Quantum advantage and CSP complexity^{*}

Lorenzo Ciardo University of Oxford lorenzo.ciardo@cs.ox.ac.uk

Abstract

Information-processing tasks modelled by homomorphisms between relational structures can witness quantum advantage when entanglement is used as a computational resource. We prove that the occurrence of quantum advantage is determined by the same type of algebraic structure (known as a minion) that captures the polymorphism identities of CSPs and, thus, CSP complexity. We investigate the connection between the minion of quantum advantage and other known minions controlling CSP tractability and width. In this way, we make use of complexity results from the algebraic theory of CSPs to characterise the occurrence of quantum advantage in the case of graphs, and to obtain new necessary and sufficient conditions in the case of arbitrary relational structures.

1 Introduction

The constraint satisfaction problem (CSP) is a computational paradigm that provides a common framework for studying the complexity of a variety of combinatorial problems. The input consists in a set of variables, a set of possible values for the variables, and a set of constraints between the variables; the goal is to determine whether there exists an assignment of values to the variables that satisfies the constraints. In its full generality, the CSP is NP-complete. Nevertheless, restricting the admissible constraints to a selected class can make the problem solvable in polynomial time. This is the case of linear equations, which can be formulated as a CSP where all constraints are linear relations between the variables. The theory of CSP complexity aims at explaining the complexity of a CSP over some specific type of constraints in terms of the structure of the constraints.

An elegant way of formalising the CSP is through the notion of homomorphisms between relational structures. A relational structure consists of a set of vertices and a set of relations on the vertices; for example, a graph is a relational structure having a single binary symmetric irreflexive relation. A homomorphism $\mathbf{X} \to \mathbf{Y}$ between two relational structures \mathbf{X} and \mathbf{Y} is a map between the vertex sets of \mathbf{X} and \mathbf{Y} that preserves all relations, in the sense that the image of a tuple in a relation of \mathbf{X} must belong to the corresponding relation of \mathbf{Y} . We can then formulate the CSP *parameterised* by \mathbf{Y} as the problem of deciding if an arbitrary structure \mathbf{X} admits a homomorphism to \mathbf{Y} . In this setting, the relations of \mathbf{X} encode the tuples of variables that appear in some constraint of the CSP, while the relations of \mathbf{Y} encode the types of admissible constraints.

^{*}This work will appear in the Proceedings of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'24). The research leading to these results was supported by UKRI EP/X024431/1.

For example, linear equations over a finite field \mathbb{F} of n elements correspond to CSP(**Y**), where the domain of **Y** is $\{1, \ldots, n\}$, and the relations of **Y** are affine hyperplanes in vector spaces over \mathbb{F} . If **Y** is the *n*-clique \mathbf{K}_n , CSP(**Y**) corresponds to another primary example of a CSP — namely, the graph *n*-colouring problem, which consists in testing whether the chromatic number of a graph **X** is at most *n*. A natural generalisation is the graph **Y**-colouring problem, which corresponds to CSP(**Y**) for some fixed graph **Y**. The complexity of graph *n*-colouring was characterised by Karp [48]: If n = 2, the problem is in P; otherwise, it is NP-complete. Hell and Nešetřil [41] extended the dichotomy by proving that graph **Y**-colouring is in P if **Y** is bipartite, and it is NP-complete otherwise. In [34], Feder and Vardi proposed the conjecture that Karp's and Hell–Nešetřil's dichotomies could in fact extend to all CSPs. Feder–Vardi's Dichotomy Conjecture inspired a major research programme in the theory of CSPs, that culminated twenty years later with its positive resolution obtained by Bulatov [21] and, independently, by Zhuk [60].

The advent of quantum computation led to the natural programme of understanding the type of advantage derived by having access to quantum resources — as opposed to classical resources — in computational tasks. In order to formalise this notion of *quantum advantage* for computational tasks consisting in constraint satisfaction problems, it shall be convenient to adopt an alternative description of homomorphisms between relational structures, in terms of winning strategies for non-local games.

For the sake of simplicity, suppose that \mathbf{X} and \mathbf{Y} are two graphs. The \mathbf{XY} -homomorphism game, introduced in [53], involves two players (Alice and Bob) and a Verifier. The Verifier sends Alice and Bob vertices x_A and x_B of \mathbf{X} , respectively. Alice and Bob (who cannot communicate during the game, but can agree on a strategy before the game starts) respond by sending the Verifier vertices y_A and y_B of \mathbf{Y} , respectively. They win the \mathbf{XY} -homomorphism game if y_A and y_B are adjacent (resp. equal) vertices of \mathbf{Y} whenever x_A and x_B are adjacent (resp. equal) vertices of \mathbf{X} . If only resources of classical type are allowed, the existence of a perfect strategy for this game — i.e., a strategy that makes Alice and Bob win with probability 1 — is equivalent to the existence of a homomorphism $\mathbf{X} \to \mathbf{Y}$. In fact, a classical perfect strategy simply consists in Alice and Bob sharing a protocol for selecting some homomorphism $f: \mathbf{X} \to \mathbf{Y}$, and responding with the vertex f(x) whenever they are sent a vertex x by the Verifier.

Suppose now that Alice and Bob have access to quantum resources, modelled by an entangled state in the space $\mathbb{H}_A \otimes \mathbb{H}_B$, where $\mathbb{H}_A \simeq \mathbb{H}_B$ are finite-dimensional (real or complex) Hilbert spaces. This state is shared by Alice and Bob; the separation between the two players corresponds to the fact that Alice (resp. Bob) can only perform measurements on \mathbb{H}_A (resp. \mathbb{H}_B). This extra resource allows the players to devise quantum-assisted strategies: Each player performs a positive operator-valued measure of their part of the state, which results in a vertex of **Y** that is then sent to the Verifier. A quantum perfect strategy over the space $\mathbb{H} = \mathbb{H}_A \simeq \mathbb{H}_B$ for the **XY**-homomorphism game is a quantum-assisted strategy that makes Alice and Bob win with probability 1. In this case, **X** and **Y** are said to be quantum homomorphic over \mathbb{H} . The extra correlation between the answers of Alice and Bob produced by the entanglement of the quantum state can result in strategies that are strictly stronger than classical strategies. In other words, it is possible for two graphs **X** and **Y** to be quantum homomorphic even when they are not homomorphic in the classical sense. In [53], the question of understanding for which graphs **Y** the notions of quantum advantage if there exists some graph **X** such that **X**

and Y are quantum homomorphic but not classically homomorphic.

Question 1 ([53]). Characterise the class of graphs Y having quantum advantage.

The notions of quantum perfect strategy, quantum homomorphism, and quantum advantage that we consider in this work extend those described above, in that they apply to arbitrary relational structures as opposed to only graphs. This extension was recently introduced in [2], and it involves a modified version of the **XY**-homomorphism game discussed above, where the roles of Alice and Bob are asymmetric: Alice receives a tuple \mathbf{x}_A in some relation of **X**, and returns a tuple \mathbf{y}_A in the corresponding relation of **Y**; Bob receives a vertex x_B of **X** and returns a vertex y_B of **Y**. They win the game if, whenever x_B appears in the tuple \mathbf{x}_A , y_B appears in the corresponding position of the tuple \mathbf{y}_B .¹

The main conceptual contribution of this work is to establish that the occurrence of quantum advantage is governed by the same type of algebraic structure that determines the complexity of constraint satisfaction problems — namely, the algebraic structure known as a minion. We start by discussing the role of minions in CSP complexity.

The mathematical framework that allowed attacking and, eventually, proving Feder-Vardi's Dichotomy Conjecture was the so-called *algebraic approach to CSPs*, whose origins can be traced back to [45, 46, 56]. The gist of this approach is that the complexity of a CSP is entirely determined by a type of identities satisfied by its higher-order symmetries, known as *polymorphisms*. A polymorphism of a relational structure \mathbf{Y} is a homomorphism $\mathbf{Y}^{\ell} \to \mathbf{Y}$, where \mathbf{Y}^{ℓ} is the ℓ -fold direct power of \mathbf{Y} . If \mathbf{Y} and \mathbf{Y}' are two structures such that the set $Pol(\mathbf{Y})$ of the polymorphisms of \mathbf{Y} is included in $Pol(\mathbf{Y}')$ — thus meaning that \mathbf{Y}' has "more symmetries" than \mathbf{Y} — there exists a specific kind of polynomial-time reduction (known as a *gadget reduction*) from $CSP(\mathbf{Y}')$ to $CSP(\mathbf{Y})$ (in symbols, $CSP(\mathbf{Y}') \leq_g CSP(\mathbf{Y})$). This gives a powerful way of comparing the complexity of different CSPs, by looking at the properties of their polymorphisms. The implication

$$\operatorname{Pol}(\mathbf{Y}) \subseteq \operatorname{Pol}(\mathbf{Y}') \quad \Rightarrow \quad \operatorname{CSP}(\mathbf{Y}') \leq_{\operatorname{g}} \operatorname{CSP}(\mathbf{Y})$$

can in fact be strengthened: As established in [11], in order for a gadget reduction to exist, it is enough for $Pol(\mathbf{Y}')$ to satisfy all identities of a certain type (known as *height-1* or *minor* identities) that are satisfied in $Pol(\mathbf{Y})$. One can then endow $Pol(\mathbf{Y})$ with a particular algebraic structure (known as *minion* or *minor-closed class* [56]) that encodes minor identities. This results in the stronger implication

$$\operatorname{Pol}(\mathbf{Y}) \to \operatorname{Pol}(\mathbf{Y}') \quad \Rightarrow \quad \operatorname{CSP}(\mathbf{Y}') \leq_{\mathrm{g}} \operatorname{CSP}(\mathbf{Y}),$$

where " \rightarrow " denotes a minor-preserving map (or *minion homomorphism*) between the minions $Pol(\mathbf{Y})$ and $Pol(\mathbf{Y}')$. Furthermore, it was shown in [11] that the converse of the implication above also holds. In this sense, the polymorphism minion precisely captures CSP complexity.

Contributions Our first contribution is to show that quantum advantage is a property of the polymorphism minion, just like CSP complexity. To that end, we construct a minion

¹If **X** and **Y** are graphs, it was noted in [2] that this definition results in a stronger notion of quantum homomorphism than the one from [53]. As we shall see in Section 6, the results of this paper apply to both definitions.

 $\mathscr{Q}_{\mathbb{H}}$ having the property that a structure **Y** has quantum advantage over \mathbb{H} if and only if $\mathscr{Q}_{\mathbb{H}} \not\rightarrow \operatorname{Pol}(\mathbf{Y})$. This quantum minion belongs to a particular class of minions identified in [28], known as *linear minions*. Its objects are tuples of orthogonal subspaces of the Hilbert space \mathbb{H} , while the minors are defined in terms of sums of linear spaces.

Using this connection, the question "Which structures have quantum advantage?" can be reformulated as "Which minions admit a homomorphism from the quantum minion?". In other words, the problem of classifying the occurrence of quantum advantage is translated into the problem of locating $\mathcal{Q}_{\mathbb{H}}$ within the preorder on the class of minions induced by the minion homomorphism relation. We prove the following three results:

(i) If \mathbb{H} has dimension 1 or 2, $\mathcal{Q}_{\mathbb{H}}$ is equivalent to the so-called *dictator minion* \mathcal{D} (also known as projection minion). To show this, we make use of the Hopf fibration of the complex projective line $\mathbb{C}\mathbf{P}^1$ into the 2-sphere S^2 , which has the property of mapping pairs of orthogonal vectors into antipodal points. Then, we consider an antipodal partition of S^2 . In this way, we can select in any element of $\mathcal{Q}_{\mathbb{H}}$ a distinguished coordinate — namely, the coordinate whose corresponding subspace is mapped by the Hopf fibration to, say, the first part of the partition. The map assigning to each element of $\mathcal{Q}_{\mathbb{H}}$ its distinguished coordinate is then shown to be a minion homomorphism from $\mathcal{Q}_{\mathbb{H}}$ to \mathcal{D} .

(*ii*) If \mathbb{H} has dimension at least 3, $\mathscr{Q}_{\mathbb{H}}$ is not equivalent to the dictator minion. Our proof of this fact works by showing that a homomorphism $\mathscr{Q}_{\mathbb{H}} \to \mathscr{D}$ would induce a Boolean measure on the subspaces of \mathbb{H} . By virtue of Gleason's Theorem [39], any measure μ on the subspaces of \mathbb{H} must be of the form $\mu(v) = \operatorname{tr}(m \cdot \operatorname{pr}_v)$ for some positive semidefinite linear operator m (where pr_v denotes the orthogonal projector onto the space v). Thus, in particular, μ must have a continuous dependence on its inputs, which contradicts the fact that μ is Boolean.

(*iii*) For any \mathbb{H} , $\mathcal{Q}_{\mathbb{H}}$ admits a homomorphism to a linear minion \mathcal{S} introduced in [28] and, independently, in [15], which captures the power of the *standard semidefinite programming* relaxation for CSPs. The result is obtained by projecting a fixed vector of \mathbb{H} onto the spaces constituting the elements of $\mathcal{Q}_{\mathbb{H}}$. This yields a minion homomorphism from $\mathcal{Q}_{\mathbb{H}}$ to a complex version $\mathcal{S}_{\mathbb{C}}$ of \mathcal{S} ; then, we show that $\mathcal{S}_{\mathbb{C}}$ and \mathcal{S} are in fact equivalent.

The minion-theoretic results on $\mathcal{Q}_{\mathbb{H}}$ listed above directly translate into results on the occurrence of quantum advantage. In particular, (i) implies that quantum advantage can only happen over Hilbert spaces of dimension 3 or higher — an instance of the known phenomenon that strong non-locality cannot be realised by two-qubit systems, see [1,18]. For such spaces, (ii) and (iii) yield new sufficient and necessary conditions for the occurrence of quantum advantage in terms of the complexity of the corresponding CSP: (ii) implies that, if CSP(**Y**) is not decided by a polynomial-time algorithm, **Y** has quantum advantage; (iii) implies that, if **Y** has quantum advantage, CSP(**Y**) has unbounded width² — where the width of a CSP is a central notion in the theory of CSP complexity, expressing the power of local-consistency techniques for its solution [9, 34].

In the case of graphs, the combination of these results completely characterises the occurrence of quantum advantage: Using Hell–Nešetřil's dichotomy for the complexity of graph CSPs (in Bulatov's reformulation [20] in terms of polymorphisms), we show that, if $\dim(\mathbb{H}) \geq 3$, a graph has quantum advantage if and only if it is non-bipartite (while, if

²We also describe an alternative derivation of this result, by showing that $\mathcal{Q}_{\mathbb{H}}$ admits a homomorphism to the *skeletal minion* introduced in [27], see Remark 26. This minion corresponds to a stronger version of the *singleton arc consistency* algorithm for CSPs [12, 24, 33] and thus, similarly to the minion \mathscr{S} of semidefinite programming, it captures bounded-width CSPs [50].

 $\dim(\mathbb{H}) \leq 2$, no graph has quantum advantage as stated above). This provides a complete answer to Question 1.

Finally, we introduce a natural generalisation of quantum advantage. A structure **Y** having quantum advantage means that the presence of a quantum homomorphism from **X** to **Y** is not sufficient to guarantee the complete classical information $\mathbf{X} \to \mathbf{Y}$. Can it still provide some *partial* classical information, in the form of a classical homomorphism $\mathbf{X} \to \mathbf{Y}'$ to a different structure \mathbf{Y}' ? If the answer is negative, we say that the pair $(\mathbf{Y}, \mathbf{Y}')$ has quantum advantage. We show that our minion framework is also able to capture this more general notion of quantum advantage. As a consequence, we connect quantum advantage for pairs of structures to the complexity of a *promise* variant of CSPs formalised in [6,8,14]. In this way, we show that the two notions of quantum advantage do not collapse, in the sense that there exist two structures **Y** and **Y**' such that both **Y** and **Y**' have quantum advantage, but the pair $(\mathbf{Y}, \mathbf{Y}')$ does not.

