP-isomorphism indistinguishability result Theorem 3.2.1.

Let $\mathcal{EQ}_2 \subset \mathcal{EQ}$ be the set of equality signatures of even arity. Schrijver [Sch08b] shows[1] that H satisfies $H\,\mathcal{EQ}_2 = \mathcal{EQ}_2$ if and only if H is a signed permutation matrix (a matrix with entries in $\{0, \pm 1\}$ and exactly one nonzero entry in each row and column). As in the argument leading to (2.3.2), $\mathrm{Holant}(\mathcal{F} \cup \mathcal{EQ}_2)$ is equivalent to $\mathrm{Holant}(\mathcal{F} \mid \mathcal{EQ}_2)$ (critically, if $=_a \in \mathcal{EQ}_2$ and $=_b \in \mathcal{EQ}_2$, then $=_{a+b-2} \in \mathcal{EQ}_2$). Then, defining $\#\mathrm{CSP}^2(\mathcal{F}) := \mathrm{Holant}(\mathcal{F} \mid \mathcal{EQ}_2)$ as $\#\mathrm{CSP}(\mathcal{F})$ restricted to instances in which every variable appears an even number of times [Cai+15; HL16], we have

Corollary 8.4.1. *Let $\mathcal{F}$ and $\mathcal{G}$ be sets of real-valued signatures. Then there is a signed permutation matrix P satisfying $\mathcal{G} = P\,\mathcal{F}$ if and only if $\mathcal{F}$ and $\mathcal{G}$ are $\#\mathrm{CSP}^2$-indistinguishable.*

Let $\mathcal{F}$ and $\mathcal{G}$ be sets of binary signatures, thought of as matrices. Any connected $\mathcal{F}$-grid Ω is a cycle. Breaking an edge of the cycle, we obtain a binary path gadget with signature matrix $\prod_{i=1}^{c} F_i$, where, depending on its orientation, each $F_i \in \mathcal{F}$ or $F_i^\top \in \mathcal{F}$. Connecting the path's two dangling edges, we reform Ω, which thus has Holant value $\mathrm{tr}\left(\prod_{i=1}^{c} F_i\right)$. Let $\Gamma_\mathcal{F}$ be the set of all finite products of matrices in $\mathcal{F}$ and $\mathcal{F}^\top := \{F^\top \mid F \in \mathcal{F}\}$. Define $\Gamma_\mathcal{G}$ similarly and, for a word $w \in \Gamma_\mathcal{F}$, construct $w_{\mathcal{F}\to\mathcal{G}} \in \Gamma_\mathcal{G}$ by replacing every character F or $F^\top$ in w by the corresponding G or $G^\top$, respectively. For orthogonal H, we have $H\,\mathcal{F} = \mathcal{G} \iff HF^{1,1} = G^{1,1}H$ for every $F \leftrightsquigarrow G$ (by (4.4.2)), so, in this setting, Theorem 8.1.1 is equivalent to the following real-valued case of a classical theorem from representation theory, due to Specht [Spe40] and Wiegmann [Wie61]. Grohe, Rattan, and Seppelt [GRS25] also give a combinatorial proof.

Corollary 8.4.2. *Let $\mathcal{F}, \mathcal{G} \subset \mathbb{R}^{q\times q}$. Then there is an $H \in O(q)$ such that $HF = GH$ for every $\mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}$ if and only if $\mathrm{tr}(w) = \mathrm{tr}(w_{\mathcal{F}\to\mathcal{G}})$ for every $w \in \Gamma_\mathcal{F}$.*

[1] The First Fundamental Theorem for $S_q^\pm \subset O(q)$ (the group of signed permutation matrices) states that $T(\mathbb{R}^q)^{S_q^\pm} = [\mathcal{EQ}_2]$. It follows as in the proof of Corollary 4.5.1 that $\mathrm{Stab}_O(\mathcal{EQ}_2) = S_q^\pm$. The fact that $\mathrm{Stab}(\mathcal{EQ}) = S_q \subset O(q)$ (the group of permutation matrices) similarly follows from the First Fundamental Theorem for S_q, which states that $T(\mathbb{R}^q)^{S_q} = [\mathcal{EQ}]$.

Suppose $\mathcal{F} = \{A_X\}$ and $\mathcal{G} = \{A_Y\}$ for graphs X and Y. Transform an A_X-grid Ω to a $(A_X \mid \mathcal{EQ})$-grid Ω' by inserting a dummy degree-2 vertex assigned $I = (=_2) \in \mathcal{EQ}$ between every consecutive pair of vertices in the cycle. Recall from Section 2.3 that $\text{Holant}_{\Omega'}(A_X \mid \mathcal{EQ})$ counts the number of homomorphisms from graph K to X, where K is the graph obtained from Ω' by ignoring the vertices assigned A_X. Here K is a cycle, so we have the following well-known result, an alternate formulation of this case of Corollary 8.4.2.

Corollary 8.4.3. *Let X and Y be graphs. Then there is an orthogonal matrix H satisfying $HA_X = A_Y H$ if and only if X and Y are homomorphism-indistinguishable over the class of all cycles.*

A matrix H is *pseudo-stochastic* if all of its rows and columns sum to 1. Dell, Grohe, and Rattan [DGR18] proved that graphs X and Y admit the same number of homomorphisms from all paths if and only if there is a pseudo-stochastic matrix H such that $HA_X = A_Y H$. Using Theorem 8.1.1, we combine this result with Corollary 8.4.3, which also reproduces a combinatorial explanation for the connection between pseudo-stochastic matrices and homomorphisms from paths [GRS25].

Corollary 8.4.4. *Let X and Y be graphs. Then there is a pseudo-stochastic orthogonal matrix H satisfying $HA_X = A_Y H$ if and only if X and Y are homomorphism-indistinguishable over the class of all paths and cycles.*

Proof. Consider $\text{Holant}(A_X \cup \{=_1\})$. Any $A_X \cup \{=_1\}$-grid is a disjoint union of cycles composed of signatures assigned A_X and paths with degree-2 internal vertices assigned A_X and degree-1 endpoints assigned $=_1 \in \mathcal{EQ}$. As discussed before Corollary 8.4.3, every cycle A_X-grid Ω has the same Holant value as $\Omega_{A_X \to A_Y}$ if and only if X and Y admit the same number of homomorphisms from every cycle. Similarly inserting dummy vertices assigned $=_2$ between every pair of A_X vertices in a path component, we produce a $(A_X \mid \mathcal{EQ})$-grid whose Holant value equals the number of homomorphisms to X from the underlying path. Thus X and Y admit the same number of homomorphisms from all cycles and all paths if and only if $A_X \cup \{=_1\}$ and $A_Y \cup \{=_1\}$ are Holant-indistinguishable. By Theorem 8.1.1, this is equivalent to the existence of an orthogonal H satisfying $HA_X = A_Y H$ and $H(=_1) = (=_1)$. The vector form of $=_1$ is the all-ones vector, so $H(=_1) = (=_1)$ if and only if the rows of H sum to 1 and (since $H(=_1) = (=_1) \iff H^\top(=_1) = (=_1)$) the columns of H sum to 1. $\square$

8.5 Odeco signature sets

In this section, we derive a consequence of Theorem 8.1.1 that is not a counting indistinguishability theorem, but a combinatorial characterization of signatures that are simultaneously 'diagonalizable'.

Definition 8.5.1 ($\mathcal{GEQ}$, odeco)**.** Define the set of *general equality* (or *weighted equality*) signatures on domain $[q]$ as $\mathcal{GEQ} = \{=_n^{\mathbf{a}} | \; n \geq 0, \mathbf{a} \in \mathbb{R}^{[q]}\}$, where $=_n^{\mathbf{a}}$ is the symmetric n-ary signature defined by, for $\mathbf{x} \in [q]^n$,

$$(=_n^{\mathbf{a}})(\mathbf{x}) = \begin{cases} a_i & x_1 = \ldots = x_n = i \\ 0 & \text{otherwise.} \end{cases}$$

A set $\mathcal{F}$ of symmetric signatures is *orthogonally decomposable*, or *odeco*, if there is an $H \in O(q)$ such that $H\,\mathcal{F} \subset \mathcal{GEQ}$.

The term "odeco" was coined by Robeva [Rob16] to refer to individual symmetric tensors (signatures) orthogonally transformable to a general equality. A binary $\mathcal{GEQ}$ signature has a diagonal signature matrix, so the spectral theorem states that every (real) symmetric binary signature F is odeco (recall from (4.4.2) that the matrix form of $(H \cdot F)$ is $HF^{1,1}H^\top$). Any nonzero edge assignment for a connected $\mathcal{GEQ}$-gadget $\mathbf{K}$ must assign all edges, including dangling edges, the same domain element, so, if $\mathbf{K}$ has arity n and is composed of vertices assigned signatures with weights $\mathbf{a}^1, \ldots, \mathbf{a}^p$, then $\mathbf{K}$ has signature $=_n^{\mathbf{a}^1 \bullet \ldots \bullet \mathbf{a}^p} \in \mathcal{GEQ}$, where $\bullet$ denotes entrywise product. In particular, if $\mathbf{K}$ is a $\mathcal{GEQ}$-grid Ω, then $\text{Holant}_\Omega = \sum_{i=1}^q (\mathbf{a}^1 \bullet \ldots \bullet \mathbf{a}^p)_i$. Thus, if $\mathcal{F}$ is odeco, then the computational problem Holant$(\mathcal{F})$ is polynomial-time tractable via orthogonal Holographic transformation.

Definition 8.5.2 ($*$)**.** For symmetric signatures $F_1, F_2 \in \mathcal{F}$ of arity n_1 and n_2, respectively, construct the $(n_1 + n_2 - 2)$-ary signature $F_1 * F_2 \in [\mathcal{F}]$ from $F_1 \otimes F_2$ by contracting an input of F_1 and an input of F_2. See Figure 8.1.

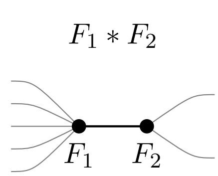

Figure 8.1: Illustrating (the gadget with signature) $F_1 * F_2$ for 6-ary F_1 and 3-ary F_2.

Since F_1 and F_2 are symmetric, $F_1 * F_2$ doesn't depend on which inputs we connect. For $\mathbf{x} \in [q]^{n_1-1}$ and $\mathbf{y} \in [q]^{n_2-1}$, we have (with vectors viewed as input lists)

$$(F_1 * F_2)(\mathbf{x}, \mathbf{y}) = \sum_{z \in [q]} F_1(\mathbf{x}, z) F_2(\mathbf{y}, z).$$

Proposition 8.5.1. *For any $H \in O(q)$, we have $HF \in \mathcal{GEQ}$ if and only if $H(F * F) \in \mathcal{GEQ}$.*

Proof. ($\Longrightarrow$): If $HF = E$ for $H \in O(q)$ and $E \in \mathcal{GEQ}$, then, by Corollary 4.4.1, $H(F * F) = (HF) * (HF) = E * E \in \mathcal{GEQ}$.

($\Longleftarrow$): Let $F \in \mathbb{R}^{[q]^n}$. Every unary signature is in $\mathcal{GEQ}$, so if $n = 1$ then we are done. If $n = 2$ then $F * F = F^2$ (a matrix product) and F is a real symmetric matrix, so if H diagonalizes F^2 then H diagonalizes F. Now assume $n \geq 3$. Let $\mathcal{GEQ} \ni E = H(F * F) = (HF) * (HF)$ for $H \in O(q)$. Suppose toward contradiction that $HF \notin \mathcal{GEQ}$, so there is a $\mathbf{x} \in [q]^n$ such that $(HF)(\mathbf{x}) \neq 0$ but $\exists i, j$ such that $x_i \neq x_j$. Assume WLOG that $i, j \neq n$. Construct $\mathbf{x}' \in [q]^{n-1}$ by deleting the nth (last) entry of $\mathbf{x}$. Then

$$E(\mathbf{x}', \mathbf{x}') = ((HF) * (HF))(\mathbf{x}', \mathbf{x}') = \sum_{z \in [q]} (HF)(\mathbf{x}', z)^2 \geq (HF)(\mathbf{x})^2 > 0,$$

contradicting $E \in \mathcal{GEQ}$. Thus $HF \in \mathcal{GEQ}$. $\square$

Theorem 8.5.1. *Let $\mathcal{F}$ be a set of real-valued symmetric signatures. The following are equivalent.*

(i) $\mathcal{F}$ is odeco.

(ii) Every connected $\mathcal{F}$-gadget has a symmetric signature.

(iii) For every $F_1, F_2 \in \mathcal{F}$, the signature $F_1 * F_2$ is symmetric.

Robeva [Rob16] conjectured the equivalence of items (i) and (iii) when $\mathcal{F}$ contains a single signature. Boralevi, Draisma, Horobeţ, and Robeva [Bor+17] confirmed this conjecture using techniques from algebraic geometry. We use Theorem 8.1.1 to give a combinatorial proof, generalized to arbitrary symmetric signature sets.

Remark 8.5.1. If $\mathcal{F}$ is a set of symmetric binary signatures, thought of as matrices, then $F_1 * F_2 = F_1 \circ F_2 = F_1 F_2$ (a matrix product) for $F_1, F_2 \in \mathcal{F}$. In general, symmetric matrices commute if and only if their product is symmetric (as if $F_1 F_2$ is symmetric then $F_1 F_2 = (F_1 F_2)^\top = F_2^\top F_1^\top = F_2 F_1$

and if F_1 and F_2 commute then $(F_1F_2)^\top = F_2^\top F_1^\top = F_2F_1 = F_1F_2$). Therefore Theorem 8.5.1 encompasses the extension of the spectral theorem which states that commuting symmetric real matrices are simultaneously diagonalizable. We use this fact in the proof below.

Proof of Theorem 8.5.1. (i) $\implies$ (ii),(iii): Suppose $H\mathcal{F} \subset \mathcal{GEQ}$ for some $H \in O(q)$. Let $K \in [\mathcal{F}]$ be the signature of a connected $\mathcal{F}$-gadget (e.g. $K = F_1 * F_2$). By Corollary 4.4.1, $H \cdot K = J$, where J is the signature of a connected $\mathcal{GEQ}$-gadget. Then $J \in \mathcal{GEQ}$, so J, and hence $K = H^{-1} \cdot J$, are symmetric.

(ii) $\implies$ (i): First, replace every non-unary odd-arity $F \in \mathcal{F}$ by $F * F$ (every unary signature is in $\mathcal{GEQ}$, so simply remove all unaries from $\mathcal{F}$). This does not change the fact that $\mathcal{F}$ satisfies item (ii), and, by Proposition 8.5.1, does not change whether $\mathcal{F}$ satisfies item (i). Thus we may assume all signatures in $\mathcal{F}$ have even arity. For $F \in \mathcal{F}$, let $\widetilde{F} \in {}_1[\mathcal{F}]_1$ be the matrix of the binary signature constructed by contracting all but one pair of inputs of F (see Figure 8.2).

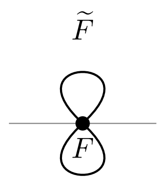

Figure 8.2: Illustrating (the gadget with signature) $\widetilde{F}$ for 6-ary F.

Since F is symmetric, $\widetilde{F}$ doesn't depend on how we pair up F's inputs. Every $\widetilde{F}$, and every composition $\widetilde{F_1} \circ \widetilde{F_2}$ for $F_1, F_2 \in \mathcal{F}$, is the signature of a connected $\mathcal{F}$-gadget, so is symmetric by assumption. Therefore, as in Remark 8.5.1, the matrices $\widetilde{F}$ for $F \in \mathcal{F}$ all commute.

Claim 8.5.1. If $\mathbf{K}$ is a connected binary $\mathcal{F}$-gadget with p vertices, assigned signatures $F_1, \dots, F_p \in \mathcal{F}$, then $M(\mathbf{K}) = \prod_{i=1}^{p} \widetilde{F}_i$.

We prove Claim 8.5.1 by induction on p. For $p = 1$, by the symmetry of F_1, every connected binary F_1-gadget with a single vertex has signature $\widetilde{F_1}$. Now suppose $p \geq 2$. $\mathbf{K}$ contains a vertex v whose removal does not disconnect $\mathbf{K}$ (take e.g. the final vertex visited by breadth first search). Assume WLOG that v is assigned signature F_p. Construct $\mathbf{K}'$ from $\mathbf{K}$ by breaking all but one edge between v and other vertices (see Figure 8.3). Each broken edge becomes two new dangling edges. Since $\mathbf{K}'$ is connected, its signature is symmetric by assumption. The number of dangling edges incident to v is odd, as v has even degree and exactly one edge to another vertex (loops on v do not affect the parity). Since $\mathbf{K}'$ has an even number of dangling edges (two plus twice the number of edges broken), there are an odd number of dangling edges incident to the other vertices of $\mathbf{K}'$. Now create a binary gadget $\mathbf{K}''$ from $\mathbf{K}'$ by arbitrarily pairing up and connecting all but one dangling edge incident to v, and similarly pairing up and connecting all but one dangling edge incident to the other vertices of $\mathbf{K}'$. We may also recover $\mathbf{K}$ from $\mathbf{K}'$ by connecting possibly different pairs of dangling edges (reforming the edges broken to create $\mathbf{K}'$) and, since the signature of $\mathbf{K}'$ is symmetric, the signature of a gadget produced by connecting dangling edges of $\mathbf{K}'$ does not depend on which pairs of dangling edges we connect (although the underlying graphs of the gadgets differ). Therefore $M(\mathbf{K}'') = M(\mathbf{K})$.

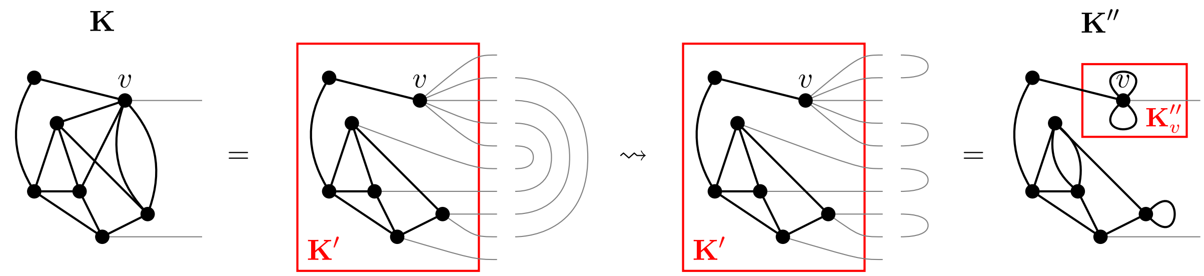

Figure 8.3: Illustrating the proof of Claim 8.5.1 when v has a single incident dangling edge in $\mathbf{K}$. The cases where v has zero or two incident dangling edges are similar.

Let $\mathbf{K}''_v$ be the subgadget of $\mathbf{K}''$ induced by v. $\mathbf{K}''_v$ has two dangling edges, one dangling in $\mathbf{K}''$, and one which, in $\mathbf{K}''$, connects v to a vertex in $\overline{\mathbf{K}''_v}$ (the complement of $\mathbf{K}''_v$), with the remaining edges incident to v paired into loops. Thus $M(\mathbf{K}''_v) = \widetilde{F_p}$. Similarly, $\overline{\mathbf{K}''_v}$ has two dangling edges, one dangling in $\mathbf{K}''$ and one incident to to v in $\mathbf{K}''$, so we recover $\mathbf{K}''$ by reconnecting the edge between $\mathbf{K}''_v$ and $\overline{\mathbf{K}''_v}$. Thus, applying induction to $\overline{\mathbf{K}''_v}$, which has $p-1$ vertices, we obtain

$$M(\mathbf{K}) = M(\mathbf{K}'') = M(\mathbf{K}''_v) \circ M(\overline{\mathbf{K}''_v}) = \widetilde{F_p} \circ \prod_{i=1}^{p-1} \widetilde{F}_i = \prod_{i=1}^{p} \widetilde{F}_i$$

(recall that the $\widetilde{F}_i$ commute). This completes the proof of Claim 8.5.1.

The matrices $\widetilde{F}$ for $F \in \mathcal{F}$ are symmetric and commute, so they are simultaneously diagonalizable: there is an $H \in O(q)$ such that, for every $F \in \mathcal{F}$, $H\widetilde{F}H^\top = (=_2^{\mathbf{a}^F})^{1,1}$ for some $\mathbf{a}^F \in \mathbb{R}^q$. Replace $\mathcal{F}$ with $H\mathcal{F}$ to assume each $\widetilde{F} = (=_2^{\mathbf{a}^F})^{1,1}$ (this does not change whether $\mathcal{F}$ is odeco). Define

$$\mathcal{G} = \{=_{\text{arity}(F)}^{\mathbf{a}^F} \mid F \in \mathcal{F}\} \subset \mathcal{GEQ},$$

a set bijective with $\mathcal{F}$. Let Ω be a $\mathcal{F}$-grid, containing signatures $F_1, \ldots, F_p$. As Ω is not a tree (all of its vertices have even degree), we can break some edge of Ω to produce a connected binary $\mathcal{F}$-gadget $\mathbf{K}$. By Claim 8.5.1,

$$M(\mathbf{K}) = \prod_{i=1}^{p} \widetilde{F}_i = \prod_{i=1}^{p} (=_2^{\mathbf{a}^{F_i}})^{1,1} = \left(=_2^{\mathbf{a}^{F_1} \bullet \ldots \bullet \mathbf{a}^{F_p}}\right)^{1,1} = M(\mathbf{K}_{\mathcal{F}\to\mathcal{G}}).$$

Now, reconnecting the dangling edges of $\mathbf{K}$, we find $\text{Holant}_{\Omega} = \text{Holant}_{\Omega_{\mathcal{F}\to\mathcal{G}}}$. Thus $\mathcal{F}$ and $\mathcal{G}$ are Holant-indistinguishable, so, by Theorem 8.1.1, there is an $H' \in O_q$ such that $H'\mathcal{F} = \mathcal{G}$. Hence $\mathcal{F}$ is odeco.

(iii) $\implies$ (ii): Let K be the n-ary signature of a connected $\mathcal{F}$-gadget $\mathbf{K}$. Every unary signature is trivially symmetric, so assume $n \geq 2$. It suffices to show that, for any fixed partial input $\mathbf{z} \in [q]^{n-2}$ to K and any $x, y \in [q]$, we have $K(x, y, \mathbf{z}) = K(y, x, \mathbf{z})$ (where we assume WLOG that x and y are the first two inputs to K by reordering the dangling edges of $\mathbf{K}$). Let u and w be the vertices of $\mathbf{K}$ incident to the first and second dangling edges of $\mathbf{K}$ (after reordering). If $u = w$ then we are done, as every $F \in \mathcal{F}$ is symmetric. Otherwise, since $\mathbf{K}$ is connected, it contains a path $P = (u = v_0, v_1, \ldots, v_{p-2}, v_{p-1} = w)$ from u to w, where v_i is assigned signature $F_i \in \mathcal{F}$, for $i \in [p]$. Let $E(P) := \{e_0, e_1, \ldots, e_{p-1}, e_p\}$ be the edges of P, including the dangling edges e_0 and e_p incident to u and w, respectively. Then e_i and e_{i+1} are inputs to F_i for all $i \in [p]$. For any fixed assignment $\sigma : E(\mathbf{K}) \setminus E(P) \to [q]$, define the matrix $F_i^\sigma \in \mathbb{R}^{q \times q}$ by $(F_i^\sigma)_{a,b} := F_i(\sigma|_{\delta(v_i)}, a, b)$ (that is, fix the inputs to F_i from edges outside of $E(P)$). Then

$$K(x, y, \mathbf{z}) = \sum_{\substack{\sigma: E(\mathbf{K}) \setminus E(P) \to [q] \\ \sigma(D) = \mathbf{z}}} \left(\prod_{v \in V(\mathbf{K}) \setminus P} F_v(\sigma|_{\delta(v)}) \right) F_P^\sigma(x, y), \tag{8.5.1}$$

where D is the ordered list of the last $n - 2$ dangling edges of $\mathbf{K}$ and

$$\begin{aligned} F_P^\sigma(x, y) &= \sum_{\substack{\phi: E(P) \to [q] \\ \phi(e_0) = x, \phi(e_p) = y}} \prod_{i=0}^{p-1} F_i(\sigma|_{\delta(v_i)}, \phi(e_i), \phi(e_{i+1})) \\ &= \sum_{\substack{\phi: E(P) \to [q] \\ \phi(e_0) = x, \phi(e_p) = y}} \prod_{i=0}^{p-1} (F_i^\sigma)_{\phi(e_i), \phi(e_{i+1})} = \left(\prod_{i=0}^{p-1} F_i^\sigma \right)_{x,y}. \end{aligned}$$

On the RHS of (8.5.1), x and y appear only in $F_P^\sigma(x, y)$. Thus it suffices to show that, for any fixed

σ, $F_P^\sigma(x,y) = F_P^\sigma(y,x)$. For any $i,j \in [p]$ and $a,b \in [q]$,

$$\begin{aligned}(F_i^\sigma F_j^\sigma)_{a,b} &= \sum_{z\in[q]} (F_i^\sigma)_{a,z}(F_j^\sigma)_{z,b} = \sum_{z\in[q]} F_i(\sigma|_{\delta(v_i)}, a, z) F_j(\sigma|_{\delta(v_j)}, b, z)\\ &= (F_i * F_j)(\sigma|_{\delta(v_i)}, a, \sigma|_{\delta(v_j)}, b) = (F_i * F_j)(\sigma|_{\delta(v_i)}, b, \sigma|_{\delta(v_j)}, a)\\ &= \sum_{z\in[q]} F_i(\sigma|_{\delta(v_i)}, b, z) F_j(\sigma|_{\delta(v_j)}, a, z) = \sum_{z\in[q]} (F_i^\sigma)_{b,z}(F_j^\sigma)_{z,a} = (F_i^\sigma F_j^\sigma)_{b,a},\end{aligned}$$

where the fourth equality uses the assumption that $F_i * F_j$ is symmetric. Thus $F_i^\sigma F_j^\sigma$ is symmetric. Both F_i^σ and F_j^σ are symmetric, as F_i and F_j are symmetric, so, as in Remark 8.5.1, F_i^σ and F_j^σ commute. Therefore

$$F_P^\sigma(x,y) = \left(\prod_{i=0}^{p-1} F_i^\sigma\right)_{x,y} = \left(\prod_{i=0}^{p-1} F_i^\sigma\right)^\top_{y,x} = \left(\prod_{i=0}^{p-1} (F_{p-1-i}^\sigma)^\top\right)_{y,x} = \left(\prod_{i=0}^{p-1} F_i^\sigma\right)_{y,x} = F_P^\sigma(y,x). \quad \square$$

Chapter 9

Planar Holant and Quantum Orthogonal Transformation

In this chapter, we prove the converse of the quantum orthogonal Holant theorem (Corollary 6.3.3). This is a generalization of Theorem 7.1.1 from planar #CSP to planar Holant, and a quantum/planar version of Theorem 8.1.1.

Theorem 9.0.1. *The following hold for bijective complex-valued signature sets $\mathcal{F}$ and $\mathcal{G}$.*

- *Suppose $\mathcal{F}$ and $\mathcal{G}$ are conjugate-closed and satisfy $F \leftrightsquigarrow G$ if and only if $\overline{F} \leftrightsquigarrow \overline{G}$. Then $\mathcal{F}$ and $\mathcal{G}$ are $Pl^{\top}$-Holant-indistinguishable if and only if there is a quantum orthogonal matrix U satisfying $U\mathcal{F} = \mathcal{G}$.*
- *$\mathcal{F}$ and $\mathcal{G}$ are $Pl^{\dagger}$-Holant-indistinguishable if and only if there is a quantum orthogonal matrix U satisfying $U\mathcal{F} = \mathcal{G}$.*

The 'easy' ($\Longleftarrow$) direction of Theorem 9.0.1 follows from the quantum orthogonal Holant theorem (Corollary 6.3.3). For the 'hard' ($\Longrightarrow$) direction, just as the lack of orbits of $\mathrm{Stab}_O(\mathcal{F})$ relative to $\mathrm{Aut}(\mathcal{F})$ forced us to apply new techniques in Chapter 8, the lack of orbits of $\mathrm{Stab}_{O^+}(\mathcal{F})$ relative to $\mathrm{Qut}(\mathcal{F})$ forces us to apply new techniques in this chapter. We apply a theorem of Seppelt and Spitzer [SS26] – based on a category-theoretic result of Neshveyev and Tuset [NT14] – originally stated for the special case of graph homomorphism counting but which directly generalizes to the Holant setting.

