

Entangled Constraint Satisfaction Problems

Reductions, Complexity, and Approximation Algorithms

by

Eric Časlav Culf

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Applied Mathematics – Quantum Information

Waterloo, Ontario, Canada, 2026

© Eric Časlav Culf 2026

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner:	Ryan O'Donnell Professor, Dept. of Computer Science, Carnegie Mellon University
Supervisor(s):	Richard Cleve Professor, Dept. of Computer Science, University of Waterloo William Slofstra Associate Professor, Dept. of Pure Mathematics, University of Waterloo
Internal Member:	Graeme Smith Associate Professor, Dept. of Applied Mathematics, University of Waterloo
Internal-External Member:	Michael Brannan Associate Professor, Dept. of Pure Mathematics, University of Waterloo
Other Member(s):	Shalev Ben-David Associate Professor, Dept. of Computer Science, University of Waterloo

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see the Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

This thesis is based on the following work:

- Eric Culf, Hamoon Mousavi, and Taro Spirig. Approximation algorithms for noncommutative CSPs. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 920–929. IEEE, 2024. ([Chapter 7](#)) [[52](#)]
- Eric Culf, Josse van Dobben de Bruyn, Matthijs Vernooij, and Peter Zeman. Existence and nonexistence of commutativity gadgets for entangled CSPs. *arXiv preprint arXiv:2509.07835*, 2025. ([Chapters 3 and 4](#)) [[47](#)]
- Eric Culf, Josse van Dobben de Bruyn, and Peter Zeman. Quantum polymorphism characterisation of commutativity gadgets in all quantum models. *arXiv preprint arXiv:2604.01408*, 2026. ([Chapter 4](#)) [[48](#)]
- Eric Culf. Almost perfect strategies for projection games are approximately tracial. *arXiv preprint arXiv:2603.14746*, 2026. ([Chapter 3](#)) [[45](#)]
- Richard Cleve, Eric Culf, and Aviv Tallor. XOR games at full tilt: The hardness of binary nonlocal games. *arXiv preprint arXiv:2607.06876*, 2026. ([Chapter 6](#)) [[40](#)]
- Eric Culf. The complexity of entangled graph colouring via polymorphisms, 2026. In preparation. ([Chapter 5](#)) [[46](#)]

Abstract

The result $\text{MIP}^* = \text{RE}$ shows that the problem of approximating the quantum value of a nonlocal game is undecidable. This thesis studies how this undecidability may be distilled to well-structured classes of nonlocal games, specifically those coming from constraint satisfaction problems (CSPs). A CSP is a decision problem where one must decide if a given set of constraints may be simultaneously satisfied; presenting it as a nonlocal game gives rise to an entangled CSP where the possible solutions are generalised to operator assignments. Many computational problems may be presented as CSPs, such as graph colouring. This work studies the various families of entangled CSP that arise due to the noncommutativity of quantum mechanics — notably non-oracular CSPs, where the players only receive variables as questions, and oracular CSPs, where the players also receive constraints.

We study how reductions between CSPs may be made quantum-sound. First, we develop the theory of commutativity gadgets, which allow any classical constraint expressible by the CSP to be expressed in the entangled setting. This allows for classical reductions between CSPs to lift to gap-preserving reductions between entangled CSPs. We fully characterise the existence of commutativity gadgets by means of entangled polymorphisms of the CSP, a type of quantum symmetry, and use this to show equivalences and separations between commutativity gadgets in different models. Second, we extend the most powerful classical reduction between CSPs, based on minion homomorphisms of the polymorphism spaces, to the setting of gapped entangled CSPs. We apply both techniques to show that entangled graph k -colouring is undecidable for all $k \geq 3$.

Famously, XOR games, which correspond to the CSP of boolean linear equations with two variables per equation, have a classical value which is **NP**-hard to approximate but a quantum value that can be approximated in polynomial time. We define the tilted XOR games, which correspond to the boolean linear CSP with one or two variables per equation. Though these are qualitatively close to XOR games, we show that their quantum value is undecidable to approximate, by quantising the classical reduction from linear system games to XOR games. This gives the first example of binary nonlocal games, where both players answer a single bit, whose quantum value is undecidable to approximate.

Finally, we study approximation algorithms for entangled CSPs. We develop an algorithm to approximate the quantum value of Max- k -Cut, the optimisation analogue of graph k -colouring. For $k = 3$ and $k = 4$, we find approximation ratios that are strictly better than for classical colouring, despite the problems being significantly harder. Our algorithm consists of a relaxation to a semidefinite program, followed by a randomised rounding using a Haar-random unitary. To analyse the algorithm, we make use of ideas from free probability to simplify the behaviour of the rounding in the limit of large dimension.

Acknowledgements

First, I would like to thank my supervisors, Richard Cleve and William Slofstra. I know I can always count on them for support or advice, or simply to listen to some half-baked idea.

I would like to thank my family: Mama, Papa, and Baka for supporting my studies since the very first day of kindergarten; my brother Marko for friendship and music; and Sausage and Vuk for the unconditional joy that only dogs know.

Next, I thank my officemates, Ben and Tiasa, for always making our office a happy place to work.

Also, I would like to thank my Masters supervisor, Anne Broadbent, for introducing me to quantum information science and supporting me ever since.

Finally, I would like to thank Archishna for her love and friendship. Nowhere is there a truer best friend, a more loving partner, or a more insightful colleague.

Dedication

*Za Baku. Mislim na tebe, kao što si uvek mislila na mene.*¹

¹For my grandmother. Thinking of you, as you have always thought of me.

Table of Contents

Examining Committee Membership	ii
Author’s Declaration	iii
Statement of Contributions	iv
Abstract	vi
Acknowledgements	vii
Dedication	viii
List of Figures	xv
List of Tables	xvi
1 Introduction	1
1.1 Summary of Results	4
1.1.1 Entangled constraint satisfaction problems	4
1.1.2 Existence of commutativity gadgets	5
1.1.3 Reductions in the absence of commutativity gadgets	10
1.1.4 Complexity of binary nonlocal games	11
1.1.5 Approximation algorithms for entangled CSPs	13
1.2 Outline	15

2	Preliminaries	16
2.1	Notation	16
2.2	Categories	17
2.3	Groups	17
2.3.1	Presentations	17
2.3.2	Representations	18
2.4	Algebras	18
2.4.1	$\ast$ - and $C^\ast$ -algebras	18
2.4.2	States and representations	20
2.4.3	von Neumann algebras	20
2.4.4	Noncommutative L^p spaces	23
2.4.5	Weighted algebras	24
2.5	Probability theory	25
2.5.1	Distributions and random variables	25
2.5.2	The Haar distribution	26
2.5.3	Random variables on the circle	26
2.6	Quantum theory	28
2.6.1	The postulates of quantum mechanics	28
2.6.2	States, operators, and measurements	30
2.6.3	Entanglement	31
2.7	Quantum spaces	32
2.7.1	Motivation	32
2.7.2	Quantum groups	32
2.8	Free probability theory	33
2.9	Complexity theory	35
2.9.1	Turing machines	35
2.9.2	Complexity classes	36

2.9.3	Reductions and hardness	36
2.10	Relational structures	37
2.10.1	Definitions	37
2.10.2	Graphs	38
2.10.3	Constraint satisfaction problems	39
2.10.4	Reductions and polymorphisms	40
2.11	Semidefinite programming	44
2.11.1	Definitions and duality	44
2.11.2	Approximation algorithms and Max-Cut	45
2.12	Nonlocal games	47
2.12.1	Games and strategies	47
2.12.2	XOR games	49
2.12.3	Synchronous and projection games	50
2.12.4	Complexity of nonlocal games	53
3	Nonlocal games for constraint satisfaction	56
3.1	Entangled constraint satisfaction problems	56
3.1.1	Relational structure homomorphisms as nonlocal games	57
3.1.2	Weighted algebra formalism	59
3.1.3	Definitions of entangled CSPs	62
3.1.4	Some equivalences based on stability	66
3.1.5	Complexity of entangled CSPs	69
3.2	Quantum relational structure homomorphisms	70
3.2.1	The structure of perfect strategies for CSP nonlocal games	70
3.2.2	Cocategories	72
3.2.3	The cocategory of relational structures	74
3.3	Almost perfect strategies for projection games are approximately tracial	78
3.3.1	Degenerate and non-degenerate projection games	78

3.3.2	Perfect quantum case	80
3.3.3	Perfect commuting operator case	82
3.3.4	Approximate case	83
3.3.5	Application to CSPs and oracularisation	90
4	Commutativity gadgets	92
4.1	Reductions based on commutativity gadgets	92
4.1.1	Definitions	92
4.1.2	Hardness from the existence of a commutativity gadget	94
4.1.3	Algebraic commutativity gadgets	103
4.1.4	Examples of commutativity gadget	104
4.2	Commutativity gadget for oracular entangled graph colouring	105
4.3	Polymorphism characterisation	107
4.3.1	Exact commutativity gadgets	108
4.3.2	Stable commutation	110
4.3.3	Robust commutativity gadgets	114
4.3.4	Implications and equivalences	118
4.4	Polymorphism algebras of the complete graphs	121
4.5	Separations	125
4.5.1	Relational structures with no commutativity gadget	126
4.5.2	Oracular and non-oracular gadgets	130
4.5.3	Separations based on solution groups	131
4.6	Complexity of deciding the existence of a commutativity gadget	139
5	Reductions beyond commutativity gadgets	141
5.1	Entangled polymorphism minions	141
5.2	Gapless reduction	144
5.3	Gap-preserving reduction	147

5.3.1	Weighted cominions	148
5.3.2	Stable minion homomorphisms	149
5.3.3	Proof of the gap-preserving reduction	151
5.4	Stable minion homomorphisms from cominion homomorphisms	153
5.5	Examples of reductions without commutativity gadgets	157
5.5.1	Hardness of entangled colouring	157
5.5.2	Hardness of homomorphism to the diamond graph	158
5.6	Equivalence of the constraint-variable and constraint-constraint frameworks	159
5.6.1	Reduction from constraint-variable to constraint-constraint	159
5.6.2	Reduction from constraint-constraint to constraint-variable	160
6	Hardness of binary games	162
6.1	Classical hardness of XOR games	162
6.2	Tilted XOR games	164
6.2.1	Definitions and basic properties	164
6.2.2	The tilted bipartite cube game	166
6.3	Gapless hardness of tilted XOR games	167
6.3.1	Optimal strategies of the tilted bipartite cube game	168
6.3.2	Reduction from 3LIN	170
6.4	Gapped hardness of tilted XOR games	171
6.4.1	Improved XOR game structure theorem	171
6.4.2	Application to the tilted cube game	173
6.4.3	Proof of the reduction	178
6.5	Hardness of approximation	180
7	Approximation algorithms for entangled CSPs	183
7.1	Graph colouring and Max-Cut problems	183
7.2	An approximation algorithm	186

7.2.1	SDP relaxation	186
7.2.2	Vector-to-unitary construction	188
7.2.3	Rounding unitaries	189
7.2.4	The algorithm	190
7.3	Analysis of the algorithm	191
7.3.1	Relative distribution	191
7.3.2	Integral formula	195
7.3.3	Application to Max- k -Cut	198
7.3.4	Minimising the ratio	202
7.4	Dimension-efficient algorithm	203
7.4.1	Generalised Weyl-Brauer group and strong isometry	204
7.4.2	The algorithm and its analysis	212
8	Discussion	216
8.1	Related work	217
8.2	Open questions	218
	References	221

List of Figures

1.1	Implications and separations for existence of commutativity gadgets	9
1.2	Comparing the hardness of XOR and tilted XOR games	12
1.3	Example of the different presentations of graph 2-colouring	14
1.4	Approximation ratios for Max- k -Cut	15
4.1	The graphs $\overline{C}_6$ and $\overline{C}_8$ presented as prisms	106
4.2	Relationships between commutativity gadget classes	120
4.3	The diamond graph D	130
4.4	Separations between commutativity gadget classes	131
6.1	Schematic representation of the tilted bipartite cube game	168
6.2	Hardness of tilted XOR games	182
7.1	The wrapped Cauchy distribution	194
7.2	The fidelity function	199

List of Tables

1.1	Sizes of known commutativity gadgets	7
7.1	Approximation ratios for classical and entangled Max- k -Cut for various k .	185
7.2	Approximation ratios for classical and entangled HMax-Lin(k) for various k .	186

Chapter 1

Introduction

Quantum mechanics has given rise to a variety of perplexing implications that have forced the reconsideration of many preconceived notions about the nature of the universe. One of these is **entanglement**, which is a form of correlation stronger than anything possible classically. A pair of particles may become maximally entangled, which implies that the result of a measurement of one particle fully prescribes the resulting state of the other, irrespective of the observable measured. When this property was first noticed by Einstein, Podolsky, and Rosen [61], it was expected to be a flaw in the theory of quantum mechanics, as entangled states violate locality. That is, measuring one system can instantaneously change the state of a distant other system, seemingly transmitting information faster than the speed of light. This would violate a key principle of relativity. They anticipated that there could be a theory underlying quantum mechanics where there are hidden variables, which remain local and could be used to recover locality.

To test the veracity of this idea, one needs a framework to compare correlations due to local hidden variables to those predicted by entangled quantum systems. The work of Bell [16] defined the first such framework, which has become known as a **Bell inequality**. This is a classically-defined scenario where one asks a correlation to satisfy certain constraints, with respect to which any classical correlation coming from a local hidden variable theory fares strictly worse than a correlation allowed by quantum mechanics. By violating a Bell inequality, one can disprove all local hidden variable theories at once. This was achieved for the first time experimentally by Aspect, Dalibard, and Roger [11], cementing entanglement as a tangible aspect of the universe.

Where does this leave locality? The standard Copenhagen interpretation of quantum mechanics sees the universe as inherently nonlocal, but still posits that information cannot

travel faster than the speed of light. This squares with the Einstein-Podolsky-Rosen paradox because measurements in quantum mechanics are also **probabilistic**: until an observer learns the measurement result of the first particle – which is constrained to travel at or below the speed of light – they see the state of the second particle as a random mixture of all possible resulting states, containing no information about the result of the first measurement. This implies that quantum correlations are **non-signalling** in the sense that the probability of a measurement result on the second system, conditioned on the observables measured on the two systems, is independent of the observable used to measure the first system. This suggests a simple resolution to the question of entanglement — perhaps all non-signalling correlations may be attained by quantum mechanics? As always, nature proves much more subtle and beautiful. The set of quantum correlations lies somewhere between the sets of classical and non-signalling correlations, but equals neither of them.

To quantify these ideas, we turn to **nonlocal games**, which provide a distillation of the properties evinced by Bell inequalities. In a nonlocal game, a referee interacts with two players (*i.e.* observers) Alice and Bob, who may share a correlation but do not communicate. The referee samples a pair of questions x and y from a predefined probability distribution, and sends them to Alice and Bob, respectively. Then, the players use their correlation to come up with answers a and b locally. Then, they each send their answer to the referee, who decides if the players win, by means of a predefined rule predicate depending on the questions and answers. The value of the correlation is the expected winning probability; and we may compare the classical, quantum, and non-signalling values, which correspond to the suprema of the values over the corresponding sets of correlations.

The prototypical example of a nonlocal game is the **CHSH game**, first studied by Clauser, Horne, Shimony, and Holt [39]. In the CHSH game, the questions are single bits 0 or 1, sampled uniformly at random; the answers are also bits; and the players win if the answers satisfy $a + b = xy$ modulo 2. This nonlocal game witnesses the separation between quantum correlations and both classical and non-signalling correlations in the following way. The classical value is equal to $\frac{3}{4}$, but there is a quantum correlation with value $\cos^2(\frac{\pi}{8}) \approx 0.85$. Further, this correlation actually attains the quantum value, by Tsirelson’s bound [114]. However, there is a non-signalling correlation, called the Popescu-Rohrlich box, that attains value 1 [103].

As seen in the above example, the quantum value of a nonlocal game can say a lot about the nature of quantum correlations. With this motivation, one may ask how hard it is to compute the value of a nonlocal game or, similarly, how much information may be gleaned by determining the value of a nonlocal game. From the perspective of computer science, this falls within the purview of complexity theory, which aims to capture how hard it is to solve a given computational problem. Specifically, a nonlocal game can be seen as a form of

multiprover interactive proof system. In a multiprover interactive proof (MIP) system, a verifier interacts with untrusted provers, who have unbounded computational power but do not communicate, to attempt to decide a problem. Under classical strategies, the lack of communication between the provers allows the verifier to interrogate them separately and corroborate their answers, allowing the verifier to decide much harder problems than they could alone, or even with a single prover. Using an MIP, a polynomial-time verifier can decide NEXP, the nondeterministic exponential-time problems [14]. By shrinking the question and answer lengths to logarithmic in the instance size, MIPs also capture the nondeterministic polynomial-time problems NP [8, 7, 56].

Nonlocal games with quantum strategies correspond to entangled multiprover interactive proof (MIP*) systems. At first, it was highly unclear whether this strengthened the class, by allowing the verifier to demand more, or weakened the class, by allowing the provers to cheat better. It was finally shown that the complexity of entangled MIP is *significantly* higher than the non-entangled version. The breakthrough result of [76] shows that a polynomial-time verifier can use an MIP* to decide RE, which is a complexity class containing even undecidable computational problems. Specifically, RE are those problems which are at most as hard as the **halting problem**, which is the problem of deciding if a computer program will halt and produce an output, given its source code. The halting problem is impossible to decide using a computer [115, 82]. This incredibly hard complexity is attained in the worst case for general nonlocal games: is it possible to restrict the problem to a smaller well-structured class of nonlocal games while preserving the hardness?

A well-structured class of decision problems whose complexity is fully understood are the **constraint satisfaction problems** (CSPs). An instance of a CSP consists of a collection of variables and associated constraints over a domain; it is an accepting instance of the decision problem if there is a map from the variables to the domain that satisfies all of the constraints. This captures a variety of relevant computational problems, such as satisfiability of boolean formulae or graph colouring. CSP complexity satisfies a dichotomy theorem, which had long been conjectured and proved only recently, stating that any CSP is either solvable in polynomial-time or NP-complete [32, 122]. This dichotomy relies heavily on the symmetry of the CSP: there is a precise sense in which the CSPs with more symmetry are easier than those with less. Also, CSPs embed naturally into MIPs, by having the verifier ask constraints and variables to the provers.

In this thesis, we study the quantum values of nonlocal games coming from CSPs, which we call **entangled CSPs**, in order to achieve a more fine-grained understanding of where the hard MIP* instances occur. We consider the problem from two directions. When and how do classical reductions between CSPs lift to reductions between the corresponding entangled CSPs? Is it possible to approximate the quantum value of an entangled CSP?

1.1 Summary of Results

We show a variety of reductions between entangled CSPs, coming from commutativity gadgets and polymorphism minion homomorphisms. We use these reductions to demonstrate hardness of the entangled versions of a variety of CSPs, notably graph colouring. Via different techniques, we also show hardness of nonlocal games where each player outputs a single bit. Finally, we develop an approximation algorithm for entangled CSPs, which we apply to graph colouring, among other problems. This section summarises these results, roughly in the order they are presented in the thesis.

1.1.1 Entangled constraint satisfaction problems

There are a few equivalent presentations of constraint satisfaction problems; we make use of the presentation via **relational structures** in this work. A relational structure consists of a set called the domain, and a collection of labelled relations of potentially different arities. A homomorphism of relational structures is a map between the domains that preserves each of the relations — note for a homomorphism between relational structures to make sense, the labels of the relations need to be the same for the two structures, which we assure by fixing a set of labels called the relational signature. In this setting, the CSP parametrised by a relational structure A is the problem of deciding whether there exists a homomorphism from a given relational structure X to A . X is an accepting instance (yes instance) if there is a homomorphism $X \rightarrow A$ and a rejecting instance (no instance) if not. Here, a constraint corresponds to a single element of a relation of X . We can also define a **gapped CSP** as the promise problem where we guarantee that the fraction of constraints that may be simultaneously satisfied in a rejecting instance is upper-bounded.

An important class of CSPs arises from graphs. A graph gives rise to a relational structure with a single binary relation; and a homomorphism of graphs as relational structures is simply a graph homomorphism. Graph n -colouring is the graph homomorphism problem to the complete graph K_n . Many other important decision problems may be phrased as CSPs, such as 3SAT and satisfiability of linear relations.

There a variety of ways to present an instance of a CSP as a nonlocal game. These reduce to two relevant cases: a nonlocal game where the players are asked to satisfy variables, or a nonlocal game where the players are asked to satisfy whole constraints. The second case is called the **oracular** framework and, in contrast, the first case is called the non-oracular framework. In either of the two cases, the winning condition is that the answered assignments to the variables must satisfy the associated constraints, and be consistent between the

players. With respect to classical strategies, there is no significant difference between the frameworks — the existence of a perfect strategy in either framework corresponds exactly to the existence of a relational structure homomorphism. However, if we consider quantum strategies, the two frameworks diverge, requiring us to consider them separately. This is because, in the oracular framework, variables that belong to the same relation must be measured together on the same system, inducing additional commutation relations between their observables.

Considering the framework as well as the class of strategies – quantum, quantum approximate, or commuting operator – gives rise to a variety of different entangled CSPs for each relational structure. We show first that there exists a relational structure for which both the quantum and quantum approximate gapped entangled CSPs are **RE**-complete and the commuting-operator gapped entangled CSP is **coRE**-complete, in both the oracular and non-oracular frameworks. This problem forms the basis of all the hardness results based on reductions between CSPs. Next, we study the algebraic structure of perfect strategies for entangled CSPs. Notably, we can see perfect strategies for CSP nonlocal games as a quantum generalisation of homomorphisms between the relational structures. Finally, we show that near-perfect strategies for CSP nonlocal games imply near-perfect strategies with respect to a tracial (*i.e.* maximally entangled) state, so we may reduce to the case of tracial strategies, for entangled CSPs with perfect completeness.

1.1.2 Existence of commutativity gadgets

Commutativity gadgets provide a way of lifting classical reductions between CSPs to quantum-sound reductions between their entangled analogues. A commutativity gadget for a relational structure A is an instance G of the CSP of A along with a pair of variables from G (called distinguished variables), such that for any pair of elements of A , there exists a quantum homomorphism from G to A where the distinguished variables are sent to those elements; and for every quantum homomorphism from G to A , the observables associated to the distinguished variables commute. That is, a commutativity gadget provides a way to restrict a pair of variables to be commuting while imposing no further restrictions.

A reduction between CSPs is local in the sense that the reduction proceeds by replacing every constraint of the original instance with a conjunction of constraints. Therefore, a classical reduction can be lifted to a reduction of entangled CSPs by placing a commutativity gadget between all the pairs of variables in the new instance that come from the same constraint of the original instance. This forces all of the local assignments to satisfy the classical constraint, guaranteeing soundness of the reduction. As such, if a CSP is **NP**-complete and admits a commutativity gadget, the associated (gapped) quantum entangled

CSP must be RE-complete; and the associated commuting operator entangled CSP must be coRE-complete.

The idea of a commutativity gadget was introduced by Ji, who found commutativity gadgets for 3-colouring and 1-in-3-SAT [75]. [49] shows that all NP-complete boolean CSPs admit a commutativity gadget: in fact, every non-TVF CSP (where there is a relation that does not constrain a pair of variables at all) admits a commutativity gadget consisting of a single relation in the oracular framework, and boolean TVF CSPs admit commutativity gadgets in both the oracular and non-oracular models. [38] shows that the odd cycles and the Siggers digraph (a directed graph on 3 vertices) admit commutativity gadgets. In this work, we show that all the complete graphs on $n \geq 3$ vertices admit commutativity gadgets generalising the construction of [75], in the oracular framework.

On the other hand, there are strong constraints on which CSPs admit commutativity gadgets. We show that non-classical **quantum symmetries** of the relational structure preclude the existence of any commutativity gadget. This connects reductions between entangled CSPs to the theory of quantum groups (in the sense of Woronowicz [121]). An illustrative example is the case of four-colouring, where these are the symmetries of the graph K_4 . Classically, the symmetries of K_4 are the permutations of the vertices, as for K_3 . Beyond this, it is known that K_4 has a non-classical quantum endomorphism [119]; in contrast, every quantum endomorphism of K_3 is classical. As noted above, however, this no-go result only applies in the non-oracular framework — every complete graph admits a commutativity gadget in the oracular framework. Using this noncommutative quantum homomorphism, we can show that K_4 does not admit a commutativity gadget.

Theorem 1.1.1. K_4 does not admit a (non-oracular) commutativity gadget.

In the proof, for clarity, we write quantum homomorphisms as collections of PVMs, whereas later in the work we present them as representations of a homomorphism algebra — these two presentations are wholly equivalent.

Proof. Suppose on the contrary that K_4 admits a commutativity gadget G with distinguished variables $x, y \in G$. First, since there exists a non-classical quantum homomorphism $\{\Pi_b^a\}_{a,b \in K_4} : K_4 \rightarrow K_4$, there exist $a_1, b_1, a_2, b_2 \in K_4$ such that $\Pi_{b_1}^{a_1}$ and $\Pi_{b_2}^{a_2}$ do not commute. Next, the relations on the gadget cannot constrain the values x and y may be assigned to, so there must exist a quantum homomorphism $\{P_a^z\}_{z \in G, a \in K_4} : G \rightarrow K_4$ such that $P_a^x = \delta_{a,a_1}$ and $P_a^y = \delta_{a,a_2}$. Consider the composition of P with Π , which is the quantum homomorphism $G \rightarrow K_4$ defined as

$$(\Pi \circ P)_a^z = \sum_{b \in K_4} P_b^z \otimes \Pi_a^b.$$

This is in fact a well-defined quantum homomorphism. Furthermore, we find that

$$\begin{aligned}(\Pi \circ P)_{b_1}^x &= \sum_{b \in K_4} P_b^x \otimes \Pi_{b_1}^b = 1 \otimes \Pi_{b_1}^{a_1}, \\ (\Pi \circ P)_{b_2}^y &= \sum_{b \in K_4} P_b^y \otimes \Pi_{b_2}^b = 1 \otimes \Pi_{b_2}^{a_2},\end{aligned}$$

meaning that $(\Pi \circ P)_{b_1}^x$ and $(\Pi \circ P)_{b_2}^y$ do not commute. Therefore, there exists a quantum four-colouring of G that does not constrain the projectors for x and y to be commuting. So G cannot be a commutativity gadget for K_4 . ■

What does this say about the hardness of quantum four-colouring? We know that unlike for classical three-colouring, we cannot use the classical reductions to directly rewrite the constraints of any CSP in terms of it. However, this does not preclude the possibility that entangled four-colouring is undecidable. Nevertheless, since the classical reduction does not quantise, we would need a natively quantum reduction to achieve this.