Related work Understanding the power of entanglement as a resource in non-local games (like the homomorphism game described above) is a well-established research programme in quantum information and theoretical computer science. A prominent example of this line of work is the "MIP^{*} = RE" Theorem, recently proved in [47], that characterises the class MIP^{*} of languages decided by a classical polynomial-time verifier interacting with multiple non-communicating computationally unbounded players sharing a finite-dimensional entangled state, as the class RE of recursively enumerable languages. In the same line of works, we mention [43, 44, 49, 55].

Quantum perfect strategies in the special case of the graph *n*-colouring problem (i.e., when \mathbf{Y} is the *n*-clique) and the corresponding notion of *quantum chromatic number* were considered in [7,23,29,35,36,52,54,59]. Boolean CSPs (i.e., CSP(\mathbf{Y}) with \mathbf{Y} being a Boolean relational structure) that are unsatisfiable but admit compatible assignments of self-adjoint unitary operators were characterised in [4]. The relation between this notion of CSP satisfiability and quantum perfect strategies for non-local games was described in [30], in the Boolean setting.

Unlike their non-promise variant, the complexity of promise CSPs remains largely undiscovered: Even the complexity of the promise version of graph *n*-colouring — known as *approximate graph colouring* [37] — is unknown in general. Since the introduction of a minion-based algebraic approach to promise CSPs in [8], the theory of minions has proved important in the investigation of this type of problems. For example, minions were used in [15, 16, 27, 28, 32] to describe the power of relaxation techniques for promise CSPs, and in [25, 26, 51] to investigate the complexity and tractability under certain relaxation models of approximate graph colouring. Transferring the geometric techniques we develop in this work for describing the structure of the quantum minion to other types of minions used in the theory of promise CSP complexity appears to be an interesting direction for future investigation.

Finally, we note that the dual version of Question 1 — characterise the structures **X** for which there exists some **Y** such that **X** and **Y** are quantum but not classically homomorphic — corresponds to the so-called *left-hand side* CSP (LCSP), where **X** is a fixed parameter and **Y** is the input. For a single **X**, LCSP is always trivially decidable in polynomial time. Hence, LCSP complexity questions are only studied in the setting where the left-hand side is an infinite class of structures, and their analysis is not captured by the algebraic approach [31, 40]. Therefore, our methods are inapplicable for characterising this dual type of quantum advantage.

2 Preliminaries

By \mathbb{N} , \mathbb{R} , and \mathbb{C} we denote the sets of positive integer numbers, real numbers, and complex numbers, respectively. For $\ell \in \mathbb{N}$, $[\ell]$ is the set $\{1, 2, \ldots, \ell\}$. For $i \leq \ell \in \mathbb{N}$, $\mathbf{e}_{i;\ell}$ is the *i*-th standard unit vector of length ℓ (i.e., the *i*-th entry of $\mathbf{e}_{i;\ell}$ is 1 and all other entries are 0). Given $\ell \in \mathbb{N}$ and a set S, S^{ℓ} denotes the set of tuples of elements of S of length ℓ . For a complex matrix M, we denote the transpose and the conjugate transpose of M by the symbols M^{\top} and M^* , respectively.

2.1 Finite-dimensional Hilbert spaces

Throughout this work, we let \mathbb{H} denote a finite-dimensional Hilbert space (i.e., a finitedimensional complex or real vector space equipped with an inner product $\langle \cdot, \cdot \rangle$). By dim (\mathbb{H}) we denote the dimension of \mathbb{H} , while $\mathbf{0}_{\mathbb{H}}$ indicates the zero vector in \mathbb{H} . To avoid trivial cases, we require $\mathbb{H} \neq \{\mathbf{0}_{\mathbb{H}}\}$. Two subsets v, w of \mathbb{H} are orthogonal (in symbols, $v \perp w$) if $\langle \mathbf{v}, \mathbf{w} \rangle = 0$ for each $\mathbf{v} \in v$, $\mathbf{w} \in w$. By v^{\perp} we denote the *orthogonal complement* of v; i.e., the set of all vectors $\mathbf{w} \in \mathbb{H}$ such that $\langle \mathbf{v}, \mathbf{w} \rangle = 0$ for each $\mathbf{v} \in v$. The set of all linear maps from \mathbb{H} to itself shall be denoted by $\operatorname{End}_{\mathbb{H}}$. Given $p, p' \in \operatorname{End}_{\mathbb{H}}, p+p'$ denotes their sum, pp' (or sometimes, for typographical convenience, $p \cdot p'$ denotes their product or composition, while [p,p'] = pp' - p'p is their commutator. The adjoint of $p \in End_{\mathbb{H}}$ is the unique map $p^* \in End_{\mathbb{H}}$ satisfying $\langle \mathbf{h}, p^*(\mathbf{h}') \rangle = \langle p(\mathbf{h}), \mathbf{h}' \rangle$ for each $\mathbf{h}, \mathbf{h}' \in \mathbb{H}$. p is self-adjoint if $p^* = p$, and it is an orthogonal projector if it is self-adjoint and idempotent (i.e., if $p^2 = p = p^*$). Proj_H denotes the set of all orthogonal projectors in $\operatorname{End}_{\mathbb{H}}$, while $L_{\mathbb{H}}$ denotes the set of all linear subspaces of \mathbb{H} . The function rg_{\bullet} sending a linear map $p \in End_{\mathbb{H}}$ to its range $rg_p \subseteq \mathbb{H}$ establishes a bijection between $\operatorname{Proj}_{\mathbb{H}}$ and $L_{\mathbb{H}}$. We denote the inverse of this bijection by $\operatorname{pr}_{\bullet}$ (i.e., for any $v \in L_{\mathbb{H}}$, $\mathrm{pr}_v \in \mathrm{Proj}_{\mathbb{H}}$ is the orthogonal projector onto v). By $\mathrm{id}_{\mathbb{H}}$ and $\mathbb{O}_{\mathbb{H}}$ we denote the identity and zero maps on \mathbb{H} , respectively.

2.2 Relational structures

A signature σ is a finite set of relation symbols R, each with its arity $\operatorname{ar}(R) \in \mathbb{N}$. A relational structure \mathbf{X} with signature σ (in short, a σ -structure) consists of a set X (the domain³ of \mathbf{X}) and a relation $R^{\mathbf{X}} \subseteq X^{\operatorname{ar}(R)}$ for each symbol $R \in \sigma$. If σ contains a unique symbol R of arity 2, and the relation $R^{\mathbf{X}}$ is symmetric and irreflexive (i.e., $(x, x') \in R^{\mathbf{X}}$ implies that $(x', x) \in R^{\mathbf{X}}$ and $x \neq x'$), \mathbf{X} is called a (simple) graph. Two relational structures \mathbf{X}, \mathbf{Y} are similar if they have the same signature. In this case, a homomorphism from \mathbf{X} to \mathbf{Y} is a map $f: X \to Y$ that preserves all relations; i.e., $f(\mathbf{x}) \in R^{\mathbf{Y}}$ for each $\mathbf{x} \in R^{\mathbf{X}}$, where f is applied entrywise to the entries of \mathbf{x} . We denote the existence of a homomorphism from \mathbf{X} to \mathbf{Y} by the expression $\mathbf{X} \to \mathbf{Y}$.

Given a σ -structure \mathbf{Y} , the constraint satisfaction problem parameterised by \mathbf{Y} (in short, $\mathrm{CSP}(\mathbf{Y})$) is the following computational problem: Given as input a σ -structure \mathbf{X} , output YES if $\mathbf{X} \to \mathbf{Y}$, and No if $\mathbf{X} \not\to \mathbf{Y}$.⁴ For $\ell \in \mathbb{N}$, we denote by \mathbf{Y}^{ℓ} the ℓ -th direct power of \mathbf{Y} ; i.e., \mathbf{Y}^{ℓ} is the σ -structure whose domain is Y^{ℓ} and whose relations are defined as follows:

 $^{^{3}}$ Unless otherwise stated, the domains of all relational structures appearing in this work shall be assumed to be finite.

⁴This is the *decision* version of CSP(**Y**). In the *search* version, the problem is to find an explicit homomorphism $f : \mathbf{X} \to \mathbf{Y}$ for any input **X** such that $\mathbf{X} \to \mathbf{Y}$. The two versions of CSP are equivalent up to polynomial-time reductions [22].

Given $R \in \sigma$ and any $\ell \times \operatorname{ar}(R)$ matrix M whose rows are tuples in $R^{\mathbf{Y}}$, the columns of M form a tuple in $R^{\mathbf{Y}^{\ell}}$. A *polymorphism* of \mathbf{Y} (of arity ℓ) is a homomorphism from \mathbf{Y}^{ℓ} to \mathbf{Y} . By $\operatorname{Pol}(\mathbf{Y})$ (resp. $\operatorname{Pol}^{(\ell)}(\mathbf{Y})$) we denote the set of all polymorphisms (resp. all polymorphisms of arity ℓ) of \mathbf{Y} .

2.3 Quantum advantage

Given two σ -structures **X** and **Y**, consider a family $\{p_{x,y}\}_{x \in X, y \in Y}$ of orthogonal projectors in $\operatorname{Proj}_{\mathbb{H}}$. For each $R \in \sigma$, $\mathbf{x} \in X^{\operatorname{ar}(R)}$, and $\mathbf{y} \in Y^{\operatorname{ar}(R)}$, define $p_{\mathbf{x},\mathbf{y}} = p_{x_1,y_1} \cdot p_{x_2,y_2} \cdot \ldots \cdot p_{x_{\operatorname{ar}(R)},y_{\operatorname{ar}(R)}}$. Consider the following three conditions:

- (Q1) $\sum_{y \in Y} p_{x,y} = \operatorname{id}_{\mathbb{H}}$ for each $x \in X$;
- (Q2) $[p_{x_i,y}, p_{x_i,y'}] = \mathbb{O}_{\mathbb{H}}$ for each $R \in \sigma$, $\mathbf{x} \in R^{\mathbf{X}}$, $i, j \in [\operatorname{ar}(R)]$, and $y, y' \in Y$;
- (Q3) $p_{\mathbf{x},\mathbf{y}} = \mathbf{0}_{\mathbb{H}}$ for each $R \in \sigma$, $\mathbf{x} \in R^{\mathbf{X}}$, and $\mathbf{y} \in Y^{\operatorname{ar}(R)} \setminus R^{\mathbf{Y}}$.

Following [2], we say that \mathbf{X} and \mathbf{Y} are quantum homomorphic over \mathbb{H} (and write $\mathbf{X} \to_{\mathbb{H}} \mathbf{Y}$) if there exists a family $\{p_{x,y}\}_{x \in X, y \in Y}$ having the three properties listed above. It was shown in [2] that $\mathbf{X} \to_{\mathbb{H}} \mathbf{Y}$ if and only if there exists a quantum perfect strategy for the $\mathbf{X}\mathbf{Y}$ -homomorphism game mentioned in the Introduction. Observe that $\mathbf{X} \to \mathbf{Y}$ always implies $\mathbf{X} \to_{\mathbb{H}} \mathbf{Y}$. Indeed, if $f: \mathbf{X} \to \mathbf{Y}$ is a homomorphism, letting $p_{x,y} = \mathrm{id}_{\mathbb{H}}$ if y = f(x) and $p_{x,y} = \mathbb{O}_{\mathbb{H}}$ otherwise yields a proper quantum homomorphism. We say that \mathbf{Y} has quantum advantage over \mathbb{H} if the converse is not true; i.e., if there exists a σ -structure \mathbf{X} such that $\mathbf{X} \to_{\mathbb{H}} \mathbf{Y}$ but $\mathbf{X} \not\to \mathbf{Y}$. In terms of the $\mathbf{X}\mathbf{Y}$ -homomorphism game, this means that Alice and Bob can win the game using a quantum-assisted strategy, but they cannot using a classical strategy.⁵

2.4 Minions

A minion \mathscr{M} is the disjoint union of non-empty sets $\mathscr{M}^{(\ell)}$ for $\ell \in \mathbb{N}$, equipped with operations $(\cdot)_{/\pi} : \mathscr{M}^{(\ell)} \to \mathscr{M}^{(\ell')}$ (known as *minor operations*) for each pair of integers $\ell, \ell' \in \mathbb{N}$ and each map $\pi : [\ell] \to [\ell']$, which must satisfy the following two conditions:

- $M_{/id} = M$ for each $\ell \in \mathbb{N}$ and $M \in \mathscr{M}^{(\ell)}$, where id is the identity map on the set $[\ell]$;
- $(M_{/\pi})_{/\tilde{\pi}} = M_{/\tilde{\pi}\circ\pi}$ for each $\ell, \ell', \ell'' \in \mathbb{N}, M \in \mathscr{M}^{(\ell)}, \pi : [\ell] \to [\ell'], \text{ and } \tilde{\pi} : [\ell'] \to [\ell''].^6$

Given $M \in \mathcal{M}$, the index ℓ for which $M \in \mathcal{M}^{(\ell)}$ is referred to as the *arity* of M. A subminion of a minion \mathcal{M} is a non-empty subset of \mathcal{M} that is closed under the minor operations of \mathcal{M} . In this work, we shall consider two concrete types of minions, described in the two examples below.

⁵We point out that, if **X** and **Y** are graphs, the definition of quantum homomorphism from [2] is more restrictive than the one given in [53], in that the latter does not enforce the commutativity condition (Q2) among the projectors. Nevertheless, in the case of graphs, our results hold for both notions of quantum homomorphisms, see Remark 30.

⁶Equivalently, we may define a minion as a functor from the skeleton category of non-empty finite sets to the category of non-empty sets. The definition of minion as we present it here was introduced in [16], and it is an abstraction of the *minor-closed classes of functions* considered in [56] to study a Galois connection between functions and relations. In fact, the latter objects are precisely the function minions of Example 2.

Example 2 (Function minions). Take two sets S and T. For any integer $\ell \in \mathbb{N}$, consider the set $\mathscr{F}_{S,T}^{(\ell)}$ of all functions from S^{ℓ} to T. Given $f \in \mathscr{F}_{S,T}^{(\ell)}$ and $\pi : [\ell] \to [\ell']$, we let $f_{/\pi} \in \mathscr{F}_{S,T}^{(\ell')}$ be the function defined as follows:

$$f_{/\pi}: S^{\ell'} \to T$$
$$(s_1, \dots, s_{\ell'}) \mapsto f(s_{\pi(1)}, \dots, s_{\pi(\ell)}).$$

The set $\mathscr{F}_{S,T}$ consisting in the disjoint union of all sets $\mathscr{F}_{S,T}^{(\ell)}$ for $\ell \in \mathbb{N}$ is easily seen to be a minion. A *function minion* (over the sets S and T) is any subminion of $\mathscr{F}_{S,T}$.

Example 3 (Linear minions). Let \mathcal{A} be an abelian monoid — i.e., a set equipped with a binary operation " $+_{\mathcal{A}}$ " that is associative and commutative, and admits an identity element " $0_{\mathcal{A}}$ ". Following [28], we let a *linear minion* over \mathcal{A} be a non-empty set \mathcal{M} of matrices having entries in \mathcal{A} that is closed under the following elementary row operations:

- swapping two rows;
- replacing two rows with their sum;
- inserting an extra zero row.

(Here, the sum is the entrywise application of the " $+_{\mathcal{A}}$ " operation, while the zero row is a tuple of copies of " $0_{\mathcal{A}}$ ".) We let $\mathscr{M}^{(\ell)}$ be the set of all matrices in \mathscr{M} having exactly ℓ -many rows. Moreover, given a matrix $M \in \mathscr{M}^{(\ell)}$ and a map $\pi : [\ell] \to [\ell']$, we define $M_{/\pi}$ as the matrix whose *i*-th row, for $i \in [\ell']$, is the sum of all rows of M having index in $\pi^{-1}(i)$; i.e.,

$$\mathbf{e}_{i;\ell'}^\top M_{/\pi} = \sum_{j \in \pi^{-1}(i)} \mathbf{e}_{j;\ell}^\top M.$$

Starting from M, one can realise $M_{/\pi}$ through a sequence of the elementary row operations listed above. Therefore, $M_{/\pi} \in \mathcal{M}^{(\ell')}$. Furthermore, this choice for the minor operations guarantees that \mathcal{M} is a minion.