9.1 Planar weak isomorphism

The following definition generalizes the concept of weak isomorphism introduced in [MR20] from planar homomorphism counting to $\mathrm{Pl}^\dagger$-Holant. A Bi-Holant version will also play a role in Chapter 10. Recall that $\leftrightsquigarrow$, defined for $\mathrm{Pl}^\dagger$-Holant and $\mathrm{Pl}^\top$-Holant in Section 6.2, is a multiset bijection between the tensors of quantum $\mathcal{F}$ and $\mathcal{G}$ gadgets, but is not necessarily well-defined as a linear map between the vector subspaces ${}_\ell([\mathcal{F}]_{\mathrm{Pl}^\dagger})_r$ and ${}_\ell([\mathcal{G}]_{\mathrm{Pl}^\dagger})_r$ of ${}_\ell\mathcal{V}_r$, because it is possible that two different quantum $\mathrm{Pl}^\top$-Holant($\mathcal{F}$) gadgets $\mathbf{K}$ and $\mathbf{J}$ have the same tensor (that is, $M(\mathbf{K}) = M(\mathbf{J})$), but the corresponding $\mathrm{Pl}^\top$-Holant($\mathcal{F}$) gadgets $\mathbf{K}_{\mathcal{F}\to\mathcal{G}}$ and $\mathbf{J}_{\mathcal{F}\to\mathcal{G}}$ do not (or vice-versa).

Definition 9.1.1 ($\mathrm{Pl}^\dagger$-weakly-isomorphic, $\mathrm{Pl}^\top$-weakly-isomorphic)**.** Bijective $\mathcal{F}, \mathcal{G} \subset \mathcal{S}(\mathbb{C}^q)$ are *$Pl^\dagger$-weakly-isomorphic* if they satisfy any of the following equivalent conditions:

(1) For every $\ell, r \geq 0$, $\leftrightsquigarrow$ is an invertible linear map, called a *$Pl^\dagger$-weak-isomorphism*, between ${}_\ell([\mathcal{F}]_{\mathrm{Pl}^\dagger})_r$ and ${}_\ell([\mathcal{G}]_{\mathrm{Pl}^\dagger})_r$.

(2) Every $[\mathcal{F}]_{\mathrm{Pl}^\dagger} \ni F \leftrightsquigarrow G \in [\mathcal{G}]_{\mathrm{Pl}^\dagger}$ satisfy $F = 0 \iff G = 0$.

(3) $M(\mathbf{K}) = M(\mathbf{J}) \iff M(\mathbf{K}_{\mathcal{F}\to\mathcal{G}}) = M(\mathbf{J}_{\mathcal{F}\to\mathcal{G}})$ for every pair of quantum $\mathrm{Pl}^\dagger$-Holant($\mathcal{F}$) gadgets $\mathbf{K}$ and $\mathbf{J}$.

Define $\mathrm{Pl}^\top$-weakly-isomorphic analogously.

Proof. (1) $\implies$ (2) is a standard property of invertible linear maps. (2) $\implies$ (3) because $\mathbf{K} - \mathbf{J}$ corresponds to the quantum gadget $\mathbf{K}_{\mathcal{F}\to\mathcal{G}} - \mathbf{J}_{\mathcal{F}\to\mathcal{G}}$, so

$$[\mathcal{F}]_{\mathrm{Pl}^\dagger} \ni M(\mathbf{K}) - M(\mathbf{J}) = M(\mathbf{K} - \mathbf{J}) \leftrightsquigarrow M(\mathbf{K}_{\mathcal{F}\to\mathcal{G}} - \mathbf{J}_{\mathcal{F}\to\mathcal{G}}) = M(\mathbf{K}_{\mathcal{F}\to\mathcal{G}}) - M(\mathbf{J}_{\mathcal{F}\to\mathcal{G}}) \in [\mathcal{G}]_{\mathrm{Pl}^\dagger}.$$

(3) $\implies$ (2) by putting $F = M(\mathbf{K})$ and letting $\mathbf{J}$ be any quantum gadget with global coefficient 0. Finally, (3) and (2) imply that $\leftrightsquigarrow$ is a well-defined linear map and is invertible, respectively. □

Lemma 9.1.1. *The following hold for $\mathcal{F}, \mathcal{G} \subset \mathcal{S}(\mathbb{C}^q)$.*

- *If $\mathcal{F}$ and $\mathcal{G}$ are conjugate-closed and $Pl^\top$-Holant-indistinguishable, then $\mathcal{F}$ and $\mathcal{G}$ are $Pl^\top$-weakly-isomorphic.*
- *If $\mathcal{F}$ and $\mathcal{G}$ are $Pl^\dagger$-Holant-indistinguishable, then $\mathcal{F}$ and $\mathcal{G}$ are $Pl^\dagger$-weakly-isomorphic.*

Proof. Recall from Theorem 6.3.1 that $[\mathcal{F}]_{\text{Pl}^\dagger} = \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\dagger}$ and, under the conjugate-closed assumption, $[\mathcal{F}]_{\text{Pl}^\top} = \langle \mathcal{F}, \supset \rangle_{+,\circ,\otimes,\dagger}$ as well. Given this, we now proceed with the proof for $\text{Pl}^\dagger$-Holant, and the proof for $\text{Pl}^\top$-Holant is similar. Assume $\mathcal{F}$ and $\mathcal{G}$ are not $\text{Pl}^\dagger$-weakly-isomorphic, so WLOG there is a $[\mathcal{F}]_{\text{Pl}^\dagger} \ni K \leftrightsquigarrow 0 \in [\mathcal{G}]_{\text{Pl}^\dagger}$ with $K \neq 0$. By Lemma 4.2.1, we may assume $K \in {}_n([\mathcal{F}]_{\text{Pl}^\dagger})_0$ for some n. Then $K^\dagger \in {}_0([\mathcal{F}]_{\text{Pl}^\dagger})_n$, so

$$[\mathcal{F}]_{\text{Pl}^\dagger} \ni K^\dagger \circ K = \langle K^\dagger, K \rangle \leftrightsquigarrow \langle K', 0 \rangle \in [\mathcal{G}]_{\text{Pl}^\dagger},$$

where $[\mathcal{F}]_{\text{Pl}^\dagger} \ni K^\dagger \leftrightsquigarrow K' \in [\mathcal{G}]_{\text{Pl}^\dagger}$. But $\langle K^\dagger, K \rangle = \|K\| \neq 0 = \langle K', 0 \rangle$ so at least one pair of corresponding $\text{Pl}^\dagger$-Holant($\mathcal{F}$) and $\text{Pl}^\dagger$-Holant($\mathcal{G}$) signature grids in the quantum signature grids with values $\langle K^\dagger, K \rangle$ and $\langle K', 0 \rangle$ must have different $\text{Pl}^\dagger$-Holant values (recall Proposition 4.1.1). Therefore $\mathcal{F}$ and $\mathcal{G}$ are not $\text{Pl}^\dagger$-Holant-indistinguishable. □

Lemma 9.1.1 is related to the fact that $[\mathcal{F}]_{\text{Pl}^\dagger}$ is *quantum-nonvanishing* – see Proposition 10.3.2 and Corollary 10.5.1 below. An analogous version of Lemma 9.1.1 does not hold for $\text{Pl}^\top$-weak-isomorphism for non-conjugate-closed sets. For example, recall from Section 8.1 that the unary domain-2 signature $[1, i]$ is Holant-indistinguishable from 0, and hence $\text{Pl}^\top$-Holant-indistinguishable from 0, as all signature grids are planar.

9.2 The quantum orthogonal converse

The converse of the quantum orthogonal Holant theorem is due to a minor generalization of a result of Seppelt and Spitzer [SS26]. Before presenting this result, we first introduce a few new concepts. An *easy quantum group* $\mathcal{Q}$ [BS09] is a compact matrix quantum group satisfying $C_\mathcal{Q} \subset \langle \mathcal{EQ} \rangle_{+,\circ,\otimes,\dagger}$. Examples include the quantum orthogonal group O_q^+, with $C_{O_q^+} = \langle \supset \rangle_{+,\circ,\otimes,\dagger}$, and the quantum symmetric group S_q^+, with $C_{S_q^+} = \langle \supset, =_3 \rangle_{+,\circ,\otimes,\dagger}$. An *orthogonal easy quantum group* is an easy quantum group that is a quantum subgroup of O_q^+, in that its intertwiner space contains $\supset$. Easy quantum groups have an easy representation theory defined by partitions, which are a combinatorial interpretation of $\mathcal{EQ}$ gadgets (the connected components of an $\mathcal{EQ}$ gadget partitions the set of its inputs into subsets that must take the same value). Given a graph X and an orthogonal easy quantum group $\mathcal{Q}$ with $C_\mathcal{Q} = \langle M_1, \ldots, M_k \rangle_{+,\circ,\otimes,\dagger} \subset \langle \mathcal{EQ} \rangle_{+,\circ,\otimes,\dagger}$, Seppelt and Spitzer [SS26, Definition 37] define $\mathcal{Q}(X)$ to be the unique quantum subgroup of O_q^+ with $C_{\mathcal{Q}(X)} =$

$\langle A_X, M_1, \ldots, M_k\rangle_{+,\circ,\otimes,\dagger}$ – whose existence is guaranteed by Tannaka-Krein duality (Theorem 6.4.1). For example, if $\mathcal{Q}$ is generated by $M_1 = \supset$, then $\mathcal{Q}(X) = O_q^+(X) = \mathrm{Stab}_{O_q^+}(X)$, and if $\mathcal{Q}$ is generated by $M_1 = \supset$ and $M_2 = (=_3)$, then $\mathcal{Q}(X) = S_q^+(X) = \mathrm{Qut}(X)$. We generalize this definition to arbitrary signature sets $\mathcal{F} \subset \mathcal{S}(\mathbb{C}^q)$ in exactly the same way that we generalized $\mathrm{Qut}(X)$ to $\mathrm{Qut}(\mathcal{F})$.

Definition 9.2.1 ($\mathcal{Q}(\mathcal{F})$)**.** For an orthogonal easy quantum group $\mathcal{Q}$ with intertwiner space $C_{\mathcal{Q}} = \langle M_1, \ldots, M_k\rangle_{+,\circ,\otimes,\dagger} \subset \langle \mathcal{EQ}\rangle_{+,\circ,\otimes,\dagger}$, define $\mathcal{Q}(\mathcal{F})$ to be the unique quantum subgroup of O_q^+ with $C_{\mathcal{Q}(\mathcal{F})} = \langle \mathcal{F}, M_1, \ldots, M_k\rangle_{+,\circ,\otimes,\dagger}$.

Seppelt and Spitzer [SS26, Definition 40], following Mančinska and Roberson [MR20], define, for graphs X and Y, a weak isomorphism between $\mathcal{Q}(X)$ and $\mathcal{Q}(Y)$ to be a map $\Phi : C_{\mathcal{Q}(X)} \to C_{\mathcal{Q}(Y)}$ satisfying the following properties:

1. $\Phi(A_X) = A_Y$ and $\Phi(M_i) = M_i$ for all $i \in [k]$,
2. $\Phi(T \circ T') = \Phi(T) \circ \Phi(T')$, $\Phi(T \otimes T') = \Phi(T) \otimes \Phi(T')$, and $\Phi(T^\dagger) = \Phi(T)^\dagger$ for all $T, T' \in C_{\mathcal{Q}(X)}$,
3. The restriction of Φ to ${}_\ell(C_{\mathcal{Q}(X)})_r$ is a bijective linear map to ${}_\ell(C_{\mathcal{Q}(Y)})_r$ for every ℓ, r.

Theorem 9.2.1 ([SS26, Corollary 63])**.** *Let $\mathcal{Q}$ be an orthogonal easy quantum group and let X and Y be graphs. If there is a weak isomorphism Φ between $\mathcal{Q}(X)$ and $\mathcal{Q}(Y)$, then there exists a quantum orthogonal matrix U with $UA_X = A_Y U$, such that $U^{\otimes \ell} M = M U^{\otimes r}$ for all $M \in {}_\ell(C_{\mathcal{Q}})_r$.*

The proof of Theorem 9.2.1 is not combinatorial but category-theoretic, applying a result of Neshveyev and Tuset on fiber functors between C^*-tensor categories [NT14]. In particular, it treats A_X and A_Y as arbitrary corresponding elements of the intertwiner spaces of $C_{\mathcal{Q}(X)}$ and $C_{\mathcal{Q}(Y)}$, and does not use the fact that they are graph adjacency matrices. Therefore, as observed in [SS26, Section 7.1], the same proof goes through if we add any bijective $\mathcal{F}, \mathcal{G} \subset \mathcal{S}(\mathbb{C}^q)$ to the intertwiner space of $\mathcal{Q}$ in place of A_X and A_Y. Naturally generalizing the definition for $\mathcal{Q}(X)$ and $\mathcal{Q}(Y)$, a weak isomorphism between $\mathcal{Q}(\mathcal{F})$ and $\mathcal{Q}(\mathcal{G})$ is a map $\Phi : C_{\mathcal{Q}(\mathcal{F})} \to C_{\mathcal{Q}(\mathcal{G})}$ satisfying the following properties:

1. $\Phi(F) = G$ for every $\mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}$ and $\Phi(M_i) = M_i$ for all $i \in [k]$,

2. $\Phi(T \circ T') = \Phi(T) \circ \Phi(T')$, $\Phi(T \otimes T') = \Phi(T) \otimes \Phi(T')$, and $\Phi(T^\dagger) = \Phi(T)^\dagger$ for all $T, T' \in C_{\mathcal{Q}(\mathcal{F})}$,

3. The restriction of Φ to ${}_\ell(C_{\mathcal{Q}(\mathcal{F})})_r$ is a bijective linear map to ${}_\ell(C_{\mathcal{Q}(\mathcal{G})})_r$ for every ℓ, r.

Theorem 9.2.2 ([SS26]). *Let $\mathcal{Q}$ be an orthogonal easy quantum group and let $\mathcal{F}, \mathcal{G} \subset \mathcal{S}(\mathbb{C}^q)$ be bijective. If there is a weak isomorphism Φ between $\mathcal{Q}(\mathcal{F})$ and $\mathcal{Q}(\mathcal{G})$, then there exists a quantum orthogonal matrix U such that $U^{\otimes \ell} F^{\ell,r} = G^{\ell,r} U^{\otimes r}$ for every ${}_{\ell+r}\mathcal{F} \ni F \leftrightsquigarrow G \in {}_{\ell+r}\mathcal{G}$ and $U^{\otimes \ell} M = M U^{\otimes r}$ for all $M \in {}_\ell(C_{\mathcal{Q}})_r$.*

Corollary 9.2.1. *The following hold for bijective $\mathcal{F}, \mathcal{G} \subset \mathcal{S}(\mathbb{C}^q)$.*

- *Suppose $\mathcal{F}$ and $\mathcal{G}$ are conjugate-closed and satisfy $F \leftrightsquigarrow G$ iff $\overline{F} \leftrightsquigarrow \overline{G}$. If $\mathcal{F}$ and $\mathcal{G}$ are $\mathrm{Pl}^\top$-Holant-indistinguishable, then there is a quantum orthogonal matrix U satisfying $U\mathcal{F} = \mathcal{G}$.*

- *If $\mathcal{F}$ and $\mathcal{G}$ are $\mathrm{Pl}^\dagger$-Holant-indistinguishable, then there is a quantum orthogonal matrix U satisfying $U\mathcal{F} = \mathcal{G}$.*

Proof. Let $\mathcal{Q} = O_q^+$. By Theorem 6.4.2, $C_{\mathcal{Q}(\mathcal{F})} = C_{\mathrm{Stab}_{O^+}(\mathcal{F})} = [\mathcal{F}]_{\mathrm{Pl}^\dagger}$, and similarly $C_{\mathcal{Q}(\mathcal{G})} = [\mathcal{G}]_{\mathrm{Pl}^\dagger}$. If $\mathcal{F}$ and $\mathcal{G}$ are conjugate-closed, then Theorem 6.4.2 also gives $C_{\mathcal{Q}(\mathcal{F})} = [\mathcal{F}]_{\mathrm{Pl}^\top}$ and $C_{\mathcal{Q}(\mathcal{G})} = [\mathcal{G}]_{\mathrm{Pl}^\top}$. Assuming previous results, the proofs of both points are now the same; we give the proof for $\mathrm{Pl}^\dagger$-Holant. We claim that $\Phi := \leftrightsquigarrow$ is a weak isomorphism between $\mathcal{Q}(\mathcal{F})$ and $\mathcal{Q}(\mathcal{G})$. By Lemma 9.1.1, the map $\leftrightsquigarrow$ is a $\mathrm{Pl}^\dagger$-weak-isomorphism between $\mathcal{F}$ and $\mathcal{G}$, so is a bijective linear map between ${}_\ell([\mathcal{F}]_{\mathrm{Pl}^\dagger})_r = {}_\ell(C_{\mathcal{Q}(\mathcal{F})})_r$ and ${}_\ell([\mathcal{G}]_{\mathrm{Pl}^\dagger})_r = {}_\ell(C_{\mathcal{Q}(\mathcal{G})})_r$. Thus $\leftrightsquigarrow$ satisfies item 3 of the definition of a weak isomorphism between $\mathcal{Q}(\mathcal{F})$ and $\mathcal{Q}(\mathcal{G})$. By definition, $\leftrightsquigarrow$ maps every generator $F \in \mathcal{F}$ to the corresponding $G \in \mathcal{G}$, and $\supset \leftrightsquigarrow \supset$, satisfying item 1. Item 2 follows from the correspondence between quantum gadget and tensor operations, along with Proposition 6.2.1: if $T, T' \in [\mathcal{F}]_{\mathrm{Pl}^\dagger}$ are the tensors of quantum $\mathrm{Pl}^\dagger$-Holant($\mathcal{F}$) gadgets $\mathbf{K}$ and $\mathbf{K}'$, then

$$\Phi(T) \circ \Phi(T') = M(\mathbf{K}_{\mathcal{F}\to\mathcal{G}}) \circ M(\mathbf{K}'_{\mathcal{F}\to\mathcal{G}}) = M(\mathbf{K}_{\mathcal{F}\to\mathcal{G}} \circ \mathbf{K}'_{\mathcal{F}\to\mathcal{G}}) = M((\mathbf{K} \circ \mathbf{K}')_{\mathcal{F}\to\mathcal{G}}) = \Phi(T \circ T')$$

and similarly for $\otimes$ and $\dagger$. Therefore Theorem 9.2.2 gives a quantum orthogonal matrix U satisfying $U^{\otimes \ell} F^{\ell,r} = G^{\ell,r} U^{\otimes r}$ for every ${}_{\ell+r}\mathcal{F} \ni F \leftrightsquigarrow G \in {}_{\ell+r}\mathcal{G}$, which, by (6.3.3), is equivalent to $U\mathcal{F} = \mathcal{G}$. $\square$

Now Theorem 9.0.1 follows from Corollary 6.3.3 and Corollary 9.2.1.

Chapter 10

Bipartite Holant, Orbit Closure Intersection, and Invertible Transformation

This chapter is based on [CY26], which is joint work with Jin-Yi Cai.

10.1 Introduction

Chapters 5, 7, 8, and 9 all successfully proved theorems of the following form: if $\mathcal{F}$ and $\mathcal{G}$ are indistinguishable – by #CSP, $\text{Pl}^{\dagger}$-#CSP, Holant, or $\text{Pl}^{\dagger}$-Holant – then there is always some matrix – permutation, quantum permutation, orthogonal, or quantum orthogonal – transforming $\mathcal{F}$ to $\mathcal{G}$. In this final chapter, we attempt the same for Bi-Holant and matrices in GL_q, but will not be as successful. An analogous result to the four earlier chapters would be a converse to Valiant's bipartite Holant theorem (Theorem 2.4.1; recall that the orthogonal Holant theorem is a special case), as conjectured by Xia [Xia10]: if $\mathcal{F} \,|\, \mathcal{F}'$ and $\mathcal{G} \,|\, \mathcal{G}'$ are Bi-Holant-indistinguishable, then there is a $T \in \text{GL}_q$ such that $T(\mathcal{F} \,|\, \mathcal{F}') = (\mathcal{G} \,|\, \mathcal{G}')$. However, it is now known that this converse is false. Instead, we must settle for two near-converses. On the bright side, though, the two near-converses are more interesting than a true converse of the bipartite Holant theorem, as they develop new connections between counting problems and other areas of mathematics and complexity theory.

There are two related reasons why the converse of the bipartite Holant theorem fails, whereas the indistinguishability theorems of earlier chapters succeed. First, the groups S_q, S_q^+, O_q, and O_q^+ whose invariant theory governs the behavior of #CSP and, for real-valued sets, Pl-#CSP, Holant, and Pl-Holant quantum gadgets are *compact* (S_q and O_q in terms of the usual Euclidean topology, S_q^+ and O_q^+ in terms of being compact matrix quantum groups), while GL_q is not compact. Similarly, the converse of the (quantum) orthogonal Holant theorem does not hold for complex-valued and non-conjugate-closed sets, and the complex orthogonal group O_q is not compact.

The second, more directly relevant reason is that, in the settings in which the earlier chapters' theorems apply – #CSP($\mathcal{F}$), Pl†-#CSP, and Pl†-Holant for complex-valued $\mathcal{F}$, Pl$^\top$-#CSP and Pl$^\top$-Holant for conjugate-closed $\mathcal{F}$, and Holant for real $\mathcal{F}$ – the quantum gadget spaces admit a nondegenerate inner product (in this chapter's terminology, they are *quantum-nonvanishing*): given any quantum gadget tensor $K \in {}_\ell\mathcal{V}(\mathbb{C}^q)_r$, we can always find another quantum gadget tensor $\widehat{K} \in {}_r\mathcal{V}(\mathbb{C}^q)_\ell$ such that $\langle K, \widehat{K} \rangle \neq 0$. In fact, in every case $\widehat{K} = K^\dagger$ (for Pl†-Holant, recall the proof of Lemma 9.1.1; for #CSP, as shown in Theorem 5.1.1, closure under entrywise product means we can effectively assume sets are conjugate-closed). The same holds for Holant($\mathcal{F}$) with conjugate-closed $\mathcal{F}$, and indeed in this chapter we generalize the orthogonal converse to this setting.

Returning to the converse of the bipartite Holant theorem, Cai, Guo, and Williams [CGW16, Section 4.3] discovered the following Boolean-domain counterexample.

Example 10.1.1. Consider the arity-4 signature $F' = [f_0, f_1, f_2, f_3, f_4] = [a, b, 1, 0, 0]$, where f_i is the value of F' on inputs of Hamming weight i and a and b are not both 0. Define $G' = [0, 0, 1, 0, 0]$ and $(\neq_2) = [0, 1, 0]$ similarly. Define $\mathcal{F} \,|\, \mathcal{F}' = (\neq_2)|F'$ and $\mathcal{G} \,|\, \mathcal{G}' = (\neq_2)|G'$. In an $(\neq_2)|F'$-grid Ω, the $\neq_2$ signatures in the left bipartition force any nonzero edge assignment σ to assign 0 to exactly half of the edges and 1 to the other half. Also, σ must provide every $[a, b, 1, 0, 0]$ in the right bipartition no more 1 than 0 inputs. If σ provides any $[a, b, 1, 0, 0]$ strictly fewer 1 than 0 inputs (to obtain a or b), it must provide a different $[a, b, 1, 0, 0]$ strictly more 1 than 0 inputs to preserve the 0/1 balance, and becomes zero. Hence $(\neq_2)|F'$ is indistinguishable from $(\neq_2)|G'$.

However, there is no $T \in \mathrm{GL}_2$ satisfying $T \cdot (\neq_2) = (\neq_2)$ and $[a, b, 1, 0, 0]T^{-1} = [0, 0, 1, 0, 0]$. To see this, if $T = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$, then, with $(\neq_2)^{2,0} = (0,1,1,0)^\top$,

$$T^{\otimes 2}\begin{bmatrix}0\\1\\1\\0\end{bmatrix} = \begin{bmatrix} a^2 & ab & ba & b^2 \\ ac & ad & bc & bd \\ ca & cb & da & db \\ c^2 & cd & dc & d^2 \end{bmatrix}\begin{bmatrix}0\\1\\1\\0\end{bmatrix} = \begin{bmatrix}0\\1\\1\\0\end{bmatrix} \iff ab = cd = 0 \text{ and } ad + bc = 1,$$

which implies that $a = d = 0$ or $b = c = 0$. If $b = c = 0$ then T simply rescales a signature's entries, and if $a = d = 0$ then T rescales a signature's entries and exchanges the roles of 0 and 1, which has the effect of reversing the entries in the $[\cdot, \dots, \cdot]$ notation. Thus $[0,0,1,0,0]T = [0,0,*,0,0] \neq [a,b,1,0,0]$.

The two main results of this chapter are motivated by two observations about Example 10.1.1. We discuss these results in the following two subsections.

10.1.1 Orbit closure intersection.

The GL_q-*orbit* of a finite set $\mathcal{F}$ of tensors is the set $\{T \cdot \mathcal{F} \mid T \in \mathrm{GL}_q\}$, where T acts simultaneously on the tensors in $\mathcal{F}$, in our setting by holographic transformation. Therefore the converse of the Holant theorem would state that, if $\mathcal{F} \,|\, \mathcal{F}'$ and $\mathcal{G} \,|\, \mathcal{G}'$ are Holant-indistinguishable, then the GL_q-orbits of $\mathcal{F} \,|\, \mathcal{F}'$ and $\mathcal{G} \,|\, \mathcal{G}'$ intersect and hence are equal. A weaker and often better-behaved notion is that of orbit *closure* intersection (Euclidean closure, for $\mathbb{K} = \mathbb{C}$). There has been much research in recent years on the computational complexity of orbit intersection and orbit closure intersection for various actions of a linear-algebraic group H [GQ23; Che+24; GQ25; IQ23; DM20; All+18; Gar+20; Acu+23; LW25] with connections to geometric complexity theory [Lan17], including border rank with applications to matrix multiplication [BI11], and polynomial identity testing.

Several such works [DM20; Gar+20; IQ23; Acu+23; LW25] apply a theorem (Theorem 10.2.1 below) from geometric invariant theory which states that the H-orbit closures of $\mathcal{F}$ and $\mathcal{G}$ intersect if and only if $\mathcal{F}$ and $\mathcal{G}$ are indistinguishable by all H-invariant polynomials (i.e. every such polynomial takes the same value on inputs $\mathcal{F}$ and $\mathcal{G}$). The authors of [Acu+23] study a family of vertex-regular tensor networks from quantum physics called PEPS networks, which admit a variant of holographic transformation called a *gauge transformation.* A PEPS signature set $\mathcal{F}$ has common arity $2n$, with inputs paired into n dimensions (with possibly distinct domains) and only allows contractions between inputs in the same dimension. Since every GL_q-invariant polynomial is a linear combination of contractions of PEPS networks, $\mathcal{F}$ and $\mathcal{G}$ are indistinguishable by all PEPS networks if and only if the GL_q-orbit closures of $\mathcal{F}$ and $\mathcal{G}$ intersect [Acu+23, Theorem 4.11]. Lysikov and Walter [LW25] define the complexity class **TOCI** of orbit closure intersection problems, showing that it contains **GI** (all problems reducible to graph isomorphism).

Returning to Example 10.1.1, observe that

$$\begin{bmatrix} \epsilon^{-1} & 0 \\ 0 & \epsilon \end{bmatrix}^{\otimes 2} (\neq_2) = (\neq_2) \quad \text{and} \quad [a, b, 1, 0, 0] \begin{bmatrix} \epsilon & 0 \\ 0 & \epsilon^{-1} \end{bmatrix}^{\otimes 4} = [a\epsilon^4, b\epsilon^2, 1, 0, 0] \xrightarrow[\epsilon \to 0]{} [0, 0, 1, 0, 0],$$

so $\left[\begin{smallmatrix} \epsilon^{-1} & 0 \\ 0 & \epsilon \end{smallmatrix}\right] \in \mathrm{GL}_2$ take $\mathcal{F} \mid \mathcal{F}'$ arbitrarily close to $\mathcal{G} \mid \mathcal{G}'$ as $\epsilon \to 0$. The following result extends this to *any* Bi-Holant-indistinguishable $\mathcal{F}$ and $\mathcal{G}$: the converse of Theorem 2.4.1 holds up to orbit closure.