We also use the no-go theorem to find an example of a relational structure that does not admit a commutativity gadget in either the oracular or non-oracular frameworks. This is the diamond graph D , which is the complete graph K_4 with one edge removed. To prove this result, we extend a result of Schmidt [106], which constructs quantum automorphisms of graphs from disjoint classical automorphisms, to work in the setting of (non-invertible) endomorphisms.

Structure	Oracular	Non-oracular
Non-TVF	$\text{ar}(R)$	?
TVF Boolean	16	16
Siggers digraph	19683	19683
K_n	$2n$	✗
C_{2n+1}	$(2n+1)^{(2n+1)^2}$	$(2n+1)^{(2n+1)^2}$
Diamond graph	✗	✗

Table 1.1: Sizes of commutativity gadgets for the relational structures where they are known to exist, in terms of the number of variables. The question mark indicates a case where it is unknown whether a commutativity gadget exists or not, and the ✗ mark indicates a case where it is known that there is no commutativity gadget. $\text{ar}(R)$ is the arity of one of the relations of the structure.

Recent work of Ciardo, Joubert, and Mottet [38] has extended this no-go theorem to use not just homomorphisms, but also **polymorphisms**, which are homomorphisms from a finite number of copies of the relational structure to the structure itself. Surprisingly, extended to this setting, the converse of the result also holds: a relational structure admits a commutativity gadget if and only if all of its quantum polymorphisms are classical. The converse is also constructive in the sense that if a relational structure A admits a commutativity gadget, then it admits a commutativity gadget on $|A|^{|A|^2}$ vertices. Of course, it may be possible to find significantly smaller or simpler commutativity gadgets, for example the oracular gadget we find for K_n . The relational structures known to admit commutativity gadgets, as well as the sizes of the gadgets, are displayed in Table 1.1.

We extend this characterisation in two ways. First, we extend it to different models of entanglement. Apart from quantum commutativity gadgets, we also get quantum approximate and commuting operator gadgets, which apply directly to different models of entangled CSP. On top of this, we consider algebraic gadgets, which work at the level of the entangled CSP C^* -algebra, and are useful in showing existence of all other types of gadget. Second, we extend the characterisation to work in the setting of robust commutativity gadgets, which facilitate reductions between gapped entangled CSPs. This requires a notion of stable commutativity of the entangled polymorphism algebras under approximate representations. Knowing this characterisation allows us to find a variety of relationships between the different types of commutativity gadget. Notably, the existence of a non-oracular gadget always implies the existence of an oracular gadget; robust quantum and quantum approximate gadgets are equivalent to exact quantum approximate gadgets; and robust and exact commuting operator gadgets are equivalent. This is illustrated in Fig. 1.1

On the other hand, the characterisation also allows us to show separations between different classes of commutativity gadget. First, as noted earlier, the fact that the complete graphs on ≥ 4 vertices admit an oracular but not a non-oracular commutativity gadget separates these two models of commutativity gadget. Next, we also show a separation between quantum and quantum approximate commutativity gadgets. To do so, we embed a hyperlinear solution group which has no non-trivial finite-dimensional representations into the polymorphisms of a relational structure. We construct this group by taking the solution group with no noncommutative finite-dimensional representations constructed by Slofstra [109], and using a construction using the magic square game and a nice solution group presentation of the dihedral group to eliminate the finite-dimensional representations. Then, to embed this solution group into the polymorphisms, we adapt a construction of van Dobben de Bruyn, Roberson, and Schmidt [58] to realise the solution group as the entangled endomorphisms of a coloured graph. To allow the endomorphisms to describe all the polymorphisms, we make use of a construction based on a complete graph. We give




Figure 1.1: An overview of the implications and separations for existence of commutativity gadgets in different models. An arrow from one model to another indicates that existence of a commutativity gadget in the first model implies existence of a commutativity gadget in the second. Models that are equivalent in this sense are grouped together in shaded areas. Red lines between the areas indicate separations, meaning that there are CSPs which admit a commutativity gadget in the weaker model but not in the stronger model. The dashed red line indicates a separation that is conditional on the existence of a non-hyperlinear group.

a complete description of the entangled polymorphisms of a complete graph, and show that all of the polymorphisms are essentially unary in the sense that it decomposes as a direct sum of quantum automorphisms of the input copies. Using this description, placing a complete graph as an additional relation in our constructed relational structure allows to control all the polymorphisms using only the endomorphisms. Under the assumption of a non-hyperlinear group, we can use an analogous construction to separate non-oracular quantum approximate and commuting operator commutativity gadgets as well. These separations are also illustrated in Fig. 1.1.

Finally, we investigate the complexity of deciding the existence of a commutativity gadget. Using the same construction, we can reduce the problem of deciding whether a linear system game has a perfect strategy in a specific quantum model, known to be undecidable by [109, 110], to the problem of deciding if a given relational structure has a commutativity gadget. This allows us to show, in the non-oracular framework, that deciding if a relational structure admits a quantum or quantum approximate commutativity gadget is undecidable, and deciding if a relational structure admits a commuting operator or algebraic commutativity gadget is **coRE**-hard.

1.1.3 Reductions in the absence of commutativity gadgets

The polymorphisms of a relational structure can also be used to access a stronger form of reduction than that made possible by commutativity gadgets. These are the reductions based on **minion homomorphisms**, which are mappings between the spaces of (quantum or classical) polymorphisms that preserve some compositions. Classically, if there is a minion homomorphism from the polymorphisms of A to the polymorphisms of B , then the CSP of B reduces to the CSP of A . It turns out that such reductions are sufficient to capture the hardness of all classically NP-complete CSPs. Ciardo, Joubert, and Mottet [38] lift this form of reduction to the quantum setting, showing that if there is a minion homomorphism from the entangled polymorphisms of A to the entangled polymorphisms of B , then the gapless entangled CSP of B reduces to the gapless entangled CSP of A .

We extend this result to the setting of gap-preserving reductions, allowing us to show undecidability for gapped entangled CSPs with no commutativity gadget. To do so, we develop a theory of **stable homomorphisms** between minions, which allow us quantify how well approximate satisfaction of constraints is preserved. This allows us to show that if there exists a stable minion homomorphism from the entangled polymorphisms of A to the entangled polymorphisms of B , then there is a reduction from the gapped entangled CSP of B to the gapped entangled CSP of A .

To allow us to apply this result to specific CSPs, we show a simpler sufficient condition for the existence of a stable minion homomorphism. We introduce an algebraic dual notion of a minion called a **comin** — the noncommutative polymorphism algebras provide examples. We show that a homomorphism between polymorphism cominions always implies a stable minion homomorphism.

We apply these results to show hardness in a few important cases. First, using our complete characterisation of the (non-oracular) polymorphism algebras of K_n , we show that entangled n -colouring is RE-complete, by means of a comin homomorphism to the polymorphism comin of 3SAT, which is classical. This reduction also allows us to establish hardness for entangled graph homomorphism to the diamond graph D , by using the fact that it is homomorphism equivalent with K_3 . These two examples give us that we can nevertheless recover hardness for the CSPs that we have found to fail to admit commutativity gadgets.

As a final application of this minion-based reduction, we show an equivalence between oracular entangled CSPs where the players only receive constraints, and those where the players receive both constraints and variables. This is done by noting that the polymorphism cominions of these two models are equal, giving a trivial stable minion homomorphism.


1.1.4 Complexity of binary nonlocal games

For general nonlocal games, determining the quantum value is significantly harder than determining the classical value. The situation is reversed for XOR games, which correspond to 2LIN, the CSP of linear equations with two variables per equation over the boolean field. Cleve, Høyer, Toner, and Watrous [41] show that the quantum value of an XOR game can be approximated in polynomial time, whereas the classical value is NP-complete to approximate [72]. This is due to a result of Tsirelson [114] which allows the optimisation over quantum strategies for an XOR game to be expressed as a semidefinite program (SDP).


The classical hardness of XOR games (*i.e.* 2LIN) is ascertained via a reduction from 3LIN with imperfect completeness, due to Håstad [72]. The hardness of 3LIN with completeness $1 - \varepsilon$ for every $\varepsilon > 0$ was also shown in [72]. The reduction from 3LIN to 2LIN works by replacing each relation from a 3LIN instance with a gadget that can be satisfied with probability $\frac{3}{4}$ if and only if the original equation is satisfied. Recently, Tallar and Vidick [112] have shown that Håstad’s proof of the hardness of 3LIN can be made quantum-sound, hence giving hardness of approximation for entangled 3LIN with imperfect completeness. If the reduction from 3LIN to 2LIN were also quantum-sound, that would imply that XOR games are undecidable to approximate. As noted above, this is well known to not be the case, and hence there must be some issue in quantising this reduction.

We studied the question of why the quantisation of the reduction from 3LIN to 2LIN fails, and found that it hinges on one simple step. In the reduction of [72], there is a variable that is shared between all the gadgets, and Håstad makes the classically-innocuous assumption that this variable takes value 0 in any assignment. This is a valid assumption as linear equations in two boolean variables are invariant under negation. However, for quantum assignments, the variables are represented by operators; since the operators may not commute, this variable cannot be set to 0 while preserving the remaining operators as binary observables.


To remedy this problem, we introduce the notion of **tilted XOR games**, which is the generalisation of XOR games where one of the questions is fixed to always require output 0. An example of such a nonlocal game known in the literature is the tilted CHSH Bell inequality [2]. Under classical strategies, tilted XOR games are equivalent to XOR games. However, for quantum strategies, they differ wildly. We show that the reduction of [72] gives in fact a quantum-sound reduction from 3LIN to tilted XOR games. This provides a class of nonlocal games that is qualitatively very close to XOR games but captures all the hardness of general nonlocal games. In particular, this provides the first known example of **binary games** – nonlocal games where each of the players answers a single bit – that are undecidable. We study the ranges of completeness and soundness parameters where tilted




(a) Hardness of classical XOR games.



(b) Hardness of entangled XOR games.



(c) Hardness of classical tilted XOR games.



(d) Hardness of entangled tilted XOR games.

Figure 1.2: Hardness of XOR and tilted XOR games for different values of the completeness and soundness parameters. The green regions indicate problems solvable in polynomial time, the red and purple regions indicate hard instances, the white regions indicate problems where the complexity is unknown, and the gray regions indicates $s > c$ where the problem makes no sense. Note that the classical complexities of XOR and tilted XOR games are identical, and the RE-complete region is not to scale (though it is the interior of a triangle with positive area).

XOR games are easy or hard to approximate. This is illustrated and compared to XOR games in [Fig. 1.2](#).

1.1.5 Approximation algorithms for entangled CSPs

An alternate way to approach the question of hardness for computational problems is to ask how well they may be approximated. For classical CSPs, there is a well-developed theory of **approximation algorithms**. An approximation algorithm takes as input an instance of the CSP in question, and outputs in polynomial time an assignment to the variables. The quality of the approximation is given by the comparing the number of constraints satisfied by this assignment to the maximum number of constraints satisfiable by any assignment. We call the worst case ratio of these two quantities the **approximation ratio** of the algorithm.

More formally, we will only require the approximation algorithm to output the value of the approximation, though approximation algorithms usually work by finding an assignment and then computing its value — however, in the entangled setting, this distinction is important as the operator solution may have dimension too high to be presented in polynomial time, but its value may be much easier to compute.

The most significant approximation algorithm is the **Goemans-Williamson algorithm** for Max-Cut, which is the optimisation problem associated to graph 2-colouring [\[67\]](#). The name Max-Cut comes by noting that an assignment of colours to the vertices of a graph can be seen as a cut that splits the vertices into two subsets, where maximising the number of edges cut is the same as maximising the number of colouring constraints satisfied. This comparison is illustrated in [Fig. 1.3](#). The algorithm proceeds by first relaxing the optimisation problem to an SDP – which can be solved to sufficient precision in polynomial time – and then rounding the SDP solution to an actual assignment. Importantly, the rounding step is randomised. This gives an approximation ratio of ≈ 0.878 . It turns out that, assuming the unique games conjecture, the optimal approximation algorithm for any CSP has this form [\[104\]](#); and the Goemans-Williamson algorithm is in fact optimal for Max-Cut [\[80\]](#).

In the quantum world, since 2-colouring is a special case of 2LIN, it is easy to approximate by [\[114, 41\]](#). Hence, that means that entangled Max-Cut has approximation ratio 1. Nevertheless, we do not need to stray far from 2-colouring to find an interesting entangled CSP that is hard to approximate. In this work, we study approximation algorithms for Max- k -Cut with $k \geq 3$, which is the optimisation problem associated to graph k -colouring. We construct an approximation algorithm inspired by the Goemans-Williamson algorithm, but which provides an operator assignment rather than a classical assignment. First, we




Figure 1.3: An example of graph 2-colouring / Max-Cut seen as (non-perfect) colouring, homomorphism, and cut.

make use of an SDP relaxation. This is the same SDP used as a relaxation of the classical problem by Goemans and Williamson [66] and Newman [94]; this means that it is easy to compare the classical approximation ratio to the entangled approximation ratio. Next, we need to round the SDP solution, consisting of vectors, to an operator solution — which we express in terms of the observables, so it consists of order- k unitaries. First, we turn the vectors into operators using a **vector-to-unitary construction** based on the Weyl-Brauer operators due to Tsirelson [114]; then we apply a randomised rounding to transform these operators into order- k unitaries. The rounding works by multiplying the operators by a random unitary and then rounding the eigenvalues to the nearest k -th roots of unity. We show that this approximation algorithm has approximation ratio comparable to the best known classical algorithm, and in fact has a higher approximation ratio for $k = 3, 4$. Note that for $k \geq 4$, the best known classical algorithm, due to de Klerk, Pasechnik, and Warners [54], uses a different SDP relaxation than the one we use; perhaps there is an algorithm based on this SDP relaxation that improves the approximation ratio for all k . The approximation ratios of our algorithm and of [54] are compared in Fig. 1.4.

To analyse our algorithm, we make use of a new technique we call the **relative distribution**. Since we are interested in inner products of random unitaries, we keep track of the products of the eigenvalues as a distribution weighted by the inner products of the corresponding eigenspace projections. Then, the value of the approximation algorithm can be expressed as an integral with respect to this distribution. Using techniques from free probability theory, we show that the relative distribution simplifies significantly in the limit of large dimension. Using this simplification, we are able to compute the approximation ratio exactly for all k .

Since we are also interested in bounding the dimension where the operator assignments live, we modify our algorithm by making use of an algebraic approach that captures the




Figure 1.4: Approximation ratios for Max- k -Cut, comparing our algorithm (red) to the best known classical algorithm (blue). The values are comparable as every classical strategy is a quantum strategy.

same properties as the relative distribution. We generalise the vector-to-unitary construction to better approximate higher-order unitaries, by studying a **generalised Weyl-Brauer group**. Then, we use this to show that the same approximation ratio can be attained to precision ε in dimension $2^{O(n/\varepsilon)}$, where the instance size is n .

1.2 Outline

In [Chapter 2](#), we discuss the technical preliminaries needed for this work. In [Chapter 3](#), we define entangled CSPs and study their basic structure. In [Chapter 4](#), we study reductions based on commutativity gadgets, prove the polymorphism characterisation, and use this to show relationships and separations between different classes of commutativity gadget. In [Chapter 5](#), we study reductions between CSPs based on minion homomorphisms of the entangled polymorphisms. In [Chapter 6](#), we prove the reduction from linear system games to tilted XOR games, and use this to study hardness of approximation for binary nonlocal games. In [Chapter 7](#), we construct and analyse approximation algorithms for entangled CSPs. Finally, in [Chapter 8](#), we discuss the main results and finish with some open questions.

Chapter 2

Preliminaries

In this chapter, we provide the notation and technical preliminaries used for the results of this thesis. Sections 2.1, 2.4.1 to 2.4.3, 2.4.5, 2.5.1, 2.6, 2.9, 2.10 and 2.12 are used throughout. Section 2.2 is used in Section 3.2 and Chapter 5. Section 2.3 is used in Chapters 4 and 7. Section 2.4.4 is used in Section 3.3. Sections 2.5.2, 2.5.3 and 2.8 are used in Chapter 7. Section 2.7 is used in Chapter 3. Section 2.11 is used in Chapters 6 and 7.

2.1 Notation

We denote standard sets using set-constructor notation as $\{x_i \mid i \in I\}$; labelled sets as $\{x_i\}_{i \in I}$; and ordered sets or tuples as $(x_i)_{i \in I}$ (inheriting an order from I). For a natural number n , write the standard n -element set $[n] = \{1, 2, \dots, n\}$, and the order- n cyclic group $\mathbb{Z}_n = \{0, \dots, n-1\}$. To disambiguate, we often denote elements of a cartesian power A^n with bold face $\mathbf{a} = (a_1, \dots, a_n)$, and elements of a double cartesian power $(A^m)^n$ with blackboard bold face $\mathbf{a} = (\mathbf{a}_1, \dots, \mathbf{a}_n) = ((a_{11}, \dots, a_{1m}), \dots, (a_{n1}, \dots, a_{nm}))$. We identify tuples in A^n with functions in $A^{[n]}$, so that $\mathbf{a}(i) = a_i$.

We denote the set of $n \times n$ complex matrices by $\mathbb{M}_n = \mathcal{B}(\mathbb{C}^n)$. The group of invertible matrices in $\mathbb{M}_n$ is denoted $\mathcal{GL}_n$, and the subgroup of unitary matrices is denoted $\mathcal{U}_n$. The standard trace on $\mathbb{M}_n$ is Tr , and the normalised trace is denoted $\text{tr}_n = \frac{1}{n} \text{Tr}$, where the subscript is dropped if clear from context. For $T \in \mathbb{M}_n$, the operator norm is $\|T\|$, the **Frobenius norm** is $\|T\|_F = \sqrt{\text{Tr}(T^*T)}$ and the **little Frobenius norm** is $\|T\|_f = \sqrt{\text{tr}(T^*T)}$.

Write the imaginary unit $i = \sqrt{-1}$. We write u_X to mean the uniform distribution on a finite set X , and drop the subscript when the set is implied from context.

2.2 Categories

A **category** $\mathcal{C}$ consists of a class of **objects** $\text{ob}(\mathcal{C})$ and for every pair $a, b \in \text{ob}(\mathcal{C})$, a class of **morphisms** $\text{mor}_{\mathcal{C}}(a, b)$. The morphisms admit a composition operation $\circ$ such that for all $f \in \text{mor}_{\mathcal{C}}(a, b)$ and $g \in \text{mor}_{\mathcal{C}}(b, c)$, $g \circ f \in \text{mor}_{\mathcal{C}}(a, c)$. The composition is associative $h \circ (g \circ f) = (h \circ g) \circ f$ for all $f \in \text{mor}_{\mathcal{C}}(a, b)$, $g \in \text{mor}_{\mathcal{C}}(b, c)$, and $h \in \text{mor}_{\mathcal{C}}(c, d)$; and admits an identity morphism $\text{id}_a \in \text{mor}_{\mathcal{C}}(a, a)$ such that $f \circ \text{id}_a = \text{id}_b \circ f = f$ for all $f \in \text{mor}_{\mathcal{C}}(a, b)$. A category is **small** if $\text{ob}(\mathcal{C})$ and all $\text{mor}_{\mathcal{C}}(a, b)$ are sets. Where clear from context, we write $a \in \mathcal{C}$ to mean $a \in \text{ob}(\mathcal{C})$, $\text{mor}(a, b)$ to mean $\text{mor}_{\mathcal{C}}(a, b)$, and $f : a \rightarrow b$ to mean $f \in \text{mor}_{\mathcal{C}}(a, b)$.

An **endomorphism** is a morphism $f : a \rightarrow a$, and we denote the set of endomorphisms by $\text{end}(a) = \text{mor}(a, a)$. An **isomorphism** is a morphism $f : a \rightarrow b$ with an inverse morphism $g : b \rightarrow a$ such that $f \circ g = \text{id}_b$ and $g \circ f = \text{id}_a$, and we denote the set of isomorphisms by $\text{iso}(a, b)$. An **automorphism** is a morphism that is both an isomorphism and an endomorphism, and we write $\text{aut}(a) = \text{iso}(a, a)$.

A **covariant functor** F from a category $\mathcal{C}$ to a category $\mathcal{D}$ is a map, denoted $F : \mathcal{C} \rightarrow \mathcal{D}$, which to each object $a \in \mathcal{C}$ assigns an object $F(a) \in \mathcal{D}$, and to each morphism $f : a \rightarrow b$ assigns a morphism $F(f) : F(a) \rightarrow F(b)$, such that $F(\text{id}_a) = \text{id}_{F(a)}$ and $F(g \circ f) = F(g) \circ F(f)$. Two categories $\mathcal{C}$ and $\mathcal{D}$ are **isomorphic**, denoted $\mathcal{C} \cong \mathcal{D}$, if there exist functors $F : \mathcal{C} \rightarrow \mathcal{D}$ and $G : \mathcal{D} \rightarrow \mathcal{C}$ such that $F \circ G = \text{Id}_{\mathcal{D}}$ and $G \circ F = \text{Id}_{\mathcal{C}}$, where Id denotes the identity functor. A **contravariant functor** F is defined identically except that $F(g \circ f) = F(f) \circ F(g)$, *i.e.* the order of the arrows is reversed.

2.3 Groups

2.3.1 Presentations

The **free group** generated by a set S , denoted FS , is the set of words in $S \sqcup S^{-1}$ subject to the relations $ss^{-1} = s^{-1}s = \varepsilon$ for all $s \in S$, where ε is the empty word. Given a finite set S and a finite set R of words in $S \sqcup S^{-1}$, the **finitely-presented group** with generators S and relations R , denoted $\langle S : R \rangle$, is the quotient of the free group

FS by the normal subgroup $\langle\langle R \rangle\rangle$ generated by R . The **free product** of two finitely presented groups $G_1 = \langle S_1 : R_1 \rangle$ and $G_2 = \langle S_2 : R_2 \rangle$ is the finitely-presented group $G_1 * G_2 = \langle S_1 \sqcup S_2 : R_1 \sqcup R_2 \rangle$. If G_1 and G_2 contain H_1 and H_2 isomorphic to H , then the **amalgamated product** of G_1 and G_2 over H , denoted $G_1 *_H G_2$, is the quotient of $G_1 * G_2$ by the normal subgroup generated by the relations $\varphi_1(h)^{-1}\varphi_2(h)$ for all $h \in H$, where $\varphi_i : H \rightarrow H_i$ are the isomorphisms. If H is finitely presented, then $G_1 *_H G_2$ remains finitely presented. The natural homomorphisms $G_1 \rightarrow G_1 *_H G_2$ and $G_2 \rightarrow G_1 *_H G_2$ are injective [77].

2.3.2 Representations

A (finite-dimensional $\mathbb{C}$ -linear) **representation** of a group G is a group homomorphism $\pi : G \rightarrow \mathcal{GL}_n$ for some n . A representation is **unitary** if $\text{im } \pi \subseteq \mathcal{U}_n$. Two representations π and ρ are **equivalent** if there exists an invertible intertwiner operator T such that $T\pi(g) = \rho(g)T$ for all $g \in G$. A representation is **irreducible** if the only subspaces of $\mathbb{C}^n$ invariant under the action of the representation are the trivial subspaces 0 and $\mathbb{C}^n$. We call a group **linear** if it admits a faithful (*i.e.* injective) representation.

If G is a finite group, then G is linear and every representation is equivalent to a unitary representation. Also, every representation decomposes as a direct sum of irreducible representations. Further, G has a finite number of inequivalent irreducible representations, called **irreps**, and their number is equal to the number of conjugacy classes of G .

For $S \subseteq G$ and $\varepsilon > 0$, an (S, ε) -**approximate representation** is a function $\phi : S \rightarrow \mathcal{U}_n$ such that for all $g, h \in G$ such that $g, h, gh \in S$, $\|\phi(g)\phi(h) - \phi(gh)\|_f < \varepsilon$. We call a group **hyperlinear** if for all $g \in G$, there exists a constant $r_g > 0$, satisfying that for every $\varepsilon > 0$ and finite set $S \subseteq G \setminus \{1\}$ there exists an (S, ε) -approximate representation ϕ of G such that $\|\phi(g) - 1\|_f \geq r_g$ for all $g \in S$ [36].

2.4 Algebras

The main references for this section are [97, 21, 30].

2.4.1 *- and C^* -algebras

A unital ***-algebra** is an algebra $\mathcal{A}$ over $\mathbb{C}$ (or $\mathbb{R}$) equipped with an antilinear involution $x \mapsto x^*$ such that $1^* = 1$ and $(xy)^* = y^*x^*$. Denote the (real) subspace of hermitian elements

by $\mathcal{A}_h = \{x \in \mathcal{A} \mid x^* = x\}$. A ***-positive cone** over $\mathcal{A}$ is a set $\mathcal{A}_+ \subseteq \mathcal{A}_h$ such that $1 \in \mathcal{A}_+$, $x + y \in \mathcal{A}_+$ for all $x, y \in \mathcal{A}_+$, and $axa^* \in \mathcal{A}_+$ for all $x \in \mathcal{A}_+$ and $a \in \mathcal{A}$. A *-positive cone induces an order on $\mathcal{A}_h$ via $x \geq y$ iff $x - y \in \mathcal{A}_+$. We say $\mathcal{A}_+$ is **archimedean** if for all $x \in \mathcal{A}$, there exists $R \in \mathbb{R}$ such that $x^*x \leq R1$. Following [97], a *-algebra equipped with an archimedean *-positive cone is called a **semi-pre- C^* -algebra**. A standard choice of *-positive cone is the **sum-of-squares cone** given by $\mathcal{A}_+ = \{\sum_{i=1}^n x_i^* x_i \mid n \in \mathbb{N}, x_i \in \mathcal{A}\}$ — we will always make this choice. As a shorthand, we often write $|a|_2^2 = a^*a$.

Given a set of generators S and relations R , we denote the ***-algebra generated by S subject to R** as $\mathbb{C}\langle S : R \rangle = F(S)/\langle\langle R \rangle\rangle$, where $F(S)$ is the **free algebra** with generators S and $\langle\langle R \rangle\rangle$ is the two-sided ideal of $F(S)$ generated by R . Note that a finitely-generated *-algebra is archimedean if and only if all its generators are bounded.