Let \mathcal{M} and \mathcal{N} be two minions. A minion homomorphism $\xi : \mathcal{M} \to \mathcal{N}$ is a map from the underlying set of \mathcal{M} to the underlying set of \mathcal{N} that

- preserves the arity i.e., $\xi(M) \in \mathcal{N}^{(\ell)}$ if $M \in \mathcal{M}^{(\ell)}$;
- preserves the minors i.e., $\xi(M_{/\pi}) = \xi(M)_{/\pi}$ if $M \in \mathscr{M}^{(\ell)}$ and $\pi : [\ell] \to [\ell']^{.7}$

Given a relational structure \mathbf{Y} , the set $\text{Pol}(\mathbf{Y})$ of all polymorphisms of \mathbf{Y} is a subset of $\mathscr{F}_{Y,Y}$ (as defined in Example 2), it is non-empty (since it contains the identity homomorphism), and it is closed under the minor operations of $\mathscr{F}_{Y,Y}$; hence, it is a (function) minion. Polymorphism minions capture the complexity of CSPs, in the sense of the following result.

Theorem 4 ([11]). Let **Y** and **Y'** be two relational structures. If there exists a minion homomorphism $\operatorname{Pol}(\mathbf{Y}) \to \operatorname{Pol}(\mathbf{Y}')$, then $\operatorname{CSP}(\mathbf{Y}')$ reduces in polynomial time to $\operatorname{CSP}(\mathbf{Y})$.⁸

⁷Equivalently, a minion homomorphism $\xi: \mathcal{M} \to \mathcal{N}$ is a natural transformation from \mathcal{M} to \mathcal{N} .

⁸In fact, a log-space reduction from $CSP(\mathbf{Y}')$ to $CSP(\mathbf{Y})$ exists.

3 The quantum minion

In this section, we describe a new minion that captures quantum advantage. The geometric properties of this minion shall be used in the later sections to obtain information on the occurrence of quantum advantage.

Recall that $L_{\mathbb{H}}$ denotes the set of all linear subspaces of \mathbb{H} . Given $v, w \in L_{\mathbb{H}}$, we let $v \boxplus w$ denote their sum as linear spaces; i.e., $v \boxplus w = \{\mathbf{v} + \mathbf{w} \mid \mathbf{v} \in v, \mathbf{w} \in w\}$. Let $\mathcal{L}_{\mathbb{H}}$ denote the set $L_{\mathbb{H}}$ equipped with the operation " \boxplus ", and observe that $\mathcal{L}_{\mathbb{H}}$ is an abelian monoid, with the subspace $\{\mathbf{0}_{\mathbb{H}}\}$ as the identity element. For $\ell \in \mathbb{N}$, consider the set

$$\mathcal{Q}_{\mathbb{H}}^{(\ell)} = \left\{ \begin{array}{cc} q_1 \\ \vdots \\ q_\ell \end{array} \right| \quad \text{such that} \quad \begin{array}{c} q_i \in L_{\mathbb{H}} \text{ for each } i, \\ q_i \perp q_j \text{ for each } i \neq j, \\ & & & \\ \vdots \\ i \in [\ell] \end{array} \right\}.$$

We define the quantum minion $\mathcal{Q}_{\mathbb{H}}$ as the disjoint union of the sets $\mathcal{Q}_{\mathbb{H}}^{(\ell)}$ for $\ell \in \mathbb{N}$.

Lemma 5. $Q_{\mathbb{H}}$ is a linear minion over the abelian monoid $\mathcal{L}_{\mathbb{H}}$.⁹

Proof. All elements of $\mathcal{Q}_{\mathbb{H}}$ are vectors (i.e., matrices having a single column) of spaces in $L_{\mathbb{H}}$. The statement immediately follows by observing that $\mathcal{Q}_{\mathbb{H}}$ is closed under the three elementary row operations of Example 3.

Using the quantum minion, we will show that quantum advantage is a property of the polymorphism minion of relational structures. Formally, we will prove the next result.

Theorem 6. A relational structure **Y** has quantum advantage over \mathbb{H} if and only if $\mathbb{Q}_{\mathbb{H}} \not\rightarrow \operatorname{Pol}(\mathbf{Y})$.

The proof of Theorem 6 makes use of two ingredients: the *free structures* from [8] and the *minion tests* from [28].

Let \mathscr{M} be a minion and let \mathbf{Y} be a σ -structure. On a high level, the free structure $\mathbb{F}_{\mathscr{M}}(\mathbf{Y})$ of \mathscr{M} generated by \mathbf{Y} is obtained by simulating the relations of \mathbf{Y} on a domain consisting of elements of \mathscr{M} . Formally, following [8], we let $\mathbb{F}_{\mathscr{M}}(\mathbf{Y})$ be the (potentially infinite) σ -structure whose domain is the set $\mathscr{M}^{(n)}$ (where n is the domain size of \mathbf{Y}) and whose relations are described as follows. For each symbol $R \in \sigma$, let m be the cardinality of $R^{\mathbf{Y}}$. A tuple $(M_1, \ldots, M_{\operatorname{ar}(R)})$ of elements of $\mathscr{M}^{(n)}$ belongs to $R^{\mathbb{F}_{\mathscr{M}}(\mathbf{Y})}$ if and only if there exists some $\widetilde{M} \in \mathscr{M}^{(m)}$ such that $M_i = \widetilde{M}_{/\pi_i}$ for each $i \in [\operatorname{ar}(R)]$, where $\pi_i : [m] \to [n]$ is the function mapping each tuple $\mathbf{y} \in R^{\mathbf{Y}}$ to its *i*-th entry y_i .¹⁰ Notice that the domain of $\mathbb{F}_{\mathscr{M}}(\mathbf{Y})$ can be infinite even when the domain of \mathbf{Y} is finite. In particular, this is the case for $\mathscr{M} = \mathscr{Q}_{\mathbb{H}}$ (except in the cases where the dimension of \mathbb{H} or the domain size of \mathbf{Y} equals 1).

⁹In fact, $\mathcal{Q}_{\mathbb{H}}$ belongs to a particular class of linear minions known as *conic minions*, introduced in [28] (see also [32]).

¹⁰Here we are implicitly identifying Y with [n] and $R^{\mathbf{Y}}$ with [m], in order to avoid introducing extra notation for the bijections between those sets.

The next result, whose proof is deferred to Appendix A, shows that quantum homomorphisms correspond to classical homomorphisms to the free structure of the quantum minion.¹¹

Proposition 7. Let \mathbf{X}, \mathbf{Y} be similar relational structures. Then $\mathbf{X} \to_{\mathbb{H}} \mathbf{Y}$ if and only if $\mathbf{X} \to \mathbb{F}_{\mathcal{Q}_{\mathbb{H}}}(\mathbf{Y})$.

Next, we describe the notion of minion tests from [28]. Let \mathcal{M} be a minion, let σ be a signature, and let Σ be the class of all (finite) σ -structures. Consider the function

$$\begin{aligned} \text{Test}_{\mathscr{M}} &: \Sigma^2 \to \{ \text{Yes}, \text{No} \} \\ & (\mathbf{X}, \mathbf{Y}) \mapsto \begin{cases} \text{Yes} & \text{if } \mathbf{X} \to \mathbb{F}_{\mathscr{M}}(\mathbf{Y}) \\ \text{No} & \text{otherwise.} \end{cases} \end{aligned}$$

Fix a σ -structure \mathbf{Y} . It is not hard to verify that $\mathbf{Y} \to \mathbb{F}_{\mathscr{M}}(\mathbf{Y})$ for every \mathscr{M} . Hence, $\mathbf{X} \to \mathbf{Y}$ implies $\operatorname{Test}_{\mathscr{M}}(\mathbf{X}, \mathbf{Y}) = \operatorname{YES}$. We say that $\operatorname{Test}_{\mathscr{M}}$ solves $\operatorname{CSP}(\mathbf{Y})$ if the converse implication also holds; i.e., if, for every σ -structure \mathbf{X} , $\operatorname{Test}_{\mathscr{M}}(\mathbf{X}, \mathbf{Y}) = \operatorname{YES}$ implies $\mathbf{X} \to \mathbf{Y}$. As stated next, this can be reformulated in terms of minion homomorphisms.

Theorem 8 ([28]). Let \mathcal{M} be a minion and let \mathbf{Y} be a relational structure. Then $\text{Test}_{\mathcal{M}}$ solves $\text{CSP}(\mathbf{Y})$ if and only if $\mathcal{M} \to \text{Pol}(\mathbf{Y})$.

We can now prove Theorem 6.

Proof of Theorem 6. By virtue of Proposition 7, the existence of a quantum homomorphism over \mathbb{H} can be viewed as a minion test for the quantum minion $\mathcal{Q}_{\mathbb{H}}$. More precisely, given two σ -structures \mathbf{X} and $\mathbf{Y}, \mathbf{X} \to_{\mathbb{H}} \mathbf{Y}$ if and only if $\text{Test}_{\mathcal{Q}_{\mathbb{H}}}(\mathbf{X}, \mathbf{Y}) = \text{YES}$. Therefore, \mathbf{Y} has quantum advantage over \mathbb{H} precisely when $\text{Test}_{\mathcal{Q}_{\mathbb{H}}}$ does not solve $\text{CSP}(\mathbf{Y})$. The result then follows from Theorem 8.

4 $Q_{\mathbb{H}}$ and the dictator minion

By virtue of Theorem 6, in order to investigate the occurrence of quantum advantage, one needs to understand the structure of the quantum minion $\mathcal{Q}_{\mathbb{H}}$. "Minion homomorphism" is a binary, reflexive, and transitive relation; thus, it induces a preorder on the class of all minions. It shall be crucial to obtain information on where $\mathcal{Q}_{\mathbb{H}}$ is located within this preorder. We shall see that, when dim $(\mathbb{H}) \leq 2$, $\mathcal{Q}_{\mathbb{H}}$ is a *least element* of the preorder, as it is equivalent to a specific minion — known as the dictator minion — that admits a homomorphism to all minions. This implies that no structure has quantum advantage over such Hilbert spaces. Things get more interesting when dim $(\mathbb{H}) \geq 3$. In this case, we shall see that $\mathcal{Q}_{\mathbb{H}}$ is essentially different from the dictator minion. Through known results on the complexity of CSPs and the theory of their polymorphism minions, this separation shall lead to a sufficient condition for the occurrence of quantum advantage.

¹¹The idea of expressing quantum homomorphisms as classical homomorphisms to a modified structure already appeared in [2, 53]. More precisely, in the case of graphs, the structure $\mathbb{F}_{\mathcal{Q}_{\mathbb{H}}}(\mathbf{Y})$ of Proposition 7 corresponds to the graph $M(\mathbf{Y}, d)$ from [53, Theorem 2.9], where $d = \dim(\mathbb{H})$. Moreover, for arbitrary relational structures, it corresponds to the structure $\mathcal{Q}_d \mathbf{Y}$, where \mathcal{Q}_d is the *quantum monad* introduced in [2]. In our setting, however, it is crucial to formulate Proposition 7 in terms of the free structure of a minion. This allows us to use the notion of minion tests and, through Theorem 8, to express quantum advantage as a property of the polymorphism minion (Theorem 6), thus linking it to the complexity of the corresponding CSP.

We start by defining the *dictator minion* as the linear minion \mathscr{D} over the abelian monoid $(\mathbb{R}, +)$ consisting of all standard unit vectors $\mathbf{e}_{i;\ell}$ for $1 \leq i \leq \ell \in \mathbb{N}$.¹² The next simple result shows that \mathscr{D} is a least element of the minion homomorphism preorder.¹³

Lemma 9. $\mathcal{D} \to \mathcal{M}$ for every minion \mathcal{M} .

Proof. Fix an element $M \in \mathcal{M}^{(1)}$ (which must exist as minions are non-empty by definition). For $i \leq \ell \in \mathbb{N}$, consider the map $\pi_{i,\ell} : [1] \to [\ell]$ given by $1 \mapsto i$. It is straightforward to check that the function $\mathbf{e}_{i;\ell} \mapsto M_{/\pi_{i,\ell}}$ yields a minion homomorphism from \mathcal{D} to \mathcal{M} .

We shall establish the following characterisation of when the quantum minion admits a homomorphism (and, thus, is homomorphically equivalent) to the dictator minion.

Theorem 10. $\mathcal{Q}_{\mathbb{H}} \to \mathcal{D}$ if and only if dim $(\mathbb{H}) \leq 2$.

Our proofs of the two directions of Theorem 10 are both geometric. This is not surprising, as the quantum minion $\mathcal{Q}_{\mathbb{H}}$ has a natural geometric interpretation.

The "if" part (Proposition 12) is established as follows. First, we show that every 2dimensional Hilbert space can be partitioned into two parts in a way that any pair of nonzero orthogonal vectors in the space is separated by the partition. In order to prove this fact, we consider the Hopf fibration of the complex projective line \mathbb{CP}^1 into the 2-sphere¹⁴ — which has the property of mapping orthogonal vectors to antipodal points of the sphere. We can then partition the Hilbert space by taking the preimage under the Hopf fibration of an antipodal partition of the 2-sphere.

Lemma 11. If dim(\mathbb{H}) = 2, there exists $C \subset \mathbb{H}$ such that, given any two nonzero orthogonal vectors $\mathbf{v}, \mathbf{w} \in \mathbb{H}$, exactly one of \mathbf{v} and \mathbf{w} lies in C.

A self-contained proof of Lemma 11 is given in Appendix A. Once such a partition $\{C, \mathbb{H} \setminus C\}$ is fixed, we can select in each pair of 1-dimensional orthogonal subspaces of \mathbb{H} a distinguished element — namely, the subspace generated by a vector in C. As we shall see next, the property of the partition guarantees that the distinguished element is preserved under taking minors, thus yielding a proper homomorphism from $\mathcal{Q}_{\mathbb{H}}$ to \mathcal{D} .

Proposition 12. If $\dim(\mathbb{H}) \leq 2$, $\mathcal{Q}_{\mathbb{H}} \to \mathcal{D}$.

Proof. Given $\mathbf{q} \in \mathcal{Q}_{\mathbb{H}}$ of some arity ℓ , an *essential coordinate* of \mathbf{q} is an index $i \in [\ell]$ for which $q_i \neq \{\mathbf{0}_{\mathbb{H}}\}$. We denote the set of essential coordinates of \mathbf{q} by $\operatorname{sec}(\mathbf{q})$. Note that the cardinality of $\operatorname{sec}(\mathbf{q})$ is at most $\dim(\mathbb{H}) \leq 2$. For any 1-dimensional subspace $v \in L_{\mathbb{H}}$, fix a nonzero vector $g(v) \in v$. Let also $C \subset \mathbb{H}$ be the set constructed in Lemma 11.

We define $\xi : \mathcal{Q}_{\mathbb{H}} \to \mathcal{D}$ as follows. Given $\mathbf{q} \in \mathcal{Q}_{\mathbb{H}}$ of arity ℓ , if \mathbf{q} has a unique essential coordinate i, we let $\xi(\mathbf{q}) = \mathbf{e}_{i;\ell}$. Otherwise, \mathbf{q} has two essential coordinates $i \neq j$ — in which case, it must be that $\dim(\mathbb{H}) = 2$ and $\dim(q_i) = \dim(q_j) = 1$. If $g(q_i) \in C$, we let $\xi(\mathbf{q}) = \mathbf{e}_{i;\ell}$,

¹²Equivalently, we may define \mathscr{D} as a function minion over some finite set with at least two elements, consisting of functions that project any input tuple onto one fixed coordinate — whence the alternative name *projection minion*, which is common in the CSP literature (see [8]). Yet another description of \mathscr{D} is the following: For $\ell \in \mathbb{N}, \ \mathscr{D}^{(\ell)} = [\ell]$; for $\pi : [\ell] \to [\ell']$ and $i \in \mathscr{D}^{(\ell)}, i_{/\pi} = \pi(i)$.

¹³The minion homomorphism preorder also has a greatest element: The linear minion (over any abelian monoid) whose only ℓ -ary element, for each $\ell \in \mathbb{N}$, is the all-zero vector of length ℓ .