Theorem (Approximate converse, Theorem 10.2.4)**.** *Finite sets $\mathcal{F} \mid \mathcal{F}'$ and $\mathcal{G} \mid \mathcal{G}'$ are Holant-indistinguishable if and only if the* GL_q*-orbit closures of $\mathcal{F} \mid \mathcal{F}'$ and $\mathcal{G} \mid \mathcal{G}'$ intersect.*

It follows that the problem of testing whether $\mathcal{F} \mid \mathcal{F}'$ and $\mathcal{G} \mid \mathcal{G}'$ are Holant-indistinguishable is decidable. A special case is a characterization of vanishing sets which applies to any set on any domain: $\mathcal{F} \mid \mathcal{F}'$ is vanishing if and only $\mathcal{F} \mid \mathcal{F}'$ is in the *null cone*, meaning its GL_q-orbit closure contains the set of all-0 signatures. This greatly generalizes the symmetric Boolean-domain characterization of [CGW16]. Furthermore, general null cone membership testing is a well-studied problem, including for actions of GL_q [Bür+19]. In Section 10.5, we apply the approximate converse and results of [LW25] to show that the problem of Holant-indistinguishability is **TOCI**-complete and **GI**-hard.

We also resolve an open problem from homomorphism indistinguishability. One notable graph class $\mathfrak{G}$ whose homomorphism indistinguishability relation had, since the seminal 2010 work of Dvořák [Dvo10], eluded any full characterization is the graphs of bounded degree (on the other hand, the complexity classification of homomorphism counting from bounded-degree graphs has been achieved [GCD23; CG20]). Roberson [Rob22] showed that homomorphism indistinguishability from graphs of degree at most d define distinct relations strictly weaker than isomorphism on the set of graphs for distinct d, but did not characterize them further. By expressing bounded-degree graph homomorphism as a bipartite Holant problem, we obtain as a corollary of the approximate converse the first characterization of its indistinguishability relation.

10.1.2 Quantum-nonvanishing bipartite sets

Cai, Guo and Williams discovered Example 10.1.1 while studying *vanishing* signature sets, those sets which are Holant-indistinguishable from 0 (more precisely, the appropriate all-0 set). Reasoning similarly to Example 10.1.1, $(\neq_2 | \, [a, b, 0, 0, 0])$ is vanishing. Our second near-converse of the Holant theorem shows that that *any* counterexample $\mathcal{F} \,|\, \mathcal{F}'$ to the converse of the Holant theorem is due to the presence of a signature that vanishes in the context of $\mathcal{F} \,|\, \mathcal{F}'$. Say, roughly, that $\mathcal{F} \,|\, \mathcal{F}'$ is *quantum-nonvanishing* if $\mathcal{F} \,|\, \mathcal{F}'$ cannot produce a quantum gadget that causes every $\mathcal{F} \,|\, \mathcal{F}'$-grid containing it to have Holant value 0. This generalizes the concept of the annihilator of the quantum labeled graph algebra [FLS07; CG22].

Theorem (Conditional converse, Theorem 10.3.2)**.** *If $\mathcal{F} \,|\, \mathcal{F}'$ and $\mathcal{G} \,|\, \mathcal{G}'$ are Holant-indistinguishable and quantum-nonvanishing, then there is a holographic transformation between $\mathcal{F} \,|\, \mathcal{F}'$ and $\mathcal{G} \,|\, \mathcal{G}'$.*

The proof of this theorem uses an invariant-theoretic characterization due to Derksen and Makam [DM23] of the wheeled PROP $\langle \mathcal{F} \,|\, \mathcal{F}' \rangle$ for quantum-nonvanishing $\mathcal{F} \,|\, \mathcal{F}'$, which generalizes the orthogonal duality Theorem 4.5.3 used to prove the #CSP and Holant indistinguishability theorems in chapters Chapter 5 and Chapter 8. The proof of the conditional converse in many ways resembles the proof of the latter Theorem 8.1.1, with Derksen and Makam's theorem playing exactly the same role as Theorem 4.5.3. However, the same domain-induction approach fails because subsignatures of $\mathcal{F} \,|\, \mathcal{F}'$ do not necessarily inherit the quantum-nonvanishing property. Instead, we use Derksen and Makam's theorem to initially split the problem into two subdomains, then gradually refine these subdomains by holographic transformations until quantum-nonvanishing forces $\mathcal{F} \,|\, \mathcal{F}' = \mathcal{G} \,|\, \mathcal{G}'$. We use similar techniques to prove Theorem 10.4.2, a variant of the second main theorem for quantum-nonvanishing sets $\mathcal{F}$ and $\mathcal{G}$ of matrices: every product of matrices in $\mathcal{F}$ has the same trace as the corresponding product in $\mathcal{G}$ if and only if $\mathcal{F}$ and $\mathcal{G}$ are simultaneously similar. The proof of this result is 'constructive' in the sense that the recovered transformation between $\mathcal{F}$ and $\mathcal{G}$ is composed of Jordan decompositions of quantum-$\mathcal{F}$-gadget-realizable matrices, and of these matrices themselves (although the gadgets are obtained nonconstructively). The proof of the second main theorem is similarly 'constructive' except for an initial application of Derksen and Makam's theorem.

In Section 10.5, we use the conditional converse to show that, while homomorphism indistinguishability of graphs F and G over graphs of any bounded degree is not in general equivalent to isomorphism, homomorphism indistinguishability over graphs of maximum degree at most three is equivalent to isomorphism for F and G with invertible adjacency matrices.

For the reader's convenience, Table 10.1 lists Holant concepts in this chapter on the left, and other names for these concepts on the right.

Signature F	Tensor
Signature set $\mathcal{F}$ (finite)	Tensor tuple
Signature grid Ω	Tensor network
Quantum gadget signature algebra $\langle \mathcal{F} \rangle$	Sub-wheeled PROP of $\mathcal{V}$
$\mathcal{F}$ is quantum-nonvanishing	Wheeled PROP $\langle \mathcal{F} \rangle$ is simple
$\mathcal{F}$ is (Bi-Holant-)vanishing	$\mathcal{F}$ is in the null cone

Table 10.1: Dictionary of terminologies

10.2 The approximate converse

In this section, let $\mathbb{K} = \mathbb{C}$ and assume all tensor sets are finite. We prove the first main theorem of this chapter, Theorem 10.2.4.

Definition 10.2.1 ($\operatorname{GL}_q \mathcal{F}$, $\overline{\operatorname{GL}_q \mathcal{F}}$)**.** The GL_q-*orbit* $\operatorname{GL}_q \mathcal{F}$ of a finite $\mathcal{F} \subset \mathcal{V}$ is $\{T \mathcal{F} \mid T \in \operatorname{GL}_q\}$. If $\mathcal{F} = \{F_1, \dots, F_m\}$ with $F_i \in {}_{\ell_i}\mathcal{V}_{r_i}$, then view $\mathcal{F}$ as an element of the finite-dimensional $\mathbb{C}$-vector space $V := \bigoplus_{i=1}^m {}_{\ell_i}\mathcal{V}_{r_i}$. Then $\operatorname{GL}_q \mathcal{F} \subset V$, so define the GL_q-*orbit closure* $\overline{\operatorname{GL}_q \mathcal{F}}$ of $\mathcal{F}$ as the closure of $\operatorname{GL}_q \mathcal{F}$ in the standard Euclidean topology. Equivalently $\mathcal{G} \in V$ is in $\overline{\operatorname{GL}_q \mathcal{F}}$ if, for every $\epsilon > 0$, there is a $T_\epsilon \in \operatorname{GL}_q$ such that $\|T_\epsilon \mathcal{F} - \mathcal{G}\| < \epsilon$ (using the standard Euclidean norm on V).

Definition 10.2.2 ($\mathbb{C}[\mathcal{X}]$)**.** Let $\mathcal{X}$ be a finite set of domain-q variable-valued tensors (that is, tensors whose underlying signature is variable-valued). For every $X \in \mathcal{X}$ of left arity ℓ and right arity r (corresponding to a complex-valued tensor in ${}_\ell\mathcal{V}_r$) and $\mathbf{a} \in [q]^\ell, \mathbf{b} \in [q]^r$ we introduce a variable $x_{\mathbf{a},\mathbf{b}}$. Define $\mathbb{C}[\mathcal{X}]$ to be the ring of polynomials $\mathbb{C}[\{x_{\mathbf{a},\mathbf{b}} : X \in \mathcal{X}, \mathbf{a} \in [q]^{\ell(X)}, \mathbf{b} \in [q]^{r(X)}\}]$.

Equivalently, $\mathbb{C}[\mathcal{X}] \cong \mathbb{C}[V]$ is the coordinate ring of the vector space V from Definition 10.2.1 (where $\mathcal{X}$ is bijective with $\mathcal{F}$). For variable-valued $\mathcal{X}$ and $\mathcal{X}$-grid Ω, $\text{Bi-Holant}_\Omega(\mathcal{X})$ is a polynomial in the entries of $\mathcal{X}$. Evaluating this polynomial at $\mathcal{F}$ for $\mathbb{C}$-valued $\mathcal{F}$ bijective with $\mathcal{X}$ (by substituting

$F(\mathbf{a},\mathbf{b})$ for $x_{\mathbf{a},\mathbf{b}}$ with $\mathcal{F} \ni F \leftrightsquigarrow X \in \mathcal{X}$) yields $\text{Bi-Holant}_\Omega(\mathcal{F}) \in \mathbb{C}$. Figure 10.1 shows an example on the Boolean domain with $\mathcal{X} = \{X, Y\}$ for binary covariant X and unary contravariant Y.

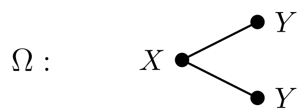

Figure 10.1: $\text{Holant}_\Omega = x_{(\cdot,00)}y^2_{(0,\cdot)} + x_{(\cdot,01)}y_{(0,\cdot)}y_{(1,\cdot)} + x_{(\cdot,10)}y_{(1,\cdot)}y_{(0,\cdot)} + x_{(\cdot,11)}y^2_{(1,\cdot)}$, with the four monomials corresponding to the edge assignments $00, 01, 10, 11$, respectively.

Define an action of GL_q on $\mathbb{C}[\mathcal{X}]$ as follows. For $T \in \text{GL}_q$ and $p \in \mathbb{C}[\mathcal{X}]$, construct $Tp \in \mathbb{C}[\mathcal{X}]$ by substituting every variable $x_{\mathbf{a},\mathbf{b}}$ with the $(\mathbf{a},\mathbf{b})$-entry of $T^{-1} \cdot X$. Equivalently,

$$(Tp)(\mathcal{F}) = p(T^{-1}\mathcal{F}) \tag{10.2.1}$$

for $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$ bijective with $\mathcal{X}$. Then define

$$\mathbb{C}[\mathcal{X}]^{\text{GL}_q} := \{p \in \mathbb{C}[\mathcal{X}] \mid Tp = p \text{ for every } T \in \text{GL}_q\}$$

to be the set of polynomials invariant under this action. The following theorem from geometric invariant theory, stated in this form in [DM22, Theorem 2.3], [DK15, Corollary 2.3.8], shows that the GL_q-orbit closures of $\mathcal{F}$ and $\mathcal{G}$ intersect if and only if $\mathcal{F}$ and $\mathcal{G}$ are indistinguishable by all GL_q-invariant polynomials.

Theorem 10.2.1 (Mumford, Fogarty, and Kirwan [MFK94])**.** *Let $\mathcal{F}, \mathcal{G} \subset \mathcal{V}(\mathbb{C}^q)$ be bijective with $\mathcal{X}$. Then $\overline{\text{GL}_q \mathcal{F}} \cap \overline{\text{GL}_q \mathcal{G}} \neq \varnothing$ if and only if $p(\mathcal{F}) = p(\mathcal{G})$ for every $p \in \mathbb{C}[\mathcal{X}]^{\text{GL}_q}$.*

More generally, Theorem 10.2.1 applies to any reductive algebraic group in place of GL_q acting on any vector space V over an algebraically closed field (although for fields other than $\mathbb{C}$ we must define $\overline{\text{GL}_q \mathcal{F}}$ as the Zariski, instead of Euclidean, closure). Accompanying Theorem 10.2.1 is a result of Hilbert (see [Der01]), which implies that it suffices to check finitely many (with the exact number depending on the arity profile of $\mathcal{X}$) polynomial invariants to ensure orbit closure intersection.

Theorem 10.2.2 (Hilbert [Hil90])**.** *The $\mathbb{C}$-algebra $\mathbb{C}[\mathcal{X}]^{\text{GL}_q}$ is finitely generated.*

To convert the condition in Theorem 10.2.1 from polynomial indistinguishability to Bi-Holant indistinguishability, we apply the following minor generalization of Weyl's Polynomial First Fundamental Theorem for GL_q [Wey66; GW09] more suited to our purpose. The proof applies an argument similar to [LW25, Theorem 4.23 and Lemma 4.26] (see also [Acu+23, Proposition 4.13]).

Theorem 10.2.3. *For variable-valued tensor set $\mathcal{X} = \{X_1, \ldots, X_m\}$ on domain $[q]$,*

$$\mathbb{C}[\mathcal{X}]^{\mathrm{GL}_q} = \mathrm{span}\{\text{Bi-Holant}_\Omega(\mathcal{X}) : \mathcal{X}\text{-}grid\ \Omega\}$$

Proof. The $\supset$ inclusion follows from (10.2.1), Theorem 4.1.1, and the fact that two polynomials which take the same value on every point must be identical.

For the $\subset$ inclusion, let $p \in \mathbb{C}[\mathcal{X}]^{\mathrm{GL}_q}$. Split p into a sum $p = \sum_{d_1,\ldots,d_m \geq 0} p_{\mathbf{d}}$ of multihomogeneous polynomials, where d_i is the total degree of the entries of X_i in $p_{\mathbf{d}}$ (and only finitely many $p_{\mathbf{d}}$ are nonzero). Since the action of GL_q replaces each variable $(x_i)_{\mathbf{a},\mathbf{b}}$ with a linear polynomial in the entries of the same tensor X_i, it preserves the multihomogeneous degree of each $p_{\mathbf{d}}$. Therefore each $p_{\mathbf{d}} \in \mathbb{C}[\mathcal{X}]^{\mathrm{GL}_q}$, and it suffices to find an Ω such that $\text{Bi-Holant}_\Omega(\mathcal{X}) = p_{\mathbf{d}}$. Let X_i have left arity ℓ_i and right arity r_i. Each $p_{\mathbf{d}}$ corresponds to a tensor $A_{\mathbf{d}} \in W^*$, where, up to reordering of factors,

$$W := \bigotimes_{i=1}^{m} \mathrm{Sym}^{d_i}\left({}_{\ell_i}\mathcal{V}_{r_i}\right) \subset (\mathbb{C}^q)^{\otimes \sum_i \ell_i d_i} \otimes ((\mathbb{C}^q)^*)^{\otimes \sum_i r_i d_i} \tag{10.2.2}$$

(here $\mathrm{Sym}^n(V)$ denotes the space of symmetric tensors in $V^{\otimes n}$) such that, for every complex-valued $\mathcal{F} = \{F_1, \ldots, F_m\}$ bijective with $\mathcal{X}$ and $\bigotimes_{i=1}^m F_i^{\otimes d_i} \in W$, we have

$$p_{\mathbf{d}}(\mathcal{F}) = \left\langle A_{\mathbf{d}},\ \bigotimes_{i=1}^{m} F_i^{\otimes d_i} \right\rangle. \tag{10.2.3}$$

For example, if $q = 5$, $\mathcal{X} = \{X, Y, Z\}$, $(\ell_1, \ell_2, \ell_3) = (0, 3, 1)$, $(r_1, r_2, r_3) = (2, 0, 1)$, and $p_{1,2,1} = x_{(\cdot,45)} y_{(124,\cdot)} y_{(555,\cdot)} z_{(3,4)} = x_{(\cdot,45)} y_{(555,\cdot)} y_{(124,\cdot)} z_{(3,4)}$, then

$$\begin{aligned} A_{1,2,1} = \frac{1}{2}\big(&(e_4 \otimes e_5) \otimes (e_1^* \otimes e_2^* \otimes e_4^*) \otimes (e_5^* \otimes e_5^* \otimes e_5^*) \otimes (e_3^* \otimes e_4) \\ &+ (e_4 \otimes e_5) \otimes (e_5^* \otimes e_5^* \otimes e_5^*) \otimes (e_1^* \otimes e_2^* \otimes e_4^*) \otimes (e_3^* \otimes e_4)\big). \end{aligned}$$

Now with $\bigotimes_{i=1}^m (X_i)^{\otimes d_i}$ having left arity $\sum_i \ell_i d_i$ and right arity $\sum_i r_i d_i$, (10.2.3) is equivalent to

$$p_{\mathbf{d}} = \left\langle A_{\mathbf{d}},\ \bigotimes_{i=1}^{m} X_i^{\otimes d_i} \right\rangle. \tag{10.2.4}$$

Furthermore, for any $T \in \mathrm{GL}_q$,

$$T p_{\mathbf{d}} = p_{\mathbf{d}}(T^{-1}\mathcal{X}) = \left\langle A_{\mathbf{d}},\ \bigotimes_{i=1}^{m} (T^{-1} \cdot X_i)^{\otimes d_i} \right\rangle = \left\langle T \cdot A_{\mathbf{d}},\ \bigotimes_{i=1}^{m} (X_i)^{\otimes d_i} \right\rangle$$

so the map $p_{\mathbf{d}} \mapsto A_{\mathbf{d}}$ is GL_q-equivariant. With $p_{\mathbf{d}} \in \mathbb{C}[\mathcal{X}]^{\mathrm{GL}_q}$, it follows that $A_{\mathbf{d}} \in \mathcal{V}(\mathbb{C}^q)^{\mathrm{GL}_q}$ (up to the reordering of factors in (10.2.2), which doesn't affect this invariance), so, by Theorem 4.3.1, $A_{\mathbf{d}} \in \langle \varnothing \rangle$ is the tensor of a wire gadget. Now (10.2.4) says that $p_{\mathbf{d}}$ is a full contraction consisting only of wires and tensors in $\mathcal{X}$, which is Bi-Holant$_\Omega(\mathcal{X})$ for some $\mathcal{X}$-grid Ω. □

Theorem 10.2.4 (The approximate converse)**.** *Finite tensor sets $\mathcal{F}, \mathcal{G} \subset \mathcal{V}(\mathbb{C}^q)$ are Bi-Holant-indistinguishable if and only if $\overline{\mathrm{GL}_q \mathcal{F}} \cap \overline{\mathrm{GL}_q \mathcal{G}} \neq \varnothing$.*

Proof. The ($\Rightarrow$) direction follows from Theorem 10.2.1 and Theorem 10.2.3. ($\Leftarrow$) follows from Theorem 4.1.1 and the fact that Bi-Holant$_\Omega(\mathcal{F})$ is a polynomial, hence continuous, function in $\mathcal{F}$. □

Combining the approximate converse with Mumford and Hilbert's theorems gives the following.

Corollary 10.2.1. *The problem of determining whether two finite $\mathcal{F}, \mathcal{G} \subset \mathcal{V}(\mathbb{C}^q)$ are Bi-Holant-indistinguishable is decidable.*

See Sections Section 10.5.2 and Section 10.6.2 for some discussion of the complexity of Bi-Holant-indistinguishability. Say $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$ is *Bi-Holant-vanishing* if it is Bi-Holant-indistinguishable from the set of all-0 tensors. Since Holant is a special case of Bi-Holant, this notion captures both Bi-Holant and Holant vanishing. Say $\mathcal{F}$ is in the *null cone* if $0 \in \overline{\mathrm{GL}_q \mathcal{F}}$.

Corollary 10.2.2. *Finite $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$ is Bi-Holant-vanishing if and only if $\mathcal{F}$ is in the null cone.*

10.3 Quantum-nonvanishing

Definition 10.3.1 ($\mathcal{F}$-nonvanishing, Quantum-nonvanishing)**.** Say $K \in {}_\ell\langle \mathcal{F} \rangle_r$ is *$\mathcal{F}$-nonvanishing* if it satisfies any of the following equivalent conditions.

(1) There is a $\widehat{K} \in {}_r\langle \mathcal{F} \rangle_\ell$ such that $\langle K, \widehat{K} \rangle \neq 0$, or

(2) there is an $\langle \mathcal{F} \rangle$-grid Ω containing K such that Bi-Holant$(\Omega) \neq 0$, or

(3) there is an $\mathcal{F} \cup \{K\}$-grid Ω containing K such that Bi-Holant$(\Omega) \neq 0$.

Then say $\mathcal{F} \subset \mathcal{V}$ is (ℓ, r)-*quantum-nonvanishing* if every nonzero $K \in {}_\ell\langle\mathcal{F}\rangle_r$ is $\mathcal{F}$-nonvanishing (equivalently, the bilinear form $\langle\cdot,\cdot\rangle : {}_\ell\langle\mathcal{F}\rangle_r \times {}_r\langle\mathcal{F}\rangle_\ell \to \mathbb{K}$ is nondegenerate), and $\mathcal{F}$ is *quantum-nonvanishing* if it is (ℓ, r)-quantum-nonvanishing for every (ℓ, r).

Proof. (1) $\Longrightarrow$ (2) because $\langle K, \widehat{K}\rangle$ is the Holant value of an $\langle\mathcal{F}\rangle$-grid containing K, and (2) $\Longrightarrow$ (1) because, given Ω, let $\widehat{K}$ be the tensor of the $\langle\mathcal{F}\rangle$-gadget formed by removing a vertex assigned K from Ω, leaving its formerly incident edges dangling. (3) $\Longrightarrow$ (2) because every $\mathcal{F} \cup \{K\}$-grid is an $\langle\mathcal{F}\rangle$-grid, and (2) $\Longrightarrow$ (3) because expanding as quantum-$\mathcal{F}$-gadgets the other tensors in the $\langle\mathcal{F}\rangle$ grid Ω containing K yields a quantum $\mathcal{F} \cup \{K\}$-grid with each term containing K, at least one of which has nonzero Holant value. □

The following theorem of Derksen and Makam on wheeled PROPs implies that, if $\mathcal{F}$ is quantum-nonvanishing ("simple" in Derksen and Makam's terminology), then there is a subgroup $\mathrm{Stab}(\mathcal{F}) \subset \mathrm{GL}_q$ such that every tensor in $\mathcal{V}$ invariant under the action of $\mathrm{Stab}(\mathcal{F})$ is realizable as a quantum-$\mathcal{F}$-gadget tensor.

Theorem 10.3.1 ([DM23, Theorem 6.2, Proposition 6.5, Corollary 6.6])**.** *A set $\mathcal{F} \subset \mathcal{V}(\mathbb{K}^q)$ is quantum-nonvanishing if and only if $\langle\mathcal{F}\rangle = \mathcal{V}^{\mathrm{Stab}(\mathcal{F})}$ for some reductive subgroup $\mathrm{Stab}(\mathcal{F}) \subset \mathrm{GL}_q$. Furthermore, if these conditions hold, then $\langle\mathcal{F}\rangle$ is finitely generated.*

We will not define a reductive subgroup of GL_q here; it suffices to say that GL_q, U_q, and O_q are reductive. By the proof of Proposition 10.5.1 below, which shows that the relevant sets are quantum-nonvanishing, Theorem 10.3.1 generalizes Theorem 4.5.3, itself a generalization of Corollary 4.5.2, the key result underlying the proofs of the #CSP and Holant indistinguishability theorems in chapters Chapter 5 and Chapter 8.

In Section 10.4.2, we use Theorem 10.3.1 to prove the following theorem.

Theorem 10.3.2 (The conditional converse)**.** *Let $\mathcal{F} \mid \mathcal{F}'$ and $\mathcal{G} \mid \mathcal{G}'$ be quantum-nonvanishing. Then $\mathcal{F} \mid \mathcal{F}'$ and $\mathcal{G} \mid \mathcal{G}'$ are Holant-indistinguishable if and only if there is a $T \in \mathrm{GL}_q$ such that $T(\mathcal{F} \mid \mathcal{F}') = \mathcal{G} \mid \mathcal{G}'$.*

Theorem 10.3.2 implies that any $\mathcal{F} \mid \mathcal{F}'$ and $\mathcal{G} \mid \mathcal{G}'$ serving as a counterexample to the converse of the Holant theorem cannot both be quantum-nonvanishing. In Example 10.1.1, $\mathcal{F} \mid \mathcal{F}'$ is quantum-vanishing. To see this, consider the quantum $\mathcal{F} \mid \mathcal{F}'$-gadget $4\,\mathbf{K}_1 - \mathbf{K}_2$ shown in Figure 10.1. Reason-

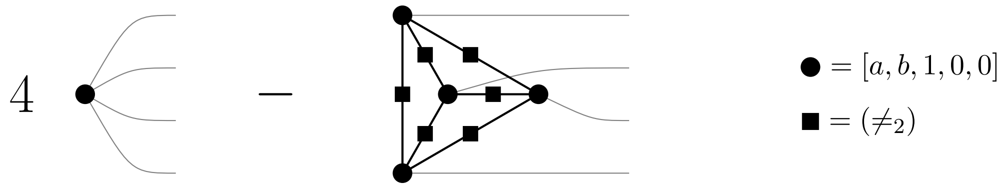

Figure 10.1: A quantum gadget $4\,\mathbf{K}_1 - \mathbf{K}_2$ with signature $K = [4a - p_1(a,b), 4b - p_2(b), 0, 0, 0]$

ing as in Example 10.1.1, the symmetric gadget $\mathbf{K}_2$ has signature $[p_1(a,b), p_2(b), 4, 0, 0]$ for polynomials p_1 in a and b and p_2 in b, so the signature of $4\,\mathbf{K}_1 - \mathbf{K}_2$ is $K := [4a-p_1(a,b), 4b-p_2(b), 0, 0, 0] \in {}_0\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_4$. But, in any $(\neq_2) \mid F', K$-grid Ω containing K, every nonzero assignment is forced to assign K strictly fewer 1s than 0s, so must assign strictly more 1s than 0s to another $[a, b, 1, 0, 0]$ or K, which then evaluates to 0. Therefore K is $\mathcal{F} \,|\, \mathcal{F}'$-vanishing (if a, b are such that $K = 0$, Theorem 10.3.2 asserts that some nonzero quantum gadget must be $\mathcal{F} \,|\, \mathcal{F}'$-vanishing).

Observe that the $\mathcal{F} \,|\, \mathcal{F}'$-vanishing K corresponds to $0 \in \langle \mathcal{G} \,|\, \mathcal{G}' \rangle$. This motivates the following definition, a wheeled PROP version of Definition 9.1.1 (with the same equivalence argument).

Definition 10.3.2. Bijective $\mathcal{F}, \mathcal{G} \subset \mathcal{V}$ are (ℓ, r)-*weakly isomorphic* if they satisfy any of the following equivalent conditions:

1. $\leftrightsquigarrow$ is an invertible linear map, called a *$Pl^\dagger$-weak-isomorphism*, between ${}_\ell\langle \mathcal{F} \rangle_r$ and ${}_\ell\langle \mathcal{G} \rangle_r$.

2. Every ${}_\ell\langle \mathcal{F} \rangle_r \ni F \leftrightsquigarrow G \in {}_\ell\langle \mathcal{G} \rangle_r$ satisfy $F = 0 \iff G = 0$.

3. $M(\mathbf{K}) = M(\mathbf{J}) \iff M(\mathbf{K}_{\mathcal{F}\to\mathcal{G}}) = M(\mathbf{J}_{\mathcal{F}\to\mathcal{G}})$ for every pair of (ℓ, r)-quantum Bi-Holant($\mathcal{F}$) gadgets $\mathbf{K}$ and $\mathbf{J}$.

Then $\mathcal{F}$ and $\mathcal{G}$ are *weakly isomorphic* if they are (ℓ, r)-weakly-isomorphic for all $\ell, r \geq 0$.

Weak isomorphism generalizes indistinguishability in the following sense.