A semi-pre- C^* -algebra is a **pre- C^* -algebra** if $x^*x \leq \varepsilon 1$ for all $\varepsilon > 0$ implies that $x = 0$ in $\mathcal{A}$. If $\mathcal{A}$ is a pre- C^* -algebra, we can define a norm on it by

$$\|x\| = \sup \{ \|\pi(x)\| \mid \pi : \mathcal{A} \rightarrow \mathcal{B}(H) \text{ is a } *- \text{homomorphism} \} = \inf \{ \lambda \in \mathbb{R}_{\geq 0} \mid x^*x \leq \lambda^2 1 \}.$$

This norm satisfies the properties $\|xy\| \leq \|x\|\|y\|$ and $\|x^*x\| = \|xx^*\| = \|x\|^2$. Note that this norm can be defined on a semi-pre- C^* -algebra as well by restricting to the positive *-homomorphisms; however, it is only a seminorm in this context. See [97] for details. A **C^* -algebra** is a pre- C^* -algebra that is complete with respect to this norm. In a C^* -algebra $\mathcal{A}$, there is a unique *-positive cone given by $\mathcal{A}_+ = \{x^*x \mid x \in \mathcal{A}\}$.

Given a semi-pre- C^* -algebra $\mathcal{A}$, we can map it to a pre- C^* -algebra by taking the quotient by the infinitesimal ideal $I(\mathcal{A}) = \{x \in \mathcal{A} \mid x^*x \leq \varepsilon 1 \ \forall \varepsilon > 0\}$. We call the completion of $\mathcal{A}/I(\mathcal{A})$ with respect to the norm above the **universal C^* -algebra** of $\mathcal{A}$, and denote it $C_u^*(\mathcal{A})$. Given a set of generators S and a set of relations R , the **C^* -algebra generated by S subject to R** is $C^*\langle S : R \rangle = C_u^*(\mathbb{C}\langle S : R \rangle)$.

Given a (finitely-presented) group G , the **group *-algebra** of G is the *-algebra $\mathbb{C}[G]$ generated by e_g for $g \in G$, subject to the relations $e_g^* = e_{g^{-1}}$ and $e_g e_h = e_{gh}$ for all $g, h \in G$. The *-homomorphisms of $\mathbb{C}[G]$ into $\mathcal{B}(H)$ for a Hilbert space H are in bijective correspondence with the unitary representations of G in $\mathcal{B}(H)$. We can see $\mathbb{C}[G]$ as a semi-pre- C^* -algebra with respect to the sum-of-squares cone. Then, the **group C^* -algebra** of G is the universal C^* -algebra $C^*G = C_u^*(\mathbb{C}[G])$.

The **tensor product** of *-algebras $\mathcal{A}$ and $\mathcal{B}$ is denoted $\mathcal{A} \otimes \mathcal{B}$, or sometimes $\mathcal{A} \odot \mathcal{B}$ in order to emphasize that we are not taking a completion. We implicitly use the isomorphism $\mathbb{C} \otimes \mathcal{A} \cong \mathcal{A} \otimes \mathbb{C} \cong \mathcal{A}$. For C^* -algebras $\mathcal{A}$ and $\mathcal{B}$, there are a variety of (generally inequivalent) ways to complete $\mathcal{A} \odot \mathcal{B}$ to a C^* -algebra. Notably, we make use of the **minimal tensor**

product (or min-tensor product), which is the completion $\mathcal{A} \otimes_{\min} \mathcal{B}$ of $\mathcal{A} \odot \mathcal{B}$ with respect to the norm $\|x\|_{\min} = \|(\pi_A \otimes \pi_B)(x)\|$ for π_A and π_B faithful $*$ -representations of $\mathcal{A}$ and $\mathcal{B}$, respectively.

2.4.2 States and representations

A **state** on a semi-pre- C^* -algebra $\mathcal{A}$ is a linear map $\rho : \mathcal{A} \rightarrow \mathbb{C}$ such that $\rho(1) = 1$ and $\rho(x) \geq 0$ for all $x \in \mathcal{A}_+$. We say ρ is **tracial** if $\rho(xy) = \rho(yx)$ for all $x, y \in \mathcal{A}$; and we say ρ is **faithful** if $\rho(x^*x) = 0$ implies $x = 0$ for all $x \in \mathcal{A}$. The existence of a faithful state automatically guarantees that $\mathcal{A}$ is a pre- C^* -algebra, since the infinitesimal ideal is trivial. Tracial states are usually denoted by τ . Any state induces a seminorm $\|a\|_\rho = \sqrt{\rho(a^*a)}$ called the **ρ -norm**. ρ is faithful if and only if $\|\cdot\|_\rho$ is a norm.

A **representation** of a $*$ -algebra $\mathcal{A}$ is a $*$ -homomorphism $\pi : \mathcal{A} \rightarrow \mathcal{B}(H)$ for a Hilbert space H . A vector $v \in H$ is called **cyclic** if $(\text{im } \pi)v$ is dense in H ; a vector $v \in H$ is **separating** if $\pi(x)v = 0$ if and only if $x = 0$. We say π is **irreducible** if H has no nontrivial subspaces invariant under the action of π . A **character** is a state that is a $*$ -representation.

Every state ρ on a $*$ -algebra induces a **Gelfand-Naimark-Segal (GNS) representation**, which is a representation $\pi : \mathcal{A} \rightarrow \mathcal{B}(H)$ such that there is a unit vector $\psi \in H$ satisfying $\rho(x) = \langle \psi | \pi(x)\psi \rangle$. There exists a **minimal** GNS representation in the sense that ψ may be chosen to be cyclic; all such minimal representations are unitarily equivalent. If ρ is a faithful tracial state, then ψ is cyclic and separating in the minimal GNS representation. A state is called **finite-dimensional** if it admits a finite-dimensional GNS representation; a tracial state is called **Connes-embeddable** if it admits a GNS representation where the von Neumann algebra generated by $\pi(\mathcal{A})$ is Connes-embeddable (see [Section 2.4.3](#)). In this context, it is often useful to restrict the image of π to the von Neumann algebra generated by $\pi(\mathcal{A})$, and restrict the state to that von Neumann algebra.

2.4.3 von Neumann algebras

A **von Neumann algebra** is a unital $*$ -subalgebra $\mathcal{M} \subseteq \mathcal{B}(H)$ for a Hilbert space H that is closed under the **weak operator topology (WOT)**, defined via convergence as $(x_\lambda)_\lambda \rightarrow x$ iff $\langle \psi | x_\lambda \phi \rangle \rightarrow \langle \psi | x \phi \rangle$ for all $\psi, \phi \in H$. Equivalently, $\mathcal{M}$ is a unital $*$ -subalgebra of $\mathcal{B}(H)$ equal to its bicommutant $\mathcal{M}''$, where the **commutant** of $S \subseteq \mathcal{B}(H)$ is $S' = \{a \in \mathcal{B}(H) \mid ax = xa \ \forall x \in S\}$; or that $\mathcal{M}$ is closed under the **strong operator topology (SOT)**, defined via convergence as $(x_\lambda)_\lambda \rightarrow x$ iff $\|(x_\lambda - x)\psi\| \rightarrow 0$ for all $\psi \in H$. In particular $\mathcal{M}$ is a C^* -algebra.

A **weight** on a von Neumann algebra is a map $\tau : \mathcal{M}_+ \rightarrow [0, \infty]$ satisfying $\tau(x + y) = \tau(x) + \tau(y)$ and $\tau(\lambda x) = \lambda\tau(x)$ for all $x, y \in \mathcal{M}_+$ and $\lambda \geq 0$ (where we take $0 \cdot \infty = 0$). A weight is **tracial** if $\tau(x^*x) = \tau(xx^*)$ for all $x \in \mathcal{M}$, in which case it is called a **trace**. A weight is **normal** if for every SOT-convergent increasing net $(x_\lambda)_\lambda$ in $\mathcal{M}_+$, $\lim_\lambda \tau(x_\lambda) = \tau(\lim_\lambda x_\lambda)$. A weight is **faithful** if $\tau(x) = 0$ iff $x = 0$. A weight is **finite** if $\tau(1) < \infty$; in this case $\infty \notin \text{im } \tau$ and τ can be extended to a positive linear functional on all $\mathcal{M}$. A state is a finite weight satisfying $\tau(1) = 1$; for a state on a von Neumann algebra, being normal is equivalent to having a GNS representation that is continuous in the strong operator topology. A weight is **semifinite** if for every $x \in \mathcal{M}_+$, there exists an increasing net $(x_\lambda)_\lambda$ in $\mathcal{M}_+$ such that $\lim_\lambda x_\lambda = x$ in the strong operator topology and $\tau(x_\lambda) < \infty$ for all λ .

A von Neumann algebra is a **factor** if $\mathcal{M} \cap \mathcal{M}' = \mathbb{C} \cdot 1$; any von Neumann algebra can be decomposed as a direct integral of factors. The finite-dimensional factors are exactly the matrix algebras M_n . A von Neumann algebra is **finite** if there exists a faithful tracial state on $\mathcal{M}$. An important example of a finite von Neumann algebra is the **hyperfinite type II₁ factor** $\mathcal{R}$. One way of constructing $\mathcal{R}$ is by taking a closure of the algebra $\mathcal{A} = \bigoplus_{n=1}^{\infty} M_{2^n}/I$, where I is the two-sided ideal generated by $x \otimes 1 - x$ for all x . Within $\mathcal{A}$, this provides an embedding of M_{2^n} into $M_{2^{n+1}}$ for all n . To find the right notion of closure, note first that there is a tracial state τ on $\mathcal{A}$ given by the normalised trace on each component. We can complete $\mathcal{A}$ as a Hilbert space H with respect to the inner product $\langle x, y \rangle = \tau(x^*y)$. Then, $\mathcal{A} \subseteq B(H)$ via the action of left multiplication, and we can take the closure in the strong operator topology to get $\mathcal{R}$. A finite von Neumann algebra is **Connes-embeddable** if there is an injective $*$ -homomorphism $\mathcal{M} \rightarrow \mathcal{R}^\omega$, where $\mathcal{R}^\omega$ is an ultrapower of the hyperfinite II₁ factor.

To define $\mathcal{R}^\omega$, we need the concept of an ultrafilter. An **ultrafilter** ω on the natural numbers is a collection of subsets of $\mathbb{N}$ such that $\emptyset \notin \omega$; if $A, B \in \omega$, then $A \cap B \in \omega$; and for all $A \subseteq \mathbb{N}$, either $A \in \omega$ or $A^c \in \omega$. Note that these properties also imply that if $A \in \omega$ and $B \supseteq A$, then $B \in \omega$ as well. An ultrafilter is called **principal** if it contains a singleton, and **free** otherwise. The existence of a free ultrafilter follows from the axiom of choice. Ultrafilters allow us to define a robust notion of limit. Given an ultrafilter ω , we say x is the **ultralimit** of a sequence $(x_n)_n$ if for each $\varepsilon > 0$, $\{n \in \mathbb{N} \mid |x - x_n| < \varepsilon\} \in \omega$; we denote this as $\lim_{n \rightarrow \omega} x_n = x$ or $x_n \xrightarrow{\omega} x$. It holds that if the ultralimit exists, it is unique; and every bounded sequence admits an ultralimit. Further, if ω is free and $(x_n)_n$ converges in the usual way, then $\lim_{n \rightarrow \omega} x_n = \lim_{n \rightarrow \infty} x_n$. Given a sequence of finite von Neumann algebras $(\mathcal{M}_n, \tau_n)$, their **ultraproduct** is the space of bounded sequences (a_n) such that $a_n \in \mathcal{M}_n$, modulo the equivalence relation $(a_n) \sim (b_n)$ if $\|a_n - b_n\|_{\tau_n} \xrightarrow{\omega} 0$. This space, denoted $\prod_{n \rightarrow \omega} \mathcal{M}_n$ or $\prod_\omega \mathcal{M}_n$ if the index is clear from context, is a finite von Neumann algebra with tracial state $\tau_\omega(a) = \lim_{n \rightarrow \omega} \tau_n(a_n)$ for a representative (a_n) of a . If $\mathcal{M}_n = \mathcal{M}$

for each n , then we write $\prod_{\omega} \mathcal{M}_n = \mathcal{M}^{\omega}$, and call it the **ultrapower** of $\mathcal{M}$. Note that ultraproducts of separable von Neumann algebras are not in general separable.

It is known that $\mathcal{R}^{\omega} \cong \mathcal{R}^{\omega'}$ for all free ultrafilters if and only if the continuum hypothesis holds (see [68]). However, we can usually make use of weaker relations between different ultraproducts: below, we provide a proof of the (folklore) result that quantifies the sense in which we can treat $\mathcal{R}^{\omega}$ as ultraproduct of matrix algebras of increasing dimension.

Lemma 2.4.1. Let ω be an ultrafilter on $\mathbb{N}$.

- (i) For any sequence of integers $(k(n))_n$, there is an embedding $\prod_{\omega} \mathbb{M}_{k(n)} \hookrightarrow \mathcal{R}^{\omega}$.
- (ii) There exists an ultrafilter ω' on $\mathbb{N}$ and a sequence $(k(n))_n$ such that there is an embedding $\mathcal{R}^{\omega} \hookrightarrow \prod_{\omega'} \mathbb{M}_{k(n)}$.

Proof. (i) By a standard result, there is a trace-preserving embedding $\iota_n : \mathbb{M}_n \rightarrow \mathcal{R}$ for all n [91]. Define the map $\iota : \prod_{\omega} \mathbb{M}_{k(n)} \rightarrow \mathcal{R}^{\omega}$ by taking $\iota(a)$ to be the element with representative $(\iota_{k(n)}(a_n))$, for (a_n) a representative of a . First, ι is well-defined: if $(a_n) \sim (b_n)$, then, $\|a_n - b_n\|_{\text{tr}_{k(n)}} \xrightarrow{\omega} 0$; as ι_n is trace-preserving, then $\|\iota_{k(n)}(a_n) - \iota_{k(n)}(b_n)\|_{\tau} \xrightarrow{\omega} 0$. Therefore, ι is a $*$ -homomorphism. Next, suppose the image of $\iota(a) = 0$ in $\mathcal{R}^{\omega}$. Then, letting (a_n) be a representative of a , $0 = \lim_{n \rightarrow \omega} \tau(\iota_{k(n)}(a_n)^* \iota_{k(n)}(a_n)) = \lim_{n \rightarrow \omega} \text{tr}_{k(n)}(a_n^* a_n)$, so $a = 0$ in $\prod_{\omega} \mathbb{M}_{k(n)}$. So ι is injective.

- (ii) $\mathcal{R}$ is constructed as the strong operator topology closure of $\bigoplus_{n=1}^{\infty} \mathbb{M}_{2^n}$, modulo the identification $x \otimes 1 = x$. In particular, for all $x \in \mathcal{R}$, there exists a sequence (x_m) such that $x_m \in \mathbb{M}_{2^m}$ and $x_m \rightarrow x$ with respect to the trace norm; via the axiom of choice, we can choose such a sequence for every element of $\mathcal{R}$. Fix any free ultrafilter v on $\mathbb{N}$. Using [36, Definition III.2.3], consider $\mathcal{M} = \prod_{(n,m) \rightarrow \omega \times v} \mathbb{M}_{2^m}$. Since $\mathbb{N} \times \mathbb{N} \cong \mathbb{N}$, $\omega \times v$ can be seen as a free ultrafilter on $\mathbb{N}$, and hence we intend to construct an embedding of $\mathcal{R}^{\omega}$ in $\mathcal{M}$. Define the map $\phi : \mathcal{R}^{\omega} \rightarrow \mathcal{M}$ as $\phi(a) = [(a_{n,m})_{n,m}]$ where (a_n) is a representative of a and $(a_{n,m})_m$ is the sequence converging to a_n chosen above. First, ϕ is well-defined, as for any $(a_n) \sim (b_n)$,

$$\lim_{(n,m) \rightarrow \omega \times v} \|a_{n,m} - b_{n,m}\|_{\text{tr}_{2^m}} = \lim_{n \rightarrow \omega} \lim_{m \rightarrow v} \|a_{n,m} - b_{n,m}\|_{\text{tr}_{2^m}} = \lim_{n \rightarrow \omega} \|a_n - b_n\|_{\tau} = 0.$$

Also, ϕ is a $*$ -homomorphism. For example, if $a, b \in \mathcal{R}^{\omega}$ then $\phi(a + b) = \phi(a) + \phi(b)$ as

$$\lim_{(n,m) \rightarrow \omega \times v} \|(a_n + b_n)_m - a_{n,m} - b_{n,m}\|_{\text{tr}_{2^m}} = \lim_{n \rightarrow \omega} \|a_n + b_n - a_n - b_n\|_{\tau} = 0.$$

ϕ preserves the identity, products, and the adjoint in the same way. To finish, we show ϕ is injective. Suppose $\phi(a) = 0$. Then, $0 = \lim_{(n,m) \rightarrow \omega \times v} \|a_{nm}\|_{\text{tr}_{2m}} = \lim_{n \rightarrow \omega} \|a_n\|_{\tau} = 0$, so $a = 0$. $\blacksquare$

2.4.4 Noncommutative L^p spaces

Let $\mathcal{M} \subseteq \mathcal{B}(H)$ be a von Neumann algebra. From Tomita-Takesaki theory [111], there exists a Hilbert space $c(H)$ and a unique (up to isomorphism) von Neumann algebra $c(\mathcal{M}) \subseteq \mathcal{B}(c(H))$, called the **core** of $\mathcal{M}$, equipped with a normal faithful semifinite trace τ , a continuous group homomorphism $\theta : \mathbb{R} \rightarrow \text{Aut}(c(\mathcal{M}))$, and an embedding $\iota : \mathcal{M} \hookrightarrow c(\mathcal{M})$, satisfying $\text{im } \iota = \{x \in c(\mathcal{M}) \mid \theta_s(x) = x \ \forall s \in \mathbb{R}\}$ and $\tau \circ \theta_s = e^{-s}\tau$.

An **unbounded operator** on a Hilbert space H is a linear map $x : D(x) \rightarrow H$, where $D(x)$ is a subspace of H , called the **domain** of x . x is **densely-defined** if $D(x)$ is dense in H ; and x is **closed** if the graph $\{(\psi, x\psi) \mid \psi \in D(x)\}$ is closed in $H \oplus H$. Let $\mathcal{N} \subseteq \mathcal{B}(H)$ be a von Neumann algebra with a normal faithful semifinite trace τ . An unbounded operator y on H is **affiliated** with $\mathcal{N}$ if $yx \subseteq xy$ for all $x \in \mathcal{N}'$. We call a closed densely-defined y affiliated with $\mathcal{N}$ **τ -measurable** if for all $\delta > 0$, there exists a projection $p \in \mathcal{N}$ such that $\text{im } p \subseteq D(y)$ and $\tau(1 - p) < \delta$. The set of affiliated operators is a $*$ -algebra with respect to the strong sum and product, *i.e.* standard sums and products of operators are closable and we take the strong versions of the operations to be their respective closures. It admits a Hausdorff topology with respect to which $\mathcal{N}$ is dense.

For $p \geq 1$, define $L^p(\mathcal{M})$ to be the set of τ -measurable operators affiliated to $c(\mathcal{M})$ that are $e^{-s/p}$ -eigenvectors of θ_s for all s . The map $x \mapsto \|x\|_p = \tau(\chi_{(1,\infty)}(|x|))^{1/p}$ is a norm with respect to which $L^p(\mathcal{M})$ is complete. We have natural identifications $L^\infty(\mathcal{M}) \cong \mathcal{M}$, $L^1(\mathcal{M}) \cong \mathcal{M}_*$, the predual. For $x \in L^1(\mathcal{M})$, define $\text{tr}(x) = \varphi(1)$, where $\varphi \in \mathcal{M}_*$ is the functional corresponding to x . tr is a linear functional on $L^1(\mathcal{M})$ satisfying $\text{tr}(xy) = \text{tr}(yx)$ for $x \in L^p(\mathcal{M})$ and $y \in L^q(\mathcal{M})$ with $\frac{1}{p} + \frac{1}{q} = 1$. We can write the p -norm as $\|x\|_p = \text{tr}(|x|^p)^{1/p}$; it follows that for positive $x \in L^1(\mathcal{M})$, we have that $\text{tr}(x) = \tau(\chi_{(1,\infty)}(x))$. We can define the inner product $\langle x, y \rangle = \text{tr}(x^*y)$ on $L^2(\mathcal{M})$, which makes it into a Hilbert space. See [71, 113] for a full exposition.

Example 2.4.2. To illustrate the construction, consider the case $\mathcal{M} = \mathbb{M}_n$. In this case, we can simply take $c(H) = L^2(\mathbb{R}) \otimes \mathbb{C}^n$ and $c(\mathcal{M}) = L^\infty(\mathbb{R}) \otimes \mathbb{M}_n$, the space of bounded measurable matrix-valued functions on the real line. In this case, $\theta_s(x)(t) = x(t - s)$ and the trace $\tau(x) = \int_{-\infty}^{\infty} \text{Tr}(x(t))e^{-t}dt$, where Tr is the usual trace on $\mathbb{M}_n$. The unbounded operators affiliated with $c(\mathcal{M})$ are the measurable functions $\mathbb{R} \rightarrow \mathbb{M}_n$; and all of these are τ -measurable. $L^p(\mathcal{M})$ consists of those τ -measurable x such that $x(t) = x(0)e^{t/p}$; hence in

this case $L^p(\mathcal{M})$ is isomorphic to $\mathbb{M}_n$ for all p . The isomorphism $L^1(\mathcal{M}) \rightarrow \mathcal{M}_*$ is given by $x \mapsto \text{Tr}(\cdot x(0))$, and hence $\text{tr}(x) = \text{Tr}(x(0))$. The p -norms in this case reduce to the Schatten norms. The same holds for any finite von Neumann algebra.

2.4.5 Weighted algebras

A **weighted algebra** [87] is a pair $(\mathcal{A}, \mu)$ where $\mathcal{A}$ is a $*$ -algebra and $\mu : \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}$ is a finitely-supported function, called the **weight function**. The **defect polynomial** of a weighted algebra is

$$D(\mathcal{A}, \mu) = \sum_{a \in \mathcal{A}} \mu(a) a^* a.$$

Given a tracial state $\tau : \mathcal{A} \rightarrow \mathbb{C}$, the **defect** of τ is

$$\text{def}(\tau; \mu) = \tau(D(\mathcal{A}, \mu)) = \sum_{a \in \mathcal{A}} \mu(a) \|a\|_{\tau}^2.$$

We denote it $\text{def}(\tau)$ if the weight function is clear from context. Weighted algebras capture approximate representations of $\mathcal{A}/\langle\langle \text{supp}(\mu) \rangle\rangle$ and $\text{def}(\tau)$ captures how much the GNS representation of the tracial state τ deviates from an exact representation of this algebra.

Since we consider only the representations of weighted algebras induced by tracial states, the natural notion of order is weaker than that induced by the sums of squares. We say $a, b \in \mathcal{A}$ are **cyclically equivalent** if there exist $c_1, d_1, \dots, c_k, d_k \in \mathcal{A}$ such that $a - b = \sum_i [c_i, d_i]$; we denote this by $a \approx b$. Write $a \gtrsim b$ if there exist $s_1, \dots, s_k \in \mathcal{A}$ such that $a - b \approx \sum_i s_i^* s_i$.

A **C -homomorphism** $\alpha : (\mathcal{A}, \mu) \rightarrow (\mathcal{B}, \nu)$ is a $*$ -homomorphism $\alpha : \mathcal{A} \rightarrow \mathcal{B}$ such that $\alpha(D(\mathcal{A}, \mu)) \lesssim C \cdot D(\mathcal{B}, \nu)$. If there is a C -homomorphism $\alpha : (\mathcal{A}, \mu) \rightarrow (\mathcal{B}, \nu)$, then for every tracial state τ on $\mathcal{B}$, there exists a tracial state $\tau' = \tau \circ \alpha$ on $\mathcal{A}$ such that $\text{def}(\tau') \leq C \text{def}(\tau)$.

The following inequality is often useful in the context of weighted algebras.

Lemma 2.4.3. Let $\mathcal{A}$ be a $*$ -algebra equipped with the sum-of-squares cone. For any $a_1, \dots, a_k \in \mathcal{A}$,

$$\left| \sum_{i=1}^k a_i \right|^2 \leq k \sum_{i=1}^k |a_i|^2.$$

Further, by taking $a_i = 1$ for all i , we can easily see that the bound is tight.

Proof. Let $\omega = e^{2\pi i/k}$ be a primitive k -th root of unity. Then, we upper bound

$$\left| \sum_{i=1}^k a_i \right|^2 \leq \sum_{j=0}^{k-1} \left| \sum_{i=1}^k \omega^{(i-1)j} a_i \right|^2 = \sum_{j=0}^{k-1} \sum_{i,i'=1}^k \omega^{(i-i')j} a_{i'}^* a_i = k \sum_{i=1}^k |a_i|^2,$$

where the inequality is due to the fact $\left| \sum_{i=1}^k a_i \right|^2$ is the term $\left| \sum_{i=1}^k \omega^{(i-1)j} a_i \right|^2$ with $j = 0$ and all the terms are positive; the first equality is by expanding the square; and the second equality is because $\sum_{j=0}^{k-1} \omega^{ij} = 0$ unless $i = 0$. ■

2.5 Probability theory

The main reference for this section is [13].

2.5.1 Distributions and random variables

For a **measurable space** $(\Omega, \mathcal{S})$, we call the set Ω the sample space and the σ -algebra $\mathcal{S}$ on Ω the event space. If Ω is a topological space, there is a canonical choice of σ -algebra to give it a measurable structure, the **Borel algebra** $\mathcal{B}(\Omega)$; in this context, we omit reference to the σ -algebra. A **probability space** is $(\Omega, \mathcal{S}, \mu)$ where μ is a measure on $(\Omega, \mathcal{S})$ such that $\mu(\Omega) = 1$. We call μ a **(probability) distribution** on Ω . If Ω is a finite set, then we always take $\mathcal{S}$ to be the entire power set, and we identify μ with a function $\mu : \Omega \rightarrow [0, 1]$ such that $\sum_{\omega \in \Omega} \mu(\omega) = 1$ via $\mu(\omega) = \mu(\{\omega\})$.

Given a measurable space $(M, \mathcal{F})$, an **M -valued random variable** on $(\Omega, \mathcal{S}, \mu)$ is a measurable function $X : \Omega \rightarrow M$, *i.e.* a function for which $X^{-1}(E) \in \mathcal{S}$ for all $E \in \mathcal{F}$. This induces a distribution on M where $\mu_X(E) = \mu(X^{-1}(E))$ for all $E \in \mathcal{F}$; this is called the **distribution** of X . We also use probability notation for this distribution, writing $\Pr[X \in E] = \mu_X(E)$. For any function $f : M \rightarrow \mathbb{C}$, the composition $f \circ X$ is also a random variable, which we denote $f(X)$. If M is a real or complex vector space, the **expectation** of X is given by the integral $\mathbb{E}X = \int_{\Omega} X(\omega) d\mu(\omega) = \int_M x d\mu_X(x)$ when defined. Since a probability measure is finite, the expectation of any bounded random variable is defined. Two random variables X, Y are **independent** if for all measurable E, F we have $\Pr[X \in E \wedge Y \in F] = \Pr[X \in E] \Pr[Y \in F]$.