¹⁴This construction results in the object that is sometimes called the *Riemann sphere* or *Bloch sphere*.

otherwise, we let $\xi(\mathbf{q}) = \mathbf{e}_{j;\ell}$. Since $\langle g(q_i), g(q_j) \rangle = 0$, ξ is well defined by virtue of the property of C. We claim that ξ yields a minion homomorphism from $\mathcal{Q}_{\mathbb{H}}$ to \mathcal{D} .

The facts that $\xi(\mathbf{q}) \in \mathcal{D}$ for each $\mathbf{q} \in \mathcal{Q}_{\mathbb{H}}$ and that ξ preserves the arity directly follow from the definition. To show that ξ preserves the minors, take $\pi : [\ell] \to [\ell']$ and $\mathbf{q} \in \mathcal{Q}_{\mathbb{H}}^{(\ell)}$. Observe that

$$\sec(\mathbf{q}_{/\pi}) = \{\pi(i) : i \in \sec(\mathbf{q})\}.$$

If $sec(\mathbf{q}) = \{i\}$, it follows that

$$\xi(\mathbf{q}_{/\pi}) = \mathbf{e}_{\pi(i);\ell'} = (\mathbf{e}_{i;\ell})_{/\pi} = \xi(\mathbf{q})_{/\pi}.$$

If $\operatorname{sec}(\mathbf{q}) = \{i, j\}$ with $i \neq j$, let $k \in \{i, j\}$ be such that $\xi(\mathbf{q}) = \mathbf{e}_{k;\ell}$. There are two cases. If $\pi(i) = \pi(j)$, we have that $\operatorname{sec}(\mathbf{q}_{/\pi}) = \{\pi(k)\}$, and it follows that $\xi(\mathbf{q}_{/\pi}) = \mathbf{e}_{\pi(k);\ell'}$. If $\pi(i) \neq \pi(j), \mathbf{q}_{/\pi}$ is the vector in $\mathcal{Q}_{\mathbb{H}}^{(\ell')}$ whose $\pi(i)$ -th entry is q_i , whose $\pi(j)$ -th entry is q_j , and all of whose other entries are $\{\mathbf{0}_{\mathbb{H}}\}$. Hence, also in this case it holds that $\xi(\mathbf{q}_{/\pi}) = \mathbf{e}_{\pi(k);\ell'}$. This means that, in either case, we have

$$\xi(\mathbf{q}_{/\pi}) = \mathbf{e}_{\pi(k);\ell'} = (\mathbf{e}_{k;\ell})_{/\pi} = \xi(\mathbf{q})_{/\pi}.$$

Therefore, the identity $\xi(\mathbf{q}_{/\pi}) = \xi(\mathbf{q})_{/\pi}$ is satisfied in all cases, and it follows that ξ is a minion homomorphism.

As a direct consequence, we obtain the following.

Corollary 13. No relational structure has quantum advantage over \mathbb{H} if dim $(\mathbb{H}) \leq 2$.

Proof. If dim(\mathbb{H}) ≤ 2 , Lemma 9 and Proposition 12 yield

$$\mathcal{Q}_{\mathbb{H}} \to \mathcal{D} \to \operatorname{Pol}(\mathbf{Y})$$

for any relational structure \mathbf{Y} , and the conclusion immediately follows from Theorem 6. \Box

Corollary 13 is an instance of a more general phenomenon in quantum information, observed in [18] in the context of pseudo-telepathy games (see also [1]): In order for an *n*-qubit system to exhibit strong non-locality with any finite number of local measurements, *n* must be at least 3. We now turn to the "only if" part of Theorem 10 (Proposition 15). Our proof builds on a deep result by Gleason on measures on the subspaces of Hilbert spaces — which we state next, after introducing the necessary terminology.

The trace $\operatorname{tr}(p)$ of a linear map $p \in \operatorname{End}_{\mathbb{H}}$ is the trace of a matrix representation of p in any basis of \mathbb{H} . We say that p is positive semidefinite (and write $p \succeq \mathbb{O}_{\mathbb{H}}$) if p is self-adjoint and $\langle p(\mathbf{h}), \mathbf{h} \rangle \geq 0$ for each $\mathbf{h} \in \mathbb{H}$. A measure on the subspaces of \mathbb{H} is a function $\mu : L_{\mathbb{H}} \to \mathbb{R}_{\geq 0}$ such that, for each collection $\{v_1, \ldots, v_\ell\}$ of mutually orthogonal subspaces of \mathbb{H} , it holds that

$$\mu\left(\bigoplus_{i\in[\ell]}v_i\right) = \sum_{i\in[\ell]}\mu(v_i).$$
(1)

Theorem 14 ([39]). Let μ be a measure on the subspaces of a finite-dimensional (real or complex) Hilbert space \mathbb{H} of dimension at least three. Then there exists a positive semidefinite linear map $m \in \operatorname{End}_{\mathbb{H}}$ such that

$$\mu(v) = \operatorname{tr}(m \cdot \operatorname{pr}_v) \tag{2}$$

for each $v \in L_{\mathbb{H}}$.¹⁵

Proposition 15 is established by showing that a homomorphism from the quantum minion to the dictator minion would yield a well-defined, Boolean measure μ on the subspaces of \mathbb{H} . However, if the dimension of \mathbb{H} is at least 3, Theorem 14 implies that μ must have a continuous dependence on its inputs and, in particular, it must assume intermediate values between 0 and 1, thus yielding a contradiction.

Proposition 15. If $\dim(\mathbb{H}) \geq 3$, $\mathcal{Q}_{\mathbb{H}} \not\to \mathcal{D}$.

Proof. Suppose, for the sake of contradiction, that there exists a homomorphism $\xi : \mathcal{Q}_{\mathbb{H}} \to \mathcal{D}$. We associate with any linear subspace $v \in L_{\mathbb{H}}$ the vector

$$\widehat{v} = \begin{bmatrix} v \\ v^{\perp} \end{bmatrix} \in \mathcal{Q}_{\mathbb{H}}^{(2)}.$$

Recall that $\mathscr{D}^{(2)} = \{\mathbf{e}_{1;2}, \mathbf{e}_{2;2}\}$, and consider the function $\mu : L_{\mathbb{H}} \to \mathbb{R}_{\geq 0}$ defined by

$$\mu(v) = \begin{cases} 1 & \text{if } \xi(\widehat{v}) = \mathbf{e}_{1,2} \\ 0 & \text{if } \xi(\widehat{v}) = \mathbf{e}_{2,2} \end{cases}$$

for each $v \in L_{\mathbb{H}}$. We claim that μ is a measure on the subspaces of \mathbb{H} . Take a collection $\{v_1, \ldots, v_\ell\}$ of mutually orthogonal subspaces of \mathbb{H} , and let $w = \bigoplus_{i \in [\ell]} v_i$.

Suppose first that $\mu(w) = 1$, which means that $\xi(\widehat{w}) = \mathbf{e}_{1,2}$. Take the vector

$$\mathbf{q} = \begin{bmatrix} v_1 & v_2 & \dots & v_\ell & w^\perp \end{bmatrix}^\top \in \mathcal{Q}_{\mathbb{H}}^{(\ell+1)},$$

and consider the map $\pi : [\ell + 1] \to [2]$ such that $\pi(i) = 1$ if $i \in [\ell]$, and $\pi(\ell + 1) = 2$. Note that $\mathbf{q}_{/\pi} = \widehat{w}$; using that ξ is a homomorphism, we get

$$\mathbf{e}_{1;2} = \xi(\widehat{w}) = \xi(\mathbf{q}_{/\pi}) = \xi(\mathbf{q})_{/\pi},\tag{3}$$

whence it follows that $\xi(\mathbf{q}) = \mathbf{e}_{j;\ell+1}$ for some $j \in [\ell]$. Consider now, for $k \in [\ell]$, the map $\pi_k : [\ell+1] \to [2]$ such that $\pi_k(k) = 1$ and $\pi_k(i) = 2$ for each $i \in [\ell+1] \setminus \{k\}$. Observe that $\mathbf{q}_{/\pi_k} = \hat{v_k}$. Therefore,

$$\xi(\hat{v}_k) = \xi(\mathbf{q}_{/\pi_k}) = \xi(\mathbf{q})_{/\pi_k} = (\mathbf{e}_{j;\ell+1})_{/\pi_k} = \mathbf{e}_{\pi_k(j);2}.$$
(4)

It follows from the above that $\mu(v_k) = 1$ precisely when $\pi_k(j) = 1$. In other words, $\mu(v_j) = 1$ and $\mu(v_k) = 0$ for each $k \neq j$. Therefore,

$$\sum_{i \in [\ell]} \mu(v_i) = 1 = \mu(w) = \mu\left(\bigoplus_{i \in [\ell]} v_i\right),$$

¹⁵We point out that Gleason's Theorem holds in the more general setting of separable, not necessarily finite-dimensional Hilbert spaces \mathbb{H} . In the general formulation, μ needs to be defined on all closed subspaces of \mathbb{H} , and it needs to satisfy the condition (1) for each countable collection of mutually orthogonal subspaces whose linear span is closed.

as needed.

Suppose now that $\mu(w) = 0$. In this case, the equation (3) turns into

$$\mathbf{e}_{2;2} = \xi(\widehat{w}) = \xi(\mathbf{q}_{/\pi}) = \xi(\mathbf{q})_{/\pi},$$

which means that $\xi(\mathbf{q}) = \mathbf{e}_{\ell+1;\ell+1}$. Then, for each $k \in [\ell]$, the equation (4) becomes

$$\xi(\hat{v}_k) = \xi(\mathbf{q}_{/\pi_k}) = \xi(\mathbf{q})_{/\pi_k} = (\mathbf{e}_{\ell+1;\ell+1})_{/\pi_k} = \mathbf{e}_{\pi_k(\ell+1);2} = \mathbf{e}_{2;2},$$

thus showing that $\mu(v_k) = 0$ for each $k \in [\ell]$. It follows that

$$\sum_{i \in [\ell]} \mu(v_i) = 0 = \mu(w) = \mu\left(\bigoplus_{i \in [\ell]} v_i\right).$$

Hence, μ is a measure on the subspaces of \mathbb{H} , as claimed. We can then invoke Theorem 14 and find a linear map $m \in \operatorname{End}_{\mathbb{H}}$ such that $m \succeq \mathbb{O}_{\mathbb{H}}$ and $\mu(v) = \operatorname{tr}(m \cdot \operatorname{pr}_{v})$ for each $v \in L_{\mathbb{H}}$. In particular,

$$\mu(\mathbb{H}) = \operatorname{tr}(m \cdot \operatorname{pr}_{\mathbb{H}}) = \operatorname{tr}(m \cdot \operatorname{id}_{\mathbb{H}}) = \operatorname{tr}(m).$$
(5)

We now claim that $\mu(\mathbb{H}) = 1$. Indeed, if $\mu(\mathbb{H}) = 0$, it would follow from (5) that $\operatorname{tr}(m) = 0$; since *m* is positive semidefinite, this would mean that $m = \mathbb{O}_{\mathbb{H}}$; then, (2) would force $\mu(v) = 0$ for each $v \in L_{\mathbb{H}}$. On the other hand, letting $\tau : [2] \to [2]$ be the map $1 \mapsto 2, 2 \mapsto 1$, we would find

$$\xi(\widehat{\{\mathbf{0}_{\mathbb{H}}\}}) = \xi(\widehat{\mathbb{H}}_{/\tau}) = \xi(\widehat{\mathbb{H}})_{/\tau} = (\mathbf{e}_{2;2})_{/\tau} = \mathbf{e}_{1;2}$$

and, thus, $\mu({\mathbf{0}_{\mathbb{H}}}) = 1$, a contradiction.

Next, let $d = \dim(\mathbb{H})$, let $\{\mathbf{v}_1, \ldots, \mathbf{v}_d\}$ be an orthonormal basis of eigenvectors for m, and let $\lambda_1, \ldots, \lambda_d$ be the corresponding (real, nonnegative) eigenvalues of m. It follows from (5) that

$$1 = \mu(\mathbb{H}) = \operatorname{tr}(m) = \sum_{i \in [d]} \lambda_i.$$

Therefore, letting v_i be the linear span of \mathbf{v}_i , we have

$$\mu(v_i) = \operatorname{tr}(m \cdot \operatorname{pr}_{v_i}) = \lambda_i$$

for each $i \in [d]$. On the other hand, the range of μ is included in $\{0, 1\}$ by construction. We deduce that the spectrum of m consists of the values 1 (with multiplicity 1) and 0 (with multiplicity d-1). In other words, $m = \operatorname{pr}_{v_j}$ for some $j \in [d]$. Choose an index $k \neq j$, and let w be the linear span of the vector $\mathbf{v}_j + \mathbf{v}_k$. Using that m is self-adjoint, we obtain

$$\mu(w) = \operatorname{tr}(m \cdot \operatorname{pr}_w) = \sum_{i \in [d]} \langle m \cdot \operatorname{pr}_w(\mathbf{v}_i), \mathbf{v}_i \rangle$$
$$= \sum_{i \in [d]} \langle \operatorname{pr}_w(\mathbf{v}_i), m(\mathbf{v}_i) \rangle = \sum_{i \in [d]} \langle \operatorname{pr}_w(\mathbf{v}_i), \operatorname{pr}_{v_j}(\mathbf{v}_i) \rangle$$
$$= \langle \operatorname{pr}_w(\mathbf{v}_j), \mathbf{v}_j \rangle.$$

Since an orthonormal basis for w is given by the single vector $\mathbf{w} = \frac{1}{\sqrt{2}}(\mathbf{v}_j + \mathbf{v}_k)$, we have that

$$\operatorname{pr}_w(\mathbf{v}) = \langle \mathbf{v}, \mathbf{w} \rangle \mathbf{w}$$

for each $\mathbf{v} \in \mathbb{H}$. Observe that $\langle \mathbf{w}, \mathbf{v}_j \rangle = \frac{1}{\sqrt{2}}$. As a consequence,

$$\mu(w) = \langle \langle \mathbf{v}_j, \mathbf{w} \rangle \mathbf{w}, \mathbf{v}_j \rangle = \frac{1}{2},$$

a contradiction.

Proposition 15 yields a separation of the quantum minion from the dictator minion. Together with Theorem 6, it implies that \mathbf{Y} has quantum advantage over any Hilbert space of dimension at least 3 whenever $\text{Pol}(\mathbf{Y}) \to \mathcal{D}$. In turn, the algebraic theory of CSPs allows formulating the latter condition in terms of the complexity of CSP(\mathbf{Y}).

Since homomorphisms between polymorphism minions provide polynomial-time reductions between CSPs (Theorem 4) and since the dictator minion is homomorphic to any minion (Lemma 9), if $Pol(\mathbf{Y}) \rightarrow \mathcal{D}$ it must hold that $CSP(\mathbf{Y})$ is NP-complete. It follows from the CSP Dichotomy Theorem proved in [21,60] (in its reformulation in terms of polymorphism minions from [11]) that the converse implication is also true.

Theorem 16 ([11,21,60]). Let **Y** be a relational structure. Then $CSP(\mathbf{Y})$ is NP-complete if $Pol(\mathbf{Y}) \rightarrow \mathcal{D}$, and it is tractable in polynomial time otherwise.

As a consequence, we obtain the following sufficient condition for the occurrence of quantum advantage.

Corollary 17. Let **Y** be a relational structure such that $CSP(\mathbf{Y})$ is not tractable in polynomial time. If dim(\mathbb{H}) ≥ 3 , **Y** has quantum advantage over \mathbb{H} .

Proof. If $\text{CSP}(\mathbf{Y})$ is not in P, it follows from Theorem 16 that $\text{Pol}(\mathbf{Y}) \to \mathcal{D}$. If $\dim(\mathbb{H}) \geq 3$, we must then have $\mathcal{Q}_{\mathbb{H}} \not\to \text{Pol}(\mathbf{Y})$, as otherwise composing the homomorphisms would yield $\mathcal{Q}_{\mathbb{H}} \to \mathcal{D}$, contradicting Proposition 15. Hence, the conclusion follows from Theorem 6. \Box

Equivalently, Corollary 17 may be phrased as follows: If $CSP(\mathbf{Y})$ is not in P and $\dim(\mathbb{H}) \geq 3$, there must exist some \mathbf{X} such that $\mathbf{X} \to_{\mathbb{H}} \mathbf{Y}$ but $\mathbf{X} \not\to \mathbf{Y}$.