Proposition 10.3.1. *$\mathcal{F}$ and $\mathcal{G}$ are $(0,0)$-weakly-isomorphic iff they are Bi-Holant-indistinguishable.*

Proof. We have ${}_0\langle \mathcal{F} \rangle_0 = \text{span}\{\text{Bi-Holant}_{\mathcal{F}}(\Omega) : \mathcal{F}\text{-grid } \Omega\} \subset \mathbb{K}$. So if $\mathcal{F}$ and $\mathcal{G}$ are indistinguishable, then every scalar in ${}_0\langle \mathcal{F} \rangle_0$ equals the corresponding scalar in ${}_0\langle \mathcal{G} \rangle_0$, hence $\mathcal{F}$ and $\mathcal{G}$ are $(0,0)$-weakly-isomorphic. Conversely, suppose there is an $\mathcal{F}$-grid Ω such that $\text{Bi-Holant}_{\mathcal{F}}(\Omega) \neq$

Bi-Holant$_{\mathcal{G}}(\Omega_{\mathcal{F}\to\mathcal{G}})$. In both $\langle\mathcal{F}\rangle$ and $\langle\mathcal{G}\rangle$, the vertexless loop $\bigcirc$ has Holant value $q \in \mathbb{K}$. Therefore

$$0 = \text{Bi-Holant}_{\mathcal{F}}\left(\Omega - \frac{\text{Bi-Holant}_{\mathcal{F}}(\Omega)}{q}\cdot\bigcirc\right) \leftrightsquigarrow \text{Bi-Holant}_{\mathcal{G}}\left(\Omega_{\mathcal{F}\to\mathcal{G}} - \frac{\text{Bi-Holant}_{\mathcal{F}}(\Omega)}{q}\cdot\bigcirc\right) \neq 0,$$

so $\mathcal{F}$ and $\mathcal{G}$ are not $(0,0)$-weakly-isomorphic. □

We also have a quasi-generalization of Lemma 9.1.1, whose proof implicitly uses the fact that sets in the context of $\text{Pl}^\dagger$-Holant are always quantum-nonvanishing.

Proposition 10.3.2. *If $\mathcal{F}$ and $\mathcal{G}$ are Bi-Holant-indistinguishable and (ℓ,r)-quantum-nonvanishing, then $\mathcal{F}$ and $\mathcal{G}$ are (ℓ,r)-weakly-isomorphic.*

Proof. Assume $\mathcal{F}$ and $\mathcal{G}$ are not (ℓ,r)-weakly-isomorphic, so WLOG there is a ${}_\ell\langle\mathcal{F}\rangle_r \ni K \leftrightsquigarrow 0 \in {}_\ell\langle\mathcal{G}\rangle_r$ with $K \neq 0$. By indistinguishability, every $\mathcal{F}\cup\{K\}$-grid Ω containing K satisfies Bi-Holant(Ω) = Bi-Holant$(\Omega_{\mathcal{F}\cup\{K\}\to\mathcal{G}\cup\{0\}}) = 0$ because $\Omega_{\mathcal{F}\cup\{K\}\to\mathcal{G}\cup\{0\}}$ contains 0. Therefore K is $\mathcal{F}$-vanishing, so $\mathcal{F}$ is (ℓ,r)-quantum-vanishing. □

The proof of the conditional converse Theorem 10.3.2, like the proof of the orthogonal converse special case (Theorem 8.1.1), makes heavy use of the subdomain restriction results from Section 8.2.2. The following result will prove useful in that vein.

Proposition 10.3.3. *If $\mathcal{F}$ is (ℓ,r)-quantum-nonvanishing and $I_X^\uparrow \in \langle\mathcal{F}\rangle$, then $\langle\mathcal{F}\rangle_X$ is (ℓ,r)-quantum-nonvanishing.*

Proof. Let $F \in {}_\ell\langle\langle\mathcal{F}\rangle_X\rangle_r$. By Proposition 8.2.2, $F^\uparrow \in {}_\ell\langle\mathcal{F}\rangle_r$. Since $\mathcal{F}$ is (ℓ,r)-quantum-nonvanishing, there is a $\widehat{F^\uparrow} \in {}_r\langle\mathcal{F}\rangle_\ell$ such that $\langle F^\uparrow, \widehat{F^\uparrow}\rangle \neq 0$. But $F^\uparrow$ is only supported on X, so

$$0 \neq \langle F^\uparrow, \widehat{F^\uparrow}\rangle = \langle F^\uparrow|_X, \widehat{F^\uparrow}|_X\rangle = \langle F, \widehat{F^\uparrow}|_X\rangle.$$

Thus $\widehat{F^\uparrow}|_X \in \langle\mathcal{F}\rangle_X$ witnesses that F is $\langle\mathcal{F}\rangle_X$-nonvanishing, so $\langle\mathcal{F}\rangle_X$ is (ℓ,r)-quantum-nonvanishing.

□

The proof of Theorem 10.3.2 also mirrors the use of $\oplus$ in the proof of the orthogonal converse. We will need the following results.

Proposition 10.3.4. *Assume $\mathcal{F}$ and $\mathcal{G}$ are Bi-Holant-indistinguishable and let $\langle\mathcal{F}\rangle \ni F \leftrightsquigarrow G \in \langle\mathcal{G}\rangle$ and $K \in \langle\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle\rangle$. Then*

$$\langle K, F \oplus G\rangle = \langle K|_{\mathcal{F}}, F\rangle + \langle K|_{\mathcal{G}}, G\rangle = 2\langle K|_{\mathcal{F}}, F\rangle.$$

Proof. In each nonzero term of $\langle K, F \oplus G\rangle$, either all inputs to both K and $F \oplus G$ are from $V(\mathcal{F})$, or all inputs to both K are $F \oplus G$ are from $V(\mathcal{G})$, giving the first equality. The second equality follows from indistinguishability and Proposition 5.2.1. □

The following lemma seems at first as though it should follow from a simple argument along the lines of the preceding propositions, but actually requires overcoming a subtle difficulty concerning the difference between $\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle$ and $\langle\mathcal{F} \oplus \mathcal{G}\rangle$. Every signature in $\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle$, such as $(F_1 \otimes F_2) \oplus (G_1 \otimes G_2)$, is zero on mixed inputs from $\mathcal{F}$ and $\mathcal{G}$. On the other hand, $(F_1 \oplus G_1) \otimes (F_2 \oplus G_2) \in \langle\mathcal{F} \oplus \mathcal{G}\rangle$, being disconnected, could be nonzero on inputs from $V(\mathcal{F})$ to the first factor and $V(\mathcal{G})$ to the second and vice-versa.

Lemma 10.3.1. *Assume $\mathcal{F}$ and $\mathcal{G}$ are Bi-Holant-indistinguishable. Then $\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle$ is quantum-nonvanishing if and only if $\mathcal{F}$ and $\mathcal{G}$ are both quantum-nonvanishing.*

Proof. ($\Longrightarrow$): We will show that $\mathcal{F}$ is quantum-nonvanishing; the proof for $\mathcal{G}$ is similar. Let $F \in \langle\mathcal{F}\rangle$ be nonzero, and $F \leftrightsquigarrow G$. Since $F \oplus G \in \langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle$, the quantum-nonvanishing of $\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle$ guarantees the existence of a $K \in \langle\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle\rangle$ such that, by Proposition 10.3.4, $0 \neq \langle K, F \oplus G\rangle = 2\langle K|_{\mathcal{F}}, F\rangle$. Proposition 5.2.1 asserts that $K|_{\mathcal{F}} \in \langle\mathcal{F}\rangle$, so $K|_{\mathcal{F}}$ witnesses that F is $\mathcal{F}$-nonvanishing.

($\Longleftarrow$): Assume $\mathcal{F}$ and $\mathcal{G}$ are quantum-nonvanishing, and let $0 \neq K \in \langle\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle\rangle$ be the tensor of a quantum $\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle$-gadget $\mathbf{K}$. First suppose that $K|_{\mathcal{F}} \neq 0$. By Proposition 5.2.1, $K|_{\mathcal{F}} \in \langle\mathcal{F}\rangle$, so by the quantum-nonvanishing of $\mathcal{F}$ there is a $\widehat{F} \in \langle\mathcal{F}\rangle$ such that $\langle K|_{\mathcal{F}}, \widehat{F}\rangle \neq 0$. Then, letting $\widehat{F} \leftrightsquigarrow \widehat{G}$, Proposition 10.3.4 gives $\langle K, \widehat{F} \oplus \widehat{G}\rangle = 2\langle K|_{\mathcal{F}}, \widehat{F}\rangle \neq 0$, so $\widehat{F} \oplus \widehat{G} \in \langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle$ witnesses that K is $\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle$-nonvanishing.

If $K|_{\mathcal{F}} = 0$ then $K|_{\mathcal{G}} = 0$ as well by Propositions 5.2.1 and 10.3.2. Since $K \neq 0$, there is a nontrivial partition of the inputs of K into $X_1 \sqcup X_2$ such that the block $K|_{X_1 \leftarrow V(\mathcal{F}), X_2 \leftarrow V(\mathcal{G})}$ of K (in which inputs in X_1 are restricted to $V(\mathcal{F})$ and inputs in X_2 are restricted to $V(\mathcal{G})$) is nonzero. Let $\mathbf{K} = \mathbf{M} + \sum_{i=1}^{j} c_i \, \mathbf{J}_i$, where each $\mathbf{J}_i$ is a $\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle$-gadget composed of two components $\mathbf{J}_{i,1}$

and $\mathbf{J}_{i,2}$, not necessarily themselves connected but disconnected from each other, such that the dangling edges of $\mathbf{J}_i$ indexed by X_1 (resp. X_2) are incident to $\mathbf{J}_{i,1}$ (resp. $\mathbf{J}_{i,2}$), and $\mathbf{M}$ is the quantum-$\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle$-gadget composed of all terms of $\mathbf{K}$ in which there is a path between some input indexed by X_1 and some input indexed by X_2. Hence the tensor M of $\mathbf{M}$ satisfies

$$M|_{X_1 \leftarrow V(\mathcal{F}), X_2 \leftarrow V(\mathcal{G})} = M|_{X_1 \leftarrow V(\mathcal{G}), X_2 \leftarrow V(\mathcal{F})} = 0. \tag{10.3.1}$$

By reordering the left dangling edges and right dangling edges of $\mathbf{K}$, which does not change whether K is $\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle$-nonvanishing, we may assume $\mathbf{J}_i = \mathbf{J}_{i,1} \otimes \mathbf{J}_{i,2}$, so their tensors satisfy

$$J_i|_{X_1 \leftarrow V(\mathcal{F}), X_2 \leftarrow V(\mathcal{G})} = J_{i,1}|_{\mathcal{F}} \otimes J_{i,2}|_{\mathcal{G}} \text{ and } J_i|_{X_1 \leftarrow V(\mathcal{G}), X_2 \leftarrow V(\mathcal{F})} = J_{i,1}|_{\mathcal{G}} \otimes J_{i,2}|_{\mathcal{F}}. \tag{10.3.2}$$

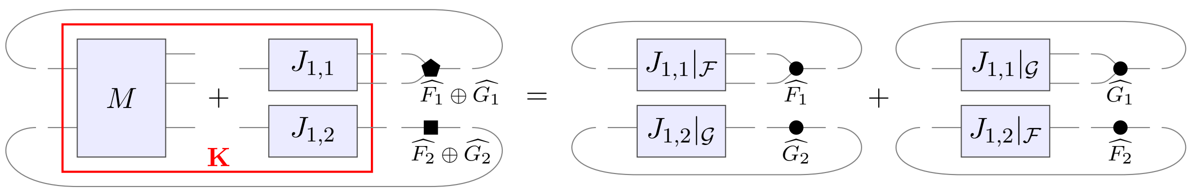

Figure 10.2: Illustrating (10.3.3) for $\mathbf{K} = \mathbf{M} + \mathbf{J}_{1,1} \otimes \mathbf{J}_{1,2}$.

For any $\langle\mathcal{F}\rangle \ni \widehat{F_1}, \widehat{F_2} \leftrightsquigarrow \widehat{G_1}, \widehat{G_2} \in \langle\mathcal{G}\rangle$ of appropriate shape (see Figure 10.2), reasoning similar to Proposition 10.3.4, with the assumption that $K|_{\mathcal{F}} = K|_{\mathcal{G}} = 0$ and (10.3.1) and (10.3.2), gives

$$\begin{aligned}
&\langle K, (\widehat{F_1} \oplus \widehat{G_1}) \otimes (\widehat{F_2} \oplus \widehat{G_2})\rangle \\
&\quad = \langle K|_{\mathcal{F}}, \widehat{F_1} \otimes \widehat{F_2}\rangle + \langle K|_{\mathcal{G}}, \widehat{G_1} \otimes \widehat{G_2}\rangle \\
&\qquad + \langle K|_{X_1 \leftarrow V(\mathcal{F}), X_2 \leftarrow V(\mathcal{G})}, \widehat{F_1} \otimes \widehat{G_2}\rangle + \langle K|_{X_1 \leftarrow V(\mathcal{G}), X_2 \leftarrow V(\mathcal{F})}, \widehat{G_1} \otimes \widehat{F_2}\rangle \\
&\quad = \sum_{i=1}^{j} c_i \langle J_{i,1}|_{\mathcal{F}}, \widehat{F_1}\rangle\langle J_{i,2}|_{\mathcal{G}}, \widehat{G_2}\rangle + \sum_{i=1}^{j} c_i \langle J_{i,1}|_{\mathcal{G}}, \widehat{G_1}\rangle\langle J_{i,2}|_{\mathcal{F}}, \widehat{F_2}\rangle \\
&\quad = 2\left\langle \sum_{i=1}^{j} c_i J_{i,1}|_{\mathcal{F}} \otimes J_{i,2}|_{\mathcal{G}},\ \widehat{F_1} \otimes \widehat{G_2} \right\rangle.
\end{aligned} \tag{10.3.3}$$

Each $J_{i,2}|_{\mathcal{G}} \in \langle\mathcal{G}\rangle$, which is closed under linear combinations, so we may successively eliminate any $J_{i,1}|_{\mathcal{F}}$ which is linearly dependent on the other $J_{i',1}|_{\mathcal{F}}$ to obtain

$$0 \neq K|_{X_1 \leftarrow V(\mathcal{F}), X_2 \leftarrow V(\mathcal{G})} = \sum_{i=1}^{j} c_i J_{i,1}|_{\mathcal{F}} \otimes J_{i,2}|_{\mathcal{G}} = \sum_{i=1}^{j'} c'_i E_i \otimes H_i \tag{10.3.4}$$

for $H_1, \dots H_{j'} \in \langle\mathcal{G}\rangle$ and linearly independent $E_1, \dots, E_{j'} \in \langle\mathcal{F}\rangle$. Substituting into (10.3.3) gives

$$\langle K, (\widehat{F_1} \oplus \widehat{G_1}) \otimes (\widehat{F_2} \oplus \widehat{G_2})\rangle = 2\left\langle \sum_{i=1}^{j'} c'_i E_i \otimes H_i,\ \widehat{F_1} \otimes \widehat{G_2} \right\rangle = \left\langle 2\sum_{i=1}^{j'} c'_i \langle H_i, \widehat{G_2}\rangle E_i,\ \widehat{F_1} \right\rangle. \tag{10.3.5}$$

Some $c'_i H_i \neq 0$ by (10.3.4), so quantum-nonvanishing of $\mathcal{G}$ gives a $\widehat{G_2}$ such that $c'_i\langle H_i, \widehat{G_2}\rangle \neq 0$. Hence, by linear independence, $0 \neq 2\sum_{i=1}^{j'} c'_i\langle H_i, \widehat{G_2}\rangle E_i \in \langle\mathcal{F}\rangle$, so by (10.3.5) and quantum-nonvanishing of $\mathcal{F}$, there is an $\widehat{F_1}$ such that $\langle K, (\widehat{F_1}\oplus\widehat{G_1})\otimes(\widehat{F_2}\oplus\widehat{G_2})\rangle \neq 0$. This $(\widehat{F_1}\oplus\widehat{G_1})\otimes(\widehat{F_2}\oplus\widehat{G_2}) \in \langle\langle\mathcal{F}\rangle\oplus\langle\mathcal{G}\rangle\rangle$ witnesses that K is $\langle\mathcal{F}\rangle\oplus\langle\mathcal{G}\rangle$-nonvanishing. □

10.4 The conditional converse

In this section, we prove the conditional converse Theorem 10.3.2, as well as Theorem 10.4.2, a similar result for sets of matrices $\mathcal{F}, \mathcal{G} \subset {}_1\mathcal{V}(\mathbb{K}^q)_1$. Throughout, let $\mathbb{K}$ be an algebraically closed field of characteristic zero.

10.4.1 Simultaneous similarity

In this subsection, we consider $\mathcal{F} \subset {}_1\mathcal{V}(\mathbb{K}^q)_1$, a set of *mixed binary* tensors with one left and one right input. Thinking of $\mathcal{F}$ as generators of a wheeled PROP, we always assume $I \in \mathcal{F}$. We also view $\mathcal{F}$ as a set of matrices in $\mathbb{K}^{q\times q}$, and for $T \in \mathrm{GL}_q$, $T\mathcal{F} = \{TFT^{-1} \mid F \in \mathcal{F}\}$ is simultaneous conjugation of the matrices in $\mathcal{F}$ by T.

Definition 10.4.1 ($\Gamma_{\mathcal{F}}$)**.** Let $\Gamma_{\mathcal{F}}$ be the set of all finite products of matrices in $\mathcal{F}$.

Every Bi-Holant $\mathcal{F}$-grid is a disjoint union of cycles, each of which defines a word $w \in \Gamma_{\mathcal{F}}$ and has value $\mathrm{tr}(w)$. Note that bipartiteness prevents transposing matrices in $\mathcal{F}$ when constructing w (this would require connecting two left or two right edges), as is allowed in non-bipartite Holant for a set of binary signatures. If transpose is allowed and $\mathbb{K} = \mathbb{R}$ or $\mathbb{K} = \mathbb{C}$, then the indistinguishability relation is always simultaneous similarity by an real or complex orthogonal matrix, respectively (the $\mathbb{R}$ case is Corollary 8.4.2; for the $\mathbb{C}$ case, see [Jin15, Corollary 2.3, Theorem 2.4]). Instead, the only conclusion we can immediately draw from indistinguishability in the bipartite setting is that every $\mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}$ have the same multiset of eigenvalues, as, by a standard argument using Newton's identities for symmetric polynomials, this is equivalent to $\mathrm{tr}(F^k) = \mathrm{tr}(G^k)$ for every $k \geq 0$. Two matrices are similar if and only if they have the same Jordan normal form, so any $\mathcal{F} = \{F\}$ and $\mathcal{G} = \{G\}$ for F and G with identical spectrum but different Jordan normal forms provide a counterexample to the converse of the Bi-Holant theorem. If F is not diagonalizable, then put F in Jordan normal form and write $F = \widetilde{F} + N$ for diagonal $\widetilde{F}$ and nilpotent N. We first make the following well-known observation [LW25]. Since they have the same multiset of eigenvalues, F and $\widetilde{F}$ are Bi-Holant indistinguishable. Therefore, if $\mathbb{K} = \mathbb{C}$, their GL_q-orbit closures intersect by the approximate converse. Indeed, the invertible matrices $\mathrm{diag}(\epsilon^q, \epsilon, \dots, 1)$ transform F arbitrarily close to $\widetilde{F}$ as $\epsilon \to 0$ – e.g.

$$\begin{bmatrix} \epsilon^2 & 0 & 0 \\ 0 & \epsilon & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \lambda & 1 & 0 \\ 0 & \lambda & 1 \\ 0 & 0 & \lambda \end{bmatrix} \begin{bmatrix} \epsilon^{-2} & 0 & 0 \\ 0 & \epsilon^{-1} & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} \lambda & \epsilon & 0 \\ 0 & \lambda & \epsilon \\ 0 & 0 & \lambda \end{bmatrix} \xrightarrow[\epsilon \to 0]{} \begin{bmatrix} \lambda & 0 & 0 \\ 0 & \lambda & 0 \\ 0 & 0 & \lambda \end{bmatrix}.$$

Second, the minimal polynomial p of $\widetilde{F}$ divides but does not equal the minimal polynomial of F, so $p(F) \neq 0 = p(\widetilde{F})$. Since $\langle F \rangle \ni p(F) \leftrightsquigarrow p(\widetilde{F}) = 0 \in \langle \widetilde{F} \rangle$, it follows from indistinguishability (as in the proof of Proposition 10.3.2) that $p(F)$ is $\{F\}$-vanishing, so $\{F\}$ is $(1,1)$-quantum-vanishing. Thus any $(1,1)$-quantum-nonvanishing $\{F\}$ and $\{G\}$ are diagonalizable, so, for any such pair, indistinguishability does imply similarity. Theorem 10.4.2 below generalizes this statement to simultaneous similarity. Note that $(1,1)$-quantum-nonvanishing does not necessarily imply full quantum-nonvanishing at all arities. So, instead of Theorem 10.3.1, our proof uses the following theorem of Kaplansky (see also [RY15, Theorem 2.1]). Say that $F \in \mathbb{K}^{q \times q}$ has *singleton spectrum* if F has (up to multiplicity) only one distinct eigenvalue.

Theorem 10.4.1 (Kaplansky [Kap72])**.** *Suppose $\mathcal{A} \subset \mathbb{K}^{q \times q}$ is closed under matrix product and every $A \in \mathcal{A}$ has singleton spectrum. Then $\mathcal{A}$ is simultaneously triangularizable under some $T \in \mathrm{GL}_q$.*

Say $\mathcal{F} \subset \mathcal{V}(\mathbb{K}^q)$ is $(1,1)$-*trivial* if ${}_1\langle \mathcal{F} \rangle_1 \subset \mathrm{span}(I)$ (i.e. is as small as possible, as the wire gadget is always present).

Corollary 10.4.1. *Let $\mathcal{F} \subset \mathbb{K}^{q \times q}$ be $(1,1)$-quantum-nonvanishing. If every $F \in {}_1\langle \mathcal{F} \rangle_1$ has singleton spectrum, then $\mathcal{F}$ is $(1,1)$-trivial.*

Proof. Applying Kaplansky's theorem to ${}_1\langle \mathcal{F} \rangle_1$, we may transform $\mathcal{F}$ so that every matrix in ${}_1\langle \mathcal{F} \rangle_1$ is upper triangular, with constant diagonal. This does not change whether $\mathcal{F}$ is quantum-vanishing. Suppose $\mathcal{F} \ni F \notin \mathrm{span}(I)$, with constant λ on the diagonal. Then $F - \lambda I \in \langle \mathcal{F} \rangle$ is nonzero and strictly upper triangular, so $(F - \lambda I)F'$ is strictly upper triangular for every $F' \in {}_1\langle \mathcal{F} \rangle_1$. But every connected $\langle\mathcal{F}\rangle$-grid containing $F-\lambda I$ is a cycle formed by a contraction between $F-\lambda I$ and a path with tensor $F' \in {}_1\langle\mathcal{F}\rangle_1$, with Holant value $\operatorname{tr}((F-\lambda I)F') = 0$. Therefore $F-\lambda I$ is $\mathcal{F}$-vanishing, contradicting $(1,1)$-quantum-nonvanishing. $\square$

Any $\mathcal{F}$ failing the condition of Kaplansky's theorem satisfies the condition of the following domain separation lemma, somewhat analogous to the proof of Lemma 8.3.2.

Lemma 10.4.1. *Let $\mathcal{F}, \mathcal{G} \subset \mathcal{V}(\mathbb{K}^q)$ be Bi-Holant-indistinguishable and $(1,1)$-quantum-nonvanishing. Either $\mathcal{F}$ and $\mathcal{G}$ are $(1,1)$-trivial and ${}_1\langle\mathcal{F}\rangle_1 = {}_1\langle\mathcal{G}\rangle_1$, or there is a nontrivial partition $(X, \overline{X})$ of $[q]$ and $T, U \in \mathrm{GL}_q$ such that $\langle T\,\mathcal{F}\rangle \ni I_X^{\uparrow}, I_{\overline{X}}^{\uparrow} \leftrightsquigarrow I_X^{\uparrow}, I_{\overline{X}}^{\uparrow} \in \langle U\,\mathcal{G}\rangle$.*

Furthermore, suppose there are ${}_1\langle\mathcal{F}\rangle_1 \ni F \leftrightsquigarrow G \in {}_1\langle\mathcal{G}\rangle_1$ that do not have singleton spectrum and have block forms $\left[\begin{smallmatrix} * & 0 \\ * & 0 \end{smallmatrix}\right]$ and $\left[\begin{smallmatrix} * & * \\ 0 & 0 \end{smallmatrix}\right]$ respectively, with the first block indexed by $\Delta \subset [q]$. Then we may choose, under the same blocks, $T = \left[\begin{smallmatrix} * & 0 \\ * & * \end{smallmatrix}\right]$ and $U = \left[\begin{smallmatrix} * & * \\ 0 & * \end{smallmatrix}\right]$ so that $X = [x] \subset \Delta$.

Proof. If $\mathcal{F}$ and $\mathcal{G}$ are $(1,1)$-trivial, then for every ${}_1\langle\mathcal{F}\rangle_1 \ni F = \lambda_F I \leftrightsquigarrow \lambda_G I = G \in {}_1\langle\mathcal{G}\rangle_1$, we have $q\lambda_F = \operatorname{tr}(F) = \operatorname{tr}(G) = q\lambda_G$, hence $\lambda_F = \lambda_G$, so $F = G$. Thus ${}_1\langle\mathcal{F}\rangle_1 = {}_1\langle\mathcal{G}\rangle_1$.

Otherwise, Corollary 10.4.1 asserts that there are ${}_1\langle\mathcal{F}\rangle_1 \ni F \leftrightsquigarrow G \in {}_1\langle\mathcal{G}\rangle_1$ such that one of F or G does not have singleton spectrum. By indistinguishability, $\operatorname{tr}(F^k) = \operatorname{tr}(G^k)$ for every $k \geq 0$. Thus F and G have the same multiset of eigenvalues. In particular, F and G share some eigenvalue λ with the same (algebraic) multiplicity. We claim that F and G must have the same minimal polynomial. Otherwise, suppose WLOG that the minimal polynomial of F does not divide the minimal polynomial p_G of G. By Proposition 10.3.2, $\mathcal{F}$ and $\mathcal{G}$ are $(1,1)$-weakly-isomorphic, but $p_G(F) \neq 0 = p_G(G)$ and ${}_1\langle\mathcal{F}\rangle_1 \ni p_G(F) \leftrightsquigarrow p_G(G) \in {}_1\langle\mathcal{G}\rangle_1$, a contradiction.

Choose T and U to be the bases under which F and G are in Jordan normal form, respectively. Then, since λ has the same multiplicity in F and G, we can define $X \subset [q]$ such that $F|_{\overline{X}}$ and $G|_{\overline{X}}$ are the union of the λ-blocks of F and G, respectively. Since F and G do not have singleton spectrum, $X \subset [q]$ is nontrivial. Then choose sufficiently large r such that

$$(F-\lambda I)^r|_{\overline{X}} = (G-\lambda I)^r|_{\overline{X}} = 0. \tag{10.4.1}$$

Then ${}_1\langle\mathcal{F}\rangle_1 \ni (F-\lambda I)^r \leftrightsquigarrow (G-\lambda I)^r \in {}_1\langle\mathcal{G}\rangle_1$ are both supported only on X, so it follows as above from $(1,1)$-weakly-isomorphic that $(F-\lambda I)^r|_X$ and $(G-\lambda I)^r|_X \in \mathbb{K}^{X\times X}$ have the same minimal polynomial p. Furthermore, $(F-\lambda I)^r|_X$ and $(G-\lambda I)^r|_X$ have no 0-eigenvalues, so p has a nonzero constant term cI_X. Expanding $p - cI_X$ removes all instances of I_X, so we can view $p - cI_X$ as a polynomial on full $q \times q$ matrices. Now, by (10.4.1),

$$_1\langle \mathcal{F} \rangle_1 \ni I_X^{\uparrow} = -\frac{1}{c}(p - cI_X)\big((F - \lambda I)^r\big) \leftrightsquigarrow -\frac{1}{c}(p - cI_X)\big((G - \lambda I)^r\big) = I_X^{\uparrow} \in {_1\langle \mathcal{G} \rangle_1}$$

and $_1\langle \mathcal{F} \rangle_1 \ni I_{\overline{X}}^{\uparrow} = I - I_X^{\uparrow} \leftrightsquigarrow I - I_X^{\uparrow} = I_{\overline{X}}^{\uparrow} \in {_1\langle \mathcal{G} \rangle_1}$.