If M is in fact a measure space with its own measure ν — for example if M is a finite-dimensional real vector space with the Lebesgue measure Λ — then we can compare this to the measure μ_X . If μ_X is absolutely continuous with respect to ν (for all $E \in \mathcal{F}$, $\nu(E) = 0$

implies $\mu_X(E) = 0$), then we can define the Radon-Nikodym derivative of μ_X with respect to ν , which we refer to as the **probability density function (PDF)** $p_X = \frac{d\mu_X}{d\nu} : M \rightarrow \mathbb{R}_{\geq 0}$, or more explicitly $\mu_X(E) = \int_E p_X(x) d\nu(x)$. More generally, we can also consider the PDF of any distribution μ on M as the Radon-Nikodym derivative $p_\mu = \frac{d\mu}{d\nu}$, assuming μ is absolutely continuous with respect to ν .

A sequence of random variables (X_n) **converges in distribution** to a random variable X if for all measurable $E \subseteq M$,

$$\lim_{n \rightarrow \infty} \Pr[X_n \in E] = \Pr[X \in E].$$

This is equivalent to $\mathbb{E} f(X_n) \rightarrow \mathbb{E} f(X)$ for all bounded continuous functions $f : M \rightarrow \mathbb{R}$.

It is not necessary to make reference to an underlying sample space to define and work with a random variable, so we generally avoid it. Instead, we define random variables according to a distribution. For a distribution μ on a measurable space M , we say a random variable $X : \Omega \rightarrow M$ is distributed according to μ if $\mu_X = \mu$. We denote this $X \sim \mu$.

2.5.2 The Haar distribution

A **topological group** is a group G equipped with a topology with respect to which the inverse and product operations are continuous. Linear groups inherit the topology of the matrix algebra in which they embed; an important example is $\mathcal{U}_n$, which is also compact as a topological space. For any compact topological group G , there exists a unique probability distribution η_G on $(G, \mathcal{B}(G))$, called the **Haar measure** (or Haar distribution), such that for all $E \subseteq G$ Borel-measurable and $g \in G$,

$$\eta_G(E) = \eta_G(gE).$$

It follows from the uniqueness and compactness that $\eta_G(Eg) = \eta_G(E)$ and $\eta_G(E^{-1}) = \eta_G(E)$ as well. It is also possible to define a Haar measure which is no longer a probability distribution on groups which are only locally compact, but we will not need that generalisation here.

2.5.3 Random variables on the circle

We will make use of random variables X on the circle $[0, 2\pi) \cong \mathbb{R}/2\pi\mathbb{Z} \cong S^1$. The **characteristic function** of X is the function

$$\chi_X : \mathbb{Z} \rightarrow \mathbb{C}, \quad \chi_X(n) = \mathbb{E}(e^{inX}).$$

If μ_X is absolutely continuous with respect to the Lebesgue measure, then we have $\chi_X = \hat{p}_X$, the Fourier series of the PDF. The **Fourier series** of an integrable function $f : [0, 2\pi) \rightarrow \mathbb{C}$ is the function $\hat{f} : \mathbb{Z} \rightarrow \mathbb{C}$ defined as

$$\hat{f}(n) = \int_0^{2\pi} f(\theta) e^{in\theta} d\theta.$$

The characteristic function is defined analogously for distributions. Due to Parseval's theorem, if X is a random variable on $[0, 2\pi)$ with square-integrable PDF and $f : [0, 2\pi) \rightarrow \mathbb{C}$ is a square-integrable function, we have

$$\int_0^{2\pi} f(\theta) d\mu_X(\theta) = \frac{1}{2\pi} \sum_{n \in \mathbb{Z}} \hat{f}(-n) \chi_X(n). \quad (2.1)$$

Convergence properties also translate to the characteristic function.

Theorem 2.5.1 (Lévy's continuity theorem on the circle). A sequence of random variables (X_n) on $[0, 2\pi)$ converges in distribution if and only if the sequence of characteristic functions (χ_{X_n}) converges pointwise.

This theorem is a consequence of the Stone-Weierstrass theorem.

To finish this section, we mention two examples of distributions on the circle which will be important. For any $\theta_0 \in [0, 2\pi)$ the **Dirac delta distribution** is the distribution

$$\delta_{\theta_0}(E) = \begin{cases} 1 & \theta_0 \in E, \\ 0 & \text{otherwise.} \end{cases} \quad (2.2)$$

The characteristic function of δ_{θ_0} is $\chi_{\delta_{\theta_0}}(n) = e^{in\theta_0}$.

The **Wrapped Cauchy Distribution** with peak position $\theta_0 \in [0, 2\pi)$ and scale factor $\gamma > 0$ is the distribution $\mathcal{W}(\theta_0, \gamma)$ with PDF

$$p_{\mathcal{W}(\theta_0, \gamma)}(\theta) = \frac{\sinh(\gamma)}{2\pi(\cosh(\gamma) - \cos(\theta - \theta_0))}, \quad (2.3)$$

with respect to the Lebesgue measure on $[0, 2\pi)$. The characteristic function of $\mathcal{W}(\theta_0, \gamma)$ is

$$\chi_{\mathcal{W}(\theta_0, \gamma)}(n) = e^{-|n|\gamma + in\theta_0}.$$

For $\lambda \in \mathbb{C}$ with $|\lambda| \leq 1$ we define the distribution Δ_λ as $\Delta_\lambda = \mathcal{W}(\arg \lambda, -\ln |\lambda|)$ when $0 < |\lambda| < 1$; as $\Delta_\lambda = \delta_{\arg \lambda}$ when $|\lambda| = 1$; and as the uniform distribution $p_{\Delta_\lambda}(\theta) = \frac{1}{2\pi}$ when $\lambda = 0$. Using this definition, when λ tends to λ_0 , Δ_λ converges to Δ_{λ_0} in distribution. This has characteristic function $\chi_{\Delta_\lambda}(n) = \lambda^n$ for $n \geq 0$ and $\chi_{\Delta_\lambda}(n) = (\lambda^*)^{-n}$ for $n < 0$.

2.6 Quantum theory

The main references for this section are [96, 26].

2.6.1 The postulates of quantum mechanics

Quantum physics is distinguished from classical physics by two fundamental properties: physical quantities can take discrete values, rather than a continuum; and both matter and radiation can behave as both waves and particles. The wave-particle duality is well-represented by the de Broglie relations $E = \hbar\omega$ and $\mathbf{p} = \hbar\mathbf{k}$, relating the particlelike properties of energy and momentum to the wavelike properties of frequency and wavenumber, via the reduced Planck constant [53].

The de Broglie relations can be used to derive the Schrödinger equation [107], which we will use to motivate the postulates of quantum mechanics, in a simple way. A plane wave, with well-defined wavenumber and frequency, is expressed as $\psi(\mathbf{r}, t) = \exp(i(\mathbf{k} \cdot \mathbf{r} - \omega t))$. Due to the de Broglie relations, this should correspond to a particle with well-defined momentum and energy, which for a nonrelativistic particle of mass m in a potential V are related by $\frac{p^2}{2m} + V = E$. To extract the momentum and energy from the wavefunction ψ , one can act by the differential operators $\hat{\mathbf{p}} = \hbar\hat{\mathbf{k}} = -i\hbar\nabla$ and $\hat{E} = \hbar\hat{\omega} = i\hbar\frac{\partial}{\partial t}$. Putting these into the classical relationship between momentum and energy gives the Schrödinger equation

$$-\frac{\hbar^2}{2m}\nabla^2\psi + V\psi = i\hbar\frac{\partial}{\partial t}\psi.$$

Though we have derived the Schrödinger equation using arguments from plane waves, we will in the ensuing discussion be assuming that the wavefunctions ψ we consider vanish as $r \rightarrow \infty$, in order for the position of the particle to not be indeterminate.

Now, consider the solution space of the Schrödinger equation. Since the Schrödinger equation is a linear differential equation, the solution space is closed under linear combinations, and hence it is a vector space. Next, it is easy to check that the Schrödinger equation implies that the integral $\int_{\mathbb{R}^3} |\psi(\mathbf{r}, t)|^2 d^3r$ is invariant in time. We assume this is bounded, since we do not want the position or momentum of the particle to tend to infinity. This gives a norm on the solution space, and this is the norm induced by the inner product $\langle\psi|\phi\rangle = \int_{\mathbb{R}^3} \psi(\mathbf{r}, t)^* \phi(\mathbf{r}, t) d^3r$. Since no physical process can change the norm of the wavefunction, we may assume the norm is normalised to 1 — an innocent assumption, but one that will underlie the probabilistic interpretation of quantum mechanics. This gives rise to the first postulate.

Postulate 1: State. An isolated quantum system is represented by a complex Hilbert space H , and a state of the system corresponds to a unit vector in this space.

As is customary in quantum mechanics, we use **bra-ket notation** when working with quantum states. Given a unit vector $\psi \in H$, we write it as a **ket** $|\psi\rangle$ to indicate that it is to be considered as a quantum state. A vector in the dual space H^* is written as a **bra** $\langle\psi|$, corresponding to the linear form $\langle\psi|(|\phi\rangle) = \langle\psi|\phi\rangle$. The inner product on H can be seen as a bra acting on a ket, or a **braket** $\langle\psi|\phi\rangle$. Finally, the **ketbra** $|\psi\rangle\langle\phi|$ is the rank-one linear map $|\psi\rangle\langle\phi|(|\chi\rangle) = \langle\phi|\chi\rangle|\psi\rangle$.

Writing the Hamiltonian operator $\hat{H} = -\frac{\hbar^2}{2m}\nabla^2 + V$, the Schrödinger equation takes the simple form $\frac{\partial}{\partial t}|\psi\rangle = -\frac{i}{\hbar}\hat{H}|\psi\rangle$. This admits the unique solution $|\psi(t)\rangle = \exp(-\frac{i}{\hbar}\hat{H}t)|\psi(0)\rangle$. Since $\hat{H}$ is hermitian, $\exp(-\frac{i}{\hbar}\hat{H}t)$ is unitary, giving the second postulate.

Postulate 2: Evolution. The time evolution of an isolated quantum system is represented by a unitary operator.

If an observer measures a property of a quantum state, this no longer constitutes an evolution of the quantum system in isolation. The Hamiltonian $\hat{H}$ constitutes the prototypical example of an **observable** — a quantity that can be measured. This is a hermitian operator, so its eigenvalues are real (and hence can correspond to real quantities) and if it has a discrete spectrum, its eigenvectors form an orthonormal basis. Further, from the previous postulate, the eigenvectors of $\hat{H}$ are stationary under time-evolution, up to phase. Hence they are the states with well-defined energy, and the eigenvalues of $\hat{H}$ are the energy values that can be measured. This is the origin of the quantisation of the energy levels as discrete values. A general state $|\psi\rangle$ is a superposition over the eigenstates of the Hamiltonian $\sum_{\lambda} c_{\lambda}|\psi_{\lambda}\rangle$, but since an energy measurement will always give one of the eigenvalues λ , the post-measurement state should be $|\psi_{\lambda}\rangle$. The eigenvalue measured is sampled randomly, with probability of measuring λ given by the Born rule as $|c_{\lambda}|^2$ [24, 65]. The physical interpretation of how the state maps discontinuously to an eigenstate differs according to the interpretation of quantum mechanics: in the Copenhagen interpretation, this is the collapse of the wavefunction, and is treated as an axiom; in a Bayesian interpretation, measurement is the act of learning something new that was already a property of the system; or in an open systems interpretation, a measurement can be seen as a unitary on a larger system; among others. These properties together give the third postulate.

Postulate 3: Measurement. An observable on a quantum system H is a Hermitian operator A on H . The possible measurement outcomes are the eigenvalues of A , the

probability of measuring eigenvalue a is $\langle \psi | \Pi_a | \psi \rangle$ where Π_a is the projector on the eigenspace of a , and the post-measurement state is $\frac{\Pi_a |\psi\rangle}{\|\Pi_a |\psi\rangle\|}$.

Given two quantum systems H and K in states $|\psi\rangle$ and $|\phi\rangle$, their joint system is also a quantum system, so it should be a Hilbert space containing a state corresponding to the pair $(|\psi\rangle, |\phi\rangle)$. Superpositions of states in the joint system are possible, and they should be compatible with local superpositions on each of the systems. These are exactly the properties satisfied by the tensor product of Hilbert spaces, giving the fourth postulate.

Postulate 4: Composition. The joint state of two quantum systems H and K is described by a unit vector in $H \otimes K$.

This postulate can fail for some complex quantum systems. The tensor product model guarantees that every observable on H commutes with every observable on K , and hence they can be simultaneously measured. However, in the context of infinite-dimensional von Neumann algebras, there are algebras of operators that commute, but cannot take on a tensor product form. Such algebras capture the observables of some quantum systems coming from quantum field theory [69] and many-body physics [28]. As we will discuss in Section 2.12, this distinction between notions of composition gives rise to different ways of defining quantum correlations between systems.

2.6.2 States, operators, and measurements

A **density operator** $\rho \in \mathcal{B}(H)$ is a positive semidefinite operator with trace 1. Density operators allow us to synthesise classical probability with quantum states in the following way. Consider an ensemble of quantum states $|\psi_i\rangle$ sampled with probability p_i . Then, the density operator of this ensemble is $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$. The density operator satisfies the property that, for any operator M , $\text{Tr}(M\rho) = \sum_i p_i \langle\psi_i|M|\psi_i\rangle$, so the measurement probabilities of the ensemble can be captured by ρ . Further, by diagonalising a density operator, we see that every density operator is the density operator of an ensemble of states. With this interpretation, we call a quantum state $|\psi\rangle$ or a rank-one density operator $|\psi\rangle\langle\psi|$ a **pure state**, and a density operator ρ of rank > 1 a **mixed state**. Further, the density operators are in bijective correspondence with the states on $\mathcal{B}(H)$, seen as a C^* -algebra. A density operator ρ corresponds to the state $A \mapsto \text{Tr}(A\rho)$.

It is very useful to consider measurements independently from the observables that generate them. A **projection-valued measurement (PVM)** with n outcomes is a

collection of projections $\{P_i\}_{i=1}^n \subseteq \mathcal{B}(H)$ such that $\sum_i P_i = 1$. In a PVM, we have that $P_i P_j = \delta_{i,j} P_i$.

PVMs admit a useful generalisation. A **positive operator-valued measurement (POVM)** with n outcomes is a collection of positive operators $\{P_i\}_{i=1}^n \subseteq \mathcal{B}(H)$ such that $\sum_i P_i = 1$. POVMs can be constructed as convex combinations of PVMs, but not all POVMs arise in this way. Generally, it is not obvious when POVMs can be reduced to PVMs. We make use of two techniques to do so. The first is **Naimark's dilation theorem** [92]. This theorem states that, given a POVM $\{P_i\}_{i=1}^n \subseteq \mathcal{B}(H)$, there exists a PVM $\{\Pi_i\}_{i=1}^n \subseteq \mathcal{B}(H \otimes \mathbb{C}^n)$ and a state $|\psi\rangle \in \mathbb{C}^n$ such that $P_i = (1 \otimes \langle\psi|) \Pi_i (1 \otimes |\psi\rangle)$. This theorem allows a POVM to be generically transformed into a PVM, but does not preserve the Hilbert space. The second technique is a rounding result due to de la Salle [55, Theorem 2]. Let $\mathcal{M}$ be a von Neumann algebra, let $\{P_i\}_{i=1}^n \subseteq \mathcal{M}$ be a POVM, and let ρ be a normal state on $\mathcal{M}$. If $\sum_i \rho(P_i^2) \geq 1 - \varepsilon$, then there exists a PVM $\{\Pi_i\}_{i=1}^n \subseteq \mathcal{M}$ such that $\sum_i \rho((P_i - \Pi_i)^2) \leq 9\varepsilon$. This allows the POVM to be rounded to a nearby PVM while preserving the space, but does not work on every POVM.

2.6.3 Entanglement

A state $|\psi\rangle \in H \otimes K$ is called **separable** if there exist $|u\rangle \in H$ and $|v\rangle \in K$ such that $|\psi\rangle = |u\rangle \otimes |v\rangle$. Otherwise, $|\psi\rangle$ is called **entangled**. Entangled states do not admit a local description as a pure state on either system.

Any entangled state can be written in terms of its **Schmidt decomposition**. Given a state $|\psi\rangle \in H \otimes K$, there exist orthonormal sets $\{|h_i\rangle\}_{i=1}^n \subseteq H$ and $\{|k_i\rangle\}_{i=1}^n \subseteq K$ and coefficients $1 \geq p_1 \geq \dots \geq p_n > 0$ such that $\sum_i p_i = 1$ and

$$|\psi\rangle = \sum_{i=1}^n \sqrt{p_i} |h_i\rangle \otimes |k_i\rangle.$$

The **Schmidt rank** of $|\psi\rangle$ is $n \leq \min\{\dim H, \dim K\}$. $|\psi\rangle$ is separable if and only if it has Schmidt rank 1.

Given a finite-dimensional Hilbert space H with orthonormal basis $\{|i\rangle\}_{i=1}^n$, the **canonical maximally entangled state** is the state $|\phi_H^+\rangle = \frac{1}{\sqrt{n}} \sum_{i=1}^n |i\rangle \otimes |i\rangle \in H \otimes H$. More generally, we call a state **maximally entangled** if it is unitarily equivalent to the canonical maximally entangled state by local unitaries. The state $|\phi_H^+\rangle$ satisfies the property that $(U \otimes \bar{U})|\phi_H^+\rangle = |\phi_H^+\rangle$ for any unitary U . In particular, this implies that for any measurement made on the first subsystem, there exists a measurement on the second subsystem that always predicts the measurement result. This gives rise to a stronger form of correlation than what can be achieved classically.

2.7 Quantum spaces

The main reference for this section is [121].

2.7.1 Motivation

For a compact space X , the set of continuous functions $X \rightarrow \mathbb{C}$, denoted $C(X)$, is a commutative $*$ -algebra via pointwise addition, multiplication, and complex conjugation. Further, $C(X)$ is a C^* -algebra via the norm $\|f\| = \max \{|f(x)| \mid x \in X\}$. The mapping $X \mapsto C(X)$ is in fact a contravariant functor from the category of compact topological spaces to the category of unital commutative C^* -algebras. **Gelfand duality** is the statement that this functor is a bijection, if homeomorphic topological spaces and isomorphic C^* -algebras are identified.

By analogy, we can see a noncommutative C^* -algebra $\mathcal{A}$ as the algebra of continuous functions on a virtual **compact quantum space** X , and denote it $\mathcal{A} = C(X)$. A compact quantum space Y is a **quantum subspace** of X if there is a surjective $*$ -homomorphism $C(X) \rightarrow C(Y)$. A natural way to construct a quantum subspace is with respect to a set of relations $R \subseteq C(X)$: then, the surjection is simply the quotient by the two-sided ideal generated by the relations.

2.7.2 Quantum groups

The group structure of a compact topological group G induces additional structure on the algebra of continuous functions $C(G)$, which can be expressed in terms of $*$ -homomorphisms acting on $C(G)$. The group product induces a $*$ -homomorphism $\Delta : C(G) \rightarrow C(G \times G) \cong C(G) \otimes_{\min} C(G)$ via $\Delta(\phi)(g, h) = \phi(gh)$; the identity element of G induces a $*$ -homomorphism $\epsilon : C(G) \rightarrow \mathbb{C}$ as $\epsilon(\phi) = \phi(1)$; and the inverse induces a $*$ -homomorphism $S : C(G) \rightarrow C(G)$ as $S(\phi)(g) = \phi(g^{-1})$. Encoded in this way, group properties may be generalised to quantum spaces.

A **compact quantum group (CQG)** [121] is a quantum space G equipped with a $*$ -homomorphism $\Delta : C(G) \rightarrow C(G) \otimes_{\min} C(G)$ called the **coproduct**, satisfying the relation $(\text{id} \otimes \Delta)\Delta = (\Delta \otimes \text{id})\Delta$, called **coassociativity**; and such that $(\text{id} \otimes C(G))\Delta(C(G))$ and $(C(G) \otimes \text{id})\Delta(C(G))$ are dense in $C(G) \otimes_{\min} C(G)$. When working with compact quantum groups, we assume all tensor products of C^* -algebras are min-tensor-products, and suppress the subscript. A **quantum subgroup** H of a CQG G is a CQG that is a quantum subspace and $\Delta_H q = (q \otimes q)\Delta_G$ where $q : C(G) \rightarrow C(H)$ is the quotient map.

In this work, we need only compact matrix quantum groups. A **compact matrix quantum group (CMQG)** is a CQG where $C(G)$ is generated by generators u_{ij} for $1 \leq i, j \leq n$ and $\Delta(u_{ij}) = \sum_{k=1}^n u_{ik} \otimes u_{kj}$. Conversely, if Δ defined in this way is a $*$ -homomorphism and the matrices $u = (u_{ij})$ and $u^T = (u_{ji})$ are invertible, then it automatically satisfies the two conditions on the coproduct of a CQG, inducing a CQG structure. Further, on the dense subalgebra generated by the polynomials in the u_{ij} , we can define a **counit** $*$ -homomorphism $\epsilon : C(G) \rightarrow \mathbb{C}$ by $\epsilon(u_{ij}) = \delta_{ij}$ and an **antipode** anti-automorphism $S : C(G) \rightarrow C(G)$ satisfying $S(S(x^*)^*) = x$ where defined, by $S(u_{ij}) = u_{ji}^*$. In general, the counit and antipode may be unbounded maps. These satisfy the conditions of a **Hopf algebra**, *i.e.* $(\text{id} \otimes \epsilon)\Delta = (\epsilon \otimes \text{id})\Delta = \text{id}$ and $\mu(\text{id} \otimes S)\Delta = \mu(S \otimes \text{id})\Delta = 1 \cdot \epsilon$, where $\mu : C(G) \otimes C(G) \rightarrow C(G)$ is the coalgebra homomorphism representing the product $\mu(x \otimes y) = xy$.

A quantum group G **acts** on a C^* -algebra $\mathcal{A}$ if there exists a $*$ -homomorphism $\beta : \mathcal{A} \rightarrow \mathcal{A} \otimes C(G)$ such that $(\beta \otimes \text{id})\beta = (\text{id} \otimes \Delta)\beta$ and $(1 \otimes C(G))\beta(\mathcal{A})$ is dense in $\mathcal{A} \otimes C(G)$ [102]. The second condition implies that $(\text{id} \otimes \epsilon)\beta = \text{id}$ where defined. The **quantum permutation group** S_n^+ of n elements is the maximal quantum group acting on $C([n])$ that preserves the uniform distribution in the sense that $(\tau \otimes \text{id})\beta = 1 \cdot \tau$ where $\tau(f) = \frac{1}{n} \sum_{i=1}^n f(i)$. This is a CMQG with algebra of functions generated by p_{ij} with $1 \leq i, j \leq n$ subject to the relations $p_{ij}^2 = p_{ij}^* = p_{ij}$ for all i, j , and $\sum_j p_{ij} = \sum_j p_{ji} = 1$ for all i [119]. Because the generators of S_n^+ are hermitian and form a unitary matrix, the counit and antipode are bounded and hence defined everywhere. In the same way, we can define the quantum permutation group of any finite set X as $S_X^+ \cong S_{|X|}^+$.

2.8 Free probability theory

The main reference for this section is [95].

A **noncommutative probability space** is a unital algebra $\mathcal{A}$ equipped with a linear functional $\tau : \mathcal{A} \rightarrow \mathbb{C}$ such that $\tau(1) = 1$, called the **expectation**. We will always assume, as is commonly done, that $\mathcal{A}$ is a C^* -algebra and that τ is a faithful tracial state. Elements of a noncommutative probability space are called **noncommutative random variables**. As a first example, it is worth noting that commutative probability theory, as seen in the previous section, fits as a special case of noncommutative probability theory. In fact, given a probability space $(\Omega, \mathcal{S}, \mu)$, its set of bounded, complex-valued random variables corresponds to the noncommutative probability space $\mathcal{A} = L^\infty(\Omega, \mu)$ with $\tau(X) = \int X(\omega) d\mu(\omega) = \mathbb{E} X$. Another example of a noncommutative probability space is $(\mathbb{M}_d(\mathbb{C}), \text{tr})$.

A **mixed *-moment** of a collection of noncommutative random variables $a_1, \dots, a_n \in \mathcal{A}$ is an expectation of the form $\tau(a_{i_1}^{e_1} \cdots a_{i_m}^{e_m})$, where $i_j \in [n]$ and $e_j \in \{1, *\}$. For a single element $a \in \mathcal{A}$, the ***-moments** of a are the mixed *-moments of a . For normal $a \in \mathcal{A}$, *i.e.* $aa^* = a^*a$, the *-moments are simply $\tau((a^*)^k a^l)$ for integers $k, l \geq 0$. In this case, the idea of a **probability distribution** for random variables generalizes quite directly. There exists a unique measure μ_a supported on a compact subset of $\mathbb{C}$ such that all the *-moments can be written as

$$\tau((a^*)^k a^l) = \int (z^*)^k z^l d\mu_a(z).$$

Knowing the *-moments, the expectation of any element of $C^*\langle a \rangle$, the unital C^* -subalgebra generated by a , can be computed.

Subalgebras $\mathcal{A}_1, \dots, \mathcal{A}_n \subseteq \mathcal{A}$ are **freely independent** (or **free** for short) if, for any $i_1, \dots, i_m \in [n]$ such that $i_j \neq i_{j+1}$ and $a_j \in \mathcal{A}_{i_j}$ such that $\tau(a_j) = 0$, we have that $\tau(a_1 \cdots a_m) = 0$. We say that noncommutative random variables $a_1, \dots, a_n \in \mathcal{A}$ are free if the C^* -subalgebras they generate are free. The usual commutative notion of independence guarantees the above property only in the case where all the i_j are distinct, which can be supposed due to commutativity. If the *-moments of a set of free random variables are known, then any of the mixed *-moments can be computed.

A collection of sequences $(a_{1,n})_n, \dots, (a_{l,n})_n$ of noncommutative random variables such that $a_{i,n} \in (\mathcal{A}_n, \tau_n)$ **converges in *-distribution** to a collection of noncommutative random variables $a_1, \dots, a_l \in (\mathcal{A}, \tau)$ if all the mixed *-moments converge, *i.e.*

$$\lim_{n \rightarrow \infty} \tau_n(a_{i_1,n}^{e_1} \cdots a_{i_m,n}^{e_m}) = \tau(a_{i_1}^{e_1} \cdots a_{i_m}^{e_m}).$$

We write this as $a_{1,n}, \dots, a_{l,n} \rightarrow a_1, \dots, a_l$ (*-dist). If $(a_{1,n})_n, \dots, (a_{l,n})_n$ converge in *-distribution to noncommutative random variables that are free, we say that the sequences are **asymptotically free**.