5 $\mathcal{Q}_{\mathbb{H}}$ and the bounded-width minion

As we have seen, no structure has quantum advantage over Hilbert spaces of dimension 1 or 2. If the dimension is at least 3, Corollary 17 provides a sufficient condition for the occurrence of quantum advantage, in terms of the complexity of $CSP(\mathbf{Y})$. In this section, we establish a necessary condition, also formulated in terms of a complexity notion of $CSP(\mathbf{Y})$ — namely, its width.

The concept of width is central in the theory of constraint satisfaction. On a high level, the width of a CSP expresses the power of local-consistency techniques for its solvability. More formally, given two similar relational structures \mathbf{X} and \mathbf{Y} and some fixed $k \in \mathbb{N}$, the *k*-consistency algorithm provides a heuristic to check if $\mathbf{X} \to \mathbf{Y}$, by testing for the existence of a non-empty family of partial homomorphisms from substructures of \mathbf{X} of size at most *k* to \mathbf{Y} , that is closed under restriction and under extension up to size *k*. If such a family exists, \mathbf{X}

and \mathbf{Y} are said to be k-consistent.¹⁶ As long as k is a constant, there exists a polynomial-time procedure (in the size of \mathbf{X}) to check whether \mathbf{X} and \mathbf{Y} are k-consistent. If $\mathbf{X} \to \mathbf{Y}$, a family of partial homomorphisms as described above always exists. We say that $\text{CSP}(\mathbf{Y})$ has width k if the k-consistency algorithm always detects unsatisfiable instances; i.e., if \mathbf{X} and \mathbf{Y} are k-consistent precisely when $\mathbf{X} \to \mathbf{Y}$. If that is the case, this method yields a polynomial-time algorithm for the solution of $\text{CSP}(\mathbf{Y})$. CSPs that are solvable in some constant width $k \in \mathbb{N}$ are known as *bounded-width* CSPs, and constitute an important fragment of all tractable CSPs.

The goal of this section is to prove that the width of a CSP parameterised by a structure having quantum advantage over some finite-dimensional Hilbert space must be unbounded (Corollary 25). Our proof is in some sense specular to the argument in the previous section. Indeed, this time we need to show that the quantum minion $\mathcal{Q}_{\mathbb{H}}$ is located "left enough" in the minion homomorphism preorder, in that it always admits a homomorphism to one specific minion — recently introduced in [28] and, independently, in [15] — that is known to capture bounded width. Then, the result will follow from Theorem 6. We shall conclude the section by outlining an alternative argument to obtain the same result, that makes use of a different minion introduced in [27].

We start by defining the minion for bounded width.

Definition 18 ([15,28]). S is the linear minion over the abelian monoid $(\mathbb{R}, +)$ whose elements are all real matrices M of finite but arbitrary size, such that MM^{\top} is a diagonal matrix of trace 1.

The fact that the linear minion \mathscr{S} captures bounded width follows from its connection to a CSP algorithm known as the *standard semidefinite programming* relaxation (SDP). Essentially, SDP relaxes a given CSP by considering a semidefinite program whose variables are real vectors having constrained inner products. We now present the formal description of SDP, following [58]. Let **X** and **Y** be two σ -structures. We consider a real vector $\mathbf{u}_{x,y}$ for each $x \in X, y \in Y$, and a real number $v_{R,\mathbf{x},\mathbf{y}}$ for each $R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}$, and $\mathbf{y} \in R^{\mathbf{Y}}$. We also fix a real vector \mathbf{u}_0 of norm 1. Consider the system

$$\begin{split} \sum_{\mathbf{y}\in R^{\mathbf{Y}}} v_{R,\mathbf{x},\mathbf{y}} &= 1 & (R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}) \\ v_{R,\mathbf{x},\mathbf{y}} &\geq 0 & (R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{y} \in R^{\mathbf{Y}}) \\ \langle \mathbf{u}_{x_{i},y}, \mathbf{u}_{0} \rangle &= \sum_{\substack{\mathbf{y}\in R^{\mathbf{Y}} \\ y_{i} = y}} v_{R,\mathbf{x},\mathbf{y}} & \begin{pmatrix} R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \\ i \in [\operatorname{ar}(R)], y \in Y \end{pmatrix} \\ \langle \mathbf{u}_{x_{i},y}, \mathbf{u}_{x_{j},y'} \rangle &= \sum_{\substack{\mathbf{y}\in R^{\mathbf{Y}} \\ y_{i} = y, y_{j} = y'}} v_{R,\mathbf{x},\mathbf{y}} & \begin{pmatrix} R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \\ i, j \in [\operatorname{ar}(R)], y, y' \in Y \end{pmatrix}. \end{split}$$

We say that SDP solves $CSP(\mathbf{Y})$ if $\mathbf{X} \to \mathbf{Y}$ whenever the system above admits a feasible solution such that all vectors $\mathbf{u}_{x,y}$ belong to a real vector space of some finite, large-enough dimension.

Theorem 19 ([15,28]). Given a relational structure \mathbf{Y} , SDP solves $CSP(\mathbf{Y})$ if and only if $\mathcal{S} \to Pol(\mathbf{Y})$.

¹⁶We refer to [9] for the formal definition of the k-consistency algorithm.

In fact, the theorem above corresponds to the fact that SDP has precisely the same power as Test_{\mathcal{S}} (i.e., the minion test associated with the minion \mathcal{S} ; cf. [15, § 6.3.2]). The following result characterises the class of CSPs solved by SDP in terms of their width.

Theorem 20 ([10]). Given a relational structure \mathbf{Y} , SDP solves $CSP(\mathbf{Y})$ if and only if $CSP(\mathbf{Y})$ has bounded width.

As a consequence of the two results stated above, Corollary 25 would follow if we show that the quantum minion admits a homomorphism to \mathcal{S} . To that end, it shall be useful to introduce a complex version of \mathcal{S} .

Definition 21. $\mathscr{S}_{\mathbb{C}}$ is the linear minion over the abelian monoid $(\mathbb{C}, +)$ whose elements are all complex matrices M of finite but arbitrary size, such that MM^* is a diagonal matrix of trace 1.

It turns out that S and $S_{\mathbb{C}}$ are in fact the same minion, up to homomorphic equivalence.

Proposition 22. S and $S_{\mathbb{C}}$ are homomorphically equivalent.

Proof. The inclusion map yields a homomorphism $\mathcal{S} \to \mathcal{S}_{\mathbb{C}}$. To build a homomorphism in the opposite direction, take a matrix $M \in \mathcal{S}_{\mathbb{C}}^{(\ell)}$ for some $\ell \in \mathbb{N}$, and let M_R and M_I denote the real and imaginary parts of M, respectively; i.e., M_R and M_I are real matrices of the same size as M, such that $M = M_R + iM_I$. We claim that the map $\vartheta : M \mapsto \begin{bmatrix} M_R & M_I \end{bmatrix}$ yields a minion homomorphism from $\mathcal{S}_{\mathbb{C}}$ to \mathcal{S} . Observe that

$$\vartheta(M)\vartheta(M)^{\top} = M_R M_R^{\top} + M_I M_I^{\top}.$$

Furthermore,

$$MM^* = (M_R + \mathrm{i}M_I)(M_R + \mathrm{i}M_I)^* = (M_R + \mathrm{i}M_I)(M_R^\top - \mathrm{i}M_I^\top)$$
$$= (M_R M_R^\top + M_I M_I^\top) + \mathrm{i}(M_I M_R^\top - M_R M_I^\top).$$

Since MM^* is diagonal and has trace 1, it follows that also its real part $M_R M_R^\top + M_I M_I^\top$ is diagonal and has trace 1. Hence, we conclude that $\vartheta(M) \in \mathcal{S}^{(\ell)}$. The fact that ϑ preserves the arity is clear. To show that it preserves the minors, we simply observe that, for $\pi : [\ell] \to [\ell']$,

$$\vartheta(M_{/\pi}) = \begin{bmatrix} (M_{/\pi})_R & (M_{/\pi})_I \end{bmatrix} = \begin{bmatrix} (M_R)_{/\pi} & (M_I)_{/\pi} \end{bmatrix}$$
$$= \begin{bmatrix} M_R & M_I \end{bmatrix}_{/\pi} = \vartheta(M)_{/\pi},$$

as required.

We shall need the following basic properties of orthogonal projectors. Recall that, for $p \in \operatorname{End}_{\mathbb{H}}$, rg_p denotes the range of p.

Lemma 23 ([42]). Take $p, p' \in \operatorname{Proj}_{\mathbb{H}}$.

- $pp' \in \operatorname{Proj}_{\mathbb{H}}$ if and only if $[p, p'] = \mathbb{O}_{\mathbb{H}}$.
- $p + p' \in \operatorname{Proj}_{\mathbb{H}}$ if and only if $pp' = p'p = \mathbb{O}_{\mathbb{H}}$ if and only if $\operatorname{rg}_p \perp \operatorname{rg}_{p'}$. In this case, $\operatorname{rg}_{p+p'} = \operatorname{rg}_p \boxplus \operatorname{rg}_{p'}$.

We now show that the quantum minion (over any \mathbb{H}) homomorphically maps to \mathcal{S} . Essentially, the proof works by fixing an arbitrary vector in \mathbb{H} and projecting it onto the spaces constituting the entries of the elements of $\mathcal{Q}_{\mathbb{H}}$. This yields a homomorphism $\mathcal{Q}_{\mathbb{H}} \to \mathcal{S}_{\mathbb{C}}$ which, by Proposition 22, is enough to establish the result.

Theorem 24. $Q_{\mathbb{H}} \to S$.

Proof. Fix an orthonormal basis B of \mathbb{H} , and let $d = \dim(\mathbb{H})$. Given any $\mathbf{v} \in \mathbb{H}$, we shall let $r(\mathbf{v})$ be the column vector of length d containing the coordinates of \mathbf{v} in the basis B. Fix also an arbitrary vector $\mathbf{w} \in \mathbb{H}$ of norm 1. For any $\mathbf{q} \in Q_{\mathbb{H}}$ of some arity ℓ , consider the complex $\ell \times d$ matrix M defined by

$$M^* \mathbf{e}_{i;\ell} = r(\mathrm{pr}_{q_i}(\mathbf{w}))$$

for each $i \in [\ell]$. We claim that the assignment $\xi : \mathbf{q} \mapsto M$ yields a homomorphism from $\mathcal{Q}_{\mathbb{H}}$ to $\mathcal{S}_{\mathbb{C}}$.

First of all, we show that $M \in \mathcal{S}_{\mathbb{C}}$. For $i, j \in [\ell]$, the (i, j)-th entry of MM^* is

$$\mathbf{e}_{i;\ell}^* M M^* \mathbf{e}_{j;\ell} = (r(\mathrm{pr}_{q_i}(\mathbf{w})))^* r(\mathrm{pr}_{q_j}(\mathbf{w})) = \langle \mathrm{pr}_{q_j}(\mathbf{w}), \mathrm{pr}_{q_i}(\mathbf{w}) \rangle.$$

Since $q_i \perp q_j$ whenever $i \neq j$, we deduce that MM^* is a diagonal matrix. Therefore, we can write its trace as

$$\begin{aligned} \operatorname{tr}(MM^*) &= \sum_{i,j \in [\ell]} \mathbf{e}_{i;\ell}^* MM^* \mathbf{e}_{j;\ell} = \sum_{i,j \in [\ell]} \left\langle \operatorname{pr}_{q_j}(\mathbf{w}), \operatorname{pr}_{q_i}(\mathbf{w}) \right\rangle \\ &= \left\| \sum_{i \in [\ell]} \operatorname{pr}_{q_i}(\mathbf{w}) \right\|^2. \end{aligned}$$

Furthermore, by Lemma 23,

$$\sum_{i \in [\ell]} \operatorname{pr}_{q_i}(\mathbf{w}) = \operatorname{pr}_{\boxplus_{i \in [\ell]} q_i}(\mathbf{w}) = \operatorname{pr}_{\mathbb{H}}(\mathbf{w}) = \mathbf{w}.$$

It follows that $\operatorname{tr}(MM^*) = \|\mathbf{w}\|^2 = 1$, as required to conclude that $M \in \mathcal{S}_{\mathbb{C}}$.

We are left to show that ξ is a minion homomorphism. The fact that it preserves the arity directly follows from its definition. To prove that it preserves the minors, take a map $\pi : [\ell] \to [\ell']$. We need to show that $\xi(\mathbf{q}_{/\pi}) = \xi(\mathbf{q})_{/\pi}$. Pick an index $j \in [\ell']$, and observe that

$$(\xi(\mathbf{q}_{/\pi}))^* \mathbf{e}_{j;\ell'} = r(\mathrm{pr}_{(\mathbf{q}_{/\pi})_j}(\mathbf{w})) = r(\mathrm{pr}_{\boxplus_{i\in\pi^{-1}(j)}q_i}(\mathbf{w}))$$
$$= r\left(\sum_{i\in\pi^{-1}(j)} \mathrm{pr}_{q_i}(\mathbf{w})\right) = \sum_{i\in\pi^{-1}(j)} r(\mathrm{pr}_{q_i}(\mathbf{w}))$$
$$= \sum_{i\in\pi^{-1}(j)} (\xi(\mathbf{q}))^* \mathbf{e}_{i;\ell} = (\xi(\mathbf{q})_{/\pi})^* \mathbf{e}_{j;\ell'},$$

thus yielding $\xi(\mathbf{q}_{/\pi}) = \xi(\mathbf{q})_{/\pi}$, as required.

It follows that ξ yields a homomorphism from $\mathcal{Q}_{\mathbb{H}}$ to $\mathcal{S}_{\mathbb{C}}$, as claimed. To conclude that $\mathcal{Q}_{\mathbb{H}} \to \mathcal{S}$, we apply Proposition 22.

Corollary 25. Let \mathbf{Y} be a relational structure such that $CSP(\mathbf{Y})$ has bounded width.¹⁷ Then \mathbf{Y} does not have quantum advantage over any \mathbb{H} .

Proof. Suppose that $\text{CSP}(\mathbf{Y})$ has bounded width. Combining Theorems 19 and 20, we deduce that $\mathcal{S} \to \text{Pol}(\mathbf{Y})$. Composing the latter homomorphism with the one from Theorem 24, we find that $\mathcal{Q}_{\mathbb{H}} \to \text{Pol}(\mathbf{Y})$ for any \mathbb{H} . Then, the conclusion follows from Theorem 6.

Remark 26. We now discuss an alternative proof of the result stated as Corollary 25, that makes use of a different minion and a different algorithm (both introduced in [27]) instead of the minion \mathcal{S} and the algorithm SDP.

The algorithm is the constraint linear programming relaxation (CLP). Informally, when applied to a pair \mathbf{X} , \mathbf{Y} of σ -structures, CLP works by running multiple times the LP relaxation of the CSP, each time adding a different extra constraint that corresponds to fixing an assignment $\mathbf{x} \mapsto \mathbf{y}$, where $\mathbf{x} \in R^{\mathbf{X}}$, $\mathbf{y} \in R^{\mathbf{Y}}$, and $R \in \sigma$. If the feasible region of the LP is empty, the assignment $\mathbf{x} \mapsto \mathbf{y}$ is marked as unfeasible and removed from the space of solutions. The procedure continues until a fixed point is reached; then, the output is YES if and only if for each \mathbf{x} there is at least one feasible assignment. CLP is at least as strong as the *singleton* arc consistency relaxation introduced in [33] (see also [12, 24]); thus, in particular, it solves all bounded-width CSPs [50].