For the second claim, it suffices to show that $F = \left[\begin{smallmatrix} * & 0 \\ * & 0 \end{smallmatrix}\right]$ can be put in Jordan normal form by T of the form $\left[\begin{smallmatrix} * & 0 \\ * & * \end{smallmatrix}\right]$; the claim about G and U follows by transposed reasoning. Note that $\{e_{|\Delta|+1}, \dots, e_q\}$ is a set of linearly independent 0-eigenvectors of F. We can always choose a T whose final $q - |\Delta|$ columns are $\{e_{|\Delta|+1}, \dots, e_q\}$, giving T the desired block form and, with $\lambda = 0$ in the proof above, ensuring that $X = [x] \subset \Delta$. □

For a word $w \in \Gamma_{\mathcal{F}}$, construct $w_{\mathcal{F}\to\mathcal{G}} \in \Gamma_{\mathcal{G}}$ by replacing every character $F \in \mathcal{F}$ in w by the corresponding $G \in \mathcal{G}$. We now obtain our characterization of simultaneously similarity of quantum-nonvanishing sets of matrices. We will use Lemma 10.4.1 to partition the domain $[q]$ as finely as possible, at which point the nonzero blocks of matrices in $_1\langle \mathcal{F} \rangle_1$ and $_1\langle \mathcal{G} \rangle_1$ are invertible and can themselves be used to construct the transformation T.

Theorem 10.4.2. *Let $\mathcal{F}, \mathcal{G} \subset \mathbb{K}^{q\times q}$ be $(1,1)$-quantum-nonvanishing. Then $\operatorname{tr}(w) = \operatorname{tr}(w_{\mathcal{F}\to\mathcal{G}})$ for every word $w \in \Gamma_{\mathcal{F}}$ if and only if there is a $T \in \mathrm{GL}_q$ such that $TFT^{-1} = G$ for every $\mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}$.*

Proof. We only need ($\Longrightarrow$). The assumption is equivalent to Bi-Holant-indistinguishability of $\mathcal{F}$ and $\mathcal{G}$. So, unless we have $\mathcal{F} \subset {_1\langle \mathcal{F} \rangle_1} = {_1\langle \mathcal{G} \rangle_1} \supset \mathcal{G}$ and are already done, Lemma 10.4.1 gives a nontrivial partition $(X, \overline{X})$ of $[q]$ such that, after suitable transformations, $\langle \mathcal{F} \rangle \ni I_X^{\uparrow}, I_{\overline{X}}^{\uparrow} \leftrightsquigarrow I_X^{\uparrow}, I_{\overline{X}}^{\uparrow} \in \langle \mathcal{G} \rangle$.

In general, suppose $\mathcal{F} \ni I_{X_1}^{\uparrow}, \dots I_{X_s}^{\uparrow} \leftrightsquigarrow I_{X_1}^{\uparrow}, \dots I_{X_s}^{\uparrow} \in \mathcal{G}$ for a partition $(X_1, \dots, X_s)$ of $[q]$. We will show that every subdomain is either $(1,1)$-trivial or can be further decomposed into smaller subdomains. By Propositions 8.2.3 and 10.3.3, each $\langle \mathcal{F} \rangle_{X_i}$ and $\langle \mathcal{G} \rangle_{X_i}$ are Bi-Holant-indistinguishable and $(1,1)$-quantum-nonvanishing. If any $\langle \mathcal{F} \rangle_{X_i}$ and $\langle \mathcal{G} \rangle_{X_i}$ are $(1,1)$-nontrivial, then by Lemma 10.4.1 there are $T, U \in \mathrm{GL}(\mathbb{K}^{X_i})$ and nontrivial $Y_i \subset X_i$ such that

$$\left\langle T\, \langle \mathcal{F} \rangle_{X_i} \right\rangle \ni I_{Y_i}^{\uparrow X_i}, I_{X_i \setminus Y_i}^{\uparrow X_i} \leftrightsquigarrow I_{Y_i}^{\uparrow X_i}, I_{X_i \setminus Y_i}^{\uparrow X_i} \in \left\langle U\, \langle \mathcal{G} \rangle_{X_i} \right\rangle.$$

Define $T^{\uparrow} := I_{X_1} \oplus \dots \oplus I_{X_{i-1}} \oplus T \oplus I_{X_{i+1}} \oplus \dots \oplus I_{X_s} \in \mathrm{GL}_q$ and replace $\mathcal{F}$ with $T^{\uparrow}\, \mathcal{F}$. This replaces $\langle \mathcal{F} \rangle_{X_i}$ with $\langle T^{\uparrow}\, \mathcal{F} \rangle_{X_i} = T\, \langle \mathcal{F} \rangle_{X_i}$ (by Proposition 4.1.2) while preserving $I_{X_1}^{\uparrow}, \dots, I_{X_s}^{\uparrow}$. Now

$I_{Y_i}^{\uparrow X_i} \in \left\langle \langle \mathcal{F} \rangle_{X_i} \right\rangle$ and we still have $I_{X_i} \in \langle \mathcal{F} \rangle$, so Proposition 8.2.2 gives $I_{Y_i}^{\uparrow} = (I_{Y_i}^{\uparrow X_i})^{\uparrow} \in \langle \mathcal{F} \rangle$. Similarly, $I_{X_i \setminus Y_i}^{\uparrow} \in \langle \mathcal{F} \rangle$ and, after transforming $\mathcal{G}$ by $U^{\uparrow}$, we obtain $I_{Y_i}^{\uparrow}, I_{X_i \setminus Y_i}^{\uparrow} \in \langle \mathcal{G} \rangle$, so we have refined the partition of $[q]$ to $(X_1, \dots, Y_i, X_i \setminus Y_i, \dots, X_s)$.

Let this process stabilize at a maximal partition $(X_1, \dots, X_m)$ of $[q]$. At this point, Lemma 10.4.1 asserts that every $\langle F \rangle_{X_i}$ and $\langle G \rangle_{X_i}$ are $(1,1)$-trivial and satisfy ${}_1\left\langle \langle \mathcal{F} \rangle_{X_i} \right\rangle_1 = {}_1\left\langle \langle \mathcal{G} \rangle_{X_i} \right\rangle_1$. We proceed to inductively transform $\mathcal{F}$ into $\mathcal{G}$. Suppose that ${}_1\left\langle \, \langle \mathcal{F} \rangle_{X_1 \cup \dots \cup X_{p-1}} \, \right\rangle_1 = {}_1\left\langle \, \langle \mathcal{G} \rangle_{X_1 \cup \dots \cup X_{p-1}} \, \right\rangle_1$. Use $I_{X_1 \cup \dots \cup X_p}^{\uparrow} = \sum_{i=1}^{p} I_{X_i}^{\uparrow}$ to isolate $\langle \mathcal{F} \rangle_{X_1 \cup \dots \cup X_p}$ and $\langle \mathcal{G} \rangle_{X_1 \cup \dots \cup X_p}$; by Proposition 8.2.3 and Proposition 10.3.3, $\langle \mathcal{F} \rangle_{X_1 \cup \dots \cup X_p}$ and $\langle \mathcal{G} \rangle_{X_1 \cup \dots \cup X_p}$ are Bi-Holant-indistinguishable and $(1,1)$-quantum-nonvanishing. Every ${}_1\left\langle \, \langle \mathcal{F} \rangle_{X_1 \cup \dots \cup X_p} \, \right\rangle_1 \ni F \leftrightsquigarrow G \in {}_1\left\langle \, \langle \mathcal{G} \rangle_{X_1 \cup \dots \cup X_p} \, \right\rangle_1$ are of the form

$$\left[\begin{array}{cccc|c} \lambda_1 I_{X_1} & F_{1,2} & \dots & F_{1,p-1} & F_{1,p} \\ F_{2,1} & \lambda_2 I_{X_2} & \dots & F_{2,p-1} & F_{2,p} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ F_{p-1,1} & F_{p-1,2} & \dots & \lambda_{p-1} I_{X_{p-1}} & F_{p-1,p} \\ \hline F_{p,1} & F_{p,2} & \dots & F_{p,p-1} & \lambda_p I_{X_p} \end{array}\right] \leftrightsquigarrow \left[\begin{array}{cccc|c} \lambda_1 I_{X_1} & F_{1,2} & \dots & F_{1,p-1} & G_{1,p} \\ F_{2,1} & \lambda_2 I_{X_2} & \dots & F_{2,p-1} & G_{2,p} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ F_{p-1,1} & F_{p-1,2} & \dots & \lambda_{p-1} I_{X_{p-1}} & G_{p-1,p} \\ \hline G_{p,1} & G_{p,2} & \dots & G_{p,p-1} & \lambda_p I_{X_p} \end{array}\right] \tag{10.4.2}$$

where $F_{i,j} := F|_{X_i,X_j}$ and $G_{i,j} := G|_{X_i,X_j}$. Extending Definition 8.2.3, for $i, j \in [m]$ with $i \neq j$, let

$$F_{i,j}^{\uparrow} = I_{X_i}^{\uparrow} F I_{X_j}^{\uparrow} \in {}_1\langle \mathcal{F} \rangle_1$$

be the matrix with $F_{i,j}$ in the (X_i, X_j) block and 0 in the other blocks. Since $\mathcal{F}$ and $\mathcal{G}$ are $(1,1)$-weakly-isomorphic by Proposition 10.3.2 and $F_{i,j}^{\uparrow} \leftrightsquigarrow G_{i,j}^{\uparrow}$, we have

$$F_{i,j} = 0 \iff G_{i,j} = 0. \tag{10.4.3}$$

If any $F_{i,j} \neq 0$ then, by $(1,1)$-quantum-nonvanishing of $\mathcal{F}$, there is a $\widehat{F_{i,j}^{\uparrow}} =: \left(\widehat{F_{k,\ell}}\right)_{k,\ell \in [p]} \in {}_1\langle \mathcal{F} \rangle_1$ such that

$$\begin{aligned} 0 \neq \left\langle F_{i,j}^{\uparrow}, \widehat{F_{i,j}^{\uparrow}} \right\rangle &= \operatorname{tr}\left(F_{i,j}^{\uparrow} \widehat{F_{i,j}^{\uparrow}}\right) \\ &= \operatorname{tr}\left(\begin{bmatrix} 0 & \dots & 0 & \dots\dots & 0 \\ \vdots & & \vdots & & \vdots \\ 0 & \dots & F_{i,j} & \dots\dots & 0 \\ \vdots & & \vdots & & \vdots \\ 0 & \dots & 0 & \dots\dots & 0 \end{bmatrix} \begin{bmatrix} * & \dots\dots & * & \dots & * \\ \vdots & & \vdots & & \vdots \\ * & \dots\dots & \widehat{F_{j,i}} & \dots & * \\ \vdots & & \vdots & & \vdots \\ * & \dots\dots & * & \dots & * \end{bmatrix}\right) = \operatorname{tr}\left(F_{i,j} \widehat{F_{j,i}}\right). \end{aligned} \tag{10.4.4}$$

Note that $F_{i,j}\widehat{F_{j,i}}$ is the (X_i, X_i)-block of $F^{\uparrow}_{i,j}\widehat{F^{\uparrow}_{i,j}} \in \langle\mathcal{F}\rangle$ in (10.4.4), so $F_{i,j}\widehat{F_{j,i}} \in \langle\mathcal{F}\rangle_{X_i}$. But $\langle\mathcal{F}\rangle_{X_i}$ is $(1,1)$-trivial, so $F_{i,j}\widehat{F_{j,i}} = \lambda_{F,i,j} I_{X_i}$ for some $\lambda_{F,i,j} \neq 0$. We simultaneously have

$$
\begin{aligned}
0 \neq \operatorname{tr}\left(F^{\uparrow}_{i,j}\widehat{F^{\uparrow}_{i,j}}\right) &= \operatorname{tr}\left(\widehat{F^{\uparrow}_{i,j}}F^{\uparrow}_{i,j}\right)\\
&= \operatorname{tr}\left(\begin{bmatrix} * & \cdots\cdots & * & \cdots & * \\ \vdots & & \vdots & & \vdots \\ * & \cdots\cdots & \widehat{F_{j,i}} & \cdots & * \\ \vdots & & \vdots & & \vdots \\ * & \cdots\cdots & * & \cdots & * \end{bmatrix}\begin{bmatrix} 0 & \cdots & 0 & \cdots\cdots & 0 \\ \vdots & & \vdots & & \vdots \\ 0 & \cdots & F_{i,j} & \cdots\cdots & 0 \\ \vdots & & \vdots & & \vdots \\ 0 & \cdots & 0 & \cdots\cdots & 0 \end{bmatrix}\right) = \operatorname{tr}\left(\widehat{F_{j,i}}F_{i,j}\right),
\end{aligned}
$$

and the $(1,1)$-triviality of $\mathcal{F}_{X_j}$ gives $\widehat{F_{j,i}}F_{i,j} = \lambda'_{F,i,j} I_{X_j}$ for $\lambda'_{F,i,j} \neq 0$. Hence $F_{i,j}$ and $\widehat{F_{j,i}}$ are both left and right-invertible, so must be square, giving $|X_i| = |X_j|$ and $\lambda_{F,i,j} = \lambda'_{F,i,j}$. On the $\mathcal{G}$ side, (10.4.3) gives $G_{i,j} \neq 0$ as well, so let $\widehat{F^{\uparrow}_{i,j}} \leftrightsquigarrow \widehat{G^{\uparrow}_{i,j}} = (\widehat{G_{k,\ell}})_{k,\ell\in[p]}$. In general, if $F, \widetilde{F} \in {}_1\langle\mathcal{F}\rangle_1$, then $F_{i,j}\widetilde{F}_{j,k}$ is the (X_i, X_k)-block of $F^{\uparrow}_{i,j}\widetilde{F}^{\uparrow}_{j,k} \in {}_1\langle\mathcal{F}\rangle_1$. In particular, $F_{i,j}\widetilde{F}_{j,i} \in {}_1\big\langle\langle\mathcal{F}\rangle_{X_i}\big\rangle_1$. Then, since ${}_1\big\langle\langle\mathcal{F}\rangle_{X_i}\big\rangle_1 = {}_1\big\langle\langle\mathcal{G}\rangle_{X_i}\big\rangle_1$,

$$
F_{i,j}\widetilde{F}_{j,i} = G_{i,j}\widetilde{G}_{j,i} \text{ for every } {}_1\langle\mathcal{F}\rangle_1 \ni F, \widetilde{F} \leftrightsquigarrow G, \widetilde{G} \in {}_1\langle\mathcal{G}\rangle_1 . \tag{10.4.5}
$$

In particular,

$$
F_{i,j}\widehat{F_{j,i}} = \widehat{F_{j,i}}F_{i,j} = G_{i,j}\widehat{G_{j,i}} = \widehat{G_{j,i}}G_{i,j} = \lambda_{F,i,j} I_{|X_i|}. \tag{10.4.6}
$$

For $k < p$, fix $F^{(k)} \in {}_1\langle\mathcal{F}\rangle_1$ with $F^{(k)}_{k,p} \neq 0$, if any such $F^{(k)}$ exists, and let $F^{(k)} \leftrightsquigarrow G^{(k)}$. Define

$$
T_p = I_{X_p} \text{ and } T_k = \begin{cases} \lambda^{-1}_{F^{(k)},k,p} G^{(k)}_{k,p}\widehat{F^{(k)}_{p,k}} & \exists F' \in {}_1\langle\mathcal{F}\rangle_1 \text{ such that } F'_{k,p} \neq 0 \\ I_{X_k} & \text{otherwise} \end{cases} \in \mathbb{K}^{X_k \times X_k} \tag{10.4.7}
$$

and $T := \bigoplus_{k=1}^{p} T_k$. By (10.4.6), T is invertible and $T^{-1} = \bigoplus_{k=1}^{p} T_k^{-1}$, where

$$
T_p^{-1} = I_{X_p} \text{ and } T_k^{-1} = \begin{cases} \lambda^{-1}_{F^{(k)},k,p} F^{(k)}_{k,p}\widehat{G^{(k)}_{p,k}} & \exists F' \in {}_1\langle\mathcal{F}\rangle_1 \text{ such that } F'_{k,p} \neq 0 \\ I_{X_k} & \text{otherwise.} \end{cases} \tag{10.4.8}
$$

We claim that $TFT^{-1} = G$ for every ${}_1\big\langle\, \langle\mathcal{F}\rangle_{X_1\cup\ldots\cup X_p}\,\big\rangle_1 \ni F \leftrightsquigarrow G \in {}_1\big\langle\, \langle\mathcal{G}\rangle_{X_1\cup\ldots\cup X_p}\,\big\rangle_1$. This is equivalent to $T_i F_{i,j} T_j^{-1} = G_{i,j}$ for arbitrary $i, j \leq p$. If $F_{i,j} = 0$ then $G_{i,j} = 0$ by (10.4.3) and we are done. Otherwise, we consider several cases.

1. If $i = j$, then $F_{i,i} = G_{i,i} = \lambda_i$ by (10.4.2), so $T_i F_{i,i} T_i^{-1} = G_{i,i}$.

2. If $i \neq p = j$, then $F_{i,j} = F_{i,p} \neq 0$ implies that $T_i = \lambda^{-1}_{F^{(i)},i,p} G^{(i)}_{i,p} \widehat{F^{(i)}_{p,i}}$, so, applying (10.4.5) followed by (10.4.6),

$$T_i F_{i,p} T_p^{-1} = \lambda^{-1}_{F^{(i)},i,p} G^{(i)}_{i,p} \widehat{F^{(i)}_{p,i}} F_{i,p} I_{X_p} = \lambda^{-1}_{F^{(i)},i,p} G^{(i)}_{i,p} \widehat{G^{(i)}_{p,i}} G_{i,p} = G_{i,p}.$$

3. If $i = p \neq j$, then $F_{i,j} = F_{p,j} \neq 0$ implies that $\widehat{F_{j,p}} \neq 0$ by (10.4.6), so $T_j^{-1} = \lambda^{-1}_{F^{(j)},j,p} F^{(j)}_{j,p} \widehat{G^{(j)}_{p,j}}$ and, applying (10.4.5) followed by (10.4.6),

$$T_p F_{p,j} T_j^{-1} = \lambda^{-1}_{F^{(j)},j,p} I_{X_p} F_{p,j} F^{(j)}_{j,p} \widehat{G^{(j)}_{p,j}} = \lambda^{-1}_{F^{(j)},j,p} G_{p,j} G^{(j)}_{j,p} \widehat{G^{(j)}_{p,j}} = G_{p,j}.$$

4. If i, j, p are all distinct, then $F_{i,j} = G_{i,j}$ by induction, so if $T_i = T_j^{-1} = I$ then we are done. Otherwise, there are two possibilities.

   a) If $T_i \neq I$ then by (10.4.7) there is a $F'_{i,p} \neq 0$. Now $\widehat{F_{j,i}}$ and $F'_{i,p}$ are invertible by (10.4.6), so $\widehat{F_{j,i}} F'_{i,p} \neq 0$ is the (X_j, X_p)-block of $\widehat{F_{j,i}}^{\uparrow} (F'_{i,p})^{\uparrow}$. Thus $T_j^{-1} = \lambda_{F^{(j)},j,p} F^{(j)}_{j,p} \widehat{G^{(j)}_{p,j}}$ by (10.4.8).

   b) If $T_j^{-1} \neq I$ then by (10.4.8) there is a $F'_{j,p} \neq 0$. Now $F_{i,j}$ and $F'_{j,p}$ are invertible by (10.4.6), so $F_{i,j} F'_{j,p} \neq 0$ is the (X_i, X_p)-block of $F^{\uparrow}_{i,j} (F'_{j,p})^{\uparrow}$. Thus $T_i = \lambda_{F^{(i)},i,p} G^{(i)}_{i,p} \widehat{F^{(i)}_{p,i}}$ by (10.4.7).

   In either case, both T_i and T_j^{-1} fall into the respective first cases in (10.4.7) and (10.4.8). Therefore, by reasoning similar to (10.4.5), followed by (10.4.6),

$$T_i F_{i,j} T_j^{-1} = \lambda^{-1}_{F^{(i)},i,p} \lambda^{-1}_{F^{(j)},j,p} G^{(i)}_{i,p} \widehat{F^{(i)}_{p,i}} F_{i,j} F^{(j)}_{j,p} \widehat{G^{(j)}_{p,j}} = \lambda^{-1}_{F^{(i)},i,p} \lambda^{-1}_{F^{(j)},j,p} G^{(i)}_{i,p} \widehat{G^{(i)}_{p,i}} G_{i,j} G^{(j)}_{j,p} \widehat{G^{(j)}_{p,j}} = G_{i,j}.$$

Now, after transforming $\mathcal{F}$ by $T \oplus I_{X_{p+1}} \oplus \ldots \oplus I_{X_m}$, we have ${}_1\big\langle \langle \mathcal{F} \rangle_{X_1 \cup \ldots \cup X_p} \big\rangle_1 = {}_1\big\langle \langle \mathcal{G} \rangle_{X_1 \cup \ldots \cup X_p} \big\rangle_1$. After similar transforms at each level of the induction, we obtain $\mathcal{F} \subset {}_1\langle \mathcal{F} \rangle_1 = {}_1\langle \mathcal{G} \rangle_1 \supset \mathcal{G}$. □

10.4.2 The bipartite case

We begin with the following analogue of Lemma 8.3.1, with a similar proof.

Lemma 10.4.2. *If $\mathcal{F}, \mathcal{G} \subset \mathcal{V}(\mathbb{K}^q)$ are Bi-Holant-indistinguishable and quantum-nonvanishing, then there exists an $H \in \operatorname{Stab}(\langle \mathcal{F} \rangle \oplus \langle \mathcal{G} \rangle)$ with $H|_{\mathcal{F},\mathcal{G}} \neq 0$ or $H|_{\mathcal{G},\mathcal{F}} \neq 0$.*

Proof. By Lemma 10.3.1, Theorem 10.3.1 applies to $\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle$. Assume that every $H \in \operatorname{Stab}(\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle)$ satisfies $H|_{\mathcal{F},\mathcal{G}} = H|_{\mathcal{G},\mathcal{F}} = 0$ (i.e. is block-diagonal). Then

$$I_{\mathcal{F}} \oplus 2I_{\mathcal{G}} = \begin{bmatrix} I & 0 \\ 0 & 2I \end{bmatrix} \in \mathcal{V}(\mathbb{K}^{2q})$$

satisfies $H(I_{\mathcal{F}} \oplus 2I_{\mathcal{G}})H^{-1} = I_{\mathcal{F}} \oplus 2I_{\mathcal{G}}$ for every $H \in \operatorname{Stab}(\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle)$, so, by Theorem 10.3.1, $I_{\mathcal{F}} \oplus 2I_{\mathcal{G}} \in {}_1\langle\langle\mathcal{F}\rangle \oplus \langle\mathcal{G}\rangle\rangle_1$. But Proposition 5.2.1 gives

$$\langle\mathcal{F}\rangle \ni I_{\mathcal{F}} = (I_{\mathcal{F}} \oplus 2I_{\mathcal{G}})|_{\mathcal{F}} \leftrightsquigarrow (I_{\mathcal{F}} \oplus 2I_{\mathcal{G}})|_{\mathcal{G}} = 2I_{\mathcal{G}} \in \langle\mathcal{G}\rangle\,,$$

violating indistinguishability, as $\operatorname{tr}(I_{\mathcal{F}}) = q \neq 2q = \operatorname{tr}(2I_{\mathcal{G}})$. □

The next step of the proof of the orthogonal converse in Chapter 8 is the inductive Lemma 8.3.2. Here, without assuming that $\mathcal{F}$ and $\mathcal{G}$ contain a subdomain restrictor $I_Z^{\uparrow}$ to begin with, we cannot assume that $\mathcal{F}\,|_Z$ and $\mathcal{G}\,|_Z$ are quantum-nonvanishing to apply induction. So, for now, we do not prove a result analogous to Lemma 8.3.2, and instead prove a result analogous to the final steps of the proof of the orthogonal converse, with $I_Z^{\uparrow}$ playing the role of the diagonal matrix D.

Lemma 10.4.3. *If $\mathcal{F}\,|\,\mathcal{F}'$ and $\mathcal{G}\,|\,\mathcal{G}'$ on domain $[q]$ are Bi-Holant-indistinguishable and quantum-nonvanishing, then there exist $\varnothing \neq Z \subset [q]$ and $T_1, T_2 \in \mathrm{GL}_q$ such that, up to swapping the roles of $\mathcal{F}\,|\,\mathcal{F}'$ and $\mathcal{G}\,|\,\mathcal{G}'$, after transforming $\mathcal{F}\,|\,\mathcal{F}'$ by T_1 and $\mathcal{G}\,|\,\mathcal{G}'$ by T_2, every ${}_n\langle\mathcal{F}\,|\,\mathcal{F}'\rangle_0 \ni F \leftrightsquigarrow G \in {}_n\langle\mathcal{G}\,|\,\mathcal{G}'\rangle_0$ and ${}_0\langle\mathcal{F}\,|\,\mathcal{F}'\rangle_n \ni F' \leftrightsquigarrow G' \in {}_0\langle\mathcal{G}\,|\,\mathcal{G}'\rangle_n$ for every $n \geq 1$ satisfy*

$$(I_Z^{\uparrow})^{\otimes n} F = G \text{ and } F' = G'(I_Z^{\uparrow})^{\otimes n}. \tag{10.4.9}$$

Proof. Lemma 10.4.2 gives an $H \in \operatorname{Stab}(\langle\mathcal{F}\,|\,\mathcal{F}'\rangle \oplus \langle\mathcal{G}\,|\,\mathcal{G}'\rangle)$ with, WLOG, $H|_{\mathcal{G},\mathcal{F}} \neq 0$. Choose $T_1, T_2 \in \mathrm{GL}_q$ so that $T_2 H_{\mathcal{G},\mathcal{F}} T_1^{-1} = I_Z^{\uparrow} \in \mathbb{K}^{q\times q}$ for some $Z \subset [q]$ with $|Z| = \operatorname{rank}(H_{\mathcal{G},\mathcal{F}}) > 0$. By Theorem 10.3.1, H satisfies $H \cdot K = K$ for every $K \in \langle\langle\mathcal{F}\,|\,\mathcal{F}'\rangle \oplus \langle\mathcal{G}\,|\,\mathcal{G}'\rangle\rangle$. Transform $\mathcal{F}\,|\,\mathcal{F}'$ by T_1 and $\mathcal{G}\,|\,\mathcal{G}'$ by T_2. By Proposition 4.1.2, this replaces $\langle\langle\mathcal{F}\,|\,\mathcal{F}'\rangle \oplus \langle\mathcal{G}\,|\,\mathcal{G}'\rangle\rangle$ with

$$\big\langle (T_1 \oplus T_2)\big(\big\langle\mathcal{F}\,|\,\mathcal{F}'\big\rangle \oplus \big\langle\mathcal{G}\,|\,\mathcal{G}'\big\rangle\big)\big\rangle = (T_1 \oplus T_2)\big\langle\big\langle\mathcal{F}\,|\,\mathcal{F}'\big\rangle \oplus \big\langle\mathcal{G}\,|\,\mathcal{G}'\big\rangle\big\rangle\,.$$

Now, after the transformation by $(T_1 \oplus T_2)$,

$$\widetilde{H} := (T_1 \oplus T_2)H(T_1 \oplus T_2)^{-1} = \begin{bmatrix} T_1 & 0 \\ 0 & T_2 \end{bmatrix} \begin{bmatrix} * & * \\ H_{\mathcal{G},\mathcal{F}} & * \end{bmatrix} \begin{bmatrix} T_1^{-1} & 0 \\ 0 & T_2^{-1} \end{bmatrix} = \begin{bmatrix} * & * \\ I_Z^{\uparrow} & * \end{bmatrix}$$

satisfies $\widetilde{H} \cdot K = K$ for every $K \in \langle\langle \mathcal{F} \,|\, \mathcal{F}' \rangle \oplus \langle \mathcal{G} \,|\, \mathcal{G}' \rangle\rangle$.