The Haar distribution on the unitary group generalises naturally in the setting of free probability. We say a noncommutative random variable $u \in \mathcal{A}$ is a **Haar unitary** if u is unitary and the *-moments $\tau(u^k) = \delta_{k,0}$ for all $k \in \mathbb{Z}$. This constitutes a generalisation as a Haar-random unitary $U \sim \eta_{\mathcal{U}_d}$ can be seen as a noncommutative random variable given by the identity function in the space $\mathcal{A} = \mathbb{M}_d(L^\infty(\mathcal{U}_d, \eta_{\mathcal{U}_d}))$ with expectation $\tau(a) = \int \text{tr}(a(u)) d\eta_{\mathcal{U}_d}(u)$, which satisfies $\tau(U^k) = \delta_{k,0}$.

We make use of the following theorem.

Theorem 2.8.1 (23.13 in [95]). Let $(d_n)_n$ be a sequence of natural numbers such that $d_n \rightarrow \infty$. For all $n \in \mathbb{N}$, let $U_n \sim \eta_{\mathcal{U}_{d_n}}$ be a Haar-random unitary and let $A_n \in \mathbb{M}_{d_n}$. Suppose that the limiting *-moments $\lim_{n \rightarrow \infty} \text{tr}(A_n^{e_1} \cdots A_n^{e_m})$ for $e_i \in \{1, *\}$ all exist. Then the sequences $(U_n)_n$ and $(A_n)_n$ are asymptotically free.

2.9 Complexity theory

The main reference for this section is [6].

2.9.1 Turing machines

A **Turing machine** is a tuple $M = (k, \Sigma, S, \delta)$, where $k \in \mathbb{Z}_{\geq 0}$, called the number of work tapes; Σ is a finite set, called the set of symbols, containing a blank symbol $\square$; S is a finite set, called the set of states, containing the start state **start** and the halt state **halt**; and $\delta : S \times \Sigma^{k+2} \rightarrow S \times \Sigma^{k+2} \times \{-1, 0, 1\}^{k+2}$ is a function, called the transition function. A Turing machine describes an abstract computation as follows. The machine consists of $k + 2$ tapes (infinite strips where one symbol can be placed for every integer) of which one is the input tape, one is the output tape, and the remainder are work tapes; one head on each tape that can read and write on one position at a time; and a register that can store one of the states. The machine starts with a blank symbol on every position, except for the input tape where a word $x \in \Sigma^*$ (where Σ^* denotes the set of words in Σ) can be input, and the register in the **start** state. Then, at each step of the computation, the head on each tape i reads the symbol σ_i it is on, and the machine computes $\delta(s, \sigma_1, \dots, \sigma_{k+2}) = (s', \sigma'_1, \dots, \sigma'_{k+2}, d_1, \dots, d_{k+2})$, where s is the state in the register. Then, the machine updates the state in the register to s' , writes σ'_i at each head position i , and moves each head i in the direction given by d_i . If the state in the register ever becomes **halt**, then the machine halts. The output of the Turing machine is the word on the output tape, omitting blank symbols, when the computation halts, and is denoted $M(x)$. Since it takes a finite amount of information to specify a Turing machine, they can be represented as bit strings; importantly, there is a **universal Turing machine** that can take a bit string representation of a Turing machine as input and produce the same output as would the inputted Turing machine.

A **probabilistic Turing machine** is a Turing machine with an additional tape that starts filled with uniformly random symbols. This provides an abstract presentation of randomised models of computation.

The **runtime** of a Turing machine on input x is the number of computation steps before reaching the halt state. A function f is **computable** if there is a Turing machine M such that $f(x) = M(x)$ for all x . A function f is **T -time computable**, where $T : \mathbb{N} \rightarrow \mathbb{N}$, if there is a Turing machine M such that $f(x) = M(x)$ and the runtime of M on input x is less than $T(|x|)$. A function f is **polynomial-time computable** if it is T -time computable with $T(n) = Cn^c$ for some $C, c \geq 0$. A function f is **exponential-time computable** if it is T -time computable with $T(n) = C \cdot 2^{kn^c}$ for some $C, c, k \geq 0$.

2.9.2 Complexity classes

A **language** is a subset of bit strings $L \subseteq \{0, 1\}^*$. A Turing machine M **decides** a language L if $M(x) = 1$ if $x \in L$ and $M(x) = 0$ if $x \notin L$. A **promise problem** P is a pair of disjoint subsets $Y(P), N(P) \subseteq \{0, 1\}^*$. The elements of $Y(P) \cup N(P)$ are called the **instances** of P ; the elements of $Y(P)$ are the **yes instances** and the elements of $N(P)$ are the **no instances**. We write $x \in P$ to mean $x \in Y(P) \cup N(P)$. Every language L corresponds to the promise problem (L, L^c) , so we will work exclusively with promise problems in the following. A Turing machine M decides P if $M(x) = 1$ for all $x \in Y(P)$ and $M(x) = 0$ for all $x \in N(P)$. Note that it does not matter what M outputs for x in neither $Y(P)$ nor $N(P)$ — in fact M may not even halt. A **complexity class** is a collection of promise problems.

P is the class of all problems that can be decided in polynomial time. NP is the class of all problems P such that there exists a Turing machine M such that for all $x \in Y(P)$, there exists $y \in \{0, 1\}^*$, called a **proof**, such that $M(x, y) = 1$ and for all $x \in N(P)$ and all $y \in \{0, 1\}^*$, $M(x, y) = 0$; and M halts in polynomial time. EXP is the class of all problems that can be decided in exponential time. $NEXP$ is the exponential-time analogue of NP , defined identically except that the Turing machine may halt in exponential time. R is the class of all problems that can be decided with a Turing machine; $P \in R$ is called **decidable**, and $P \notin R$ is called **undecidable**. RE is the class of all problems P such that there exists a Turing machine M where if $x \in Y(P)$, M halts on input x , and if $x \in N(P)$, M does not halt on input x . Given a complexity class C , the **complement** is the class coC consisting of the problems of C with the yes and no instances swapped.

2.9.3 Reductions and hardness

Let P_1 and P_2 be promise problems. A **(Karp) reduction from P_1 to P_2** is a computable function f such that for all $x \in Y(P_1)$, $f(x) \in Y(P_2)$ and for all $x \in N(P_1)$, $f(x) \in N(P_2)$. The reduction is **polynomial-time** if f is polynomial-time computable. A problem P is **hard** for a complexity class C , or **C-hard**, if every problem in C reduces to P . Note that the time complexity stipulated for the reduction depends on the complexity class: for NP and $NEXP$ it is polynomial, while for RE it just has to be computable. We say P is **C-complete** if it is C -hard and $P \in C$. Note that polynomial-time reductions are used for $NEXP$ rather than exponential-time reductions since any NP -complete problem is complete for $NEXP$ under exponential-time reductions.

The **halting problem** is the problem where the yes instances are the Turing machines that halt on empty input and the no instances are the Turing machines that do not. The halting problem is RE -complete.

2.10 Relational structures

Relational structures provide a natural way to present constraint satisfaction problems. They allow arbitrary CSPs to be presented in a manner similar to graph homomorphism problems, such as graph n -colouring, which corresponds to homomorphism to the complete graph K_n .

2.10.1 Definitions

A **relation** over a set D is a subset $R \subseteq D^n$ for some $n \in \mathbb{N}$, called the **arity** of R and denoted $\text{ar}(R)$. For example, the set of edges of a graph forms a relation of arity 2. Equivalently, a relation can be seen as a boolean function $R : D^n \rightarrow \{0, 1\}$, corresponding to the characteristic function of the relation seen as a subset. A **relational structure** is a pair $A = (\mathfrak{D}(A), \mathfrak{R}(A))$ where $\mathfrak{D}(A)$ is a set called the **domain** of A and $\mathfrak{R}(A)$ is a set of relations over $\mathfrak{D}(A)$. Unless otherwise specified, we consider relational structures with finite domains and finitely many relations. When clear from context, we write A to mean $\mathfrak{D}(A)$. For a graph G , the domain is the set of vertices, and there is a single relation consisting of the set of edges. The specific examples of graphs is discussed in more detail in [Section 2.10.2](#).

A relational structure is **boolean** if $\mathfrak{D}(A) = \{0, 1\}$, the boolean field. Many important CSPs arise from boolean relational structures, such as 3SAT and boolean linear systems. These are defined formally in [Section 2.10.4](#).

To compare relational structures, we need to keep track of the relations. A **relational signature** is a set σ with an associated arity function $\text{ar} : \sigma \rightarrow \mathbb{N}$. We assume the signature to be finite unless otherwise specified. The elements of σ are called **relational symbols**. A **relational structure over σ** (or a **σ -structure**) is a relational structure A with an associated bijective mapping $\sigma \rightarrow \mathfrak{R}(A)$, $R \mapsto R^A$ such that $\text{ar}(R^A) = \text{ar}(R)$. If $\sigma \subseteq \tau$, then every σ -structure is also a τ -structure, where the relations with symbols in $\tau \setminus \sigma$ are taken as empty. If the signature of a relational structure A is not specified, we treat A as a σ_A -structure, where $\sigma_A = \mathfrak{R}(A)$ seen as a relational signature, so that $R^A = R$.

A **homomorphism** of σ -structures $f : A \rightarrow B$ is a map $f : \mathfrak{D}(A) \rightarrow \mathfrak{D}(B)$ such that for all $R \in \sigma$ and $\mathbf{a} \in R^A$, $f(\mathbf{a}) \in R^B$, where we write $f(\mathbf{a}) = (f(a_1), \dots, f(a_{\text{ar}(R)}))$. The category of σ -structures with morphisms given by the relational structure homomorphisms is denoted $\mathcal{R}(\sigma)$.

We close this section with a few ways of building new relational structures from old. Given a σ -structure A and an equivalence relation $\sim$ on $\mathfrak{D}(A)$, the **quotient** of A by $\sim$ is the σ -

structure $A/\sim$ with domain $\mathfrak{D}(A)/\sim$ and relations $R^{A/\sim} = \{([a_1], \dots, [a_{\text{ar}(R)}]) \mid \mathbf{a} \in R^A\}$. The **disjoint union** of σ -structures A and B is the σ -structure $A \sqcup B$ with domain $\mathfrak{D}(A) \sqcup \mathfrak{D}(B)$ and relations $R^{A \sqcup B} = R^A \sqcup R^B$. The **cartesian product** of A and B is the σ -structure $A \times B$ with domain $\mathfrak{D}(A) \times \mathfrak{D}(B)$ and relations

$$R^{A \times B} = \{((a_1, b_1), \dots, (a_{\text{ar}(R)}, b_{\text{ar}(R)})) \mid \mathbf{a} \in R^A, \mathbf{b} \in R^B\}.$$

Write $A^k = A \times A \times \dots \times A$. For any finite set S , we take $A^S = A^{|S|}$ with indices labelled by S .

2.10.2 Graphs

We work with finite simple undirected loop-free graphs. A **graph** $G = (V(G), E(G))$ consists of vertices and edges. $V(G)$ is a finite set, called the **vertices**; $E(G)$ is a subset of the unordered pairs of vertices $\{\{x, y\} \mid x, y \in V(G) \wedge x \neq y\}$, called the **edges**. The edges induce a symmetric irreflexive relation on $V(G)$, which we denote by $x \sim_G y$ iff $\{x, y\} \in E(G)$. We often write $x \in G$ to mean $x \in V(G)$.

A graph G may be presented as a relational structure with domain $V(G)$ and a single relation E^G . All graphs are relational structures over the signature $\sigma = \{E\}$, consisting of a single relation of arity 2, where $E^G = \{(x, y) \in V(G)^2 \mid \{x, y\} \in E(G)\}$. **Graph homomorphisms** are relational structure homomorphisms of the associated relational structure. For graphs, the cartesian product as relational structures is the same as the categorical (or tensor) product as graphs.

A graph G is **bipartite** if there exists a partition $V(G) = A \cup B$ such that $x \sim_G y$ if $x, y \in A$ or $x, y \in B$.

The **complement** of a graph G is the graph $\overline{G}$ with the same vertices and edges $x \sim_{\overline{G}} y$ if and only if $x \neq y$ and $x \not\sim_G y$.

A **graph product** is a way to define a graph on the vertex set $V(G) \times V(H)$ for graphs G and H . We work with the **categorical product** (or tensor product), denoted $G \times H$ with edges $(x, y) \sim_{G \times H} (x', y')$ iff $x \sim_G x'$ and $y \sim_H y'$. This corresponds to the cartesian product of the graphs as relational structures.

We use some standard examples of graphs. The **complete graph** on n vertices is the graph K_n with $V(K_n) = \mathbb{Z}_n$ and $x \sim_{K_n} y$ for all $x \neq y$. The **n -cycle** is the graph C_n with $V(C_n) = \mathbb{Z}_n$ and $x \sim_{C_n} y$ iff $x = y \pm 1$.

2.10.3 Constraint satisfaction problems

Let A be a σ -structure. The **constraint satisfaction problem (CSP)** of A is the promise problem $\text{CSP}(A)$ where the instances are σ -structures X , of which the yes instances are those admitting a relational structure homomorphism $X \rightarrow A$. This problem can in fact be seen as the language corresponding to the σ -structures with homomorphism to A .

There are two alternate, but equivalent, presentations of CSPs that commonly appear in the literature. We will exclusively work with the relational structure homomorphism presentation, but discuss the equivalence here for completeness.

The first alternate presentation is via **boolean formulae**. Here, seeing the relations of A as boolean functions, the instances of $\text{CSP}(A)$ are formulae of the form

$$\exists x_1, \dots, x_n \in \mathfrak{D}(A) : R_1(x_{i_{1,1}}, \dots, x_{i_{1,\text{ar}(R_1)}}) \wedge \dots \wedge R_m(x_{i_{m,1}}, \dots, x_{i_{m,\text{ar}(R_m)}}),$$

for some relations $R_i \in \mathfrak{R}(A)$ and $i_{j,k} \in [n]$, not necessarily distinct. The instance is a yes instance if the formula is true, and a no instance if not. This formulation portrays a CSP as the question of the solubility of a system of equations, and is commonly seen for boolean CSPs such as 3SAT or LIN (see below for the definitions). The presentation via boolean formulae can be seen to be equivalent to the structure homomorphism presentation by taking X to be the σ -structure with domain $\mathfrak{D}(X) = \{x_1, \dots, x_n\}$ — where the x_i are now seen as abstract symbols rather than variables over $\mathfrak{D}(A)$ — and relations $R^X = \{(x_{i_{j,1}}, \dots, x_{i_{j,\text{ar}(R)}}) \mid R_j = R^A\}$ for all $R \in \sigma$.

The second alternate presentation is via **constraint systems**. A **constraint system** $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ over an alphabet Σ consists of a set of variables X , contexts $V_i \subseteq X$, and associated constraints $C_i \subseteq \Sigma^{V_i}$. A **satisfying assignment** to S is a function $f : X \rightarrow \Sigma$ such that $f|_{V_i} \in C_i$ for all i . In this model, instances of $\text{CSP}(A)$ are constraint systems $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ over $\mathfrak{D}(A)$ where for each $i \in [m]$, there exists $R_i \in \mathfrak{R}(A)$ and a function $f_i : [\text{ar}(R_i)] \rightarrow V_i$ such that $C_i = \{f \in \Sigma^{V_i} \mid (f(f_i(1)), \dots, f(f_i(\text{ar}(R_i)))) \in R_i\}$. The instance S is a yes instance if it admits a satisfying assignment. This presentation of CSPs is commonly seen in the quantum literature, as it presents the contexts where the variables should be measured together in a clear way. The constraint system presentation is equivalent to the structure homomorphism presentation by taking the relational structure associated to an instance S to be the σ -structure Y with domain $\mathfrak{D}(Y) = X$ and relations $R^Y = \{(f_i(1), \dots, f_i(\text{ar}(R_i))) \mid R_i = R^A\}$.

These two presentations differ most significantly from the structure homomorphism presentation in that the terms of an instance are not relations but constraints. This perspective is useful. To make use of it in the structure homomorphism model, we say

that a **constraint** of a σ -structure X is a pair $(R, \mathbf{x})$ where $R \in \sigma$ and $\mathbf{x} \in R^X$. We say a function $f : \mathfrak{D}(X) \rightarrow \mathfrak{D}(A)$ **satisfies** the constraint $(R, \mathbf{x})$ if $f(\mathbf{x}) \in R^A$.

Next, we also consider gapped versions of CSPs as follows. Let $1 \geq c \geq s \geq 0$, called the **completeness** and **soundness** parameters, respectively. One may take these parameters to be functions of the instance size, but unless otherwise specified, we assume they are constants. Then, the (c, s) -**gapped CSP** of A is the promise problem $\text{CSP}_{c,s}(A)$ where the yes instances are σ -structures X such that there exists a function $f : \mathfrak{D}(X) \rightarrow \mathfrak{D}(A)$ that satisfies at least a fraction c of the constraints; and the no instances are σ -structures X such that any function $f : \mathfrak{D}(X(M)) \rightarrow \mathfrak{D}(A)$ satisfies a fraction of the constraints strictly less than s . If $c > s$, this problem is no longer a language.

Succinct versions of CSPs are also relevant for this work. The **succinct CSP** of A is the promise problem $\text{MIP}(A)$ (see [Chapter 3](#) for a clarification of the notation) where the instances are probabilistic Turing machines M that sample constraints of a σ -structure $X(M)$ in polynomial time, where the yes instances are those Turing machines for which there is a homomorphism $X(M) \rightarrow A$. The succinct presentation as a sampler over the constraints allows relational structures of exponential size to be handled in polynomial time. The (c, s) -**gapped succinct CSP** of A is the promise problem $\text{MIP}_{c,s}(A)$ where the instances are again probabilistic Turing machines M that sample constraints of $X(M)$, where yes instances are those Turing machines M such that there is a function $f : \mathfrak{D}(X(M)) \rightarrow \mathfrak{D}(A)$ such that the probability of f satisfying a constraint is at least c , and the no instances are those Turing machines M such that any function $f : \mathfrak{D}(X) \rightarrow \mathfrak{D}(A)$ has probability strictly less than s of satisfying a constraint.

2.10.4 Reductions and polymorphisms

The complexity of CSPs is well understood. **Schaefer's dichotomy theorem** [\[105\]](#) shows that every boolean CSP (*i.e.* where A is boolean) is either contained in **P** or is **NP**-complete. Furthermore, this theorem shows that the reductions between the **NP**-complete CSPs are particularly well-behaved: the reductions are **local** in the sense that every constraint of the old CSP is rewritten as a collection of constraints in the new CSP, called a **gadget**, independently of the other constraints. An analogous dichotomy theorem for all CSPs was long conjectured, and was independently completed by Bulatov [\[32\]](#) and Zhuk [\[122\]](#).

For those relational structures where $\text{CSP}(A)$ is **NP**-complete, we also find analogous hardness for the gapped and succinct problems. Due to the PCP theorem, there always exists $s < 1$ such that $\text{CSP}_{1,s}(A)$ remains **NP**-complete. Further, due to the locality of the reductions, we also find that $\text{MIP}(A)$ and $\text{MIP}_{1,s}(A)$ are **NEXP**-complete.

The prototypical example of an NP-complete CSP is 3SAT. This is the CSP with boolean relational structure 3SAT with signature $\sigma_{3\text{SAT}} = \{R_b \mid b \in \mathbb{Z}_2^3\}$, where all the symbols have arity 3, and relations $R_b^{3\text{SAT}} = \mathbb{Z}_2^3 \setminus \{b\}$. In the presentation via boolean formulae, this corresponds to satisfiability of 3-CNF formulae, which is the usual presentation of 3SAT. The hardness of 3SAT is the content of the Cook-Levin theorem.

Another example of a CSP is LIN. This is the CSP consisting of the systems of linear equations over the boolean field. The associated relational structure is the boolean structure LIN with signature $\sigma_{\text{LIN}} = \{L_{b,n} \mid n \in \mathbb{N}, b \in \{0, 1\}\}$, where $\text{ar}(L_{b,n}) = n$, and relations $L_{b,n}^{\text{LIN}} = \{\mathbf{a} \in \mathbb{Z}_2^n \mid a_1 + \dots + a_n = b\}$. Due to Gaussian elimination, $\text{CSP}(\text{LIN})$ is solvable in polynomial time. For any $n \in \mathbb{N}$, we also write $n\text{LIN}$ for the relational structure with only the relations of LIN on n variables. Note that LIN does not have finitely many relations, but $n\text{LIN}$ does.

The reductions between CSPs are presented most elegantly via the **algebraic approach**, which identifies reductions between CSPs with homomorphisms between objects related to their defining relational structure. The objects in question are the polymorphisms: the **polymorphism space** of a σ -structure A is the collection of morphism spaces $\text{mor}(A^n, A)$ for all $n \in \mathbb{N}$, denoted $\text{pol}(A)$. Polymorphisms capture generalised symmetries of the relational structure. At each level n , $\text{mor}(A^n, A)$ always contains the **projection maps** $\pi_i^n : A^n \rightarrow A$, $\mathbf{a} \mapsto a_i$ for each $i \in [n]$; but it may contain other maps. For example, every polymorphism of 3SAT is a projection; but 3LIN also has some other maps, such as the ternary minority operation $(a, b, c) \mapsto a + b + c$, as polymorphisms.

We say that a relational structure A **pp-defines** a relational structure B if for all $S \in \mathfrak{R}(B)$, there exist $n, m \in \mathbb{N}$, $R_1, \dots, R_n$ either in $\mathfrak{R}(A)$ or equal to the equality relation $\{(a, a) \mid a \in A\}$, $i_k \in [m]$ for $k \in [\text{ar}(S)]$, and $i_{jk} \in [m]$ for $j \in [n]$ and $k \in [\text{ar}(R_j)]$ such that

$$S = \{(a_{i_1}, \dots, a_{i_{\text{ar}(S)}}) \mid \mathbf{a} \in A^m \wedge (a_{i_{j1}}, \dots, a_{i_{j \text{ar}(R_j)}}) \in R_j \forall j \in [n]\}.$$

Lemma 2.10.1 ([64, 22, 23]). Let A and B be relational structures such that $\mathfrak{D}(A) = \mathfrak{D}(B)$. Then A pp-defines B if and only if $\text{pol}(A) \subseteq \text{pol}(B)$.

We include a proof of this result here since we make use of its structure later.

Proof. Suppose that A pp-defines B . Let $f \in \text{mor}(A^k, A)$ be a polymorphism of A . To see that it is a polymorphism of B , let $S \in \mathfrak{R}(B)$, and by pp-definition, there exist $n, m \in \mathbb{N}$, $R_1, \dots, R_n$ either in $\mathfrak{R}(A)$ or equal to the equality relation $\{(a, a) \mid a \in A\}$, $i_k \in [m]$ for $k \in [\text{ar}(S)]$, and $i_{jk} \in [m]$ for $j \in [n]$ and $k \in [\text{ar}(R_j)]$ such that

$$S = \{(a_{i_1}, \dots, a_{i_{\text{ar}(S)}}) \mid \mathbf{a} \in A^m \wedge (a_{i_{j1}}, \dots, a_{i_{j \text{ar}(R_j)}}) \in R_j \forall j \in [n]\}.$$

Now, let $(a_{1i_1}, \dots, a_{1i_{\text{ar}(S)}}), \dots, (a_{ki_1}, \dots, a_{ki_{\text{ar}(S)}}) \in S$. Then, for all $l \in [k]$ and $j \in [n]$, $(a_{li_{j1}}, \dots, a_{li_{j \text{ar}(R_j)}}) \in R_j$. Since f is a polymorphism of A ,

$$(f(a_{1i_{j1}}, \dots, a_{ki_{j1}}), \dots, f(a_{1i_{j \text{ar}(R_j)}}, \dots, a_{ki_{j \text{ar}(R_j)}})) \in R_j.$$

This implies that $(f(a_{1i_1}, \dots, a_{ki_1}), \dots, f(a_{1i_{\text{ar}(S)}}, \dots, a_{ki_{\text{ar}(S)}})) \in S$. Hence $f \in \text{mor}(B^k, B)$.

For the converse, let $S \in \mathfrak{R}(B)$. Enumerate the elements $S = \{\mathbf{s}_1, \dots, \mathbf{s}_{|S|}\}$. Consider the $|A|^{|S|}$ -ary relation over A

$$P = \{(x_{\mathbf{a}})_{\mathbf{a} \in A^{|S|}} \mid \mathbf{a} \mapsto x_{\mathbf{a}} \text{ is a polymorphism of } A\}.$$

A pp-defines P as

$$P = \{(x_{\mathbf{a}})_{\mathbf{a} \in A^{|S|}} \mid (x_{(a_{11}, \dots, a_{1|S|})}, \dots, x_{(a_{\text{ar}(R)1}, \dots, a_{\text{ar}(R)|S|})}) \in R \forall R \in \mathfrak{R}(A), \mathbf{a}_1, \dots, \mathbf{a}_{|S|} \in R\}.$$

Now, for $i \in [\text{ar}(S)]$ let $\mathbf{t}_i = (s_{1i}, \dots, s_{|S|i}) \in A^{|S|}$ and define S' to be the $\text{ar}(S)$ -ary relation

$$S' = \{(x_{\mathbf{t}_1}, \dots, x_{\mathbf{t}_{\text{ar}(S)}}) \mid (x_{\mathbf{a}})_{\mathbf{a} \in A^{|S|}} \in P\}.$$

It is clear that A pp-defines S' . To finish, we show that $S' = S$. The projection onto the i -th component $\pi_i : A^{|S|} \rightarrow A$ is a polymorphism, and therefore $(a_i)_{\mathbf{a} \in A^{|S|}} \in P$. As such, $(t_{1i}, \dots, t_{\text{ar}(S)i}) = \mathbf{s}_i \in S'$, giving $S \subseteq S'$. On the other hand, if $\mathbf{a} \in S'$, then there exists a $|S|$ -ary polymorphism f such that $a_i = f(\mathbf{t}_i)$. As such, since $\mathbf{s}_1, \dots, \mathbf{s}_{|S|} \in S$, we have that

$$\mathbf{a} = (f(s_{11}, \dots, s_{|S|1}), \dots, f(s_{1 \text{ar}(S)}, \dots, s_{|S| \text{ar}(S)})) \in S.$$

Hence $S' \subseteq S$ as well. ■

We can access stronger types of reduction by noting that the polymorphism space of a σ -structure is an example of a **minion**, which is a covariant functor from a skeleton category of finite sets (where a skeleton category is such that any two isomorphic objects are identified). More concretely, a minion $\mathcal{M}$ is a collection of spaces $\mathcal{M}_n$ for $n \in \mathbb{N}$, called the **levels**, such that for every function $\pi : [n] \rightarrow [m]$, there exists a mapping $\mathcal{M}_n \rightarrow \mathcal{M}_m$, $M \mapsto M_{/\pi}$ called a **minor operation**, satisfying $M_{/\text{id}} = M$ and for all $\pi : [n] \rightarrow [m]$ and $\pi' : [m] \rightarrow [k]$, $(M_{/\pi})_{/\pi'} = M_{/\pi' \circ \pi}$. $\text{pol}(A)$ can be equipped with a minion structure by taking $\text{pol}(A)_n = \text{mor}(A^n, A)$ and, for $\phi \in \text{mor}(A^n, A)$ and $\pi : [n] \rightarrow [m]$, $\phi_{/\pi}(\mathbf{a}) = \phi(\mathbf{a} \circ \pi)$ (where we identify tuples A^m and functions $A^{[m]}$). Given minions $\mathcal{M}$ and $\mathcal{N}$, a **minion homomorphism** $\xi : \mathcal{M} \rightarrow \mathcal{N}$ is a natural transformation between the functors, or concretely, a collection of maps $\xi_n : \mathcal{M}_n \rightarrow \mathcal{N}_n$ such that $\xi_m(M_{/\pi}) = \xi_n(M)_{/\pi}$. We often drop the subscript n on ξ_n . The following theorem provides the strongest way of achieving reductions between CSP languages from the structure of the polymorphisms.