We now describe the minion \mathscr{C} capturing CLP. Let M be a real matrix having ℓ -many rows. We say that M is *skeletal* if, for any $i \in [\ell]$, whenever the *i*-th row of M is not identically zero there exists a column of M equal to $\mathbf{e}_{i;\ell}$. We say that M is *stochastic* if it is entrywise nonnegative and each of its columns sums up to 1. \mathscr{C} is defined as the linear minion over the abelian monoid $(\mathbb{R}, +)$ whose elements are all real skeletal stochastic matrices having finitely many rows. It was shown in [27] that \mathscr{C} captures the power of CLP, in the sense that CLP solves $\mathrm{CSP}(\mathbf{Y})$ if and only if $\mathscr{C} \to \mathrm{Pol}(\mathbf{Y})$.¹⁸

As a consequence, an alternative way of proving Corollary 25 is by showing that $\mathscr{Q}_{\mathbb{H}} \to \mathscr{C}$ for each \mathbb{H} . Let S be the set of vectors in \mathbb{H} of norm 1. Given $\mathbf{q} \in \mathscr{Q}_{\mathbb{H}}$ of some arity ℓ , consider the matrix M whose columns are indexed by the elements of S, defined as follows: For $\mathbf{h} \in S$, the **h**-th column of M is the vector

$$\begin{bmatrix} \|\operatorname{pr}_{q_1}(\mathbf{h})\|^2 & \|\operatorname{pr}_{q_2}(\mathbf{h})\|^2 & \dots & \|\operatorname{pr}_{q_\ell}(\mathbf{h})\|^2 \end{bmatrix}^\top.$$

It is straightforward to check that M is stochastic. The fact that M is skeletal follows by observing that, if $\operatorname{pr}_{q_i}(\mathbf{h}) \neq \mathbf{0}_{\mathbb{H}}$ for some $i \in [\ell]$ and some $\mathbf{h} \in S$, it must be that $q_i \neq \{\mathbf{0}_{\mathbb{H}}\}$. Given a vector $\mathbf{h}' \in q_i \cap S$, we find that the \mathbf{h}' -th column of M is $\mathbf{e}_{i;\ell}$, as needed. Therefore, $M \in \mathscr{C}$. Moreover, the map $\mathbf{q} \mapsto M$ is easily seen to preserve the arity and the minors, thus yielding $\mathcal{Q}_{\mathbb{H}} \to \mathscr{C}$.

¹⁷It is known that a CSP is solvable in bounded width if and only if it is solvable in *sublinear* width [5]. Hence, in the statement of Corollary 25, the word "bounded" can be replaced by the word "sublinear".

¹⁸The minion \mathscr{C} described here differs from the CLP minion in [27] (which we shall denote by $\tilde{\mathscr{C}}$) in two technical aspects, that are both inessential. The first is that the matrices in $\tilde{\mathscr{C}}$ are required to have rational (as opposed to real) elements. The reason why this difference is inessential is that running CLP on a CSP instance amounts to solving multiple LPs whose constraints are linear inequalities having rational coefficients. The columns of the matrices in the minion correspond to the solutions of the LPs. For such LPs, a rational solution exists if and only if a real solution exists. Hence, considering rational- or real-valued matrices is equivalent in order to capture the power of CLP. The second difference is that each matrix in $\tilde{\mathscr{C}}$ is required to have countably many columns, that are eventually all equal. Nevertheless, using a compactness argument akin to the one in [27, § 4.2], it follows that, for any relational structure $\mathbf{Y}, \mathscr{C} \to \text{Pol}(\mathbf{Y})$ if and only if $\tilde{\mathscr{C}} \to \text{Pol}(\mathbf{Y})$.

6 Quantum advantage for graphs

In the case of graphs, the sufficient and necessary conditions for quantum advantage established in the previous two sections turn out to collapse. As a consequence, we obtain the following complete characterisation of quantum advantage for graphs, thus answering Question 1.¹⁹

Theorem 27. If dim(\mathbb{H}) ≥ 3 , a graph has quantum advantage over \mathbb{H} if and only if it is non-bipartite.

(Note that no graph has quantum advantage over \mathbb{H} if dim $(\mathbb{H}) \leq 2$, see Corollary 13.) We now show how the machinery developed in the previous sections proves Theorem 27. The following, classic result by Hell and Nešetřil classifies the complexity of graph CSPs.

Theorem 28 ([41]). Let \mathbf{Y} be a graph. Then $CSP(\mathbf{Y})$ is in P if \mathbf{Y} is bipartite, and it is NP-complete otherwise.

This is enough for proving Theorem 27 under the assumption that $P \neq NP$, via Corollaries 17 and 25. To prove Theorem 27 unconditionally, we make use of the following, revisited version of Hell–Nešetřil's Theorem by Bulatov, that we formulate in terms of polymorphism minions.

Theorem 29 ([20]). Let **Y** be a graph. Then $Pol(\mathbf{Y}) \rightarrow \mathcal{D}$ if and only if **Y** is non-bipartite.

Proof of Theorem 27. Take a graph \mathbf{Y} , and suppose that \mathbf{Y} is bipartite. If \mathbf{Y} contains no edges, the result is trivial. Otherwise, \mathbf{Y} is homomorphically equivalent to the graph \mathbf{K}_2 consisting of a single, undirected edge. It is well known that $\text{CSP}(\mathbf{K}_2)$ has bounded width [34]; hence, the same holds for $\text{CSP}(\mathbf{Y})$. By Corollary 25, we conclude that \mathbf{Y} does not have quantum advantage over any \mathbb{H} .

Suppose now that \mathbf{Y} is non-bipartite. Since we are assuming that $\dim(\mathbb{H}) \geq 3$, it follows from Theorems 10 and 29 that $\mathcal{Q}_{\mathbb{H}} \not\rightarrow \operatorname{Pol}(\mathbf{Y})$. Applying Theorem 6, we deduce that \mathbf{Y} has quantum advantage over \mathbb{H} .

Remark 30. The notion of quantum perfect strategy and quantum homomorphism from [2] that we adopt in this paper — which is applicable to arbitrary relational structures — slightly differs from the one introduced in [53] for the case of graphs. We now show that Theorem 27 also holds if we make use of that alternative definition (which we shall indicate by the initials "MR").

For two graphs \mathbf{X} and \mathbf{Y} , an MR quantum homomorphism from \mathbf{X} to \mathbf{Y} over \mathbb{H} is a family of orthogonal projectors $\{p_{x,y}\}_{x \in X, y \in Y}$ in $\operatorname{Proj}_{\mathbb{H}}$ satisfying the conditions (Q1) and (Q2) from Section 2.3. In other words, commutativity between the projectors corresponding to adjacent vertices of \mathbf{X} is not enforced in the MR definition (cf. [53, Corollary 2.2]). It easily follows from the definitions that $\mathbf{X} \to_{\mathbb{H}} \mathbf{Y}$ implies $\mathbf{X} \xrightarrow{\mathrm{MR}}_{\mathbb{H}} \mathbf{Y}$. Therefore, if a graph \mathbf{Y} has quantum advantage over \mathbb{H} , it also has MR quantum advantage over \mathbb{H} . Consequently, the fact that non-bipartite graphs have MR quantum advantage when dim(\mathbb{H}) ≥ 3 immediately follows from Theorem 27. The converse direction of the statement of the theorem in the MR setting (namely, that bipartite graphs have no MR quantum advantage) was proved in [53].

 $^{^{19}}$ We point out that the "only if" part of Theorem 27 can also be derived from the results in [53], see Remark 30.

7 Quantum advantage and promise CSPs

Let \mathbf{Y} be a relational structure and, as usual, let \mathbb{H} be a finite-dimensional Hilbert space. If \mathbf{Y} does not have quantum advantage over \mathbb{H} , any structure \mathbf{X} such that $\mathbf{X} \to_{\mathbb{H}} \mathbf{Y}$ admits a homomorphism to \mathbf{Y} . Suppose now that \mathbf{Y} does have quantum advantage over \mathbb{H} . Can we still get some partial information of classical type from the fact that a quantum homomorphism $\mathbf{X} \to_{\mathbb{H}} \mathbf{Y}$ exists? A natural way of expressing this partial information — as opposed to the complete information " $\mathbf{X} \to \mathbf{Y}$ " — is by asking that a quantum homomorphism $\mathbf{X} \to_{\mathbb{H}} \mathbf{Y}$ should at least guarantee the existence of a classical homomorphism $\mathbf{X} \to \mathbf{Y}'$ for *some* structure \mathbf{Y}' . Notice that, if this is the case, we must have $\mathbf{Y} \to \mathbf{Y}'$, since $\mathbf{Y} \to_{\mathbb{H}} \mathbf{Y}$. We now give a formal description of this more general version of quantum advantage.

Definition 31. Let **Y** and **Y'** be two σ -structures such that $\mathbf{Y} \to \mathbf{Y'}$. We say that the pair $(\mathbf{Y}, \mathbf{Y'})$ has quantum advantage over \mathbb{H} if there exists a σ -structure **X** such that $\mathbf{X} \to_{\mathbb{H}} \mathbf{Y}$ but $\mathbf{X} \not\to \mathbf{Y'}$.

In other words, $(\mathbf{Y}, \mathbf{Y}')$ does *not* have quantum advantage over \mathbb{H} if and only if a quantum homomorphism $\mathbf{X} \to_{\mathbb{H}} \mathbf{Y}$ guarantees at least a classical homomorphism $\mathbf{X} \to \mathbf{Y}'$. Since the composition of a quantum and a classical homomorphism is a quantum homomorphism, we easily see that, if $(\mathbf{Y}, \mathbf{Y}')$ has quantum advantage over \mathbb{H} , then the same must hold for both \mathbf{Y} and \mathbf{Y}' individually. We will prove that the converse implication is not true in general.

Theorem 32. If dim(\mathbb{H}) \geq 3, there exists a pair of relational structures $\mathbf{Y} \rightarrow \mathbf{Y}'$ such that \mathbf{Y} and \mathbf{Y}' , but not $(\mathbf{Y}, \mathbf{Y}')$, have quantum advantage over \mathbb{H} .

It turns out that the framework we developed in the previous sections can be smoothly readapted to capture this generalised version of quantum advantage. As a consequence, just like the occurrence of quantum advantage for a single structure is linked to CSP complexity, we shall see that the occurrence of quantum advantage for a pair of structures is linked to the complexity of a *promise* version of CSPs that we now describe.

For two σ -structures \mathbf{Y} and \mathbf{Y}' such that $\mathbf{Y} \to \mathbf{Y}'$, the *promise* CSP parameterised by \mathbf{Y} and \mathbf{Y}' (in symbols, $\text{PCSP}(\mathbf{Y}, \mathbf{Y}')$) is the following computational problem: Given as input a σ -structure \mathbf{X} , output YES if $\mathbf{X} \to \mathbf{Y}$, and No if $\mathbf{X} \neq \mathbf{Y}'$. The fact that $\mathbf{Y} \to \mathbf{Y}'$ ensures that the answer sets are disjoint. Notice that $\text{PCSP}(\mathbf{Y}, \mathbf{Y})$ is precisely $\text{CSP}(\mathbf{Y})$, so PCSPs are a generalisation of CSPs.

This framework was formally introduced in [6] and [14] to study approximability of perfectly satisfiable CSPs, but particular examples of PCSPs have been studied for a long time. For instance, given two natural numbers $n \leq n'$, PCSP($\mathbf{K}_n, \mathbf{K}_{n'}$) is the famous approximate graph colouring problem, whose complexity is not known for general n and n' — while the complexity of its non-promise version, graph n-colouring, was already classified in [48].

Unlike for CSPs, the complexity landscape of PCSPs is widely unknown. In particular, a dichotomy for PCSP complexity is not known to hold. Nevertheless, part of the machinery from the algebraic approach to CSPs does extend to the promise setting. Following [6], we let a polymorphism of arity ℓ for the pair $(\mathbf{Y}, \mathbf{Y}')$ be a homomorphism from \mathbf{Y}^{ℓ} to \mathbf{Y}' . The set $Pol(\mathbf{Y}, \mathbf{Y}')$ of all polymorphisms of $(\mathbf{Y}, \mathbf{Y}')$ is a subset of $\mathscr{F}_{Y,Y'}$ (from Example 2) and it is closed under minors. Hence, just like $Pol(\mathbf{Y})$, $Pol(\mathbf{Y}, \mathbf{Y}')$ is a (function) minion.

The next result generalises Theorem 6, by showing that the quantum minimal also captures quantum advantage for pairs of structures.

Theorem 33. Let $\mathbf{Y} \to \mathbf{Y}'$ be relational structures. Then $(\mathbf{Y}, \mathbf{Y}')$ has quantum advantage over \mathbb{H} if and only if $\mathcal{Q}_{\mathbb{H}} \not\to \operatorname{Pol}(\mathbf{Y}, \mathbf{Y}')$.

Recall the notion of minion test used in Section 3. Given a minion \mathcal{M} , we say that $\text{Test}_{\mathcal{M}}$ solves $\text{PCSP}(\mathbf{Y}, \mathbf{Y}')$ if, for every structure \mathbf{X} similar to \mathbf{Y} and \mathbf{Y}' , $\text{Test}_{\mathcal{M}}(\mathbf{X}, \mathbf{Y}) = \text{YES}$ implies $\mathbf{X} \to \mathbf{Y}'$. Theorem 8 extends to PCSPs (in fact, the following is its original formulation in [28]).

Theorem 34 ([28]). Let \mathscr{M} be a minion and let $\mathbf{Y} \to \mathbf{Y}'$ be relational structures. Then $\text{Test}_{\mathscr{M}}$ solves $\text{PCSP}(\mathbf{Y}, \mathbf{Y}')$ if and only if $\mathscr{M} \to \text{Pol}(\mathbf{Y}, \mathbf{Y}')$.

The proof of Theorem 33 is then entirely analogous to that of Theorem 6.

Proof of Theorem 33. The pair $(\mathbf{Y}, \mathbf{Y}')$ has quantum advantage over \mathbb{H} if and only if there exists some \mathbf{X} such that $\text{Test}_{\mathcal{Q}_{\mathbb{H}}}(\mathbf{X}, \mathbf{Y}) = \text{YES}$ but $\mathbf{X} \not\to \mathbf{Y}'$; i.e., if and only if $\text{Test}_{\mathcal{Q}_{\mathbb{H}}}$ does not solve $\text{PCSP}(\mathbf{Y}, \mathbf{Y}')$. Hence, the conclusion follows from Theorem 34.

The following statement is then a direct consequence of Theorems 10, 24, and 33.

Corollary 35. Let $\mathbf{Y} \to \mathbf{Y}'$ be relational structures.

- (1) If dim(\mathbb{H}) ≤ 2 , (\mathbf{Y}, \mathbf{Y}') has no quantum advantage over \mathbb{H} .
- (2) If dim(\mathbb{H}) ≥ 3 and Pol(\mathbf{Y}, \mathbf{Y}') $\rightarrow \mathcal{D}$, (\mathbf{Y}, \mathbf{Y}') has quantum advantage over \mathbb{H} .
- (3) If $\mathcal{S} \to \text{Pol}(\mathbf{Y}, \mathbf{Y}')$, $(\mathbf{Y}, \mathbf{Y}')$ has no quantum advantage over any \mathbb{H} .

How to link the result above to PCSP complexity? It was shown in [8] that Theorem 4 keeps holding in the PCSP setting. I.e., a minion homomorphism $\operatorname{Pol}(\mathbf{Y}, \mathbf{Y}') \to \operatorname{Pol}(\tilde{\mathbf{Y}}, \tilde{\mathbf{Y}}')$ induces a polynomial-time (log-space) reduction from $\operatorname{PCSP}(\tilde{\mathbf{Y}}, \tilde{\mathbf{Y}}')$ to $\operatorname{PCSP}(\mathbf{Y}, \mathbf{Y}')$. Nevertheless, some care needs to be taken when reformulating the conditions in parts (2) and (3) of Corollary 35 in terms of PCSP-complexity statements. Regarding part (2), in sharp contrast to the CSP setting, it is not true that the polymorphism minion of any NP-hard PCSP admits a homomorphism to \mathcal{D} , as detailed in [8]. Therefore, unless P = NP, we cannot extend Corollary 17 to the PCSP setting.