Let ${}_n\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_0 \ni F \leftrightsquigarrow G \in {}_n\langle \mathcal{G} \,|\, \mathcal{G}' \rangle_0$. Then $(F \otimes I) \oplus (G \otimes I) \in {}_{n+1}\langle\langle \mathcal{F} \,|\, \mathcal{F}' \rangle \oplus \langle \mathcal{G} \,|\, \mathcal{G}' \rangle\rangle_1$, so $\widetilde{H}^{\otimes n+1}((F \otimes I) \oplus (G \otimes I)) = \big((F \otimes I) \oplus (G \otimes I)\big)\widetilde{H}$, which, by Proposition 8.2.1 and (8.2.1), has $(V(\mathcal{F}), V(\mathcal{G}))$-block matrix form

$$\begin{bmatrix} * & * & \dots & * \\ \vdots & \vdots & \iddots & \vdots \\ * & * & \dots & * \\ (I_Z^{\uparrow})^{\otimes n+1} & * & \dots & * \end{bmatrix} \begin{bmatrix} F \otimes I & 0 \\ 0 & 0 \\ \vdots & \vdots \\ 0 & 0 \\ 0 & G \otimes I \end{bmatrix} = \begin{bmatrix} F \otimes I & 0 \\ 0 & 0 \\ \vdots & \vdots \\ 0 & 0 \\ 0 & G \otimes I \end{bmatrix} \begin{bmatrix} * & * \\ I_Z^{\uparrow} & * \end{bmatrix}. \tag{10.4.10}$$

The bottom left block of (10.4.10) gives the second equality in

$$\big((I_Z^{\uparrow})^{\otimes n} F\big) \otimes I_Z^{\uparrow} = (I_Z^{\uparrow})^{\otimes n+1}(F \otimes I) = (G \otimes I) I_Z^{\uparrow} = G \otimes I_Z^{\uparrow} \tag{10.4.11}$$

(see Figure 10.1), which implies that $(I_Z^{\uparrow})^{\otimes n} F = G$.

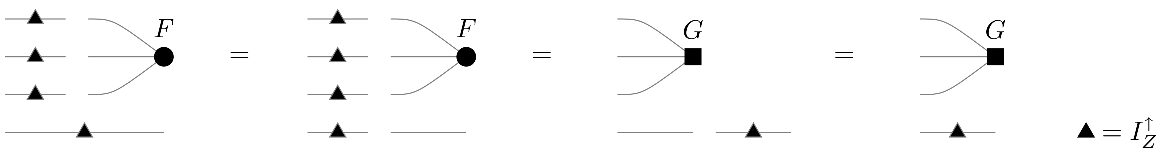

Figure 10.1: Illustrating (10.4.11) for $n = 3$.

Similarly, if ${}_0\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_n \ni F' \leftrightsquigarrow G' \in {}_0\langle \mathcal{G} \,|\, \mathcal{G}' \rangle_n$, then $(F' \otimes I) \oplus (G' \otimes I) \in {}_1\langle\langle \mathcal{F} \,|\, \mathcal{F}' \rangle \oplus \langle \mathcal{G} \,|\, \mathcal{G}' \rangle\rangle_{n+1}$, so $\widetilde{H}\,((F' \otimes I) \oplus (G' \otimes I)) = \big((F' \otimes I) \oplus (G' \otimes I)\big)\widetilde{H}^{\otimes n+1}$, or equivalently

$$\begin{bmatrix} * & * \\ I_Z^{\uparrow} & * \end{bmatrix} \begin{bmatrix} F' \otimes I & 0 & \dots & 0 & 0 \\ 0 & 0 & \dots & 0 & G' \otimes I \end{bmatrix} = \begin{bmatrix} F' \otimes I & 0 & \dots & 0 & 0 \\ 0 & 0 & \dots & 0 & G' \otimes I \end{bmatrix} \begin{bmatrix} * & * & \dots & * \\ \vdots & \vdots & \iddots & \vdots \\ * & * & \dots & * \\ (I_Z^{\uparrow})^{\otimes n+1} & * & \dots & * \end{bmatrix},$$

and the bottom left block of (10.4.2) gives

$$F' \otimes I_Z^{\uparrow} = I_Z^{\uparrow}(F' \otimes I) = (G' \otimes I)(I_Z^{\uparrow})^{\otimes n+1} = \big(G'(I_Z^{\uparrow})^{\otimes n}\big) \otimes I_Z^{\uparrow},$$

and it follows that $F' = G'(I_Z^{\uparrow})^{\otimes n}$. □

If $Z = [q]$ in Lemma 10.4.2, then, since $\mathcal{F} \subset {}_n\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_0$ corresponds to $\mathcal{G} \subset {}_n\langle \mathcal{G} \,|\, \mathcal{G}' \rangle_0$ and $\mathcal{F}' \subset {}_0\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_n$ corresponds to $\mathcal{G}' \subset {}_0\langle \mathcal{G} \,|\, \mathcal{G}' \rangle_n$, (10.4.9) already gives $\mathcal{F} \,|\, \mathcal{F}' = \mathcal{G} \,|\, \mathcal{G}'$ after the transformations by T_1 and T_2, so we are done. Otherwise, we must diverge from the proof strategy of the orthogonal converse in Chapter 8. The natural continuation along those lines would be to use Lemma 10.4.3 to add $I_Z^{\uparrow}$ to to $\mathcal{F}$ and $\mathcal{G}$ while preserving indistinguishability, then split into subdomains and apply induction. However, we cannot guarantee that these subdomains are quantum-nonvanishing. Instead, we use Lemma 10.4.3 to heavily restrict the form of $\mathcal{F} \,|\, \mathcal{F}'$ and $\mathcal{G} \,|\, \mathcal{G}'$, then use Lemma 10.4.1 to either split into subdomains or place further restrictions on ${}_1\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_1$ and ${}_1\langle \mathcal{G} \,|\, \mathcal{G}' \rangle_1$.

Proof of Theorem 10.3.2. Recall that we may assume $\mathcal{F} \mid \mathcal{F}'$ and $\mathcal{G} \mid \mathcal{G}'$ have the same domain size q, the Bi-Holant value of the vertxless loop. Lemma 10.4.3 gives $\varnothing \neq Z \subset [q]$ and T_1, T_2 such that, after replacing $\mathcal{F} \,|\, \mathcal{F}'$ with $T_1(\mathcal{F} \,|\, \mathcal{F}')$ and $\mathcal{G} \,|\, \mathcal{G}'$ with $T_2(\mathcal{G} \,|\, \mathcal{G}')$ (which preserves indistinguishability, quantum-nonvanishing and GL_q-orbits), (10.4.9) is satisfied. As mentioned in the previous paragraph, if $Z = [q]$ then we are done. Otherwise, (10.4.9) is equivalent to the statement that every $F' \in {}_0\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_n$ and $G \in {}_n\langle \mathcal{G} \,|\, \mathcal{G}' \rangle_0$ are supported only on Z, and furthermore $G|_Z = F|_Z$ for $F \leftrightsquigarrow G$ and $F'|_Z = G'|_Z$ for $F' \leftrightsquigarrow G'$. Or, assuming WLOG that $Z = [z] \subset [q]$, every ${}_n\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_0 \ni F \leftrightsquigarrow G \in {}_n\langle \mathcal{G} \,|\, \mathcal{G}' \rangle_0$ and ${}_0\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_n \ni F' \leftrightsquigarrow G' \in {}_0\langle \mathcal{G} \,|\, \mathcal{G}' \rangle_n$ have $(Z, \overline{Z})$-block form (with $\overline{Z} := [q] \setminus Z$)

$$F = \begin{bmatrix} F|_Z \\ * \\ \vdots \\ * \end{bmatrix}, G = \begin{bmatrix} F|_Z \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad \begin{matrix} F' = \begin{bmatrix} G'|_Z & 0 & \dots & 0 \end{bmatrix}, \\ G' = \begin{bmatrix} G'|_Z & * & \dots & * \end{bmatrix}. \end{matrix} \tag{10.4.12}$$

All generators (tensor in $\mathcal{F} \,|\, \mathcal{F}'$ and $\mathcal{G} \,|\, \mathcal{G}'$) are purely covariant or contravariant, so are subject to (10.4.12). Say that $\mathcal{F} \,|\, \mathcal{F}'$ and $\mathcal{G} \,|\, \mathcal{G}'$ have *skew blocks* if the purely covariant/contravariant tensors in $\langle \mathcal{F} \,|\, \mathcal{F}' \rangle$ and $\langle \mathcal{G} \,|\, \mathcal{G}' \rangle$ have zero blocks matching (10.4.12). We will use quantum-nonvanishing to force the $*$ blocks in (10.4.12) to be 0, at which point $\mathcal{F} \,|\, \mathcal{F}' = \mathcal{G} \,|\, \mathcal{G}'$.

Claim 10.4.1. Let $\mathbf{K}$ be a nontrivial (not just a wire) $(1,1)$-$\mathcal{F} \,|\, \mathcal{F}'$-gadget with tensors K and let $\widetilde{K}$ be the tensors of $\mathbf{K}_{\mathcal{F} \,|\, \mathcal{F}' \to \mathcal{G} \,|\, \mathcal{G}'}$. If $\mathcal{F} \,|\, \mathcal{F}'$ and $\mathcal{G} \,|\, \mathcal{G}'$ have skew blocks, then

$$K = \begin{bmatrix} K|_Z & 0 \\ * & 0 \end{bmatrix} \text{ and } \widetilde{K} = \begin{bmatrix} \widetilde{K}|_Z & * \\ 0 & 0 \end{bmatrix}. \tag{10.4.13}$$

Proof. Since $\mathbf{K}$ is nontrivial, it must contain at least one tensor in both $\mathcal{F}$ and $\mathcal{F}'$ to preserve covariant/contravariant balance. The right input to $\mathbf{K}$ is incident to an $F' \in \mathcal{F}'$, which by (10.4.12) is only supported on Z. Similarly, the left input to $\mathbf{K}_{\mathcal{F}\,|\,\mathcal{F}'\to\mathcal{G}\,|\,\mathcal{G}'}$ is incident to a $G \in \mathcal{G}$, which is only supported on Z. This completes the proof of Claim 10.4.1. ■

Say that $T \in \mathrm{GL}_q$ is $(Z,\overline{Z})$-*lower-triangular* if it has block form $T = \begin{bmatrix} T|_Z & 0 \\ T|_{\overline{Z},Z} & T|_{\overline{Z}} \end{bmatrix}$. Define $(Z,\overline{Z})$-*upper-triangular* similarly.

Claim 10.4.2. If $\mathcal{F}\,|\,\mathcal{F}'$ and $\mathcal{G}\,|\,\mathcal{G}'$ have skew blocks and T and U are $(Z,\overline{Z})$-lower- and upper-triangular, respectively, then $T(\mathcal{F}\,|\,\mathcal{F}')$ and $U(\mathcal{G}\,|\,\mathcal{G}')$ have skew blocks and $(T\cdot K)|_Z = T|_Z \cdot K|_Z$ for every purely covariant or contravariant $K \in \langle \mathcal{F}\,|\,\mathcal{F}'\rangle \cup \langle \mathcal{G}\,|\,\mathcal{G}'\rangle$.

Proof. The transformations T and U act on every $F \in {}_n\langle \mathcal{F}\,|\,\mathcal{F}'\rangle_0$, $F' \in {}_0\langle \mathcal{F}\,|\,\mathcal{F}'\rangle_n$, $G \in {}_n\langle \mathcal{G}\,|\,\mathcal{G}'\rangle_0$, and $G' \in {}_0\langle \mathcal{G}\,|\,\mathcal{G}'\rangle_n$ as

$$F \mapsto T^{\otimes n}F = \begin{bmatrix} (T|_Z)^{\otimes n} & 0 & \dots & 0 \\ * & * & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ * & * & \dots & * \end{bmatrix} \begin{bmatrix} F|_Z \\ * \\ \vdots \\ * \end{bmatrix} = \begin{bmatrix} (T|_Z)^{\otimes n}F|_Z \\ * \\ \vdots \\ * \end{bmatrix},$$

$$F' \mapsto F'(T^{-1})^{\otimes n} = \begin{bmatrix} F'|_Z & 0 & \dots & 0 \end{bmatrix} \begin{bmatrix} (T|_Z^{-1})^{\otimes n} & 0 & \dots & 0 \\ * & * & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ * & * & \dots & * \end{bmatrix} = \begin{bmatrix} F'|_Z(T|_Z^{-1})^{\otimes n} & 0 & \dots & 0 \end{bmatrix},$$

$$G \mapsto U^{\otimes n}G = \begin{bmatrix} (U|_Z)^{\otimes n} & * & \dots & * \\ 0 & * & \dots & * \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & * \end{bmatrix} \begin{bmatrix} G|_Z \\ 0 \\ \vdots \\ 0 \end{bmatrix} = \begin{bmatrix} (U|_Z)^{\otimes n}G|_Z \\ 0 \\ \vdots \\ 0 \end{bmatrix},$$

$$G' \mapsto G'(U^{-1})^{\otimes n} = \begin{bmatrix} G'|_Z & * & \dots & * \end{bmatrix} \begin{bmatrix} (U|_Z^{-1})^{\otimes n} & * & \dots & * \\ 0 & * & \dots & * \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & * \end{bmatrix} = \begin{bmatrix} G'|_Z(U|_Z^{-1})^{\otimes n} & * & \dots & * \end{bmatrix}. \blacksquare$$

Claim 10.4.3. Suppose $\mathcal{F}\,|\,\mathcal{F}'$ and $\mathcal{G}\,|\,\mathcal{G}'$ have skew blocks and $\langle \mathcal{F}\,|\,\mathcal{F}'\rangle \ni I_X^{\uparrow} \leftrightsquigarrow I_X^{\uparrow} \in \langle \mathcal{G}\,|\,\mathcal{G}'\rangle$, where $X = [x] \subset [z] = Z$ and $Z = X \cup \Delta$ with $|\Delta| = \delta < |Z|$. Then there are $(Z,\overline{Z})$-lower-triangular $T_\delta \in \mathrm{GL}_q$ and $(Z,\overline{Z})$-upper-triangular $U_\delta \in \mathrm{GL}_q$ such that, after transforming $\mathcal{F}\,|\,\mathcal{F}'$ by T_δ and $\mathcal{G}\,|\,\mathcal{G}'$ by U_δ, $F|_Z^{\uparrow} \in \langle \mathcal{F}\,|\,\mathcal{F}'\rangle$ for every $F \in \mathcal{F}$ and $G'|_Z^{\uparrow} \in \langle \mathcal{G}\,|\,\mathcal{G}'\rangle$ for every $G' \in \mathcal{G}'$.

Proof. We prove the claim by induction on δ. If $\delta = 0$, then $X = Z$, so we already have $F|_Z^{\uparrow} = (I_Z^{\uparrow})^{\otimes n}F \in {}_n\langle \mathcal{F}\,|\,\mathcal{F}'\rangle_0$ and $G'|_Z^{\uparrow} = G'(I_Z^{\uparrow})^{\otimes n} \in {}_0\langle \mathcal{G}\,|\,\mathcal{G}'\rangle_n$.

Otherwise $\delta > 0$. First suppose every $F \in \mathcal{F}$ and $G' \in \mathcal{G}'$ is supported on only $\overline{\Delta} = X \cup \overline{Z}$ (that is, F and G' are 0 when given any input from Δ). Then $F|_Z^{\uparrow} = F|_X^{\uparrow} = (I_X^{\uparrow})^{\otimes n} F \in {}_n\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_0$, and similarly $G'|_Z^{\uparrow} \in {}_0\langle \mathcal{G} \,|\, \mathcal{G}' \rangle_n$, so we are done. Otherwise, there is an $F \in \mathcal{F}$ or $G' \in \mathcal{G}'$ that is supported on Δ. We give a proof for the former case; the latter follows by transposed reasoning. There is an $\mathbf{a} \in [q]^n$ satisfying $F_{\mathbf{a}} \neq 0$ and $a_i \in \Delta$ for some i. Then, since $\Delta \subset \overline{X}$,

$$0 \neq F^{\Delta} := (I^{\otimes i-1} \otimes I_X^{\uparrow} \otimes I^{\otimes n-i}) F = (I^{\otimes i-1} \otimes (I - I_X^{\uparrow}) \otimes I^{\otimes n-i}) F \in {}_n\big\langle \mathcal{F} \,|\, \mathcal{F}' \big\rangle_0 .$$

Therefore, by the quantum-nonvanishing of $\mathcal{F} \,|\, \mathcal{F}'$, there is an quantum-$\mathcal{F} \,|\, \mathcal{F}'$-gadget $\mathbf{F}'$ with

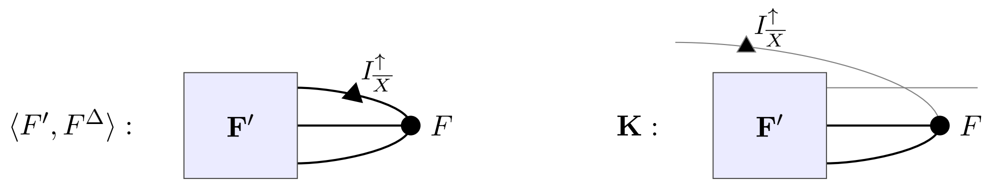

Figure 10.2: Breaking an edge of the grid $\langle F', F^{\Delta} \rangle$ to produce $\mathbf{K}$, with $n = 3$ and $i = 1$.

tensor $F' \in {}_0\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_n$ such that $\langle F', F^{\Delta} \rangle \neq 0$. View $\langle F', F^{\Delta} \rangle$ as a $\langle \mathcal{F} \,|\, \mathcal{F}' \rangle$-grid composed of F, $I_X^{\uparrow}$, and $\mathbf{F}'$ (see Figure 10.2). Breaking the edge between $I_X^{\uparrow}$ and $\mathbf{F}'$ produces a $(1,1)$-quantum-$\mathcal{F} \,|\, \mathcal{F}'$-gadget $\mathbf{K}$ with tensor $K \in {}_1\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_1$ such that $\operatorname{tr}(K) \neq 0$. The left input to $\mathbf{K}$ is incident to $I_X^{\uparrow}$ and the right input to $\mathbf{K}$ is incident to $\mathbf{F}'$, whose tensor F' is only supported on Z by skew blocks. On the $\mathcal{G} \,|\, \mathcal{G}'$ side, every term of $\mathbf{K}_{\mathcal{F} | \mathcal{F}' \to \mathcal{G} | \mathcal{G}'}$ is a nontrivial $\mathcal{G} \,|\, \mathcal{G}'$-gadget (it contains e.g. the generator $G \in \mathcal{G}$ such that $F \leftrightsquigarrow G$), so satisfies the condition of Claim 10.4.1. Thus, letting $\widetilde{K}$ be the tensor of $\mathbf{K}_{\mathcal{F} | \mathcal{F}' \to \mathcal{G} | \mathcal{G}'}$ (so $K \leftrightsquigarrow \widetilde{K}$),

$$K = \left[\begin{array}{c:c|c} 0 & 0 & 0 \\ \hdashline K|_{\Delta,X} & K|_{\Delta} & 0 \\ \hline K|_{\overline{Z},X} & K|_{\overline{Z},\Delta} & 0 \end{array}\right] \quad \text{and } \widetilde{K} = \left[\begin{array}{c:c|c} \widetilde{K}|_X & \widetilde{K}|_{X,\Delta} & \widetilde{K}|_{X,\overline{Z}} \\ \hdashline \widetilde{K}|_{\Delta,X} & \widetilde{K}|_{\Delta} & \widetilde{K}|_{\Delta,\overline{Z}} \\ \hline 0 & 0 & 0 \end{array}\right].$$

Now $\operatorname{tr}(K) \neq 0$ implies that $\operatorname{tr}(K|_{\Delta}) \neq 0$ and

$$\big\langle \mathcal{F} \,|\, \mathcal{F}' \big\rangle_{\overline{X}} \ni K|_{\overline{X}} = \left[\begin{array}{c|c} K|_{\Delta} & 0 \\ \hline K|_{\overline{Z},\Delta} & 0 \end{array}\right] \leftrightsquigarrow \left[\begin{array}{c|c} \widetilde{K}|_{\Delta} & \widetilde{K}|_{\Delta,\overline{Z}} \\ \hline 0 & 0 \end{array}\right] = \widetilde{K}|_{\overline{X}} \in \big\langle \mathcal{G} \,|\, \mathcal{G}' \big\rangle_{\overline{X}} .$$

By Proposition 8.2.3 and Proposition 10.3.3, $\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_{\overline{X}}$ and $\langle \mathcal{G} \,|\, \mathcal{G}' \rangle_{\overline{X}}$ are Bi-Holant-indistinguishable and quantum-nonvanishing. Hence $\operatorname{tr}(\widetilde{K}|_\Delta) = \operatorname{tr}(K|_\Delta) \neq 0$. Thus $K|_{\overline{X}}$ and $\widetilde{K}|_{\overline{X}}$ do not have singleton spectrum, so, by Lemma 10.4.1, there are $T = \begin{bmatrix} T|_\Delta & 0 \\ T|_{\overline{Z},\Delta} & T|_{\overline{Z}} \end{bmatrix} \in \mathrm{GL}\left(\mathbb{K}^{\overline{X}}\right)$ and $U = \begin{bmatrix} U|_\Delta & U|_{\Delta,\overline{Z}} \\ 0 & U|_{\overline{Z}} \end{bmatrix} \in \mathrm{GL}\left(\mathbb{K}^{\overline{X}}\right)$ such that, after transforming $\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_{\overline{X}}$ by T and $\langle \mathcal{G} \,|\, \mathcal{G}' \rangle_{\overline{X}}$ by U, we obtain $\left\langle \langle \mathcal{F} \,|\, \mathcal{F}' \rangle_{\overline{X}} \right\rangle \ni I_{X'}^{\uparrow \overline{X}} \leftrightsquigarrow I_{X'}^{\uparrow \overline{X}} \in \left\langle \langle \mathcal{G} \,|\, \mathcal{G}' \rangle_{\overline{X}} \right\rangle$ for some $\varnothing \neq X' \subset \Delta$. Apply the $(Z, \overline{Z})$-lower-triangular and $(Z, \overline{Z})$-upper-triangular transformations

$$I_X \oplus T = \left[\begin{array}{c:c|c} I_X & 0 & 0 \\ \hdashline 0 & T|_\Delta & 0 \\ \hline 0 & T|_{\overline{Z},\Delta} & T|_{\overline{Z}} \end{array}\right] \in \mathrm{GL}_q \text{ and } I_X \oplus U = \left[\begin{array}{c:c|c} I_X & 0 & 0 \\ \hdashline 0 & U|_\Delta & U|_{\Delta,\overline{Z}} \\ \hline 0 & 0 & U|_{\overline{Z}} \end{array}\right] \in \mathrm{GL}_q$$

to $\mathcal{F} \,|\, \mathcal{F}'$ and $\mathcal{G} \,|\, \mathcal{G}'$, respectively. This preserves $I_X^\uparrow$ and $I_{\overline{X}}^\uparrow$ in $\langle \mathcal{F} \,|\, \mathcal{F}' \rangle$ and $\langle \mathcal{G} \,|\, \mathcal{G}' \rangle$, preserves skew blocks by Claim 10.4.2, and, by the above, realizes $\left\langle \langle \mathcal{F} \,|\, \mathcal{F}' \rangle_{\overline{X}} \right\rangle \ni I_{X'}^{\uparrow \overline{X}} \leftrightsquigarrow I_{X'}^{\uparrow \overline{X}} \in \left\langle \langle \mathcal{G} \,|\, \mathcal{G}' \rangle_{\overline{X}} \right\rangle$. Now, by Proposition 8.2.2,

$$\langle \mathcal{F} \,|\, \mathcal{F}' \rangle \ni \left(I_{X'}^{\uparrow \overline{X}}\right)^\uparrow = I_{X'}^\uparrow \leftrightsquigarrow I_{X'}^\uparrow = \left(I_{X'}^{\uparrow \overline{X}}\right)^\uparrow \in \langle \mathcal{G} \,|\, \mathcal{G}' \rangle .$$

Hence $\langle \mathcal{F} \,|\, \mathcal{F}' \rangle \ni I_{X \cup X'}^\uparrow = I_X^\uparrow + I_{X'}^\uparrow \leftrightsquigarrow I_{X \cup X'}^\uparrow \in \langle \mathcal{G} \,|\, \mathcal{G}' \rangle$. We have $Z = X \cup \Delta = (X \cup X') \cup \Delta'$, where $\Delta' = \Delta \setminus X'$. This $\delta' := |\Delta'| < |\Delta| = \delta$, so, by induction (with $X := X \cup X'$), there exist $(Z, \overline{Z})$-lower- and upper-triangular transformations $T_{\delta'}$ and $U_{\delta'}$ after which $\langle \mathcal{F} \,|\, \mathcal{F}' \rangle$ and $\langle \mathcal{G} \,|\, \mathcal{G}' \rangle$ contain the desired $F|_Z^\uparrow$ and $G'|_Z^\uparrow$. In total, we have applied $T_\delta := T_{\delta'} \circ (I_X \oplus T)$ and $U_\delta := U_{\delta'} \circ (I_X \oplus U)$, which, since both components are $(Z, \overline{Z})$-lower (resp. upper)-triangular, are $(Z, \overline{Z})$-lower (resp. upper)-triangular. This completes the proof of Claim 10.4.3. ■

Unless (by Proposition 10.3.2) $\mathcal{F} \,|\, \mathcal{F}'$ and $\mathcal{G} \,|\, \mathcal{G}'$ consist only of zero tensors, there is a nonzero $F' \in \mathcal{F}'$, and there is an $\mathcal{F} \,|\, \mathcal{F}'$-grid Ω containing F' with $\mathrm{Holant}(\Omega) \neq 0$. Breaking an edge of Ω yields a nontrivial binary $\mathcal{F} \,|\, \mathcal{F}'$-gadget $\mathbf{K}$ whose tensor $K \in {}_1\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_1$ satisfies $\operatorname{tr}(K) = \mathrm{Holant}(\Omega) \neq 0$. By indistinguishability, the tensor $\widetilde{K} \in {}_1\langle \mathcal{G} \,|\, \mathcal{G}' \rangle_1$ of $\mathbf{K}_{\mathcal{F} \,|\, \mathcal{F}' \to \mathcal{G}' \,|\, \mathcal{G}'}$ has the same nonzero trace. Claim 10.4.1 asserts that K and $\widetilde{K}$ have the form (10.4.13), so $\operatorname{tr}(K|_Z) = \operatorname{tr}(\widetilde{K}|_Z) \neq 0$. Therefore K and $\widetilde{K}$ do not have singleton spectrum, so, by Lemma 10.4.1, we may transform $\mathcal{F}$ by $T = \begin{bmatrix} T|_Z & 0 \\ T|_{\overline{Z},Z} & T|_{\overline{Z}} \end{bmatrix}$ and $\mathcal{G}$ by $U = \begin{bmatrix} U|_Z & U|_{Z,\overline{Z}} \\ 0 & U|_{\overline{Z}} \end{bmatrix}$ to obtain $\langle \mathcal{F} \,|\, \mathcal{F}' \rangle \ni I_X^\uparrow \leftrightsquigarrow I_X^\uparrow \in \langle \mathcal{G} \,|\, \mathcal{G}' \rangle$ for some