Theorem 2.10.2 ([17, 37, 15]). Let A be a σ -structure and B be a τ -structure, and suppose there exists a minion homomorphism $\xi : \text{pol}(A) \rightarrow \text{pol}(B)$. Then, there exists a polynomial-time reduction from $\text{CSP}(B)$ to $\text{CSP}(A)$.

In fact, the existence of this minion homomorphism is also a necessary condition for the existence of a reduction, as long as $\text{P} \neq \text{NP}$ — however, this is very much not obvious, and follows from the fact that those relational structures that do not admit a homomorphism to 3SAT are solvable in polynomial time [32, 122]. Further, since the reduction guaranteed by this theorem is local, it can be effected in logarithmic space, and therefore lifts to a polynomial-time reduction between succinct CSPs. Due to the locality, it also preserves the existence of a constant gap for gapped CSPs, but the size of the gap may change. We include a proof of this result here, since it is instructive once we pass to the quantum setting.

Proof. Let Y be a τ -structure. Define the σ -structure $X_0 = \bigsqcup_{y \in Y} A^{\mathfrak{D}(B)} \sqcup \bigsqcup_{R \in \tau, \mathbf{y} \in R^Y} A^{R^B}$. We write elements of X_0 as tuples $(y, \mathbf{a})$ or $(R, \mathbf{y}, \mathbf{a})$ to identify the term of the disjoint union that they belong to. Next, we take the equivalence relation $\sim$ on X_0 to be defined via $(y_i, \mathbf{a}) \sim (R, \mathbf{y}, \mathbf{a} \circ \pi_i)$ for all $R \in \tau$, $\mathbf{y} \in R^Y$, $i \in [\text{ar}(R)]$, and $\mathbf{a} \in A^{\mathfrak{D}(B)}$, where π_i is the i -th projection map. Let X be the quotient σ -structure $X = X_0 / \sim$. We claim that the mapping $Y \rightarrow X$ provides a reduction from $\text{CSP}(B)$ to $\text{CSP}(A)$.

First, suppose that Y is a yes instance of $\text{CSP}(B)$. Then, there exists a homomorphism $\phi : Y \rightarrow B$. Define the map $\gamma : \mathfrak{D}(X) \rightarrow \mathfrak{D}(A)$ as $\gamma([(y, \mathbf{a})]) = \mathbf{a}(\phi(y))$ and $\gamma([(R, \mathbf{y}, \mathbf{a})]) = \mathbf{a}(\phi(\mathbf{y}))$. We need to check if this map is well-defined. First $\gamma([(R, \mathbf{y}, \mathbf{a})])$ exists, as, if $\mathbf{y} \in R^Y$, then $\phi(\mathbf{y}) \in R^B$, and hence it is in the domain of $\mathbf{a}$. Also, if $(y, \mathbf{a}) \sim (R, \mathbf{y}, \mathbf{b})$, then there exists i such that $y = y_i$ and $\mathbf{b} = \mathbf{a} \circ \pi_i$. As such, $\gamma([(R, \mathbf{y}, \mathbf{b})]) = \mathbf{a}(\pi_i(\phi(\mathbf{y}))) = \mathbf{a}(\phi(y_i)) = \gamma([(y, \mathbf{a})])$, so γ is well-defined. Next, we need to show that γ is a homomorphism. Let $S \in \sigma$ with arity n and let $\mathbf{x} \in S^X$. By definition, either $x_i = [(y, \mathbf{a}_i)]$ or $x_i = [(R, \mathbf{y}, \mathbf{a}_i)]$ for all i . In the first case, $\gamma(\mathbf{x}) = (\mathbf{a}_1(\phi(y)), \dots, \mathbf{a}_n(\phi(y))) \in S^A$, as $(\mathbf{a}_1, \dots, \mathbf{a}_n) \in S^{A^{\mathfrak{D}(B)}}$ so the projection onto one of the terms is an element of S^A . In the second case, $\gamma(\mathbf{x}) = (\mathbf{a}_1(\phi(\mathbf{y})), \dots, \mathbf{a}_n(\phi(\mathbf{y}))) \in S^A$, again as $(\mathbf{a}_1, \dots, \mathbf{a}_n) \in S^{A^{R^B}}$. So γ is a homomorphism, giving that X is a yes instance of $\text{CSP}(A)$.

For the converse, suppose that X is a yes instance of $\text{CSP}(A)$. Then, there exists a homomorphism $\gamma : X \rightarrow A$. For each $y \in Y$, define $\gamma_y : \mathfrak{D}(A^{\mathfrak{D}(B)}) \rightarrow \mathfrak{D}(A)$ as $\gamma_y(\mathbf{a}) = \gamma([(y, \mathbf{a})])$; and for each $R \in \tau$ and $\mathbf{y} \in R^Y$, define $\gamma_{R, \mathbf{y}} : \mathfrak{D}(A^{R^B}) \rightarrow \mathfrak{D}(A)$ as $\gamma_{R, \mathbf{y}}(\mathbf{a}) = \gamma([(R, \mathbf{y}, \mathbf{a})])$. By construction, these are relational structure homomorphisms, and hence polymorphisms of A . Using the minion homomorphism, we have that $\phi_y = \xi(\gamma_y) \in \text{mor}(B^{\mathfrak{D}(B)}, B)$ and $\phi_{R, \mathbf{y}} = \xi(\gamma_{R, \mathbf{y}}) \in \text{mor}(B^{R^B}, B)$. Define $\phi : \mathfrak{D}(Y) \rightarrow \mathfrak{D}(B)$ as $\phi(y) =$

$\phi_y(\text{id})$. This is well-defined, since the identity is a function from $\mathfrak{D}(B) \rightarrow \mathfrak{D}(B)$ and hence an element of $B^{\mathfrak{D}(B)}$. To finish, we show that ϕ is a relational structure homomorphism. Let $R \in \tau$ be a relation of arity n , and let $\mathbf{y} \in R^Y$. To show $\phi(\mathbf{y}) \in R^B$, we first need to see the implications of the defining equivalence relation of X on the ϕ_{y_i} and $\phi_{R,\mathbf{y}}$. Since $\pi_i : R^B \rightarrow \mathfrak{D}(B)$, $(\gamma_{R,\mathbf{y}})_{/\pi_i} \in \text{mor}(A^{\mathfrak{D}(B)}, A)$. Acting on $\mathbf{a} \in A^{\mathfrak{D}(B)}$, $(\gamma_{R,\mathbf{y}})_{/\pi_i}(\mathbf{a}) = \gamma([(R, \mathbf{y}, \mathbf{a} \circ \pi_i)]) = \gamma([(y_i, \mathbf{a})]) = \gamma_{y_i}(\mathbf{a})$. Using the properties of the minion homomorphism, this implies that $(\phi_{R,\mathbf{y}})_{/\pi_i} = \phi_{y_i}$. Then, $\phi(\mathbf{y}) = (\phi_{y_1}(\text{id}), \dots, \phi_{y_n}(\text{id})) = (\phi_{R,\mathbf{y}}(\pi_1), \dots, \phi_{R,\mathbf{y}}(\pi_n)) \in R^B$, as $(\pi_1, \dots, \pi_n) \in R^{B^{R^B}}$ — in fact, for any $\mathbf{b} \in R^B$, $(\pi_1(\mathbf{b}), \dots, \pi_n(\mathbf{b})) = \mathbf{b} \in R^B$. ■

2.11 Semidefinite programming

2.11.1 Definitions and duality

The main reference for this section is [120].

A **convex cone** in a real vector space V is a convex subset $K \subseteq V$ closed under scalar multiplication by any nonnegative scalar. Given a convex cone K in an inner product space V , the **dual cone** is the space

$$K^* = \{v \in V \mid \langle u, v \rangle \geq 0 \ \forall u \in K\}.$$

If K is a closed convex cone, then $K^{**} = K$. Let V and W be real finite-dimensional inner product spaces, let $K \subseteq V$ be a closed convex cone, let $T : V \rightarrow W$ be a real linear map, and let $a \in V$ and $b \in W$ be vectors. Then, the **conic program** $\mathcal{P}$ parametrised by (V, W, K, T, a, b) is the optimisation problem

$$\begin{aligned} & \text{maximise } \langle a, x \rangle \\ & \text{subject to } T(x) = b \\ & \quad x \in K. \end{aligned}$$

A vector x is a **feasible solution** of $\mathcal{P}$ if $x \in K$ and $T(x) = b$, and an **optimal solution** if it maximises $\langle a, x \rangle$ among the feasible solutions. The supremum over the values of $\langle a, x \rangle$ among the feasible solutions is called **optimal value**, and is denoted $\text{opt}(\mathcal{P})$. Note that the optimal value may be $\pm\infty$, in which case no optimal solution exists. A conic program is an example of a **convex optimisation problem**, and admits a nice duality property.

The **dual program** $\mathcal{P}^*$ of $\mathcal{P}$ is the optimisation problem

$$\begin{aligned} & \text{minimise } \langle b, y \rangle \\ & \text{subject to } T^*(y) - a \in K^* \\ & y \in W, \end{aligned}$$

where T^* is the adjoint of T with respect to the inner products on V and W . The dual always satisfies the property that $\text{opt}(\mathcal{P}) \leq \text{opt}(\mathcal{P}^*)$. By **Slater's theorem**, if there exists a feasible solution of $\mathcal{P}$ that is in the algebraic interior of K (the set of points x in K such that for all $v \in V$ there exists $t > 0$ such that $x + tv \in K$), then the optimal values of $\mathcal{P}$ and $\mathcal{P}^*$ are equal.

A **semidefinite program (SDP)** is a conic program where V is the real vector space of Hermitian matrices $H_d \subseteq \mathbb{M}_d$ with Hilbert-Schmidt inner product $\langle A, B \rangle = \text{Tr}(AB)$, and K is the cone of positive semidefinite matrices. Explicitly, the semidefinite program $\mathcal{P}$ parametrised by a real inner product space W , a linear map $T : H_d \rightarrow W$, a matrix $A \in H_d$, and a vector $b \in W$ is the optimisation problem

$$\begin{aligned} & \text{maximise } \text{Tr}(AX) \\ & \text{subject to } T(X) = b \\ & X \geq 0, \end{aligned}$$

Often, rather than a single linear map T with an associated vector b , the constraints of an SDP may be presented as a collection T_i of linear maps, each with an associated vector b_i . An SDP can be approximated to any precision ε in time $\text{poly}(d, \log(1/\varepsilon))$ [116].

2.11.2 Approximation algorithms and Max-Cut

A general **optimisation problem** can be expressed as the question of maximising an **objective function** obj on variables satisfying some fixed constraints. For example, deciding if an instance X of $\text{CSP}_{c,s}(A)$ is a yes instance can be phrased as the question of optimising the number of constraints satisfied by a function $f : \mathfrak{D}(X) \rightarrow \mathfrak{D}(A)$. Even if an optimisation problem cannot be solved efficiently, it may be efficiently **approximated**, which is sufficient in some cases. For example, suppose there is an efficient algorithm that outputs, given a relational structure X , a function $f : \mathfrak{D}(X) \rightarrow \mathfrak{D}(A)$ such that f satisfies a fraction ρ of the maximum number of constraints that may be satisfied by a function $\mathfrak{D}(X) \rightarrow \mathfrak{D}(A)$. Then, the problem $\text{CSP}_{c,s}(A)$ can be solved in polynomial time if $\rho c > s$, as for any yes instance, the algorithm outputs a function satisfying at least a fraction s of the constraints.

A versatile way of constructing approximation algorithms is via SDP relaxations. An **SDP relaxation** of an optimisation problem $\mathcal{P}$ is an SDP $\mathcal{Q}$ such that $\text{opt}(\mathcal{P}) \leq \text{opt}(\mathcal{Q})$. An optimal solution x of the SDP relaxation can be computed in polynomial time. To turn this into a feasible solution $\tilde{x}$ of $\mathcal{P}$, the approximation algorithm requires a second step called **rounding**. Then the **approximation ratio** of the algorithm is $\inf_{\mathcal{P}} \frac{\text{obj}(\tilde{x})}{\text{opt}(\mathcal{P})} \geq \inf_{\mathcal{P}} \frac{\text{obj}(\tilde{x})}{\text{opt}(\mathcal{Q})}$. Since the second term can actually be computed (we don't know how to find the optimal value of $\mathcal{P}$), we take this to be the approximation ratio, often denoted ρ . Assuming the Unique Games Conjecture, the optimal approximation algorithm for any CSP has this form [104].

To illustrate such an algorithm, consider the **Goemans-Williamson algorithm** [67] for the Max-Cut problem, which is the optimisation problem associated to the CSP of the graph K_2 . An instance of this problem is a graph G , and the problem consists of maximising the number of edges mapped to the edge of K_2 by a function $G \rightarrow K_2$. This can be phrased in the following way:

$$\begin{aligned} & \text{maximise} && \sum_{\{i,j\} \in E(G)} \frac{1 - x_i x_j}{2} \\ & \text{subject to} && x_i \in \{-1, +1\} \forall i \in V(G), \end{aligned}$$

since the term $\frac{1 - x_i x_j}{2} = 0$ if $x_i = x_j$ and 1 otherwise. One defines an SDP relaxation of this problem by replacing the product of the values $x_i \in \{-1, +1\}$ by the inner product of real unit vectors. In terms of the Gram matrix of these vectors, this takes the form of the SDP

$$\begin{aligned} & \text{maximise} && \sum_{\{i,j\} \in E(G)} \frac{1 - X_{i,j}}{2} \\ & \text{subject to} && X_{i,i} = 1 \forall i \\ & && X \geq 0. \end{aligned}$$

To finish, we need to round an optimal solution X of the SDP to a feasible solution of the original problem. To do so, first reexpress X as a Gram matrix $X_{i,j} = \langle v_i, v_j \rangle$. Then, sample a random unit vector r , and let $x_i = \text{sgn}\langle r, v_i \rangle$. This gives a feasible solution of the original problem, with approximation ratio

$$\rho = \inf_{u,v} \mathbb{E}_r \frac{1 - \text{sgn}(\langle r, u \rangle) \text{sgn}(\langle r, v \rangle)}{1 - \langle u, v \rangle} = \inf_{\theta \in [0, \pi]} \frac{2}{\pi} \frac{\theta}{1 - \cos \theta} \approx 0.8786,$$

which is the **Goemans-Williamson constant**.

2.12 Nonlocal games

2.12.1 Games and strategies

A (2-player) **nonlocal game** $\mathbf{G}$ consists of finite sets X and Y called Alice's and Bob's questions; collections of finite sets $\{A_x\}_{x \in X}$ and $\{B_y\}_{y \in Y}$ called Alice's and Bob's answers; a probability distribution $\mu : X \times Y \rightarrow [0, 1]$, called the **question distribution**; and a collection of values $V = \{V(a, b|x, y)\}_{x \in X, y \in Y, a \in A_x, b \in B_y}$ in $\{0, 1\}$, called the **predicate**. Alternately, we may take $A = \bigcup_{x \in X} A_x$ and $B = \bigcup_{y \in Y} B_y$ and stipulate that $V(a, b|x, y) = 0$ when $a \notin A_x$ or $b \notin B_y$. We may sometimes assume that the verifier holds on to some **side information** from a finite set S , which she uses to decide the predicate, without informing the players. Then, the question distribution is over $X \times Y \times S$ and the predicate is $V = \{V(a, b|x, y, s)\}_{x \in X, y \in Y, a \in A_x, b \in B_y, s \in S}$.

Any strategy for $\mathbf{G}$ can be presented by its **correlation**, which is a collection of values $p = \{p(a, b|x, y)\}_{x \in X, y \in Y, a \in A_x, b \in B_y}$ in $[0, 1]$ such that $(a, b) \mapsto p(a, b|x, y)$ is a probability distribution for all x, y . The **value** (or **winning probability**) of a nonlocal game $\mathbf{G}$ with respect to a correlation p is

$$\mathfrak{w}(\mathbf{G}, p) = \sum_{x \in X, y \in Y, a \in A_x, b \in B_y} \mu(x, y) V(a, b|x, y) p(a, b|x, y).$$

Different classes of strategies for nonlocal games can be captured by different classes of correlation.

- p is a **deterministic correlation** if there exist functions $f : X \rightarrow A$, $g : Y \rightarrow B$ such that $p(a, b|x, y) = \delta_{f(x), a} \delta_{g(y), b}$.
- p is a **classical correlation** if it is in the convex hull of the deterministic correlations, or more concretely if there exists a random variable Λ on a set L and functions $f : X \times L \rightarrow A$ and $g : Y \times L \rightarrow B$ such that $p(a, b|x, y) = \mathbb{E} \delta_{f(x, \Lambda), a} \delta_{g(y, \Lambda), b}$.
- p is a **quantum correlation** if there exist finite-dimensional Hilbert spaces H_A and H_B , POVMs $\{P_a^x\}_{a \in A_x} \subseteq \mathcal{B}(H_A)$ and $\{Q_b^y\}_{b \in B_y} \subseteq \mathcal{B}(H_B)$ for all $x \in X$, $y \in Y$, and a unit vector $|\psi\rangle \in H_A \otimes H_B$ such that $p(a, b|x, y) = \langle \psi | P_a^x \otimes Q_b^y | \psi \rangle$.
- p is a **quantum approximate correlation** if there exists a sequence of quantum correlations $(p_n)_n$ such that $p_n(a, b|x, y) \rightarrow p(a, b|x, y)$.

- p is a **commuting operator correlation** if there exists a Hilbert space H , POVMs $\{P_a^x\}_{a \in A_x} \subseteq \mathcal{B}(H)$ and $\{Q_b^y\}_{b \in B_y} \subseteq \mathcal{B}(H)$, and a unit vector $|\psi\rangle \in H$ such that $[P_a^x, Q_b^y] = 0$ for all x, y, a, b and $p(a, b|x, y) = \langle \psi | P_a^x Q_b^y | \psi \rangle$.
- p is a **tracial correlation** if there exists a von Neumann algebra $\mathcal{M} \subseteq \mathcal{B}(H)$, PVMs $\{P_a^x\}_{a \in A_x} \subseteq \mathcal{M}$ and $\{Q_b^y\}_{b \in B_y} \subseteq \mathcal{M}$, and a normal tracial state $\tau : \mathcal{M} \rightarrow \mathbb{C}$ such that $p(a, b|x, y) = \tau(P_a^x Q_b^y)$. We say p is a **quantum tracial correlation** if $\mathcal{M}$ may be chosen to be finite-dimensional; and p is a **quantum approximate tracial correlation** if $\mathcal{M}$ may be chosen to be Connes-embeddable.

In all of the above cases, we call the additional data needed to express the correlation the **model**. Due to Naimark's dilation theorem, every quantum and commuting operator correlation admits a model where the POVMs are projective. Every tracial correlation is a commuting-operator correlation, by taking the standard form of $\mathcal{M}$ [70]. Similarly, every quantum tracial correlation is a quantum correlation, by an analogous but simpler finite-dimensional argument; these correlations have maximally-entangled shared state $|\psi\rangle$.

For a class of correlations Q — where we use the symbol $Q = d$ for deterministic, $Q = c$ for classical, $Q = q$ for quantum, $Q = qa$ for quantum approximate, $Q = co$ for commuting operator, $Q = t$ for tracial, $Q = qat$ for quantum approximate tracial, and $Q = qt$ for quantum tracial — we denote the set of all correlations with given input and output spaces by $C_Q(A, B|X, Y)$. Write C_Q to mean the union over all $C_Q(A, B|X, Y)$. We have that

$$\begin{array}{ccccccc} C_d & \subseteq & C_c & \subseteq & C_q & \subseteq & C_{qa} & \subseteq & C_{co} \\ & & & & \cup & & \cup & & \cup \\ & & & & C_{qt} & \subseteq & C_{qat} & \subseteq & C_t. \end{array}$$

All these inclusions are strict [39, 109, 110, 76]. By definition, $C_{qa}(A, B|X, Y)$ is the closure of $C_q(A, B|X, Y)$; and $C_{qat}(A, B|X, Y)$ is the closure of $C_{qt}(A, B|X, Y)$.

The **value** of $\mathbf{G}$ with respect to a set of correlations $C_Q(A, B|X, Y)$ is

$$\mathfrak{w}_Q(\mathbf{G}) = \sup_{p \in C_Q(A, B|X, Y)} \mathfrak{w}(\mathbf{G}, p).$$

We have that $\mathfrak{w}_d(\mathbf{G}) = \mathfrak{w}_c(\mathbf{G}) \leq \mathfrak{w}_q(\mathbf{G}) = \mathfrak{w}_{qa}(\mathbf{G}) \leq \mathfrak{w}_{co}(\mathbf{G})$ and $\mathfrak{w}_{qt}(\mathbf{G}) = \mathfrak{w}_{qat}(\mathbf{G}) \leq \mathfrak{w}_t(\mathbf{G})$, where all the inequalities may be strict [39, 76]. We say a correlation $p \in C_Q(A, B|X, Y)$ is **optimal** for $\mathbf{G}$ over Q if $\mathfrak{w}(\mathbf{G}, p) = \mathfrak{w}_Q(\mathbf{G})$, and ε -**optimal** for $\mathbf{G}$ over Q if $\mathfrak{w}(\mathbf{G}, p) \geq \mathfrak{w}_Q(\mathbf{G}) - \varepsilon$. We say a correlation p is **perfect** for $\mathbf{G}$ if $\mathfrak{w}(\mathbf{G}, p) = 1$, and ε -**perfect** if $\mathfrak{w}(\mathbf{G}, p) \geq 1 - \varepsilon$.

2.12.2 XOR games

A nonlocal game $\mathbf{G} = (X, Y, A, B, \mu, V)$ is **binary** if $A = B = \{0, 1\}$. An **XOR game** is a binary nonlocal game $\mathbf{G}$ for which there exists a function $f : X \times Y \rightarrow \{0, 1\}$ such that $V(a, b|x, y) = \delta_{a+b, f(x, y)}$, where the addition is modulo 2. As for any nonlocal game, we may keep track of side information for XOR games, in which case $V(a, b|x, y, s) = \delta_{a+b, f(x, y, s)}$. The value of a quantum correlation $p(a, b|x, y) = \langle \psi | P_a^x \otimes Q_b^y | \psi \rangle$ for an XOR game $\mathbf{G}$ has a simple expression in terms of the observables $P_x = P_0^x - P_1^x$ and $Q_y = Q_0^y - Q_1^y$ for the POVMs $\{P_a^x\}_a$ and $\{Q_b^y\}_b$, respectively:

$$\begin{aligned} \mathfrak{w}(\mathbf{G}, p) &= \sum_{\substack{x \in X, y \in Y \\ a, b \in \{0, 1\}: a+b=f(x, y)}} \mu(x, y) \langle \psi | P_a^x \otimes Q_b^y | \psi \rangle \\ &= \frac{1}{4} \sum_{x \in X, y \in Y} \mu(x, y) \sum_{a \in \{0, 1\}} \langle \psi | (1 + (-1)^a P_x) \otimes (1 + (-1)^{a+f(x, y)} Q_y) | \psi \rangle \\ &= \frac{1}{2} + \frac{1}{2} \sum_{x \in X, y \in Y} (-1)^{f(x, y)} \mu(x, y) \langle \psi | P_x \otimes Q_y | \psi \rangle. \end{aligned}$$

This presentation implies that the XOR game can be fully characterised by the matrix H with entries $H_{x, y} = (-1)^{f(x, y)} \mu(x, y)$, called the **game matrix** of $\mathbf{G}$, and the strategy can be characterised by the **quantum correlation matrix** $C_{x, y} = \langle \psi | P_x \otimes Q_y | \psi \rangle$. In general, a **correlation matrix** is a matrix $C = (C_{x, y})_{x \in X, y \in Y}$ with entries in $[-1, 1]$ and the **bias** of a correlation matrix is

$$\mathfrak{b}(\mathbf{G}, C) = \sum_{x \in X, y \in Y} H_{x, y} C_{x, y}.$$

For a quantum correlation p with associated correlation matrix C , $\mathfrak{b}(\mathbf{G}, C) = 2\mathfrak{w}(\mathbf{G}, p) - 1$. The **quantum bias** of $\mathbf{G}$ is $\mathfrak{b}_q(\mathbf{G}) = 2\mathfrak{w}_q(\mathbf{G}) - 1$.

Quantum correlations for XOR games admit a useful alternate characterisation. A **vector correlation matrix** C for an XOR game $\mathbf{G}$ is a correlation matrix C such that $C_{x, y} = \langle u_x | v_y \rangle$ for some real unit vectors $|u_x\rangle, |v_y\rangle$. The **vector bias** is the supremum over the biases of all vector correlation matrices; it is denoted $\mathfrak{b}_{vect}(\mathbf{G})$. A vector correlation matrix C is **optimal** if $\mathfrak{b}(\mathbf{G}, C) = \mathfrak{b}_{vect}(\mathbf{G})$, and ε -**optimal** if $\mathfrak{b}(\mathbf{G}, C) \geq \mathfrak{b}_{vect}(\mathbf{G}) - 2\varepsilon$. We include a factor of 2 in the definition of ε -optimal bias to ensure that the value is within ε of the optimal value.