Example 36. Given a graph \mathbf{X} , let its quantum chromatic number be the quantity $\chi_{\mathbf{q}}(\mathbf{X}) = \min\{n \in \mathbb{N} : \mathbf{X} \to_{\mathbb{H}} \mathbf{K}_n\}$. Then, for $n \leq n' \in \mathbb{N}$, quantum advantage for the pair $(\mathbf{K}_n, \mathbf{K}_{n'})$ corresponds to a separation between quantum and classical chromatic numbers. Indeed, $(\mathbf{K}_n, \mathbf{K}_{n'})$ has quantum advantage over \mathbb{H} precisely when there exists a graph having quantum chromatic number $\leq n$ and classical chromatic number > n'. It was shown in [13] (see also [8]) that, given integers $3 \leq n \leq n'$, $\operatorname{Pol}(\mathbf{K}_n, \mathbf{K}_{n'}) \to \mathcal{D}$ if and only if $n' \leq 2n - 2$. Thus, if $\dim(\mathbb{H}) \geq 3$ and $n \geq 3$, Corollary 35 implies that there exist graphs having quantum chromatic number at most n and classical chromatic number at least 2n-1. However, no larger separation can be derived through Corollary 35.²⁰

²⁰We remark that this version of quantum chromatic number differs from the one introduced in [23] (which we denote by $\tilde{\chi}_{q}$), in that the latter makes use of projectors that are not necessarily commuting (cf. Remark 30). In particular, for any graph **X**, it holds that $\tilde{\chi}_{q}(\mathbf{X}) \leq \chi_{q}(\mathbf{X}) \leq \chi(\mathbf{X})$, where χ is the classical chromatic number. An exponential separation between $\tilde{\chi}_{q}$ and χ is known, see [7, 17, 19].

As for part (3) of Corollary 35, the condition $\mathcal{S} \to \text{Pol}(\mathbf{Y}, \mathbf{Y}')$ does capture PCSP solvability via SDP; also in this case, the following, more general formulation of Theorem 19 is the one given in [15, 28].

Theorem 37 ([15,28]). Let $\mathbf{Y} \to \mathbf{Y}'$ be relational structures. Then SDP solves $PCSP(\mathbf{Y}, \mathbf{Y}')$ if and only if $\mathcal{S} \to Pol(\mathbf{Y}, \mathbf{Y}')$.

As we see next, the result above allows finding a separation for the notions of quantum advantage for single structures and for pairs of structures.

Proof of Theorem 32. Let \mathbf{Z} and \mathbf{Z}' be the relational structures encoding the CSPs "1-IN-3 SAT" and "NOT-ALL-EQUAL SAT", respectively. I.e., \mathbf{Z} and \mathbf{Z}' have Boolean domain and a unique, ternary relation

$$R^{\mathbf{Z}} = \{(1,0,0), (0,1,0), (0,0,1)\},\$$
$$R^{\mathbf{Z}'} = \left\{\begin{array}{c}(1,0,0), (0,1,0), (0,0,1),\\ (0,1,1), (1,0,1), (1,1,0)\end{array}\right\}.$$

It is well known [38] that both $\text{CSP}(\mathbf{Z})$ and $\text{CSP}(\mathbf{Z}')$ are NP-complete problems, and it is straightforward to check that $\text{Pol}(\mathbf{Z}) \to \mathcal{D}$ and $\text{Pol}(\mathbf{Z}') \to \mathcal{D}$. Hence, since we are assuming $\dim(\mathbb{H}) \geq 3$, both \mathbf{Z} and \mathbf{Z}' have quantum advantage over \mathbb{H} . On the other hand, it was shown in [15] that $\text{PCSP}(\mathbf{Z}, \mathbf{Z}')$ is solved by SDP. It follows from Theorem 37 that $\mathcal{S} \to \text{Pol}(\mathbf{Z}, \mathbf{Z}')$. Hence, Corollary 35 implies that the pair $(\mathbf{Z}, \mathbf{Z}')$ has no quantum advantage over \mathbb{H} . \Box

Remark 38. The fact that the pair $(\mathbf{Z}, \mathbf{Z}')$ from the proof of Theorem 32 has no quantum advantage over any \mathbb{H} can also be obtained by using the minion \mathscr{C} from Remark 26 instead of \mathscr{S} . Indeed, also in this case, \mathscr{C} captures solvability by the CLP relaxation in the more general setting of PCSPs: It was proved in [27] that CLP solves PCSP $(\mathbf{Y}, \mathbf{Y}')$ if and only if $\mathscr{C} \to \text{Pol}(\mathbf{Y}, \mathbf{Y}')$. Since, as shown in Remark 26, $\mathscr{Q}_{\mathbb{H}} \to \mathscr{C}$ for each \mathbb{H} , part (3) of Corollary 35 also holds if we replace \mathscr{S} by \mathscr{C} . It was shown in [14] that PCSP $(\mathbf{Z}, \mathbf{Z}')$ is solved by CLP; thus, $\mathscr{C} \to \text{Pol}(\mathbf{Z}, \mathbf{Z}')$, which means that $(\mathbf{Z}, \mathbf{Z}')$ has no quantum advantage over any \mathbb{H} .

We also point out that, while a notion of width for PCSPs is known [3,8], neither of the minions \mathscr{S} and \mathscr{C} corresponds to bounded width for PCSPs. Indeed, as we have seen, $\operatorname{Pol}(\mathbf{Z}, \mathbf{Z}')$ admits a homomorphism from both \mathscr{S} and \mathscr{C} but, as shown in [3], $\operatorname{PCSP}(\mathbf{Z}, \mathbf{Z}')$ has unbounded width.

Acknowledgements

The author is grateful to David E. Roberson, Laura Mančinska, and Samson Abramsky for helpful conversations on quantum homomorphisms and quantum advantage.

References

 Samson Abramsky, Rui Soares Barbosa, Giovanni Carù, Nadish de Silva, Kohei Kishida, and Shane Mansfield. Minimum quantum resources for strong non-locality. In Mark M. Wilde, editor, 12th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2017, June 14-16, 2017, Paris, France, volume 73 of LIPIcs, pages 9:1–9:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPICS.TQC.2017.9.

- [2] Samson Abramsky, Rui Soares Barbosa, Nadish de Silva, and Octavio Zapata. The quantum monad on relational structures. In 42nd International Symposium on Mathematical Foundations of Computer Science, volume 83 of LIPIcs. Leibniz Int. Proc. Inform., pages Art. No. 35, 19. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2017.
- [3] Albert Atserias and Víctor Dalmau. Promise constraint satisfaction and width. In Proc. 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA'22), pages 1129–1153, 2022. arXiv:2107.05886, doi:10.1137/1.9781611977073.48.
- [4] Albert Atserias, Phokion G. Kolaitis, and Simone Severini. Generalized satisfiability problems via operator assignments. J. Comput. Syst. Sci., 105:171–198, 2019. doi:10.1016/J.JCSS.2019.05.003.
- [5] Albert Atserias and Joanna Ochremiak. Proof complexity meets algebra. ACM Trans. Comput. Log., 20(1):1:1–1:46, 2019. doi:10.1145/3265985.
- [6] Per Austrin, Venkatesan Guruswami, and Johan Håstad. (2+ε)-SAT is NP-hard. SIAM J. Comput., 46(5):1554–1573, 2017. doi:10.1137/15M1006507.
- [7] David Avis, Jun Hasegawa, Yosuke Kikuchi, and Yuuya Sasaki. A quantum protocol to win the graph colouring game on all Hadamard graphs. *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.*, 89-A(5):1378–1381, 2006. doi:10.1093/IETFEC/E89-A.5.1378.
- [8] Libor Barto, Jakub Bulín, Andrei A. Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. J. ACM, 68(4):28:1-28:66, 2021. arXiv:1811.00970, doi:10.1145/ 3457606.
- [9] Libor Barto and Marcin Kozik. Constraint satisfaction problems solvable by local consistency methods. J. ACM, 61(1), 2014. Article No. 3. doi:10.1145/2556646.
- [10] Libor Barto and Marcin Kozik. Robustly solvable constraint satisfaction problems. SIAM J. Comput., 45(4):1646-1669, 2016. arXiv:1512.01157, doi:10.1137/130915479.
- [11] Libor Barto, Jakub Opršal, and Michael Pinsker. The wonderland of reflections. Isr. J. Math, 223(1):363-398, Feb 2018. arXiv:1510.04521, doi:10.1007/s11856-017-1621-9.
- [12] Christian Bessiere and Romuald Debruyne. Theoretical analysis of singleton arc consistency and its extensions. Artif. Intell., 172(1):29–41, 2008. doi:10.1016/j.artint.2007.09.001.
- [13] Joshua Brakensiek and Venkatesan Guruswami. New hardness results for graph and hypergraph colorings. In Proc. 31st Conference on Computational Complexity (CCC'16), volume 50 of LIPIcs, pages 14:1–14:27. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs. CCC.2016.14.
- [14] Joshua Brakensiek and Venkatesan Guruswami. Promise constraint satisfaction: Algebraic structure and a symmetric Boolean dichotomy. SIAM J. Comput., 50(6):1663-1700, 2021. arXiv: 1704.01937, doi:10.1137/19M128212X.
- [15] Joshua Brakensiek, Venkatesan Guruswami, and Sai Sandeep. SDPs and robust satisfiability of promise CSP. In Proc. 55th Annual ACM Symposium on Theory of Computing (STOC'23), pages 609–622. ACM, 2023. arXiv:2211.08373, doi:10.1145/3564246.3585180.
- [16] Joshua Brakensiek, Venkatesan Guruswami, Marcin Wrochna, and Stanislav Živný. The power of the combined basic LP and affine relaxation for promise CSPs. SIAM J. Comput., 49:1232–1248, 2020. arXiv:1907.04383, doi:10.1137/20M1312745.
- [17] Gilles Brassard, Richard Cleve, and Alain Tapp. Cost of exactly simulating quantum entanglement with classical communication. *Physical Review Letters*, 83(9):1874, 1999.
- [18] Gilles Brassard, André Allan Méthot, and Alain Tapp. Minimum entangled state dimension required for pseudo-telepathy. *Quantum Inf. Comput.*, 5(4):275–284, 2005. doi:10.26421/QIC5.45-2.

- [19] Harry Buhrman, Richard Cleve, and Avi Wigderson. Quantum vs. classical communication and computation. In Jeffrey Scott Vitter, editor, Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998, pages 63–68. ACM, 1998. doi:10.1145/276698.276713.
- [20] Andrei A. Bulatov. H-coloring dichotomy revisited. Theor. Comput. Sci., 349(1):31-39, 2005.
 URL: https://doi.org/10.1016/j.tcs.2005.09.028, doi:10.1016/J.TCS.2005.09.028.
- [21] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In Proc. 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS'17), pages 319–330, 2017. arXiv: 1703.03021, doi:10.1109/F0CS.2017.37.
- [22] Andrei A. Bulatov, Peter Jeavons, and Andrei A. Krokhin. Classifying the complexity of constraints using finite algebras. SIAM J. Comput., 34(3):720–742, 2005. doi:10.1137/S0097539700376676.
- [23] Peter J. Cameron, Ashley Montanaro, Michael W. Newman, Simone Severini, and Andreas J. Winter. On the quantum chromatic number of a graph. *Electron. J. Comb.*, 14(1), 2007. doi:10.37236/999.
- [24] Hubie Chen, Víctor Dalmau, and Berit Grußien. Arc consistency and friends. J. Log. Comput., 23(1):87-108, 2013. doi:10.1093/logcom/exr039.
- [25] Lorenzo Ciardo and Stanislav Živný. Approximate graph colouring and crystals. In Proc. 2023 ACM-SIAM Symposium on Discrete Algorithms (SODA'23), pages 2256–2267, 2023. arXiv: 2210.08293, doi:10.1137/1.9781611977554.ch86.
- [26] Lorenzo Ciardo and Stanislav Živný. Approximate graph colouring and the hollow shadow. In Proc. 55th Annual ACM Symposium on Theory of Computing (STOC'23), pages 623–631. ACM, 2023. arXiv:2211.03168, doi:10.1145/3564246.3585112.
- [27] Lorenzo Ciardo and Stanislav Živný. CLAP: A new algorithm for promise CSPs. SIAM Journal on Computing, 52(1):1–37, 2023. arXiv:2107.05018, doi:10.1137/22M1476435.
- [28] Lorenzo Ciardo and Stanislav Živný. Hierarchies of minion tests for PCSPs through tensors. In Proc. 2023 ACM-SIAM Symposium on Discrete Algorithms (SODA'23), pages 568–580, 2023. arXiv:2207.02277, doi:10.1137/1.9781611977554.ch25.
- [29] Richard Cleve, Peter Høyer, Benjamin Toner, and John Watrous. Consequences and limits of nonlocal strategies. In 19th Annual IEEE Conference on Computational Complexity (CCC 2004), 21-24 June 2004, Amherst, MA, USA, pages 236–249. IEEE Computer Society, 2004. doi:10.1109/CCC.2004.1313847.
- [30] Richard Cleve and Rajat Mittal. Characterization of binary constraint system games. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, Automata, Languages, and Programming 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I, volume 8572 of Lecture Notes in Computer Science, pages 320–331. Springer, 2014. doi:10.1007/978-3-662-43948-7_27.
- [31] Víctor Dalmau, Phokion G. Kolaitis, and Moshe Y. Vardi. Constraint satisfaction, bounded treewidth, and finite-variable logics. In Pascal Van Hentenryck, editor, Principles and Practice of Constraint Programming - CP 2002, 8th International Conference, CP 2002, Ithaca, NY, USA, September 9-13, 2002, Proceedings, volume 2470 of Lecture Notes in Computer Science, pages 310–326. Springer, 2002. doi:10.1007/3-540-46135-3_21.
- [32] Víctor Dalmau and Jakub Opršal. Local consistency as a reduction between constraint satisfaction problems. 2023. arXiv:2301.05084.
- [33] Romuald Debruyne and Christian Bessière. Some practicable filtering techniques for the constraint satisfaction problem. In Proc. 15th International Joint Conference on Artificial Intelligence (IJCAI'97), pages 412–417. Morgan Kaufmann, 1997.

- [34] Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. SIAM J. Comput., 28(1):57–104, 1998. doi:10.1137/S0097539794266766.
- [35] Junya Fukawa, Hiroshi Imai, and François Le Gall. Quantum coloring games via symmetric SAT games. In Asian Conference on Quantum Information Science (AQIS'11), 2011.
- [36] Viktor Galliard and Stefan Wolf. Pseudo-telepathy, entanglement, and graph colorings. In *Proceedings IEEE International Symposium on Information Theory*, page 101. IEEE, 2002.
- [37] M. R. Garey and David S. Johnson. The complexity of near-optimal graph coloring. J. ACM, 23(1):43–49, 1976. doi:10.1145/321921.321926.
- [38] Michael R. Garey and David S. Johnson. Computers and Intractability: A Guide to the Theory of NP-Completeness. W.H. Freeman, 1979.
- [39] Andrew M Gleason. Measures on the closed subspaces of a Hilbert space. In The Logico-Algebraic Approach to Quantum Mechanics: Volume I: Historical Evolution, pages 123–133. Springer, 1975.
- [40] Martin Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. J. ACM, 54(1):1:1–1:24, 2007. doi:10.1145/1206035.1206036.
- [41] Pavol Hell and Jaroslav Nešetřil. On the complexity of H-coloring. J. Comb. Theory, Ser. B, 48(1):92–110, 1990. doi:10.1016/0095-8956(90)90132-J.
- [42] Leslie Hogben. Handbook of Linear Algebra. CRC Press, 2013.
- [43] Tsuyoshi Ito, Hirotada Kobayashi, and Keiji Matsumoto. Oracularization and two-prover one-round interactive proofs against nonlocal strategies. 24th Annual IEEE Conference on Computational Complexity, pages 217–228, 2008.
- [44] Tsuyoshi Ito and Thomas Vidick. A multi-prover interactive proof for NEXP sound against entangled provers. In 53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012, pages 243-252. IEEE Computer Society, 2012. doi:10.1109/F0CS.2012.11.
- [45] Peter G. Jeavons. On the algebraic structure of combinatorial problems. Theor. Comput. Sci., 200(1-2):185-204, 1998. doi:10.1016/S0304-3975(97)00230-2.
- [46] Peter G. Jeavons, David A. Cohen, and Marc Gyssens. Closure properties of constraints. J. ACM, 44(4):527–548, 1997. doi:10.1145/263867.263489.
- [47] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. MIP^{*} = RE. Commun. ACM, 64(11):131-138, 2021. doi:10.1145/3485628.
- [48] Richard M. Karp. Reducibility among combinatorial problems. In Proc. Complexity of Computer Computations, pages 85–103, 1972. doi:10.1007/978-1-4684-2001-2_9.
- [49] Julia Kempe, Hirotada Kobayashi, Keiji Matsumoto, Ben Toner, and Thomas Vidick. Entangled games are hard to approximate. SIAM J. Comput., 40(3):848–877, 2011. doi:10.1137/090751293.
- [50] Marcin Kozik. Solving CSPs using weak local consistency. SIAM J. Comput., 50(4):1263–1286, 2021. arXiv:1605.00565, doi:10.1137/18M117577X.
- [51] Andrei Krokhin, Jakub Opršal, Marcin Wrochna, and Stanislav Živný. Topology and adjunction in promise constraint satisfaction. SIAM J. Comput., 52(1):37–79, 2023. arXiv:2003.11351, doi:10.1137/20M1378223.
- [52] Laura Mančinska and David E. Roberson. Oddities of quantum colorings. Balt. J. Mod. Comput., 4(4), 2016. doi:10.22364/BJMC.2016.4.4.16.
- [53] Laura Mančinska and David E. Roberson. Quantum homomorphisms. J. Comb. Theory, Ser. B, 118:228-267, 2016. doi:10.1016/j.jctb.2015.12.009.