$\varnothing \neq X = [x] \subset Z$. By Claim 10.4.2, these transformations preserve skew blocks. Thus Claim 10.4.3 applies and we obtain T_δ and U_δ under which $F|_Z^{\uparrow} \in \langle \mathcal{F} \,|\, \mathcal{F}' \rangle$ for every $F \in \mathcal{F}$ and $G'|_Z^{\uparrow} \in \langle \mathcal{G} \,|\, \mathcal{G}' \rangle$ for every $G' \in \mathcal{G}'$. After the combined transformations $T_\delta \circ T = \left[\begin{smallmatrix} (T_\delta T)|_Z & * \\ 0 & * \end{smallmatrix}\right]$ and $U_\delta \circ U = \left[\begin{smallmatrix} (U_\delta U)|_Z & 0 \\ * & * \end{smallmatrix}\right]$, (10.4.12) becomes, by Claim 10.4.2,

$$F = \begin{bmatrix} ((T_\delta T)|_Z)^{\otimes n} \widetilde{F}|_Z \\ * \\ \vdots \\ * \end{bmatrix}, G = \begin{bmatrix} ((U_\delta U)|_Z)^{\otimes n} \widetilde{F}|_Z \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad \begin{array}{l} F' = \left[\widetilde{G}'|_Z((T_\delta T)|_Z^{-1})^{\otimes n} \quad 0 \quad \dots \quad 0\right], \\ G' = \left[\widetilde{G}'|_Z((U_\delta U)|_Z^{-1})^{\otimes n} \quad * \quad \dots \quad *\right]. \end{array} \tag{10.4.14}$$

for every ${}_n\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_0 \ni F \leftrightsquigarrow G \in {}_n\langle \mathcal{G} \,|\, \mathcal{G}' \rangle_0$ and ${}_0\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_n \ni F' \leftrightsquigarrow G' \in {}_0\langle \mathcal{G} \,|\, \mathcal{G}' \rangle_n$ (where $\widetilde{F}$ and $\widetilde{G}'$ are the pre-transformation F and G'). For $F \in \mathcal{F}$, we now have $F - F|_Z^{\uparrow} \in {}_n\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_0$, and (10.4.14) gives

$$\langle (F - F|_Z^{\uparrow}), F' \rangle = \langle (F - F|_Z^{\uparrow})|_Z, F'|_Z \rangle = \langle 0, F'|_Z \rangle = 0$$

for every $F' \in {}_0\langle \mathcal{F} \,|\, \mathcal{F}' \rangle_n$, so the quantum-nonvanishing of $\mathcal{F} \,|\, \mathcal{F}'$ implies that $F - F|_Z^{\uparrow} = 0$. Similarly, every $G' - G'|_Z^{\uparrow} = 0$. So (10.4.14) is

$$F = \begin{bmatrix} ((T_\delta T)|_Z)^{\otimes n} \widetilde{F}|_Z \\ 0 \\ \vdots \\ 0 \end{bmatrix}, G = \begin{bmatrix} (U_\delta U)|_Z^{\otimes n} \widetilde{F}|_Z \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad \begin{array}{l} F' = \left[\widetilde{F}'|_Z((T_\delta T)|_Z^{-1})^{\otimes n} \quad 0 \quad \dots \quad 0\right], \\ G' = \left[\widetilde{F}'|_Z((U_\delta U)|_Z^{-1})^{\otimes n} \quad 0 \quad \dots \quad 0\right] \end{array}$$

for every $\mathcal{F} \ni F \leftrightsquigarrow G \in \mathcal{G}$ and $\mathcal{F}' \ni F' \leftrightsquigarrow G' \in \mathcal{G}'$. After a final transformation of $\mathcal{F} \,|\, \mathcal{F}'$ by $(T_\delta T)|_Z^{-1} \oplus I_{\overline{Z}} \in \mathrm{GL}_q$ and $\mathcal{G} \,|\, \mathcal{G}'$ by $(U_\delta U)|_Z^{-1} \oplus I_{\overline{Z}} \in \mathrm{GL}_q$, we obtain $\mathcal{F} \,|\, \mathcal{F}' = \mathcal{G} \,|\, \mathcal{G}'$. □

We conclude this section by noting that Theorem 10.3.1 applies to any field $\mathbb{K}$ of characteristic 0. However, the multitude of Jordan decompositions performed – via Lemma 10.4.1 – in the proof of Theorem 10.3.2 necessitate the extra assumption that $\mathbb{K}$ is algebraically closed. Indeed, Theorem 10.3.2 does not hold without this assumption. For example, let $\mathbb{K} = \mathbb{R}$ and consider $\mathcal{F} = (=_2 \mid =_2)$ and $\mathcal{G} = (-(=_2)| - (=_2))$. Every $\mathcal{G}$-grid must contain an equal number of covariant and contravariant $-(=_2)$ tensors, hence an even number of total tensors. Therefore $\mathcal{F}$ and $\mathcal{G}$ are (Bi-)Holant-indistinguishable. Furthermore, if $K \in {}_\ell\langle \mathcal{G} \rangle_r$, then construct $\pm K^\top \in {}_r\langle \mathcal{G} \rangle_\ell$ by connecting a left-facing $-(=_2)$ to every right input of K and connecting a right-facing $-(=_2)$ to each left input of K (this exchanges the left and right inputs of K while preserving the underlying tensor up to a global $\pm$). Now $\langle K, \pm K^\top \rangle \neq 0$ (as this is effectively a contraction of K with itself), so K is $\mathcal{G}$-nonvanishing. Thus $\mathcal{G}$ and, similarly, $\mathcal{F}$, are quantum-nonvanishing. Theorem 10.3.2 guarantees the existence of a complex T (in this case $T = iI$) transforming $\mathcal{F}$ to $\mathcal{G}$, but any such T must satisfy $TT^\top = T(=_2)^{1,1}T^\top = (-(=_2))^{1,1} = -I$, which is impossible for real-valued T.

10.5 More corollaries of the two converses

In this section, we exploit the expressive power of Holant and Bi-Holant to derive novel consequences of Theorems 10.2.4 and 10.3.2. To apply 10.3.2, we need a method of showing that a set is quantum-nonvanishing. In general, this seems to be difficult, but one easy case is the canonical class of quantum-nonvanishing sets admitting a complex inner product:

Proposition 10.5.1. *If $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$ is conjugate-closed, then $(\mathcal{F}, \supset, \subset)$ is quantum-nonvanishing.*

Proof. Let $0 \neq K \in {}_\ell\langle \mathcal{F}, \supset, \subset \rangle_r$. By (4.4.4), $\langle \mathcal{F}, \supset, \subset \rangle$ is conjugate-closed, so $\overline{K} \in \langle \mathcal{F}, \supset, \subset \rangle$. Now construct $K^\dagger$ in ${}_r\langle \mathcal{F}, \supset, \subset \rangle_\ell$ by connecting $\supset$ to each right input of $\overline{K}$ and $\subset$ to each left input of $\overline{K}$. Then $\langle K, K^\dagger \rangle = \|K\|^2 \neq 0$, so $K^\dagger$ witnesses that K is $(\mathcal{F}, \supset, \subset)$-nonvanishing. See Figure 10.1a. □

We now obtain a complex generalization of Theorem 8.1.1, which we recall does not hold as stated for complex-valued signatures.

Corollary 10.5.1. *Suppose $\mathcal{F}, \mathcal{G} \subset \mathcal{S}(\mathbb{C}^q)$ are conjugate-closed. Then $\mathcal{F}$ and $\mathcal{G}$ are Holant-indistinguishable if and only if there is a $T \in O_q(\mathbb{C})$ such that $T\mathcal{F} = \mathcal{G}$.*

Proof. By definition, $[\mathcal{F}] = \langle \mathcal{F}, \supset, \subset \rangle$. Now the result follows from Proposition 10.5.1, Proposition 2.4.1, and Theorem 10.3.2 with $\mathbb{K} = \mathbb{C}$. □

10.5.1 Bounded-degree graph homomorphisms and #CSP

Define $\#\text{CSP}^{(d)}(\mathcal{F})$ to be the restriction of $\#\text{CSP}(\mathcal{F})$ to instances I in which every variable appears at most d times across all constraints [CS24]. Recall from (2.3.2) that $\#\text{CSP}(\mathcal{F}) \equiv \text{Holant}(\mathcal{F} \mid \mathcal{EQ}) \equiv \text{Holant}(\mathcal{F} \cup \mathcal{EQ})$, where each signature $=_n$ in a $\text{Holant}(\mathcal{F} \mid \mathcal{EQ})$ grid models a variable appearing exactly n times across all constraints in the corresponding $\#\text{CSP}(\mathcal{F})$ instance.

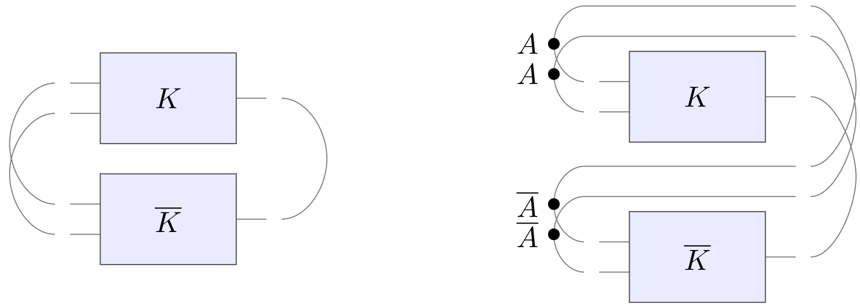

(a) The construction in Proposition 10.5.1 (b) The construction in Lemma 10.5.1

Figure 10.1: Connecting quantum gadgets with their conjugates for quantum-nonvanishing.

Therefore $\#\mathrm{CSP}^{(d)}(\mathcal{F}) \equiv \mathrm{Holant}(\mathcal{F} \,|\, \mathcal{EQ}_{\leq d})$, where $\mathcal{EQ}_{\leq d} = \{=_1, \dots, =_d\}$. By Proposition 5.1.3, a matrix $T \in \mathrm{GL}_q$ is a permutation matrix if and only if $T\{=_2, =_3\} = \{=_2, =_3\}$, suggesting that $\#\mathrm{CSP}^{(d)}$-indistinguishability should also be equivalent to isomorphim. However, the equivalence between $\mathrm{Holant}(\mathcal{F} \,|\, \mathcal{EQ})$ and $\mathrm{Holant}(\mathcal{F} \cup \mathcal{EQ})$ followed from merging adjacent $\mathcal{EQ}$ signatures in a $\mathrm{Holant}(\mathcal{F} \cup \mathcal{EQ})$ grid, and merging two $\mathcal{EQ}$ signatures of arity $\leq d$ could produce an $\mathcal{EQ}$ signature of arity $> d$. Therefore $\mathrm{Holant}(\mathcal{F} \,|\, \mathcal{EQ}_{\leq d})$ and $\mathrm{Holant}(\mathcal{F} \cup \mathcal{EQ}_{\leq d})$ are not equivalent, so $\#\mathrm{CSP}^{(d)}(\mathcal{F})$ is a truly bipartite problem, to which the the results of previous chapters do not apply. Instead must apply the conditional converse Theorem 10.3.2 to obtain the following.

Corollary 10.5.2. *For $d \geq 3$, if $\mathcal{EQ}_{\leq d} \,|\, \mathcal{F}$ and $\mathcal{EQ}_{\leq d} \,|\, \mathcal{G}$ are quantum-nonvanishing, then $\mathcal{F}$ and $\mathcal{G}$ are $\#\mathrm{CSP}^{(d)}$-indistinguishable if and only if $\mathcal{F}$ and $\mathcal{G}$ are isomorphic.*

The next lemma, which generalizes the argument in Proposition 10.5.1, gives a sufficient condition for $\mathcal{EQ}_{\leq d} \,|\, \mathcal{F}$ to be quantum-nonvanishing.

Lemma 10.5.1. *If $\mathcal{F} \subset \mathcal{V}(\mathbb{C}^q)$ is conjugate-closed and satisfies $\supset \in {}_2\langle \mathcal{F} \rangle_0$ and $A \in {}_0\langle \mathcal{F} \rangle_2$ (or vice-versa) for some A whose matrix form $A^{1,1}$ is invertible, then $\mathcal{F}$ is quantum-nonvanishing.*

Proof. See Figure 10.1b. Let $0 \neq K \in {}_\ell\langle \mathcal{F} \rangle_r$. If $\ell = 0$, then, as in Proposition 10.5.1, construct $K^\dagger \in {}_r\langle \mathcal{F} \rangle_0$ by connecting each right input of K with a copy of $\supset \in {}_2\langle \mathcal{F} \rangle_0$ and conjugating all coefficients and tensors composing K. Then $\langle K, K^\dagger \rangle \neq 0$, so K is $\mathcal{F}$-nonvanishing. Otherwise, $(A^{1,1})^{\otimes \ell} K \neq 0$ by invertibility of $A^{1,1}$, and therefore the tensor $K' \in {}_0\langle \mathcal{F} \rangle_{\ell+r}$ formed by connecting ℓ copies of A with the ℓ left inputs of K (equivalently, connecting ℓ copies of right-facing $=_2$ with the ℓ left inputs of $(A^{1,1})^{\otimes \ell} K$) is nonzero. Again, since K' is now fully right-facing, its dual $(K')^*$

is in ${}_{\ell+r}\langle\mathcal{F}\rangle_0$. The $\langle\mathcal{F}\rangle$-grid formed by contracting K' and $(K')^\dagger$ contains K and has nonzero value, so K is $\mathcal{F}$-nonvanishing. □

Corollary 10.5.3. *Let $\mathcal{F},\mathcal{G} \subset \mathcal{V}(\mathbb{C}^q)$ be conjugate-closed. For $d \geq 3$, if there exist $A_1 \in {}_0\langle\mathcal{EQ}_{\leq d}\,|\,\mathcal{F}\rangle_2$ and $A_2 \in {}_0\langle\mathcal{EQ}_{\leq d}\,|\,\mathcal{G}\rangle_2$ with invertible matrix forms, then $\mathcal{F}$ and $\mathcal{G}$ are $\#\mathrm{CSP}^{(d)}$-indistinguishable if and only if $\mathcal{F}$ and $\mathcal{G}$ are isomorphic.*

To apply these results to bounded-degree homomorphism indistinguishability, we first address a technical point regarding graphs and multigraphs.

Proposition 10.5.2. *Graphs F and G are homomorphism-indistinguishable over (simple) graphs of maximum degree at most d iff $\mathcal{EQ}_{\leq d}\,|A_F$ and $\mathcal{EQ}_{\leq d}\,|A_G$ are Holant-indistinguishable.*

Proof. First note that F and G satisfying the former condition satisfy $|V(F)| = \hom(K_1, F) = \hom(K_1, G) = |V(G)|$ (where K_1 is the isolated vertex), so $\mathcal{EQ}_{\leq d}\,|A_F$ and $\mathcal{EQ}_{\leq d}\,|A_G$ have the same domain size. The discussion around Figure 2.1 shows that $\mathrm{Holant}_{\mathcal{EQ}_{\leq d}|A_G}$ exactly captures the problem of counting homomorphisms from *multigraphs* of degree at most d to G. However, for any G, $\hom(X,G) = 0$ if X has a loop, and $\hom(X,G) = \hom(X',G)$, where X' is constructed from X by removing multiedges. Critically, X' has maximum degree at most d if X has maximum degree at most d. Therefore F and G are homomorphism-indistinguishable over simple graphs of degree at most d if and only if they are homomorphism-indistinguishable over multigraphs of degree at most d. □

Specializing Corollary 10.5.3 to $\mathcal{F} = \{A_F\}$ and $\mathcal{G} = \{A_G\}$ for invertible A_F and A_G gives:

Corollary 10.5.4. *Two graphs F and G with invertible adjacency matrices are homomorphism-indistinguishable over graphs of maximum degree at most 3 if and only if they are isomorphic.*

Finally, we show that the condition in Lemma 10.5.1 is also necessary for the homomorphism counting tensor set to be quantum-nonvanishing.

Proposition 10.5.3. *For graph G and any conjugate-closed $\mathcal{F}$ with $(=_2) \in \mathcal{F}$, the set $\mathcal{F}\,|A_G$ is quantum-nonvanishing if and only if A_G is invertible.*

Proof. The ($\Leftarrow$) direction follows from Lemma 10.5.1. Conversely, suppose A_G is singular. There are a real orthogonal matrix H, diagonal matrix D, and $X \subsetneq [q]$ such that $A_G = H^\top DH$ and $D_{xx} \neq 0 \iff x \in X$. Transforming by H does not change whether a set is quantum-nonvanishing, so it suffices to show that $H(\mathcal{F}\,|A_G) = H\,\mathcal{F} \mid D$ is quantum-vanishing. Connect a copy of the left-facing $H(=_2) = (=_2) \in H\,\mathcal{F}$ to a right-facing D, giving $D \in {}_1\langle H\,\mathcal{F} \mid D\rangle_1$. Let p be a polynomial such that $p(0) = 0$ and $p(D_{xx}) = 1$ for every $x \in X$. Then $I - p(D) \in {}_1\langle H\,\mathcal{F} \mid D\rangle_1$ is a diagonal matrix that is 1 on $[q] \setminus X$ and 0 on X. Now connecting another left-facing $=_2$ to the right input of $I - p(D)$ gives $I - p(D) \in {}_2\langle H\,\mathcal{F} \mid D\rangle_0$. This left-facing $I - p(D)$ is $(H\,\mathcal{F} \mid D)$-vanishing. Indeed, in any $\mathcal{F}, (I - p(D)) \mid D$-grid Ω containing $I - p(D)$, the left-side vertex assigned $I - p(D)$ must be adjacent to some right-side vertex assigned D. The former evaluates to 0 on inputs from X and the latter evaluates to 0 on inputs from $[q] \setminus X$, so $\text{Holant}(\Omega) = 0$. □

We also apply the approximate converse Theorem 10.2.4, which applies to all graphs F and G, instead of the conditional converse to obtain the first characterization of homomorphism-indistinguishability over graphs of bounded degree, as a special case of the characterization for bounded-degree #CSP.

Corollary 10.5.5. *Finite constraint function sets $\mathcal{F}, \mathcal{G} \subset \mathcal{S}(\mathbb{C}^q)$ are $\#\text{CSP}^{(d)}$-indistinguishable if and only if* $\overline{\text{GL}_q(\mathcal{EQ}_{\leq d}\,|\,\mathcal{F})} \cap \overline{\text{GL}_q(\mathcal{EQ}_{\leq d}\,|\,\mathcal{G})} \neq \varnothing$.

Corollary 10.5.6. *Graphs F and G on q vertices are homomorphism-indistinguishable over all graphs of maximum degree at most d if and only if* $\overline{\text{GL}_q(\mathcal{EQ}_{\leq d}\,|A_F)} \cap \overline{\text{GL}_q(\mathcal{EQ}_{\leq d}\,|A_G)} \neq \varnothing$.

Therefore the problem of determining whether two graphs are homomorphism-indistinguishable over all graphs of maximum degree at most d is decidable.

The decidability result in Corollary 10.5.6, which follows from Corollary 10.2.1, answers another open question, and is interesting because homomorphism-indistinguishability over some graph classes (e.g. quantum isomorphism for planar graphs) is undecidable.

10.5.2 Indistinguishability, TOCI, and GI

Lysikov and Walter [LW25] define the class **TOCI** of (problems reducible to) orbit closure intersection problems for actions of general linear groups on finite subsets of $\bigcup_{i=1}^{m} \mathcal{V}(\mathbb{C}^{q_i})$ (sets are allowed to contain tensors with different domains). They show that **GI** $\subset$ **TOCI** by reducing isomorphism of q-vertex graphs F and G to GL_q-orbit-intersection of $(A_F, =_3 \mid =_2)$ and $(A_G, =_3 \mid =_2)$ [LW25, Lemma 5.26 and Proposition 5.28][1]. Lysikov and Walter also show that the orbit closure intersection problem for $\mathcal{F}$ containing two left-facing ternary tensor and one right-facing binary tensor, all on the same domain, is **TOCI**-complete [LW25, Corollary 1.3]. Combining these results with Theorem 10.2.4 gives the following.

Corollary 10.5.7. *The following problem is **TOCI**-complete: Given ternary signatures F_3, F_3', G_3, G_3' and binary signatures F_2, G_2, decide whether $(F_3, F_3' \mid F_2)$ and $(G_3, G_3' \mid G_2)$ are Holant-indistinguishable.*

Corollary 10.5.8. *The following problem is **GI**-hard: Given ternary signatures F_3, G_3 and binary signatures F_2, F_2', G_2, G_2', decide whether $(F_2, F_3 \mid F_2')$ and $(G_2, G_3 \mid G_2')$ are Holant-indistinguishable.*

10.6 Discussion

10.6.1 Weak isomorphism

Recall Definition 10.3.2 of weakly isomorphic $\mathcal{F}$ and $\mathcal{G}$. By Proposition 4.1.2, if $\mathcal{F}$ and $\mathcal{G}$ are not weakly isomorphic, then there is no $T \in \mathrm{GL}_q$ transforming $\mathcal{F}$ to $\mathcal{G}$ (such a T would map a nonzero tensor to 0). It is natural in light of Propositions 10.3.1 and 10.3.2 to ask whether the converse of this statement holds. That is, can we improve the conditional converse to an equivalence between weak isomorphism of $\mathcal{F} \mid \mathcal{F}'$ and $\mathcal{G} \mid \mathcal{G}'$ and the existence of a holographic transformation T between them? Indeed, the answer to the analogous question is yes in Corollary 9.2.1. However, in that setting, and the settings of all of our indistinguishability theorems before this chapter, the signature sets in question are quantum-nonvanishing, so one of the duality theorems Theorem 4.5.1,

[1] Our framework gives a short alternative proof of this reduction. First, if $F \cong G$, then, since every permutation matrix preserves $\mathcal{EQ}$, the GL_q-orbits of $(A_F, =_3 \mid =_2)$ and $(A_G, =_3 \mid =_2)$ intersect. Conversely, the 'easy' ($\Longleftarrow$) direction of Theorem 10.2.4 asserts that $(A_F, =_3 \mid =_2)$ and $(A_G, =_3 \mid =_2)$ are Holant-indistinguishable. As in the proof of Proposition 5.1.3, left-facing $=_3$ and right-facing $=_2$ together construct all of $\mathcal{EQ}$, so $(A_F \mid \mathcal{EQ})$ and $(A_G \mid \mathcal{EQ})$ are Holant-indistinguishable – that is, F and G are homomorphism-indistinguishable. Then, by Lovász's theorem, $F \cong G$.

Theorem 6.4.1 or Theorem 10.3.1 always applies. Here, we cannot apply Theorem 10.3.1 to weakly isomorphic but quantum-vanishing $\mathcal{F}$ and $\mathcal{G}$.

Question 10.6.1. *If $\mathcal{F}$ and $\mathcal{G}$ are weakly isomorphic, is there a $T \in \mathrm{GL}_q$ such that $T\mathcal{F} = \mathcal{G}$?*

A linear-algebraic argument similar to the proof of [Xia10, Theorem 5] shows that the answer to Question 10.6.1 is yes for sets of unary signatures.

Proposition 10.6.1. *Let $\mathcal{F}, \mathcal{G} \subset {}_1\mathcal{V}_0$ and $\mathcal{F}', \mathcal{G}' \subset {}_0\mathcal{V}_1$. Then $\mathcal{F} \,|\, \mathcal{F}'$ and $\mathcal{G} \,|\, \mathcal{G}'$ are weakly isomorphic if and only if there is a $T \in \mathrm{GL}_q$ such that $T(\mathcal{F} \,|\, \mathcal{F}') = \mathcal{G} \,|\, \mathcal{G}'$.*

Proof. We only need ($\Longrightarrow$). The $(1,0)$ (resp. $(0,1)$)-weak-isomorphism $\leftrightsquigarrow$ is a bijective linear map between $\mathrm{span}(\mathcal{F})$ and $\mathrm{span}(\mathcal{G})$ (resp. $\mathrm{span}(\mathcal{F}')$ and $\mathrm{span}(\mathcal{G}')$). Let $F, G \in \mathbb{K}^{q \times k}$ be matrices whose columns form a basis for $\mathrm{span}(\mathcal{F})$ and $\mathrm{span}(\mathcal{G})$, and let $F', G' \in \mathbb{K}^{\ell \times q}$ be matrices whose rows form a basis for $\mathrm{span}(\mathcal{F}')$ and $\mathrm{span}(\mathcal{G}')$, respectively. In this unary setting, the $(1,0)$-weak-isomorphism is a holographic transformation $S \in \mathrm{GL}_q$ such that $SF = G$. Replace $\mathcal{F} \,|\, \mathcal{F}'$ with $S(\mathcal{F} \,|\, \mathcal{F}')$ so that $F = G$. By Proposition 10.3.1, $\mathcal{F} \,|\, \mathcal{F}'$ and $\mathcal{G} \,|\, \mathcal{G}'$ are indistinguishable, which implies $F'F = G'G = G'F$. If $\ell < q$, then choose a vector in $(\mathbb{K}^q \setminus \mathrm{span}(\mathcal{F}')) \cap (\mathbb{K}^q \setminus \mathrm{span}(\mathcal{G}'))$ and add this vector to both $\mathcal{F}'$ and $\mathcal{G}'$. Repeat this process $q-\ell$ times to obtain invertible $F'', G'' \in \mathbb{K}^{q \times q}$. Since $F'F = G'F$ initially, we still have

$$F''F = G''F = G''G.$$

Thus, putting $T := (G'')^{-1}F''$, we have $TF = G$ and $F''T^{-1} = G''$. The latter implies $F'T^{-1} = G'$, so $T(\mathcal{F} \,|\, \mathcal{F}') = \mathcal{G} \,|\, \mathcal{G}'$. □

A tempting potential family of counterexamples to Question 10.6.1 is $\mathcal{F} = \{F \oplus F\}$ and $\mathcal{G} = \{F \oplus 0\}$ for vanishing F. This $\mathcal{F}$ and $\mathcal{G}$ are weakly isomorphic when only considering quantum gadgets with every term connected, but disconnected gadgets cause problems similar to those addressed in the proof of Lemma 10.3.1, which seem more difficult to overcome in this setting.

10.6.2 Complexity of bounded-degree homomorphism indistinguishability

Corollaries 10.5.7 and 10.5.8 give lower bounds on the complexity of Holant-indistinguishability. However, special cases may admit faster algorithms. Corollary 10.5.6 and Theorem 10.2.3 give an algorithm for determining whether two q-vertex graphs F and G are homomorphism-indistinguishable over graphs of degree at most d: for each invariant polynomial p in the finite generating set of $\mathbb{C}[{}_1\mathcal{V}_0 \oplus \ldots \oplus {}_d\mathcal{V}_0 \oplus {}_0\mathcal{V}_2]^{\mathrm{GL}_q}$ guaranteed by Theorem 10.2.2, check if $p(\mathcal{EQ}_{\leq d} \mid A_F) = p(\mathcal{EQ}_{\leq d} \mid A_G)$. There are algorithms to compute the generating set of general $\mathbb{C}[\mathcal{X}]^{\mathrm{GL}_q}$ [DK15; Der99], and there are upper bounds on the largest degree of any such generator [Der01]. However, in general these upper bounds are exponential in the size of the tensors in $\mathcal{X}$ and in certain cases there are exponential lower bounds – see e.g. [Acu+23, Proposition 4.15]. Nevertheless, it is possible that separating $p(\mathcal{EQ}_{\leq d} \mid A_F) = p(\mathcal{EQ}_{\leq d} \mid A_G)$ requires fewer invariants than the general case.

Bibliography

[Acu+23] Arturo Acuaviva, Visu Makam, Harold Nieuwboer, David Perez-Garcia, Friedrich Sittner, Michael Walter, and Freek Witteveen. " The minimal canonical form of a tensor network ". In: *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*. Los Alamitos, CA, USA: IEEE Computer Society, Nov. 2023, pp. 328–362. arXiv: 2209.14358 [quant-ph]. URL: https://arxiv.org/abs/2209.14358.

[All+18] Zeyuan Allen-Zhu, Ankit Garg, Yuanzhi Li, Rafael Oliveira, and Avi Wigderson. "Operator scaling via geodesically convex optimization, invariant theory and polynomial identity testing". In: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2018. Los Angeles, CA, USA: Association for Computing Machinery, 2018, pp. 172–181. ISBN: 9781450355599. DOI: 10.1145/3188745.3188942.

[Ats+19] Albert Atserias, Laura Mančinska, David E. Roberson, Robert Šámal, Simone Severini, and Antonios Varvitsiotis. "Quantum and non-signalling graph isomorphisms". In: *Journal of Combinatorial Theory, Series B* 136 (2019), pp. 289–328. ISSN: 0095-8956.

[Ban05] Teodor Banica. "Quantum automorphism groups of homogeneous graphs". In: *Journal of Functional Analysis* 224 (July 2005), pp. 243–280. DOI: 10.1016/j.jfa.2004.11.002.

[BBC07] Teodor Banica, Julien Bichon, and Gaetan Chenevier. "Graphs having no quantum symmetry". In: *Annales de l'institut Fourier*. Vol. 57. 3. 2007, pp. 955–971.

[BG05] Andrei Bulatov and Martin Grohe. "The complexity of partition functions". In: *Theoretical Computer Science* 348.2 (2005), pp. 148–186. ISSN: 0304-3975. DOI: https://doi.org/10.1016/j.tcs.2005.09.011.