The vector correlation matrices are in bijective correspondence with quantum correlation matrices, and therefore $\mathfrak{b}_{vect}(G) = \mathfrak{b}_q(G)$ [114]. This characterisation via vector correlations

allows the bias of an XOR game to be expressed as an SDP [41]. Following [44, 108], let $B = \frac{1}{2} \begin{bmatrix} 0 & H \\ H^T & 0 \end{bmatrix}$. Then, the quantum bias of $\mathbf{G}$ is the optimal value of the SDP over $(|X| + |Y|) \times (|X| + |Y|)$ -matrices M

$$\begin{aligned} & \text{maximise } \text{Tr}(BM) \\ & \text{subject to } M_{i,i} = 1, \forall i \in X \sqcup Y \\ & M \geq 0. \end{aligned} \tag{2.4}$$

Here, M represents the Gram matrix of the vectors defining the correlation matrix. The dual of this SDP is the SDP over vectors $\mathbf{a} \in \mathbb{R}^X$ and $\mathbf{b} \in \mathbb{R}^Y$

$$\begin{aligned} & \text{minimise } \frac{1}{2} \sum_{x \in X} a_x + \frac{1}{2} \sum_{y \in Y} b_y \\ & \text{subject to } \frac{1}{2} \Delta(\mathbf{a}, \mathbf{b}) \geq B, \end{aligned} \tag{2.5}$$

where $\Delta(\mathbf{a}, \mathbf{b})$ is the matrix with the entries of the vectors along the diagonal, and zeroes elsewhere. For an optimal strategy, the values a_x and b_y are called the **optimal row and column biases**, respectively, and satisfy

$$\begin{aligned} \sum_{y \in Y} H_{x,y} \langle u_x | v_y \rangle &= a_x \\ \sum_{x \in X} H_{x,y} \langle u_x | v_y \rangle &= b_y. \end{aligned}$$

We make use of the following structure theorem for near-optimal XOR game strategies.

Theorem 2.12.1 ([108] Theorem 3.1). Let $\mathbf{G}$ be an XOR game, and let a_x for $x \in X$ and b_y for $y \in Y$ be the optimal row and column biases, respectively. Then, for any ε -optimal vector correlation matrix $C_{x,y} = \langle u_x | v_y \rangle$,

$$\begin{aligned} \left\| \sum_{y \in Y} H_{x,y} |v_y\rangle - a_x |u_x\rangle \right\|^2 &\leq 10\sqrt{2(|X| + |Y|)\varepsilon}, \\ \left\| \sum_{x \in X} H_{x,y} |u_x\rangle - b_y |v_y\rangle \right\|^2 &\leq 10\sqrt{2(|X| + |Y|)\varepsilon}. \end{aligned}$$

2.12.3 Synchronous and projection games

A nonlocal game $\mathbf{G} = (X, Y, \{A_x\}_x, \{B_y\}_y, \mu, V)$ is **synchronous** if $X = Y$, $A_x = B_x$, μ and V are symmetric, and $V(a, b|x, x) = 0$ if $a \neq b$. If a correlation p is perfect for $\mathbf{G}$,

then it must satisfy $p(a, b|x, x) = 0$ whenever $a \neq b$ and $\mu(x, x) > 0$. If $\mu(x, x) > 0$ for all x , then any commuting operator correlation that is perfect for $\mathbf{G}$ must be tracial and **symmetric** in the sense that $p(a, b|x, y) = \tau(P_a^x P_b^y)$. Similarly, quantum and quantum approximate correlations must be quantum tracial and quantum approximate tracial, respectively [100, 81].

This condition can be made robust. For some $\alpha > 0$, we say $\mathbf{G}$ is α -**synchronous** if $\mu(x, x) \geq \alpha \sum_{y \neq x} \mu(x, y)$ for all $x \in X$. There exists a universal constant $c > 0$ such that if $\mathfrak{w}_{co}(\mathbf{G}) \geq 1 - \varepsilon$, then $\mathfrak{w}_t(\mathbf{G}) \geq 1 - c(\varepsilon/\alpha)^{1/4}$, where the near-optimal tracial strategy found is also symmetric. The same relation holds for the quantum and quantum approximate values [117, 86].

In order to reduce to working with tracial correlations, a nonlocal game may be transformed into a synchronous game. Let $\mathbf{G} = (X, Y, \{A_x\}_x, \{B_y\}_y, \mu, V)$ be a nonlocal game and $\alpha > 0$. Then, the α -**synchronisation** of $\mathbf{G}$ is the nonlocal game $\mathbf{G}^{\alpha\text{-synch}} = (X^{\alpha\text{-synch}}, X^{\alpha\text{-synch}}, \{A_x^{\alpha\text{-synch}}\}_x, \{A_x^{\alpha\text{-synch}}\}_x, \mu^{\alpha\text{-synch}}, V^{\alpha\text{-synch}})$ with question sets $X^{\alpha\text{-synch}} = X \sqcup Y$, answer sets $A_x^{\alpha\text{-synch}} = A_x$ for $x \in X$ and $A_y^{\alpha\text{-synch}} = B_y$ for $y \in Y$, question distribution

$$\mu^{\alpha\text{-synch}}(x, y) = \begin{cases} \frac{\alpha}{2(1+\alpha)} \sum_{z \in Y} \mu(x, z) & x = y \in X \\ \frac{\alpha}{2(1+\alpha)} \sum_{z \in X} \mu(z, y) & x = y \in Y \\ \frac{\mu(x, y)}{2(1+\alpha)} & x \in X \wedge y \in Y \\ \frac{\mu(y, x)}{2(1+\alpha)} & x \in Y \wedge y \in X \\ 0 & \text{else,} \end{cases}$$

and predicate

$$V^{\alpha\text{-synch}}(a, b|x, y) = \begin{cases} \delta_{a,b} & x, y \in X \vee x, y \in Y \\ V(a, b|x, y) & x \in X \wedge y \in Y \\ V(b, a|y, x) & y \in X \wedge x \in Y. \end{cases}$$

By construction, the tracial value $\mathfrak{w}_t(\mathbf{G}^{\alpha\text{-synch}}) = \frac{\alpha + \mathfrak{w}_t(\mathbf{G})}{1+\alpha}$, with identical relations holding for the quantum tracial and quantum approximate tracial cases.

If $X = Y$, a symmetric modification of the transformation is useful. The **symmetric** α -**synchronisation** of $\mathbf{G}$ is the nonlocal game $\mathbf{G}^{\alpha\text{-symm}} = (X, X, S, \{A_x \cup B_x\}_x, \{A_x \cup B_x\}_x, \mu^{\alpha\text{-symm}}, V^{\alpha\text{-symm}})$ with side information $S = \{\text{fix}, \text{swap}, \text{cc}\}$, corresponding to a round played with the players as usual, swapped, or as a consistency check; question

distribution

$$\mu^{\alpha\text{-symm}}(x, y, s) = \begin{cases} \frac{\alpha}{2(1+\alpha)} \sum_{z \in X} (\mu(x, z) + \mu(z, x)) \delta_{x,y} & s = \text{cc} \\ \frac{\mu(x, y)}{2(1+\alpha)} & s = \text{fix} \\ \frac{\mu(y, x)}{2(1+\alpha)} & s = \text{swap}; \end{cases}$$

and predicate

$$V^{\alpha\text{-symm}}(a, b|x, y, s) = \begin{cases} \delta_{a,b} & s = \text{cc} \\ V(a, b|x, y) & s = \text{fix} \\ V(b, a|y, x) & s = \text{swap}. \end{cases}$$

We have that $\mathfrak{w}_t(\mathbf{G}^{\alpha\text{-symm}}) = \frac{\alpha + \mathfrak{w}_{t, \text{symm}}(\mathbf{G})}{1 + \alpha}$, where $\mathfrak{w}_{t, \text{symm}}(\mathbf{G})$ is the tracial value of $\mathbf{G}$ restricted to symmetric strategies, with identical relations holding for the quantum tracial and quantum approximate tracial cases.

A nonlocal game $\mathbf{G} = (X, Y, \{A_x\}_{x \in X}, \{B_y\}_{y \in Y}, \mu, V)$ is a **projection game** if for all $x \in X$ and $y \in Y$, there exists a function $p_{x,y} : A_x \rightarrow B_y$ such that $V(a, b|x, y) = \delta_{b, p_{x,y}(a)}$. Perfect quantum strategies for projection games are also known to be tracial. This property extends to the robust setting but with weaker guarantees than what is known for synchronous games [117, 98].

We can also transform arbitrary nonlocal games into projection games. The (projection) **oracularisation** of a nonlocal game $\mathbf{G} = (X, Y, \{A_x\}_{x \in X}, \{B_y\}_{y \in Y}, \mu, V)$ is the nonlocal game $\mathbf{G}^{\text{orac}} = (X \times Y, X \sqcup Y, \{A_x \times B_y\}_{(x,y) \in X \times Y}, \{A_x\}_{x \in X} \cup \{B_y\}_{y \in Y}, \mu^{\text{orac}}, V^{\text{orac}})$, where

$$\mu^{\text{orac}}((x, y), z) = \begin{cases} \frac{1}{2} \mu(x, y) & z = x \vee z = y \\ 0 & \text{else,} \end{cases}$$

and

$$V^{\text{orac}}((a, b), c|(x, y), z) = \begin{cases} 1 & V(a, b|x, y) = 1 \wedge z = x \wedge c = a \\ 1 & V(a, b|x, y) = 1 \wedge z = y \wedge c = b \\ 0 & \text{else.} \end{cases}$$

Oracularisations are clearly projection games via the projection functions $p_{(x,y),x}((a, b)) = a$ and $p_{(x,y),y}((a, b)) = b$.

The relationship between the values of $\mathbf{G}$ and $\mathbf{G}^{\text{orac}}$ is a little more complicated than in the synchronous case. A tracial correlation $p(a, b|x, y) = \tau(P_a^x Q_b^y)$ is **oracularisable**

if P_a^x and Q_b^y commute for all $a \in A_x$ and $b \in B_y$ and x, y such that $\mu(x, y) > 0$; the restriction that the PVMs need only commute for questions asked at the same time makes oracularisable strategies differ significantly from commuting operator strategies, where the commutation relation is used to guarantee independence of Alice's and Bob's subsystems — in fact, if the PVMs for an oracularisable tracial correlation commute for all questions, then the correlation is classical. In general, we only have the upper bound $w_t(\mathbf{G}^{\text{orac}}) \leq \min \left\{ 1 - \frac{1}{2}(1 - w_t(\mathbf{G}))^2, \sqrt{\frac{1+w_t(\mathbf{G})}{2}} \right\}$ [57], but if the optimal tracial strategy is oracularisable, then we also find the lower bound $w_t(\mathbf{G}^{\text{orac}}) \geq \frac{1+w_t(\mathbf{G})}{2}$. Further, if $\mathbf{G}^{\text{orac}}$ admits a perfect tracial correlation, then $\mathbf{G}$ admits a perfect oracularisable tracial correlation.

2.12.4 Complexity of nonlocal games

The main reference for this section is [118].

A more advanced way to define complexity classes than directly via Turing machines is by using **proof systems**. In an **interactive proof system**, a trusted polynomial-time verifier interacts with an untrusted but unbounded prover to decide a language. **Multiprover interactive proof systems** are a generalisation where the verifier interacts with multiple non-interacting provers. Deterministic interactive proof systems cannot decide languages outside of NP but probabilistic systems, which may not have perfect completeness or soundness, can have much greater complexity. In its most general form, a multiprover interactive proof system is parametrised by the number of provers m , the number of rounds n , the bit length of the questions ℓ_q , the bit length of the answers ℓ_a , the computation time T , and the completeness and soundness parameters $1 \geq c \geq s \geq 0$, which may in general be arbitrary functions of the instance size. A promise problem P is in $\text{MIP}_{c,s}(m, n, \ell_q, \ell_a, T)$ if there exists a probabilistic Turing machine M that runs in time T such that, on the input of an instance x , there exist provers such that the following protocol outputs 1 with probability $\geq c$ if x is a yes instance, and for any provers the following protocol outputs 1 with probability $< s$ if x is a no instance:

1. The verifier inputs x into M , sends the output $(q_1, \dots, q_m)$ to the provers, who respond with answers $(a_1, \dots, a_m)$.
2. The verifier repeats this for each of the n rounds, inputting also the transcript of the previous rounds into M .
3. At the end of the last round, the verifier inputs x and the full transcript of the interaction into M , which outputs a single bit b .

If m , n , ℓ_q , ℓ_a , and T are polynomial in $|x|$, and c and s have inverse-polynomial gap, then $\text{MIP}_{c,s}(m, n, \ell_q, \ell_a, T) \subseteq \text{NEXP}$. This is an equality in a very strong way. In fact, there exists a constant ℓ_a such that $\text{MIP}_{1,1/2}(2, 1, \text{poly}, \ell_a, \text{poly}) = \text{NEXP}$ [14]. Due to this, we will always consider $m = 2$, $n = 1$, and $T = \text{poly}$ proof systems so we only write the question and answer length in the definition of MIP classes. Further, one way of phrasing the **PCP theorem** is via the equality $\text{MIP}_{1,1/2}(\log, \ell_a) = \text{NP}$, for some constant ℓ_a [8, 7, 56].

With the restriction to two provers and one round, a problem with Turing machine M in the class $\text{MIP}_{c,s}(\ell_q, \ell_a)$ corresponds to determining whether the classical value of a nonlocal game is $\geq c$ or $< s$, where M first samples the questions and at the end decides if the players won. If ℓ_q or ℓ_a is superlogarithmic, then the game is **succinctly-presented**, in the sense that the verifier can sample the questions and verify the answers but cannot describe the whole question distribution or predicate in polynomial time.

The provers in an MIP protocol are not allowed to communicate, but may agree on a strategy beforehand. Allowing for quantum mechanics, this means that the provers may share entanglement, which was not considered in the original definitions of the MIP classes. The class is modified depending on the model of entanglement considered. If the provers are allowed to share finite-dimensional entanglement in the tensor product model, we call the class $\text{MIP}_{c,s}^*(\ell_q, \ell_a)$; if the provers are allowed to share infinite-dimensional entanglement in the commuting-operator model, we call the class $\text{MIP}_{c,s}^{\text{co}}(\ell_q, \ell_a)$. The breakthrough result of [76] shows that $\text{MIP}_{1,1/2}^*(\text{poly}, \text{poly}) = \text{RE}$, capturing even undecidable problems. The equivalence here is particularly powerful in that it is via a polynomial-time reduction from the halting problem to succinctly-presented nonlocal games. Since then, the result has been tightened somewhat. Again via a polynomial-time reduction from the halting problem, [60] shows that there exists ℓ_a constant such that $\text{MIP}_{1,1/2}^*(\text{poly}, \ell_a) = \text{RE}$, which has been further tightened in [62] to $\text{MIP}_{1,1/2}^*(\text{polylog}, \ell_a) = \text{RE}$. In all the reductions above, the yes instances are mapped to synchronous games with oracularisable quantum tracial optimal correlation. In the other direction, [93] shows that there exists ℓ_q constant such that $\text{MIP}_{1,1/2}^*(\ell_q, \text{polylog}) = \text{RE}$. Interestingly, in the classical setting, we have that $\text{MIP}_{1,1/2}(\ell_q, \text{polylog}) \subseteq \text{EXP}$, weaker than the full class. In the zero-gap setting $c = s$, the complexity increases further: $\text{MIP}_{1,1}^*(\text{poly}, \text{poly}) = \Pi_2^0$ [89]. Here, the games are again synchronous with oracularisable optimal strategy, but this strategy is quantum approximate rather than quantum. For the commuting operator model, [110] shows that there exists a constant ℓ_a such that $\text{MIP}_{1,1}^{\text{co}}(\text{poly}, \ell_a) = \text{coRE}$. This has recently been extended to the gapped setting, giving that $\text{MIP}_{1,1/2}^{\text{co}}(\text{poly}, \text{poly}) = \text{coRE}$ as well, via a polynomial-time reduction from the cohalting problem [83]. Since reductions in RE and coRE are not bounded by runtime, these results also give the same hardness for non-succinct nonlocal games, *i.e.* with log length questions and answers, albeit via an exponential-time reduction from the

halting problem (or cohalting problem).

Chapter 3

Nonlocal games for constraint satisfaction

In this chapter, we introduce the different ways of presenting constraint satisfaction problems as nonlocal games, and study their general structure. In [Section 3.1](#), we define entangled CSPs using both nonlocal games and weighted algebras, and ascertain their complexity in general. In [Section 3.2](#), we study how perfect strategies for the nonlocal games arising from constraint satisfaction correspond to a form of homomorphism of relational structures. In [Section 3.3](#), we demonstrate a strong way of rounding to tracial strategies in projection nonlocal games, and apply this to entangled CSPs.

3.1 Entangled constraint satisfaction problems

In this section, we define entangled CSPs. In [Section 3.1.1](#), we introduce the different ways of presenting homomorphisms between relational structures as nonlocal games. In [Section 3.1.2](#), we give an alternate way of viewing quantum strategies for the relational structure homomorphism games via weighted algebras. In [Section 3.1.3](#), we define the different types of entangled CSP we study. In [Section 3.1.4](#), we show some straightforward equalities between some of the entangled CSPs. In [Section 3.1.5](#), we show that hardness of approximation of nonlocal games implies analogous hardness for a general entangled CSP.

3.1.1 Relational structure homomorphisms as nonlocal games

There are several ways to express constraint satisfaction problems as nonlocal games. Fix a relational signature σ , and let A and B be σ -structures. The nonlocal games capturing homomorphisms $A \rightarrow B$ differ in how the constraints are asked to the players. They fall into two main frameworks: **non-oracular**, where the players are asked to assign values to variables; and **oracular**, where at least one of the players is asked to assign values to a whole constraint.

The first variant is the oracular game where the verifier samples a constraint of A to send to each player. They each have to respond with a tuple from the corresponding relation of B , and win if the variables in the outputs are consistent.

Definition 3.1.1. Let π be a probability distribution on pairs of constraints of A . Then, the **oracular constraint-constraint game** $G_{occ}(A, B, \pi)$ is the nonlocal game with question sets $X = Y = \bigcup_{R \in \sigma} \{R\} \times R^A$, answer sets $A_{R,a} = B_{R,a} = \mathfrak{D}(B)^{\text{ar}(R)}$, question distribution $\mu = \pi$, and predicate $V(\mathbf{b}, \mathbf{b}' | (R, \mathbf{a}), (S, \mathbf{a}')) = 1$ iff $\mathbf{b} \in R^B$, $\mathbf{b}' \in S^B$, and $b_i = b'_j$ whenever $a_i = a'_j$.

This is a synchronous game, so we can restrict to working with tracial strategies. In a gapped setting, we might have to make sure that there are sufficiently many consistency checks, which can be achieved by means of a symmetric α -synchronisation transformation. Constraint-constraint games were first introduced in [4] and further studied in [81, 99, 87]

In the second oracular variant we consider, the referee samples a constraint of A to send to Alice, and a variable from that constraint to send to Bob. Then, Alice responds with a tuple from the corresponding relation of B and Bob responds with an element of B . As before, the players win if the assignments are consistent.

Definition 3.1.2. Let π be a probability distribution on the constraints of A . Then, the **oracular constraint-variable game** $G_{ocv}(A, B, \pi)$ is the nonlocal game with question sets $X = \bigcup_{R \in \sigma} \{R\} \times R^A$ and $Y = \mathfrak{D}(A)$, answer sets $A_{R,a} = \mathfrak{D}(B)^{\text{ar}(R)}$ and $B_a = \mathfrak{D}(B)$, additional side information $S = \mathbb{N}$, question distribution $\mu((R, \mathbf{a}), a, i) = \frac{\pi(R, \mathbf{a})}{\text{ar}(R)}$ if $a_i = a$ and 0 otherwise, and predicate $V(\mathbf{b}, b | (R, \mathbf{a}), a, i) = 1$ iff $\mathbf{b} \in R^B$ and $b_i = b$.

This formulation gives a projection game, which again allows us to restrict to tracial strategies. In Section 3.3, we present a robust version of this statement. Constraint-variable games were studied in [43, 5, 1].

Generally, we will only work with the oracular constraint-variable framework, rather than the oracular constraint-constraint framework. This can be done without loss of generality

in the context of the decision problems we are interested in, since the two frameworks are equivalent. We defer the proof of this equivalence to [Section 5.6](#), since we will need to develop some more powerful tools to achieve it.

A third, non-oracular, way of presenting the constraint problem as a nonlocal game arises in the 2-CSP case, *i.e.* $\text{ar}(R) = 2$ for all $R \in \sigma$. In this game, the verifier samples constraints from A and sends each player a variable from the constraint. The players respond with variables of B , and win if the variables together satisfy the corresponding relation of B .

Definition 3.1.3. Suppose $\text{ar}(R) = 2$ for all $R \in \sigma$, and let π be a probability distribution on the constraints of A . Then, the **non-oracular game** $\mathbf{G}_{no}(A, B, \pi)$ is the nonlocal game with question sets $X = Y = \mathfrak{D}(A)$, answer sets $A_{R,a} = B_{R,a} = \mathfrak{D}(B)$, additional side information $S = \sigma$, question distribution $\mu(a, a', R) = \pi(R, (a, a'))$, and predicate $V(b, b' | a, a', R) = 1$ iff $(b, b') \in R^B$.

This is again a synchronous game, so we can restrict to working with tracial strategies, once consistency checks are added via symmetric α -synchronisation. We also see that, when the non-oracular game is defined, the constraint-variable game is exactly its oracularisation — hence the choice of notation “oracular.” The non-oracular game has been studied in the context of quantum graph homomorphisms [\[31, 63, 34, 84\]](#).

To finish this section, we note that perfect deterministic classical strategies for any of these three games, given sufficiently well supported distributions, exactly correspond to homomorphisms $A \rightarrow B$. This will become a bit more complicated for quantum strategies, as we will discuss in [Section 3.2](#).

Definition 3.1.4. We say a probability distribution π on pairs of constraints of a relational structure A is **connected** if, for every $a \in A$, the graph G_a , with vertex set given by the constraints $(R, \mathbf{a})$ of A such that $a_i = a$ for some i and edges $(R, \mathbf{a}) \sim_{G_a} (S, \mathbf{a}')$ iff $\pi((R, \mathbf{a}), (S, \mathbf{a}')) > 0$, is connected.

Lemma 3.1.5. Let A and B be σ -structures. Then, there is a homomorphism $A \rightarrow B$ if and only if one of the following conditions hold:

- (i) $\mathbf{G}_{occ}(A, B, \pi_0)$, where π_0 is connected, has a perfect deterministic symmetric strategy;
- (ii) $\mathbf{G}_{ocv}(A, B, \pi_1)$, where π_1 has full support, has a perfect deterministic strategy;
- (iii) $\mathbf{G}_{no}(A, B, \pi_2)$, where π_2 has full support, has a perfect deterministic symmetric strategy (where in this case, we assume all the relations have arity 2).

Proof. First suppose that $f : A \rightarrow B$ is a homomorphism. Then, it is direct to verify that $p(\mathbf{b}, \mathbf{b}' | (R, \mathbf{a}), (S, \mathbf{a}')) = \delta_{f(\mathbf{a}), \mathbf{b}} \delta_{f(\mathbf{a}'), \mathbf{b}'}$ is a perfect strategy for $\mathbf{G}_{occ}(A, B, \pi_0)$ (recall that $f(\mathbf{a}) = (f(a_1), \dots, f(a_{\text{ar}(R)}))$); that $p(\mathbf{b}, b | (R, \mathbf{a}), a) = \delta_{f(\mathbf{a}), \mathbf{b}} \delta_{f(a), b}$ is a perfect strategy for $\mathbf{G}_{ocv}(A, B, \pi_1)$; and that $p(b, b' | a, a') = \delta_{f(a), b} \delta_{f(a'), b'}$ is a perfect strategy for $\mathbf{G}_{no}(A, B, \pi_2)$.

Now, suppose that $p(\mathbf{b}, \mathbf{b}' | (R, \mathbf{a}), (S, \mathbf{a}')) = \delta_{F(R, \mathbf{a}), \mathbf{b}} \delta_{F(S, \mathbf{a}'), \mathbf{b}'}$ is a perfect deterministic strategy for $\mathbf{G}_{occ}(A, B, \pi_0)$. Take the function $f : \mathfrak{D}(A) \rightarrow \mathfrak{D}(B)$ to be $f(a) = F(R, \mathbf{a})_i$ for any constraint $(R, \mathbf{a})$ such that $a_i = a$. f is well-defined, since the strategy is perfect and π_0 is connected. Then, for any $R \in \sigma$ and $\mathbf{a} \in R^A$, $f(\mathbf{a}) = F(R, \mathbf{a}) \in R^B$, as the strategy is perfect. So f is a homomorphism.

Next, suppose that $p(\mathbf{b}, b | (R, \mathbf{a}), a) = \delta_{F(R, \mathbf{a}), \mathbf{b}} \delta_{f(a), b}$ is a perfect deterministic strategy for $\mathbf{G}_{ocv}(A, B, \pi_1)$. We claim that f is a homomorphism $A \rightarrow B$. In fact, since the strategy is perfect and π_1 has full support, $F(R, \mathbf{a})_i = f(a_i)$. Therefore, if $R \in \sigma$ and $\mathbf{a} \in R^A$, then $f(\mathbf{a}) = F(R, \mathbf{a}) \in R^B$.

Finally, suppose that $p(b, b' | a, a') = \delta_{f(a), b} \delta_{f(a'), b'}$ is a perfect strategy for $\mathbf{G}_{no}(A, B, \pi_2)$. We claim that f is a homomorphism $A \rightarrow B$. In fact, since the strategy is perfect and π_2 has full support, for any $R \in \sigma$ and $(a_1, a_2) \in R^A$, we have that $(f(a_1), f(a_2)) \in R^B$. ■

3.1.2 Weighted algebra formalism

A useful algebraic way to express the CSP games introduced in the previous section is by means of the weighted algebra formalism, first introduced for this purpose in [87]. Here, tracial strategies will correspond to traces on the algebra and the defect of a trace will be related to the winning probability. As before, we fix σ -structures A and B .