- [54] Laura Mančinska, Giannicola Scarpa, and Simone Severini. New separations in zero-error channel capacity through projective kochen-specker sets and quantum coloring. *IEEE Trans. Inf. Theory*, 59(6):4025–4032, 2013. doi:10.1109/TIT.2013.2248031.
- [55] Anand Natarajan and John Wright. NEEXP is contained in MIP. In David Zuckerman, editor, 60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019, pages 510-518. IEEE Computer Society, 2019. doi: 10.1109/FOCS.2019.00039.
- [56] Nicholas Pippenger. Galois theory for minors of finite functions. Discret. Math., 254(1-3):405–419, 2002. doi:10.1016/S0012-365X(01)00297-7.
- [57] Eduard Prugovecki. Quantum Mechanics in Hilbert Space. Academic Press, 1982.
- [58] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In Proc. 40th Annual ACM Symposium on Theory of Computing (STOC'08), pages 245-254, 2008. doi:10.1145/1374376.1374414.
- [59] Giannicola Scarpa and Simone Severini. Kochen–Specker sets and the rank-1 quantum chromatic number. IEEE Trans. Inf. Theory, 58(4):2524–2529, 2012. doi:10.1109/TIT.2011.2178018.
- [60] Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. J. ACM, 67(5):30:1-30:78, 2020. arXiv:1704.01914, doi:10.1145/3402029.

A Omitted proofs

Proposition (Proposition 7 restated). Let \mathbf{X}, \mathbf{Y} be similar relational structures. Then $\mathbf{X} \to_{\mathbb{H}} \mathbf{Y}$ if and only if $\mathbf{X} \to \mathbb{F}_{\hat{\mathbb{Q}}_{\mathbb{H}}}(\mathbf{Y})$.

The proof of Proposition 7 makes use of the following result.

Lemma 39 ([57]). Let $\{p_1, p_2, \ldots, p_\ell\} \subseteq \operatorname{Proj}_{\mathbb{H}}$ be a finite family of orthogonal projectors. Then $\operatorname{id}_{\mathbb{H}} - \sum_{i \in [\ell]} p_i \succeq \mathbb{O}_{\mathbb{H}}$ if and only if $p_i p_j = \mathbb{O}_{\mathbb{H}}$ for each $i \neq j \in [\ell]$.

Proof of Proposition 7. Suppose that $\mathbf{X} \to_{\mathbb{H}} \mathbf{Y}$, and take a family $\{p_{x,y}\}_{x \in X, y \in Y}$ of orthogonal projectors in $\operatorname{Proj}_{\mathbb{H}}$ witnessing it. Let *n* be the domain size of \mathbf{Y} . Consider the function *f* mapping $x \in X$ to the vector of length *n* whose *y*-th element is $\operatorname{rg}_{p_{x,y}}$ for each $y \in Y$. We claim that $f(x) \in \mathcal{Q}_{\mathbb{H}}^{(n)}$. By (Q₁), the family $\{p_{x,y}\}_{y \in Y}$ satisfies the hypothesis of Lemma 39; hence, $p_{x,y}p_{x,y'} = \mathbb{O}_{\mathbb{H}}$ for each $y \neq y' \in Y$. Then, a repeated application of the second part of Lemma 23 yields that $\operatorname{rg}_{p_{x,y'}} \perp \operatorname{rg}_{p_{x,y'}}$ for each $y \neq y' \in Y$, and

$$\prod_{y \in Y} \operatorname{rg}_{p_{x,y}} = \operatorname{rg}_{\sum_{y \in Y} p_{x,y}} = \operatorname{rg}_{\operatorname{id}_{\mathbb{H}}} = \mathbb{H}.$$

This proves the claim, and shows that $f: X \to \mathcal{Q}_{\mathbb{H}}^{(n)}$ is a well-defined function. We now claim that f yields a homomorphism from \mathbf{X} to $\mathbb{F}_{\mathcal{Q}_{\mathbb{H}}}(\mathbf{Y})$. To that end, take a symbol R of arity r in the common signature σ of \mathbf{X} and \mathbf{Y} , and consider a tuple $\mathbf{x} \in R^{\mathbf{X}}$. We need to prove that $f(\mathbf{x}) \in R^{\mathbb{F}_{\mathcal{Q}_{\mathbb{H}}}(\mathbf{Y})}$ (where $f(\mathbf{x})$ denotes the entrywise application of f to the entries of \mathbf{x}). Consider the vector \mathbf{w} of length $m = |R^{\mathbf{Y}}|$ whose \mathbf{y} -th entry, for $\mathbf{y} \in R^{\mathbf{Y}}$, is $\operatorname{rg}_{p_{\mathbf{x},\mathbf{y}}}$ (where we recall that $p_{\mathbf{x},\mathbf{y}} = p_{x_1,y_1} \cdot p_{x_2,y_2} \cdot \ldots \cdot p_{x_r,y_r}$). The condition (Q₂) guarantees that the factors in $p_{\mathbf{x},\mathbf{y}}$ are pairwise commuting; hence, by the first part of Lemma 23, we find $p_{\mathbf{x},\mathbf{y}} \in \operatorname{Proj}_{\mathbb{H}}$. Notice also that

$$\sum_{\mathbf{y}\in R^{\mathbf{Y}}} p_{\mathbf{x},\mathbf{y}} = \sum_{\mathbf{y}\in Y^r} p_{\mathbf{x},\mathbf{y}} = \left(\sum_{y\in Y} p_{x_1,y}\right) \left(\sum_{y\in Y} p_{x_2,y}\right) \dots \left(\sum_{y\in Y} p_{x_r,y}\right) = \mathrm{id}_{\mathbb{H}}^r = \mathrm{id}_{\mathbb{H}},$$

where the first equality comes from (Q_3) and the third from (Q_1) . We deduce through Lemma 39 that $p_{\mathbf{x},\mathbf{y}}p_{\mathbf{x},\mathbf{y}'} = \mathbb{O}_{\mathbb{H}}$ whenever $\mathbf{y} \neq \mathbf{y}'$. It then follows from Lemma 23 that the entries of \mathbf{w} are pairwise orthogonal spaces summing up to \mathbb{H} , which yields $\mathbf{w} \in \mathcal{Q}_{\mathbb{H}}^{(m)}$. Recall from Section 3 that, given $i \in [r]$, π_i is the map assigning to each $\mathbf{y} \in R^{\mathbf{Y}}$ its *i*-th entry y_i . We are left to show that $f(x_i) = \mathbf{w}_{/\pi_i}$ for each $i \in [r]$. Given $y \in Y$, we find

$$\sum_{\mathbf{y}\in\pi_i^{-1}(y)} p_{\mathbf{x},\mathbf{y}} = \sum_{\substack{\mathbf{y}\in R^{\mathbf{Y}}\\y_i=y}} p_{\mathbf{x},\mathbf{y}} = \sum_{\substack{\mathbf{y}\in Y^r\\y_i=y}} p_{\mathbf{x},\mathbf{y}} = \mathrm{id}_{\mathbb{H}}^{r-1} \cdot p_{x_i,y} = p_{x_i,y}.$$

It follows that the *y*-th entry of $\mathbf{w}_{/\pi_i}$ is

$$\bigoplus_{\mathbf{y}\in\pi_i^{-1}(y)}\mathrm{rg}_{p_{\mathbf{x},\mathbf{y}}}=\mathrm{rg}_{\sum_{\mathbf{y}\in\pi_i^{-1}(y)}p_{\mathbf{x},\mathbf{y}}}=\mathrm{rg}_{p_{x_i,y}},$$

which equals the y-th entry of $f(x_i)$. This proves the claim and shows that $f(\mathbf{x}) \in R^{\mathbb{F}_{\mathcal{Q}_{\mathbb{H}}}(\mathbf{Y})}$, thus establishing that f is indeed a homomorphism from \mathbf{X} to $\mathbb{F}_{\mathcal{Q}_{\mathbb{H}}}(\mathbf{Y})$.

To prove the converse implication, let $f : \mathbf{X} \to \mathbb{F}_{\mathcal{Q}_{\mathbb{H}}}(\mathbf{Y})$ be a homomorphism. Recall that, for a linear subspace $v \in L_{\mathbb{H}}$, $\mathrm{pr}_v \in \mathrm{Proj}_{\mathbb{H}}$ denotes the orthogonal projector onto v; for typographical convenience, we shall denote pr_v by $\mathrm{pr}(v)$ in the rest of this proof. For $x \in X$ and $y \in Y$, let $p_{x,y} = \mathrm{pr}(f(x)_y)$, where $f(x)_y \in L_{\mathbb{H}}$ is the y-th entry of the vector f(x). We claim that the family $\{p_{x,y}\}_{x \in X, y \in Y}$ is a witness for the existence of a quantum homomorphism from \mathbf{X} to \mathbf{Y} over \mathbb{H} .

The condition (Q₁) is easily satisfied as $\coprod_{y \in Y} f(x)_y = \mathbb{H}$. Take now $R \in \sigma$ of arity r, $\mathbf{x} \in R^{\mathbf{X}}, k, \ell \in [r]$, and $y, y' \in Y$. Since f is a homomorphism, there exists some vector $\mathbf{w} \in \mathcal{Q}_{\mathbb{H}}^{(m)}$ (where $m = |R^{\mathbf{Y}}|$) such that $f(x_i) = \mathbf{w}_{/\pi_i}$ for each $i \in [r]$. Observe that

$$f(x_i)_y = (\mathbf{w}_{/\pi_i})_y = \bigoplus_{\mathbf{y} \in \pi_i^{-1}(y)} w_{\mathbf{y}} = \bigoplus_{\substack{\mathbf{y} \in R^{\mathbf{Y}} \\ y_i = y}} w_{\mathbf{y}}.$$
 (6)

Hence, we have

$$p_{x_k,y} = \operatorname{pr}\left(\bigoplus_{\substack{\mathbf{y} \in R^{\mathbf{Y}} \\ y_k = y}} w_{\mathbf{y}}\right) = \sum_{\substack{\mathbf{y} \in R^{\mathbf{Y}} \\ y_k = y}} \operatorname{pr}(w_{\mathbf{y}})$$

and, similarly,

$$p_{x_{\ell},y'} = \sum_{\substack{\mathbf{y} \in R^{\mathbf{Y}} \\ y_{\ell} = y'}} \operatorname{pr}(w_{\mathbf{y}}).$$

Using the linearity of the commutator and the fact that the product of orthogonal projectors onto orthogonal spaces is zero (the second part of Lemma 23), we deduce that

$$[p_{x_k,y}, p_{x_\ell,y'}] = \mathbb{O}_{\mathbb{H}},$$

which shows that the condition (Q₂) holds. Take now $\mathbf{z} \in Y^r \setminus R^{\mathbf{Y}}$. By (6), we have

$$p_{x_i, z_i} = \operatorname{pr}\left(\bigoplus_{\substack{\mathbf{y} \in R^{\mathbf{Y}} \\ y_i = z_i}} w_{\mathbf{y}}\right) = \sum_{\substack{\mathbf{y} \in R^{\mathbf{Y}} \\ y_i = z_i}} \operatorname{pr}(w_{\mathbf{y}})$$

for each $i \in [r]$. Using again that the product of orthogonal projectors onto orthogonal spaces is zero and the fact that $\mathbf{z} \notin R^{\mathbf{Y}}$, we conclude that

$$p_{\mathbf{x},\mathbf{z}} = \prod_{i \in [r]} p_{x_i,z_i} = \prod_{i \in [r]} \sum_{\substack{\mathbf{y} \in R^{\mathbf{Y}} \\ y_i = z_i}} \operatorname{pr}(w_{\mathbf{y}}) = \sum_{\substack{\mathbf{y} \in R^{\mathbf{Y}} \\ \mathbf{y} = \mathbf{z}}} \operatorname{pr}(w_{\mathbf{y}}) = \mathbb{O}_{\mathbb{H}},$$

as needed to show that also the condition (Q_3) is satisfied.

Lemma (Lemma 11 restated). If dim(\mathbb{H}) = 2, there exists $C \subset \mathbb{H}$ such that, given any two nonzero orthogonal vectors $\mathbf{v}, \mathbf{w} \in \mathbb{H}$, exactly one of \mathbf{v} and \mathbf{w} lies in C.

Proof. Consider the map

$$f: \mathbb{C}^2 \to \mathbb{R} \times \mathbb{C}$$
$$(x, y) \mapsto (x\bar{x} - y\bar{y}, 2x\bar{y}).$$

Let S^3 be the subset of \mathbb{C}^2 containing all pairs $(x, y) \in \mathbb{C}^2$ such that $x\bar{x} + y\bar{y} = 1$, and let S^2 be the subset of $\mathbb{R} \times \mathbb{C}$ containing all pairs $(t, z) \in \mathbb{R} \times \mathbb{C}$ such that $t^2 + z\bar{z} = 1$. A straightforward calculation shows that f maps elements in S^3 to elements in S^2 . Moreover, given two orthogonal vectors (x, y) and (p, q) in S^3 (i.e., $x\bar{p} + y\bar{q} = 0$), notice that the matrix

$$M = \begin{bmatrix} x & p \\ y & q \end{bmatrix}$$

satisfies $M^*M = I_2$; i.e., M is unitary. As a consequence,

$$I_2 = MM^* = \begin{bmatrix} x\bar{x} + p\bar{p} & x\bar{y} + p\bar{q} \\ y\bar{x} + q\bar{p} & y\bar{y} + q\bar{q} \end{bmatrix},$$

and we deduce that f(x, y) = -f(p, q). In other words, f maps pairs of orthogonal vectors of \mathbb{C}^2 having norm 1 to antipodal vectors in S^2 .

Choose an antipodal partition of S^2 ; i.e., a partition $S^2 = U \cup V$ such that, for each $(t, z) \in S^2$, (t, z) and (-t, -z) lie in different parts. For example, we might let U be the set of vectors $(t, z) \in S^2$ such that t > 0, or t = 0 and $z_R > 0$ (where z_R is the real part of z), or t = 0 and z = i, and $V = S^2 \setminus U$. Fix an orthonormal basis B for \mathbb{H} , and let $r : \mathbb{H} \to \mathbb{C}^2$ be the function mapping elements of \mathbb{H} to the corresponding vectors of coordinates in the basis B. Finally, consider the function $g : \mathbb{H} \setminus \{\mathbf{0}_{\mathbb{H}}\} \to S^2$ mapping \mathbf{v} to the image of $\frac{\mathbf{v}}{\|\mathbf{v}\|}$ under the function $f \circ r$. We define $C = g^{-1}(U)$. If \mathbf{v}, \mathbf{w} are two nonzero orthogonal vectors in \mathbb{H} , the argument above implies that $g(\mathbf{v}) = -g(\mathbf{w})$ and, hence, exactly one of $g(\mathbf{v})$ and $g(\mathbf{w})$ lies in U. This shows that the set C has the required property.