[BI11] Peter Bürgisser and Christian Ikenmeyer. "Geometric complexity theory and tensor rank". In: *Proceedings of the forty-third annual ACM Symposium on Theory of Computing (STOC)*. 2011, pp. 509–518.

[Böh+02] Elmar Böhler, Edith Hemaspaandra, Steffen Reith, and Heribert Vollmer. "Equivalence and Isomorphism for Boolean Constraint Satisfaction". In: *Proceedings of the 16th International Workshop and 11th Annual Conference of the EACSL on Computer Science Logic*. CSL '02. Berlin, Heidelberg: Springer-Verlag, 2002, pp. 412–426. ISBN: 3540442405.

[Böh+04] Elmar Böhler, Edith Hemaspaandra, Steffen Reith, and Heribert Vollmer. "The complexity of Boolean constraint isomorphism". In: *Proceedings of the 21st Annual Symposium on Theoretical Aspects of Computer Science*. STACS '04. Berlin, Heidelberg: Springer, 2004, pp. 164–175.

[Bor+17] Ada Boralevi, Jan Draisma, Emil Horobeţ, and Elina Robeva. "Orthogonal and unitary tensor decomposition from an algebraic perspective". en. In: *Israel Journal of Mathematics* 222.1 (Oct. 2017), pp. 223–260. DOI: 10.1007/s11856-017-1588-6.

[BS09] Teodor Banica and Roland Speicher. "Liberation of orthogonal Lie groups". In: *Advances in Mathematics* 222.4 (Nov. 2009), pp. 1461–1501. ISSN: 0001-8708. DOI: 10.1016/j.aim.2009.06.009.

[Bul+12] Andrei Bulatov, Martin Dyer, Leslie Ann Goldberg, Markus Jalsenius, Mark Jerrum, and David Richerby. "The complexity of weighted and unweighted #CSP". In: *Journal of Computer and System Sciences* 78.2 (Mar. 2012), pp. 681–688. DOI: 10.1016/j.jcss.2011.12.002.

[Bul13] Andrei A. Bulatov. "The Complexity of the Counting Constraint Satisfaction Problem". In: *J. ACM* 60.5 (Oct. 2013). DOI: 10.1145/2528400.

[Bür+19] Peter Bürgisser, Cole Franks, Ankit Garg, Rafael Oliveira, Michael Walter, and Avi Wigderson. "Towards a theory of non-commutative optimization: Geodesic 1st and 2nd order methods for moment maps and polytopes". In: *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2019, pp. 845–861.

[Cai+15] Jin-Yi Cai, Zhiguo Fu, Heng Guo, and Tyson Williams. "A Holant dichotomy: is the FKT algorithm universal?" In: *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*. IEEE. 2015, pp. 1259–1276.

[CC07] Jin-Yi Cai and Vinay Choudhary. "Valiant's Holant theorem and matchgate tensors". In: *Theoretical Computer Science* 384.1 (2007), pp. 22–32.

[CC17a] Jin-Yi Cai and Xi Chen. *Complexity Dichotomies for Counting Problems*. Vol. 1. Cambridge: Cambridge University Press, 2017. DOI: 10.1017/9781107477063.002.

[CC17b] Jin-Yi Cai and Xi Chen. "Complexity of Counting CSP with Complex Weights". In: *J. ACM* 64.3 (June 2017). DOI: 10.1145/2822891.

[CC19] Jin-Yi Cai and Xi Chen. "A decidable dichotomy theorem on directed graph homomorphisms with non-negative weights". In: *Computational Complexity* 28.3 (2019), pp. 345–408.

[CCL13] Jin-Yi Cai, Xi Chen, and Pinyan Lu. "Graph Homomorphisms with Complex Values: A Dichotomy Theorem". In: *SIAM Journal on Computing* 42.3 (2013), pp. 924–1029. DOI: 10.1137/110840194.

[CCL16] Jin-Yi Cai, Xi Chen, and Pinyan Lu. "Nonnegative Weighted #CSP: An Effective Complexity Dichotomy". In: *SIAM J. Comput.* 45.6 (2016), pp. 2177–2198. DOI: 10.1137/15M1032314.

[CF22] Jin-Yi Cai and Zhiguo Fu. "Holographic Algorithm with Matchgates Is Universal for Planar #CSP over Boolean Domain". In: *SIAM Journal on Computing* 51.2 (2022). DOI: 10.1137/17M1131672.

[CFS] Jin-Yi Cai, Zhiguo Fu, and Shuai Shao. "New Planar P-time Computable Six-Vertex Models and a Complete Complexity Classification". In: *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1535–1547. DOI: 10.1137/1.9781611976465.93.

[CG20] Jin-Yi Cai and Artem Govorov. "Dichotomy for graph homomorphisms with complex values on bounded degree graphs". In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2020, pp. 1103–1111. arXiv: 2004.06620 [cs.CC]. URL: https://arxiv.org/abs/2004.06620.

[CG21] Jin-Yi Cai and Artem Govorov. "On a Theorem of Lovász that $\hom(.,H)$ Determines the Isomorphism Type of H". In: *ACM Trans. Comput. Theory* 13.2 (2021). DOI: 10.1145/3448641.

[CG22] Jin-Yi Cai and Artem Govorov. "Perfect matchings, rank of connection tensors and graph homomorphisms". In: *Combinatorics, Probability and Computing* 31.2 (2022), pp. 268–303. DOI: 10.1017/S0963548321000286.

[CGW16] Jin-Yi Cai, Heng Guo, and Tyson Williams. "A Complete Dichotomy Rises from the Capture of Vanishing Signatures". In: *SIAM Journal on Computing* 45.5 (2016), pp. 1671–1728. DOI: 10.1137/15M1049798.

[Che+24] Zhili Chen, Joshua A. Grochow, Youming Qiao, Gang Tang, and Chuanqi Zhang. "On the Complexity of Isomorphism Problems for Tensors, Groups, and Polynomials III: Actions by Classical Groups". In: *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*. Ed. by Venkatesan Guruswami. Vol. 287. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024, 31:1–31:23. ISBN: 978-3-95977-309-6. DOI: 10.4230/LIPIcs.ITCS.2024.31.

[CI25] Jin-Yi Cai and Jin Soo Ihm. "Holant* Dichotomy on Domain Size 3: A Geometric Perspective". In: *52nd International Colloquium on Automata, Languages, and Programming, ICALP 2025, July 8-11, 2025, Aarhus, Denmark*. Ed. by Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis. Vol. 334. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025, 148:1–148:18. DOI: 10.4230/LIPICS.ICALP.2025.148.

[CL09] Jin-Yi Cai and Pinyan Lu. "Holographic algorithms: The power of dimensionality resolved". In: *Theoretical Computer Science* 410.18 (2009), pp. 1618–1628. ISSN: 0304-3975. DOI: https://doi.org/10.1016/j.tcs.2008.12.047.

[CL11] Jin-Yi Cai and Pinyan Lu. "Holographic algorithms: From art to science". In: *Journal of Computer and System Sciences* 77.1 (2011), pp. 41–61. ISSN: 0022-0000. DOI: https://doi.org/10.1016/j.jcss.2010.06.005.

[CLX08] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. "Holographic algorithms by Fibonacci gates and holographic reductions for hardness". In: *2008 49th Annual IEEE Symposium on Foundations of Computer Science*. IEEE. 2008, pp. 644–653.

[CLX11] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. "Computational Complexity of Holant Problems". In: *SIAM Journal on Computing* 40.4 (2011), pp. 1101–1132. DOI: 10.1137/100814585.

[CLX13] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. "Dichotomy for Holant problems with a function on domain size 3". In: *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2013, pp. 1278–1295.

[CLX17] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. "Holographic Algorithms with Matchgates Capture Precisely Tractable Planar #CSP". In: *SIAM Journal on Computing* 46.3 (2017), pp. 853–889. DOI: 10.1137/16M1073984.

[CM23] Jin-Yi Cai and Ashwin Maran. "The Complexity of Counting Planar Graph Homomorphisms of Domain Size 3". In: *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. STOC 2023. New York, NY, USA: Association for Computing Machinery, 2023, pp. 1285–1297. ISBN: 9781450399135. DOI: 10.1145/3564246.3585173.

[CM24a] Jin-Yi Cai and Ashwin Maran. *Polynomial and analytic methods for classifying complexity of planar graph homomorphisms*. 2024. arXiv: 2412.17122 [cs.CC]. URL: https://arxiv.org/abs/2412.17122.

[CM24b] Ada Chan and William J. Martin. "Quantum isomorphism of graphs from association schemes". In: *Journal of Combinatorial Theory, Series B* 164 (2024), pp. 340–363. ISSN: 0095-8956. DOI: https://doi.org/10.1016/j.jctb.2023.09.005.

[CMY26] Jin-Yi Cai, Ashwin Maran, and Ben Young. *Planar Graph Homomorphisms: A Dichotomy and a Barrier from Quantum Groups*. 2026. arXiv: 2601.23198 [cs.CC]. URL: https://arxiv.org/abs/2601.23198.

[CS24] Jin-Yi Cai and Daniel P. Szabo. "Bounded Degree Nonnegative Counting CSP". In: *ACM Trans. Comput. Theory* 16.2 (Mar. 2024). ISSN: 1942-3454. DOI: 10.1145/3632184.

[CY24] Jin-Yi Cai and Ben Young. "Planar #CSP Equality Corresponds to Quantum Isomorphism – A Holant Viewpoint". In: *ACM Transactions on Computation Theory* 16.3 (Sept. 2024).

[CY26] Jin-Yi Cai and Ben Young. "Vanishing Signatures, Orbit Closure, and the Converse of the Holant Theorem". In: *17th Innovations in Theoretical Computer Science Conference (ITCS 2026)*. Vol. 362. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2026, 32:1–32:20. DOI: 10.4230/LIPIcs.ITCS.2026.32.

[Der01] Harm Derksen. "Polynomial bounds for rings of invariants". In: *Proceedings of the American Mathematical Society* 129.4 (2001), pp. 955–963.

[Der99] Harm Derksen. "Computation of invariants for reductive groups". In: *Advances in Mathematics* 141.2 (1999), pp. 366–384.

[DG00] Martin Dyer and Catherine Greenhill. "The complexity of counting graph homomorphisms". In: *Random Structures & Algorithms* 17.3-4 (2000), pp. 260–289. ISSN: 1042-9832.

[DGR18] Holger Dell, Martin Grohe, and Gaurav Rattan. "Lovász Meets Weisfeiler and Leman". In: *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*. Ed. by Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella. Vol. 107. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018, 40:1–40:14. DOI: 10.4230/LIPIcs.ICALP.2018.40.

[DK15] Harm Derksen and Gregor Kemper. *Computational Invariant Theory*. 2nd ed. Encyclopaedia of Mathematical Sciences. Berlin, Heidelberg: Springer, 2015. DOI: https://doi.org/10.1007/978-3-662-48422-7.

[DM20] Harm Derksen and Visu Makam. "Algorithms for orbit closure separation for invariants and semi-invariants of matrices". In: *Algebra & Number Theory* 14.10 (2020), pp. 2791–2813.

[DM22] Harm Derksen and Visu Makam. "Polystability in positive characteristic and degree lower bounds for invariant rings". In: *Journal of Combinatorial Algebra* 6.3 (2022), pp. 353–405.

[DM23] Harm Derksen and Visu Makam. "Invariant Theory and wheeled PROPs". In: *Journal of Pure and Applied Algebra* 227.9 (2023), p. 107302. ISSN: 0022-4049. DOI: https://doi.org/10.1016/j.jpaa.2022.107302.

[DR13a] Jan Draisma and Guus Regts. "Tensor invariants for certain subgroups of the orthogonal group". en. In: *Journal of Algebraic Combinatorics* 38.2 (Sept. 2013), pp. 393–405. DOI: 10.1007/s10801-012-0408-7.

[DR13b] Martin Dyer and David Richerby. "An Effective Dichotomy for the Counting Constraint Satisfaction Problem". In: *SIAM Journal on Computing* 42.3 (2013), pp. 1245–1274. DOI: 10.1137/100811258.

[Dra+12] Jan Draisma, Dion C. Gijswijt, László Lovász, Guus Regts, and Alexander Schrijver. "Characterizing partition functions of the vertex model". In: *Journal of Algebra* 350.1 (Jan. 2012), pp. 197–206. DOI: 10.1016/j.jalgebra.2011.10.030.

[DRS25] Josse van Dobben de Bruyn, David E. Roberson, and Simon Schmidt. "Asymmetric graphs with quantum symmetry". en. In: *Proceedings of the London Mathematical Society* 131.5 (2025). ISSN: 1460-244X. DOI: 10.1112/plms.70098.

[Dvo10] Zdeněk Dvořák. "On recognizing graphs by numbers of homomorphisms". en. In: *Journal of Graph Theory* 64.4 (2010), pp. 330–342. DOI: 10.1002/jgt.20461.

[FLS07] Michael Freedman, László Lovász, and Alexander Schrijver. "Reflection Positivity, Rank Connectivity, and Homomorphism of Graphs". In: *Journal of the American Mathematical Society* 20.1 (2007), pp. 37–51.

[Gar+20] Ankit Garg, Christian Ikenmeyer, Visu Makam, Rafael Oliveira, Michael Walter, and Avi Wigderson. "Search problems in algebraic complexity, GCT, and hardness of generators for invariant rings". In: *Proceedings of the 35th Computational Complexity Conference*. CCC '20. Virtual Event, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2020. ISBN: 9783959771566. DOI: 10.4230/LIPIcs.CCC.2020.12. URL: https://doi.org/10.4230/LIPIcs.CCC.2020.12.

[GCD23] Artem Govorov, Jin-Yi Cai, and Martin Dyer. "A dichotomy for bounded degree graph homomorphisms with nonnegative weights". In: *Journal of Computer and System Sciences* 132 (2023), pp. 1–15. ISSN: 0022-0000. DOI: https://doi.org/10.1016/j.jcss.2022.09.002.

[Gol+10] Leslie Ann Goldberg, Martin Grohe, Mark Jerrum, and Marc Thurley. "A Complexity Dichotomy for Partition Functions with Mixed Signs". In: *SIAM J. Comput.* 39.7 (Jan. 2010), pp. 3336–3402. DOI: 10.1137/090757496.

[GQ23] Joshua Grochow and Youming Qiao. "On the Complexity of Isomorphism Problems for Tensors, Groups, and Polynomials I: Tensor Isomorphism-Completeness". In: *SIAM Journal on Computing* 52.2 (2023), pp. 568–617. DOI: 10.1137/21M1441110.

[GQ25] Joshua A. Grochow and Youming Qiao. "On the Complexity of Isomorphism Problems for Tensors, Groups, and Polynomials IV: Linear-Length Reductions and Their Applications". In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. STOC '25. Prague, Czechia: Association for Computing Machinery, 2025, pp. 766–776. ISBN: 9798400715105. DOI: 10.1145/3717823.3718282.

[Gro21] Daniel Gromada. *Quantum symmetries of Cayley graphs of abelian groups*. 2021. DOI: 10.48550/ARXIV.2106.08787. URL: https://arxiv.org/abs/2106.08787.

[GRS25] Martin Grohe, Gaurav Rattan, and Tim Seppelt. "Homomorphism Tensors and Linear Equations". In: *Advances in Combinatorics* (Apr. 2025). DOI: 10.19086/aic.2025.4.

[GRV16] Andrew Goodall, Guus Regts, and Lluís Vena. "Matroid invariants and counting graph homomorphisms". In: *Linear Algebra and its Applications* 494 (Apr. 2016), pp. 263–273. DOI: 10.1016/j.laa.2016.01.022.

[GW09] Roe Goodman and Nolan R. Wallach. *Symmetry, Representations, and Invariants*. Graduate Texts in Mathematics. New York: Springer, 2009. DOI: https://doi.org/10.1007/978-0-387-79852-3.

[Hil90] David Hilbert. "Ueber dieTheorie der algebraischen Formen". In: *Mathematische Annalen* 36 (1890), pp. 473–534. URL: http://eudml.org/doc/157506.

[HJ93] P. de la Harpe and V. F. R. Jones. "Graph invariants related to statistical mechanical models: examples and problems". In: *J. Comb. Theory Ser. B* 57.2 (Mar. 1993), pp. 207–227. DOI: 10.1006/jctb.1993.1017.

[HL16] Sangxia Huang and Pinyan Lu. "A dichotomy for real weighted Holant problems". In: *Computational Complexity* 25 (2016), pp. 255–304.

[IQ23] Gábor Ivanyos and Youming Qiao. "On the orbit closure intersection problems for matrix tuples under conjugation and left-right actions". In: *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2023, pp. 4115–4126.

[Jin15] Naihuan Jing. "Unitary and orthogonal equivalence of sets of matrices". In: *Linear Algebra and its Applications* 481 (2015), pp. 235–242. ISSN: 0024-3795. DOI: https://doi.org/10.1016/j.laa.2015.04.036.

[Kap72] Irving Kaplansky. *Fields and Rings*. 2nd ed. Chicago Lectures in Mathematics. Chicago: University of Chicago Press, 1972.

[Kar+25] Prem Nigam Kar, David E. Roberson, Tim Seppelt, and Peter Zeman. "NPA Hierarchy for Quantum Isomorphism and Homomorphism Indistinguishability". In: *52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025)*. Ed. by Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis. Vol. 334. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025, 105:1–105:19. ISBN: 978-3-95977-372-0. DOI: 10.4230/LIPIcs.ICALP.2025.105.

[Kas61] Pieter W Kasteleyn. "The statistics of dimers on a lattice: I. The number of dimer arrangements on a quadratic lattice". In: *Physica* 27.12 (1961), pp. 1209–1225.

[Lan17] J. M. Landsberg. *Geometry and Complexity Theory*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2017.

[Liu24] Yin Liu. *A combinatorial view of Holant problems on higher domains*. 2024. arXiv: 2403.14150 [cs.CC]. URL: https://arxiv.org/abs/2403.14150.

[LMR17] Martino Lupini, Laura Mančinska, and David Roberson. "Nonlocal Games and Quantum Permutation Groups". In: *Journal of Functional Analysis* 279 (Dec. 2017). DOI: 10.1016/j.jfa.2020.108592.

[Lov06] László Lovász. "The rank of connection matrices and the dimension of graph algebras". In: *European Journal of Combinatorics* 27.6 (2006), pp. 962–970.

[Lov67] László Lovász. “Operations with structures”. In: *Acta Mathematica Hungarica* 18.3-4 (1967), pp. 321–328.

[LS09] László Lovász and Balázs Szegedy. “Contractors and connectors of graph algebras”. In: *Journal of Graph Theory* 60.1 (Jan. 2009), pp. 11–30. DOI: 10.1002/jgt.20343.

[LW18] Jiabao Lin and H. Wang. “The Complexity of Boolean Holant Problems with Nonnegative Weights”. In: *SIAM Journal on Computing* 47.3 (2018), pp. 798–828. DOI: 10.1137/17M113304X.

[LW25] Vladimir Lysikov and Michael Walter. “Complexity theory of orbit closure intersection for tensors: reductions, completeness, and graph isomorphism hardness”. In: *2025 IEEE 66th Annual Symposium on Foundations of Computer Science (FOCS)*. 2025, pp. 486–511. DOI: 10.1109/FOCS63196.2025.00027. arXiv: 2411.04639 [cs.CC]. URL: https://arxiv.org/abs/2411.04639.

[Mac65] Saunders MacLane. “Categorical algebra”. In: *Bulletin of the American Mathematical Society* 71.1 (1965), pp. 40–106.

[McC23] JP McCarthy. “Tracing the orbitals of the quantum permutation group”. In: *Archiv der Mathematik* 121.2 (2023), pp. 211–224.

[MFK94] David Mumford, John Fogarty, and Frances Kirwan. *Geometric Invariant Theory*. 3rd ed. Vol. 34. Ergebnisse der Mathematik und ihrer Grenzgebiete. Berlin, Heidelberg: Springer, 1994. ISBN: 978-3-540-56963-3.

[MMS09] M. Markl, S. Merkulov, and S. Shadrin. “Wheeled PROPs, graph complexes and the master equation”. In: *Journal of Pure and Applied Algebra* 213.4 (2009), pp. 496–535. ISSN: 0022-4049. DOI: https://doi.org/10.1016/j.jpaa.2008.08.007.

[MR20] Laura Mančinska and David E. Roberson. “Quantum isomorphism is equivalent to equality of homomorphism counts from planar graphs”. In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. 2020, pp. 661–672. DOI: 10.1109/FOCS46700.2020.00067. arXiv: 1910.06958 [quant-ph]. URL: https://arxiv.org/abs/1910.06958.

[NT14] Sergey Neshveyev and Lars Tuset. *Compact Quantum Groups and Their Representation Categories*. Vol. 20. Cours Spécialisés. Société Mathématique de France, 2014. ISBN: 9782856297773.

[Reg12] Guus Regts. “The rank of edge connection matrices and the dimension of algebras of invariant tensors”. In: *European Journal of Combinatorics* 33.6 (Aug. 2012), pp. 1167–1173. DOI: 10.1016/j.ejc.2012.01.014.

[Reg13a] Guus Regts. “A characterization of edge-reflection positive partition functions of vertex-coloring models”. en. In: *The Seventh European Conference on Combinatorics, Graph Theory and Applications*. CRM Series. Pisa: Scuola Normale Superiore, 2013, pp. 305–311. ISBN: 978-88-7642-475-5. DOI: 10.1007/978-88-7642-475-5_49.

[Reg13b] Guus Regts. “Graph parameters and invariants of the orthogonal group”. PhD thesis. Universiteit van Amsterdam, 2013.

[Rob16] Elina Robeva. “Orthogonal Decomposition of Symmetric Tensors”. In: *SIAM Journal on Matrix Analysis and Applications* 37.1 (Jan. 2016), pp. 86–102. DOI: 10.1137/140989340.

[Rob22] David E. Roberson. *Oddomorphisms and homomorphism indistinguishability over graphs of bounded degree*. 2022. arXiv: 2206.10321 [math.CO]. URL: https://arxiv.org/abs/2206.10321.

[RS23] Gaurav Rattan and Tim Seppelt. “Weisfeiler-Leman and Graph Spectra”. In: *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2023, pp. 2268–2285. DOI: 10.1137/1.9781611977554.ch87.

[RS24] David E. Roberson and Tim Seppelt. “Lasserre Hierarchy for Graph Isomorphism and Homomorphism Indistinguishability”. In: *TheoretiCS* Volume 3, 20 (Sept. 2024). ISSN: 2751-4838. DOI: 10.46298/theoretics.24.20. URL: https://theoretics.episciences.org/12321.

[RY15] Heydar Radjavi and Bamdad R. Yahaghi. “A theorem of Kaplansky revisited”. In: *Linear Algebra and its Applications* 487 (2015), pp. 268–275. ISSN: 0024-3795. DOI: https://doi.org/10.1016/j.laa.2015.09.024.

[SC20] Shuai Shao and Jin-Yi Cai. “A Dichotomy for Real Boolean Holant Problems”. In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. 2020, pp. 1091–1102. DOI: 10.1109/FOCS46700.2020.00105. arXiv: 2005.07906 [cs.CC]. URL: https://arxiv.org/abs/2005.07906.

[Sch08a] Alexander Schrijver. “Graph Invariants in the Edge Model”. In: *Building Bridges: Between Mathematics and Computer Science*. Ed. by Martin Grötschel, Gyula O. H. Katona, and Gábor Sági. Berlin, Heidelberg: Springer, 2008, pp. 487–498. DOI: 10.1007/978-3-540-85221-6_16.

[Sch08b] Alexander Schrijver. “Tensor subalgebras and first fundamental theorems in invariant theory”. In: *Journal of Algebra* 319.3 (Feb. 2008), pp. 1305–1319. DOI: 10.1016/j.jalgebra.2007.10.039.

[Sch09] Alexander Schrijver. “Graph invariants in the spin model”. In: *Journal of Combinatorial Theory, Series B* 99.2 (2009), pp. 502–511. DOI: https://doi.org/10.1016/j.jctb.2008.10.003.

[Sel10] P. Selinger. “A Survey of Graphical Languages for Monoidal Categories”. In: *New Structures for Physics*. Springer Berlin Heidelberg, 2010, pp. 289–355. ISBN: 9783642128219. DOI: 10.1007/978-3-642-12821-9_4. URL: http://dx.doi.org/10.1007/978-3-642-12821-9_4.

[Sep24] Tim Frederik Seppelt. “Homomorphism indistinguishability”. Veröffentlicht auf dem Publikationsserver der RWTH Aachen University 2025; Dissertation, RWTH Aachen University, 2024. Dissertation. Aachen: RWTH Aachen University, 2024, 1 Online–Ressource : Illustrationen. DOI: 10.18154/RWTH-2024-11629. URL: https://publications.rwth-aachen.de/record/998785.

[Slo19] William Slofstra. “The Set of Quantum Correlations is not Closed”. In: *Forum of Mathematics, Pi* 7 (2019).

[Spe40] Wilhelm Specht. “Zur Theorie der Matrizen. II.” ger. In: *Jahresbericht der Deutschen Mathematiker-Vereinigung* 50 (1940), pp. 19–23.

[SS26] Tim Seppelt and Gian Luca Spitzer. *Homomorphism Indistinguishability Relations induced by Quantum Groups*. 2026. arXiv: 2505.07922 [quant-ph]. URL: https://arxiv.org/abs/2505.07922.

[Sze07] Balázs Szegedy. “Edge coloring models and reflection positivity”. en. In: *Journal of the American Mathematical Society* 20.4 (May 2007), pp. 969–988. DOI: 10.1090/S0894-0347-07-00568-1.

[TF61] Harold NV Temperley and Michael E Fisher. “Dimer problem in statistical mechanics-an exact result”. In: *Philosophical Magazine* 6.68 (1961), pp. 1061–1063.

[Val06] Leslie G Valiant. “Accidental algorthims”. In: *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS’06)*. IEEE. 2006, pp. 509–517.

[Val08] Leslie G. Valiant. “Holographic algorithms”. In: *SIAM Journal on Computing* 37.5 (2008), pp. 1565–1594. DOI: 10.1137/070682575.

[Wan95] Shuzhou Wang. “Free products of compact quantum groups”. en. In: *Communications in Mathematical Physics* 167.3 (Feb. 1995), pp. 671–692. ISSN: 1432-0916. DOI: 10.1007/BF02101540.

[Wan98] Shuzhou Wang. “Quantum Symmetry Groups of Finite Spaces”. en. In: *Communications in Mathematical Physics* 195.1 (July 1998), pp. 195–211. ISSN: 1432-0916. DOI: 10.1007/s002200050385.

[Wey66] Hermann Weyl. *The Classical Groups: Their Invariants and Representations*. Princeton University Press, 1966. ISBN: 978-0-691-05756-9. DOI: 10.2307/j.ctv3hh48t.

[Wie61] NA Wiegmann. “Necessary and sufficient conditions for unitary similarity”. In: *Journal of the Australian Mathematical Society* 2.1 (1961), pp. 122–126.

[Wor87] S. L. Woronowicz. "Compact matrix pseudogroups". en. In: *Communications in Mathematical Physics* 111.4 (Dec. 1987), pp. 613–665. DOI: 10.1007/BF01219077.

[Wor88] S. L. Woronowicz. "Tannaka-Krein duality for compact matrix pseudogroups. TwistedSU(N) groups". en. In: *Inventiones mathematicae* 93.1 (Feb. 1988), pp. 35–76.

[Xia10] Mingji Xia. "Holographic reduction: A domain changed application and its partial converse theorems". In: *Automata, Languages and Programming: 37th International Colloquium, ICALP 2010, Bordeaux, France, July 6-10, 2010, Proceedings, Part I 37*. Springer. 2010, pp. 666–677.

[You25a] Ben Young. "Equality on all #CSP Instances Yields Constraint Function Isomorphism via Interpolation and Intertwiners". In: *The Electronic Journal of Combinatorics* 32.4 (2025). DOI: https://doi.org/10.37236/13879.

[You25b] Ben Young. "The Converse of the Real Orthogonal Holant Theorem". In: *52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025)*. Vol. 334. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025, 136:1–136:20. DOI: 10.4230/LIPIcs.ICALP.2025.136. arXiv: 2409.06911 [cs.DM]. URL: https://arxiv.org/abs/2409.06911.