Definition 3.1.6. Let π be a probability distribution on pairs of constraints of A . The **oracular constraint-constraint algebra** $A \rightarrow B$ is the weighted algebra $\text{Mor}_\pi^{occ}(A, B) = (\text{Mor}^{occ}(A, B), \mu_{occ, \pi})$ where the unital $*$ -algebra $\text{Mor}^{occ}(A, B)$ is generated by $\Phi_{R, \mathbf{a}, \mathbf{b}}$ for $R \in \sigma$, $\mathbf{a} \in R^A$, and $\mathbf{b} \in R^B$ satisfying the relations

$$\Phi_{R, \mathbf{a}, \mathbf{b}} = \Phi_{R, \mathbf{a}, \mathbf{b}}^* \quad \Phi_{R, \mathbf{a}, \mathbf{b}} \Phi_{R, \mathbf{a}, \mathbf{b}'} = \delta_{\mathbf{b}, \mathbf{b}'} \Phi_{R, \mathbf{a}, \mathbf{b}} \quad \sum_{\mathbf{b} \in R^B} \Phi_{R, \mathbf{a}, \mathbf{b}} = 1;$$

and the weight function

$$\mu_{occ, \pi}(\Phi_{R, \mathbf{a}, \mathbf{b}} \Phi_{S, \mathbf{a}', \mathbf{b}'}) = \pi((R, \mathbf{a}), (S, \mathbf{a}'))$$

if there exist $i \in [\text{ar}(R)]$ and $j \in [\text{ar}(S)]$ such that $a_i = a'_j$ but $b_i \neq b'_j$, and $\mu_{occ, \pi}(x) = 0$ for all other $x \in \text{Mor}_\pi^{occ}(A, B)$. Denote the defect polynomial of $\text{Mor}_\pi^{occ}(A, B)$ as $D_\pi^{occ}(A, B) = D(\text{Mor}_\pi^{occ}(A, B), \mu_{occ, \pi})$.

Let π be a probability distribution on the constraints of A . The **oracular constraint-variable algebra** $A \rightarrow B$ is the weighted algebra $\text{Mor}_\pi^{\text{ocv}}(A, B) = (\text{Mor}^{\text{ocv}}(A, B), \mu_{\text{ocv}, \pi})$ where the unital $*$ -algebra $\text{Mor}^{\text{ocv}}(A, B)$ is generated by $\Phi_{R, \mathbf{a}, \mathbf{b}}$ for $R \in \sigma$, $\mathbf{a} \in R^A$, and $\mathbf{b} \in R^B$ and $p_{a, b}$ for $a \in A$ and $b \in B$, satisfying the relations

$$\begin{aligned} \Phi_{R, \mathbf{a}, \mathbf{b}} &= \Phi_{R, \mathbf{a}, \mathbf{b}}^* & \Phi_{R, \mathbf{a}, \mathbf{b}} \Phi_{R, \mathbf{a}, \mathbf{b}'} &= \delta_{\mathbf{b}, \mathbf{b}'} \Phi_{R, \mathbf{a}, \mathbf{b}} & \sum_{\mathbf{b} \in R^B} \Phi_{R, \mathbf{a}, \mathbf{b}} &= 1 \\ p_{a, b} &= p_{a, b}^* & p_{a, b} p_{a, b'} &= \delta_{b, b'} p_{a, b} & \sum_{b \in B} p_{a, b} &= 1; \end{aligned}$$

and the weight function

$$\mu_{\text{ocv}, \pi}(\Phi_{R, \mathbf{a}, \mathbf{b}}(1 - p_{a_i, b_i})) = \frac{\pi(R, \mathbf{a})}{\text{ar}(R)}$$

for all $R \in \sigma$, $\mathbf{a} \in R^A$, $\mathbf{b} \in R^B$, and $i \in [\text{ar}(R)]$, and $\mu_{\text{ocv}, \pi}(x) = 0$ for all other elements $x \in \text{Mor}^{\text{ocv}}(A, B)$. Denote the defect polynomial of $\text{Mor}_\pi^{\text{ocv}}(A, B)$ as $D_\pi^{\text{ocv}}(A, B) = D(\text{Mor}_\pi^{\text{ocv}}(A, B), \mu_{\text{ocv}, \pi})$.

Let π be a probability distribution on the constraints of A . The **non-oracular algebra** $A \rightarrow B$ is the weighted algebra $\text{Mor}_\pi^{\text{no}}(A, B) = (\text{Mor}^{\text{no}}(A, B), \mu_{\text{no}, \pi})$ where the unital $*$ -algebra $\text{Mor}^{\text{no}}(A, B)$ is generated by $p_{a, b}$ for $a \in A$ and $b \in B$, satisfying the relations

$$p_{a, b} = p_{a, b}^* \quad p_{a, b} p_{a, b'} = \delta_{b, b'} p_{a, b} \quad \sum_{b \in B} p_{a, b} = 1;$$

and the weight function

$$\mu_{\text{no}, \pi}(p_{a_1, b_1} \cdots p_{a_{\text{ar}(R)}, b_{\text{ar}(R)}}) = \pi(R, \mathbf{a})$$

for all $R \in \sigma$, $\mathbf{a} \in R^A$, and $\mathbf{b} \notin R^B$, and $\mu_{\text{no}, \pi}(x) = 0$ for all other $x \in \text{Mor}^{\text{no}}(A, B)$. Denote the defect polynomial of $\text{Mor}_\pi^{\text{no}}(A, B)$ as $D_\pi^{\text{no}}(A, B) = D(\text{Mor}_\pi^{\text{no}}(A, B), \mu_{\text{no}, \pi})$.

It is direct to see that that the defining relations of $\text{Mor}^{\text{occ}}(A, B)$, $\text{Mor}^{\text{ocv}}(A, B)$, and $\text{Mor}^{\text{no}}(A, B)$ precisely say that $\{p_{a, b}\}_{b \in B}$ and $\{\Phi_{R, \mathbf{a}, \mathbf{b}}\}_{\mathbf{b} \in R^B}$ are PVMs. Hence, without the weighted algebra structure, these algebras are simply free products of finite-dimensional abelian $*$ -algebras. Endowing the algebras with the sum-of-squares cone, we see immediately that the generators are bounded, and hence they also admit a semi-pre- C^* -algebra structure.

Note that unlike the non-oracular nonlocal game, the non-oracular algebra is defined for any pair of σ -structures, not just those where the arity of all the relations is 2. Now, we show that tracial states on the weighted algebras correspond exactly to symmetric tracial strategies for the corresponding nonlocal games.

Lemma 3.1.7. Let A and B be σ -structures.

- (i) For any probability distribution π_0 on pairs of constraints of A , there exists a tracial state τ on $\text{Mor}_{\pi_0}^{\text{occ}}(A, B)$ such that $\text{def}(\tau) = \varepsilon$ if and only if there exists a symmetric tracial correlation p for $\mathbf{G}_{\text{occ}}(A, B, \pi_0)$ such that $\mathfrak{w}(\mathbf{G}_{\text{occ}}(A, B, \pi_0), p) = 1 - \varepsilon$.
- (ii) For any probability distribution π_1 on the constraints of A , there exists a tracial state τ on $\text{Mor}_{\pi_1}^{\text{ocv}}(A, B)$ such that $\text{def}(\tau) = \varepsilon$ if and only if there exists a tracial correlation p for $\mathbf{G}_{\text{ocv}}(A, B, \pi_1)$ such that $\mathfrak{w}(\mathbf{G}_{\text{ocv}}(A, B, \pi_1), p) = 1 - \varepsilon$.
- (iii) Suppose that the arity of all the relations of σ is 2. For any probability distribution π_2 on the constraints of A , there exists a tracial state τ on $\text{Mor}_{\pi_2}^{\text{no}}(A, B)$ such that $\text{def}(\tau) = \varepsilon$ if and only if there exists a symmetric tracial correlation p for $\mathbf{G}_{\text{no}}(A, B, \pi_2)$ such that $\mathfrak{w}(\mathbf{G}_{\text{no}}(A, B, \pi_2), p) = 1 - \varepsilon$.

In any of the cases, τ has a GNS representation with the same von Neumann algebra as in the model of p , so finite-dimensional tracial states correspond to quantum correlations, and Connes-embeddable tracial states correspond to quantum approximate tracial correlations.

Proof. For (i), let $p(\mathbf{b}, \mathbf{b}' | (R, \mathbf{a}), (S, \mathbf{a}')) = \tau(P_{\mathbf{b}}^{R, \mathbf{a}} P_{\mathbf{b}'}^{S, \mathbf{a}'})$ be a symmetric tracial correlation for $\mathbf{G}_{\text{occ}}(A, B, \pi_0)$. Then, define the $*$ -homomorphism $\varphi : \text{Mor}^{\text{occ}}(A, B) \rightarrow \mathcal{M}$ by $\varphi(\Phi_{R, \mathbf{a}, \mathbf{b}}) = P_{\mathbf{b}}^{R, \mathbf{a}}$. This is well-defined as $\{P_{\mathbf{b}}^{R, \mathbf{a}}\}_{\mathbf{b} \in R^B}$ is a PVM for all constraints. Now, we take the tracial state $\tau' = \tau \circ \varphi$, giving

$$\begin{aligned}
\text{def}(\tau') &= \sum_{\substack{R, S \in \sigma, \mathbf{a} \in R^A, \mathbf{a}' \in S^A \\ \mathbf{b} \in R^B, \mathbf{b}' \in S^B : a_i = a'_j \neq b_i = b'_j}} \pi_0((R, \mathbf{a}), (S, \mathbf{a}')) \|\Phi_{R, \mathbf{a}, \mathbf{b}} \Phi_{S, \mathbf{a}', \mathbf{b}'}\|_{\tau'}^2 \\
&= \sum_{\substack{R, S \in \sigma, \mathbf{a} \in R^A, \mathbf{a}' \in S^A \\ \mathbf{b} \in R^B, \mathbf{b}' \in S^B}} \pi_0((R, \mathbf{a}), (S, \mathbf{a}')) (1 - V(\mathbf{b}, \mathbf{b}' | (R, \mathbf{a}), (S, \mathbf{a}'))) \tau(P_{\mathbf{b}}^{R, \mathbf{a}} P_{\mathbf{b}'}^{S, \mathbf{a}'}) \\
&= 1 - \mathfrak{w}(\mathbf{G}_{\text{occ}}(A, B, \pi_0), p).
\end{aligned}$$

For the converse, for a tracial state τ on $\text{Mor}^{\text{occ}}(A, B)$, let $\tau(x) = \langle \psi | \varphi(x) | \psi \rangle$ be the (minimal) GNS representation. Let $\mathcal{M} = \overline{\text{im } \varphi}^{\text{SOT}}$ — by construction, the map $\tau'(T) = \langle \psi | T | \psi \rangle$ is a normal tracial state on $\mathcal{M}$. With this, we take $P_{\mathbf{b}}^{R, \mathbf{a}} = \varphi(\Phi_{R, \mathbf{a}, \mathbf{b}})$ to get the correlation $p(\mathbf{b}, \mathbf{b}' | (R, \mathbf{a}), (S, \mathbf{a}')) = \tau'(P_{\mathbf{b}}^{R, \mathbf{a}} P_{\mathbf{b}'}^{S, \mathbf{a}'})$. By the same argument as above, $\text{def}(\tau) = 1 - \mathfrak{w}(\mathbf{G}_{\text{occ}}(A, B, \pi_0), p)$.

For (ii), let $p(\mathbf{b}, \mathbf{b} | (R, \mathbf{a}), a) = \tau(P_{\mathbf{b}}^{R, \mathbf{a}} Q_a^a)$ be a tracial correlation for $\mathbf{G}_{\text{ocv}}(A, B, \pi_1)$. Then, define the $*$ -homomorphism $\varphi : \text{Mor}^{\text{ocv}}(A, B) \rightarrow \mathcal{M}$ by $\varphi(\Phi_{R, \mathbf{a}, \mathbf{b}}) = P_{\mathbf{b}}^{R, \mathbf{a}}$ and $\varphi(p_{a, b}) = Q_b^a$.

This is well-defined as $\{P_b^{R,a}\}_{b \in R^B}$ and $\{Q_b^a\}_{b \in B}$ are PVMs. Now, we take the tracial state $\tau' = \tau \circ \varphi$, giving

$$\begin{aligned} \text{def}(\tau') &= \sum_{\substack{R \in \sigma, a \in R^A \\ b \in R^B, i \in [\text{ar}(R)]}} \frac{\pi_1(R, \mathbf{a})}{\text{ar}(R)} \|\Phi_{R,a,b}(1 - p_{a_i, b_i})\|_{\tau'}^2 \\ &= \sum_{\substack{R \in \sigma, a \in R^A \\ b \in R^B, i \in [\text{ar}(R)]}} \frac{\pi_1(R, \mathbf{a})}{\text{ar}(R)} \tau(P_b^{R,a}(1 - Q_{b_i}^{a_i})) \\ &= 1 - \mathfrak{w}(\mathbf{G}_{ocv}(A, B, \pi_1), p). \end{aligned}$$

For the converse, we proceed exactly as in (i).

For (iii), let $p(b, b'|a, a') = \tau(P_b^a P_{b'}^{a'})$ be a symmetric tracial correlation for $\mathbf{G}_{no}(A, B, \pi_2)$. Then, define the $*$ -homomorphism $\varphi : \text{Mor}^{no}(A, B) \rightarrow \mathcal{M}$ by $\varphi(p_{a,b}) = P_b^a$. This is well-defined as $\{P_b^a\}_{b \in B}$ is a PVM for all $a \in A$. Now, we take the tracial state $\tau' = \tau \circ \varphi$, giving

$$\begin{aligned} \text{def}(\tau') &= \sum_{\substack{R \in \sigma, (a, a') \in R^A \\ (b, b') \notin R^B}} \pi_2(R, (a, a')) \|p_{a,b} p_{a',b'}\|_{\tau'}^2 \\ &= 1 - \sum_{\substack{R \in \sigma, (a, a') \in R^A \\ (b, b') \in R^B}} \pi_2(R, (a, a')) \tau(P_b^a P_{b'}^{a'}) \\ &= 1 - \mathfrak{w}(\mathbf{G}_{no}(A, B, \pi_2), p). \end{aligned}$$

For the converse, we proceed exactly as in (i) and (ii).

To finish, note that the constructions in all three cases preserve the underlying von Neumann algebra. ■

3.1.3 Definitions of entangled CSPs

As in the classical case, homomorphisms of relation structures give rise to decision problems. We study a few different types of problems defined in this way, which we consider in aggregate as **CSP promise problems**. We present these problems in the weighted algebra model, since this allows us to capture non-oracular 2-CSPs as well, but they can be equivalently be expressed in the nonlocal game model, where this is defined, using [Lemma 3.1.7](#), which we will discuss at the end of this section.

Definition 3.1.8. Let Q be a set of von Neumann algebras. We say a state ρ on a $*$ -algebra $\mathcal{A}$ is **over** Q if there exists a representation $\pi : \mathcal{A} \rightarrow \mathcal{B}(H)$ and a unit vector $|\psi\rangle \in H$ such that $\rho(x) = \langle \psi | x | \psi \rangle$ and $\text{im } \pi$ embeds into an element of Q .

We make use of the following sets of von Neumann algebras, which correspond to the classes of strategy indicated by the same symbol:

- **Deterministic:** $d = \{\mathbb{C}\}$
- **Classical:** $c = \{C(X) \mid X \text{ compact space}\}$
- **Quantum:** $q = \{\mathbb{M}_n(\mathbb{C}) \mid n \in \mathbb{N}\}$
- **Quantum approximate:** $qa = \{\mathcal{M} \mid \mathcal{M} \hookrightarrow \mathcal{R}^\omega\}$
- **Commuting operator:** $co = \{\mathcal{M} \mid \exists \text{ normal tracial state } \tau : \mathcal{M} \rightarrow \mathbb{C}\}$

Definition 3.1.9. Let A be a σ -structure, let Q be a set of von Neumann algebras, let $1 \geq c \geq s \geq 0$ be completeness and soundness parameters, and let $w \in \{occ, ocv, no\}$. The **Q -entangled (c, s) -CSP of A** is the promise problem $\text{CSP}_w^Q(A)_{c,s}$ with instances (X, π) where X is a σ -structure and π is a probability distribution on the constraints of X (or a connected probability distribution on pairs of constraints if $w = occ$), for which (X, π) is a yes instance if $\inf_\tau \text{def}(\tau) \leq 1 - c$, where the infimum is over tracial states on $\text{Mor}_\pi^w(X, A)$ over Q , and (X, π) is a no instance if $\inf_\tau \text{def}(\tau) > 1 - s$, where the infimum is over the same set of tracial states.

The **exact Q -entangled (c, s) -CSP of A** is the promise problem ${}^e\text{CSP}_w^Q(A)_{c,s}$ with instances (X, π) where X is a σ -structure and π is a probability distribution on the constraints of X (or a connected probability distribution on pairs of constraints if $w = occ$), for which (X, π) is a yes instance if there exists a tracial state τ on $\text{Mor}_\pi^w(X, A)$ over Q such that $\text{def}(\tau) \leq 1 - c$, and X is a no instance if $\text{def}(\tau) > 1 - s$, for any tracial state τ on $\text{Mor}_\pi^w(X, A)$ over Q .

For $w = no, ocv$, we can reduce to the case where π is uniform by repeating constraints. However, we cannot do this in the constraint-constraint case as the probability distribution may be more structured.

The additional promise that the probability distribution is connected ([Definition 3.1.4](#)) in the case $w = occ$ is necessary to allow all the local realisations of a variable to be compared. It is possible to define the problem without this promise, but this would give a

more tenuous connection to the other notions of entangled CSPs. Also, this only becomes a significant issue in the case of the succinct versions introduced below; for the case considered above, the variables can always be relabelled to go from a non-connected instance to a connected instance.

For the sets of von Neumann algebras above, we get some equivalences between the problems, notably

$$\begin{aligned}\text{CSP}_w^d(A)_{c,s} &= \text{CSP}_w^c(A)_{c,s} = {}^e\text{CSP}_w^d(A)_{c,s} = {}^e\text{CSP}_w^c(A)_{c,s} \\ \text{CSP}_w^q(A)_{c,s} &= \text{CSP}_w^{qa}(A)_{c,s} = {}^e\text{CSP}_w^{qa}(A)_{c,s} \\ \text{CSP}_w^{co}(A)_{c,s} &= {}^e\text{CSP}_w^{co}(A)_{c,s}.\end{aligned}$$

Deterministic and classical CSPs are equivalent since every trace over c is a convex combination of traces over d , which are characters. Quantum and quantum approximate CSPs are equivalent as every trace over qa is an ultralimit of traces over q . The remaining equivalences between CSPs and their exact variants are because the sets Q are closed under ultraproducts, and hence there exists a tracial state that attains the infimum. These last two points are shown in [Section 3.1.4](#). Note however that the problem ${}^e\text{CSP}_w^q(A)_{c,s}$ is distinct.

We also consider the succinct versions of the constraint satisfaction problems.

Definition 3.1.10. Let A be a σ -structure, let Q be a set of von Neumann algebras, let $1 \geq c \geq s \geq 0$ be completeness and soundness parameters, and let $w \in \{occ, ocv, no\}$. The **succinct Q -entangled (c, s) -CSP of A** is the promise problem $\text{MIP}_w^Q(A)_{c,s}$ with instances M that are Turing machines that sample constraints from a σ -structure $X(M)$ according to a probability distribution π_M (or pairs of constraints from a connected probability distribution if $w = occ$), for which M is a yes instance if $\inf_\tau \text{def}(\tau) \leq 1 - c$, where the infimum is over tracial states on $\text{Mor}_{\pi_M}^w(X(M), A)$ over Q , and M is a no instance if $\inf_\tau \text{def}(\tau) > 1 - s$, where the infimum is over the same set of tracial states.

The **succinct exact Q -entangled (c, s) -CSP of A** is the promise problem ${}^e\text{MIP}_w^Q(A)_{c,s}$ with instances M that are Turing machines that sample constraints from a σ -structure $X(M)$ according to a probability distribution π_M (or pairs of constraints from a connected probability distribution if $w = occ$), for which M is a yes instance if there exists a tracial state τ on $\text{Mor}_{\pi_M}^w(X(M), A)$ over Q such that $\text{def}(\tau) \leq 1 - c$, and M is a no instance if $\text{def}(\tau) > 1 - s$, for any tracial state τ on $\text{Mor}_{\pi_M}^w(X(M), A)$ over Q .

For these problems, we find the same equivalences as in the non-succinct case:

$$\begin{aligned}\text{MIP}_w^d(A)_{c,s} &= \text{MIP}_w^c(A)_{c,s} = {}^e\text{MIP}_w^d(A)_{c,s} = {}^e\text{MIP}_w^c(A)_{c,s} \\ \text{MIP}_w^q(A)_{c,s} &= \text{MIP}_w^{qa}(A)_{c,s} = {}^e\text{MIP}_w^{qa}(A)_{c,s} \\ \text{MIP}_w^{co}(A)_{c,s} &= {}^e\text{MIP}_w^{co}(A)_{c,s}.\end{aligned}$$

In the classical setting, this connects to CSPs as introduced in [Section 2.10.3](#). We have that $\text{MIP}_{c,s}(A) = \text{MIP}_{no}^d(A)_{c,s}$; hence the notation MIP for succinct CSPs is due to the relationship with multiprover interactive proofs.

Using [Lemma 3.1.7](#), the problems $\text{CSP}_w^Q(A)_{c,s}$, ${}^e\text{CSP}_w^Q(A)_{c,s}$, $\text{MIP}_w^Q(A)_{c,s}$, and ${}^e\text{MIP}_w^Q(A)_{c,s}$ can equivalently be expressed in terms of tracial correlations for the corresponding nonlocal game. In this setting, it is natural to ask what happens when we extend the space of correlations from tracial correlations over Q to all correlations over Q . This gives rise to game versions of these problems.

Definition 3.1.11. Let Q be a set of von Neumann algebras. We say a commuting-operator correlation $p(a, b|x, y) = \langle \psi | P_a^x Q_b^y | \psi \rangle$ is **over** Q if both the von Neumann algebras generated by the $\{P_a^x\}_{x,a}$ and $\{Q_b^y\}_{y,b}$ embed into elements of Q .

Definition 3.1.12. Let A be a σ -structure, let Q be a set of von Neumann algebras, let $1 \geq c \geq s \geq 0$ be completeness and soundness parameters, and let $w \in \{occ, ocv, no\}$. If $w = no$, we suppose that all the relational symbols in σ have arity 2.

The **Q -entangled game (c, s) -CSP of A** is the promise problem ${}_g\text{CSP}_w^Q(A)_{c,s}$ with instances (X, π) where X is a σ -structure and π is a probability distribution on the constraints of X (or a connected probability distribution on pairs of constraints if $w = occ$), for which (X, π) is a yes instance if $\sup_p \mathfrak{w}(\mathbb{G}_w(X, A, \pi), p) \geq c$, where the supremum is over correlations for $\mathbb{G}_w(X, A, \pi)$ over Q , and (X, π) is a no instance if $\sup_p \mathfrak{w}(\mathbb{G}_w(X, A, \pi), p) < s$, where the supremum is over the same set of correlations.

The **exact Q -entangled game (c, s) -CSP of A** is the promise problem ${}^e{}_g\text{CSP}_w^Q(A)_{c,s}$ with instances (X, π) where X is a σ -structure and π is a probability distribution on the constraints of X (or a connected distribution on pairs of constraints if $w = occ$), for which (X, π) is a yes instance if there exists a correlation p for $\mathbb{G}_w(X, A, \pi)$ over Q such that $\mathfrak{w}(\mathbb{G}_w(X, A, \pi), p) \geq c$, and (X, π) is a no instance if $\mathfrak{w}(\mathbb{G}_w(X, A, \pi), p) < s$, for any correlation p for $\mathbb{G}_w(X, A, \pi)$ over Q .

The **succinct Q -entangled game (c, s) -CSP of A** ${}_g\text{MIP}_w^Q(A)_{c,s}$ and the **succinct exact Q -entangled game (c, s) -CSP of A** ${}^e{}_g\text{MIP}_w^Q(A)_{c,s}$ are defined analogously as in [Definition 3.1.10](#).

For CSPs with perfect completeness $c = 1$ (or near-perfect completeness), we can reduce the CSP to the game CSP considered earlier in the constraint-constraint and non-oracular settings via the result of [\[117, 86\]](#). We will also show an analogous result for constraint-variable games in [Section 3.3](#). As such, hardness results for CSPs immediately imply hardness results for game CSPs, though potentially with a much smaller (but still

constant) completeness-soundness gap. Hence, in this context, we do not consider game CSPs independently.

Lemma 3.1.13. Let A be a σ -structure, let $0 \leq s \leq 1$, and let $Q \in \{q, co\}$. Then, there exists $0 \leq s' \leq 1$ such that $s' = 1 \iff s = 1$ for which there are polynomial time reductions from $\text{CSP}_{occ}^Q(A)_{1,s}$ to ${}_g\text{CSP}_{occ}^Q(A)_{1,s'}$, ${}^e\text{CSP}_{occ}^Q(A)_{1,s}$ to ${}^e{}_g\text{CSP}_{occ}^Q(A)_{1,s'}$, $\text{MIP}_{occ}^Q(A)_{1,s}$ to ${}_g\text{MIP}_{occ}^Q(A)_{1,s'}$, and ${}^e\text{MIP}_{occ}^Q(A)_{1,s}$ to ${}^e{}_g\text{MIP}_{occ}^Q(A)_{1,s'}$.

Now suppose also that $\text{ar}(R) \leq 2$ for all $R \in \sigma$ and the equality relation $\{(a, a) \mid a \in A\}$ is one of the relations of A . Then, there exists $0 \leq s' \leq 1$ such that $s' = 1 \iff s = 1$ for which there are polynomial time reductions from $\text{CSP}_{no}^Q(A)_{1,s}$ to ${}_g\text{CSP}_{no}^Q(A)_{1,s'}$, ${}^e\text{CSP}_{no}^Q(A)_{1,s}$ to ${}^e{}_g\text{CSP}_{no}^Q(A)_{1,s'}$, $\text{MIP}_{no}^Q(A)_{1,s}$ to ${}_g\text{MIP}_{no}^Q(A)_{1,s'}$, and ${}^e\text{MIP}_{no}^Q(A)_{1,s}$ to ${}^e{}_g\text{MIP}_{no}^Q(A)_{1,s'}$.

Proof. Let $s' = 1 - \left(\frac{1-s}{2c}\right)^4$, where c is the universal constant from [86]. For an instance of one of the CSPs, we use symmetric synchronisation to modify the probability distribution to be symmetric and 1-synchronous, and take this as the instance of the corresponding game CSP. The probability distribution is made symmetric by exchanging the roles of Alice and Bob with probability $\frac{1}{2}$; then it is made 1-synchronous by sending both players Alice's question with probability $\frac{1}{2}$. This is a polynomial-time procedure for succinct games. If the tracial value of the original game is 1, the tracial value remains 1; if the tracial value is $< s$, then the tracial value becomes $< \frac{1+s}{2}$, as symmetric tracial strategies always succeed on consistency check questions. Now, if the tracial value was 1, the value must also be 1, either with $Q = q$ or $Q = co$. If the tracial value is $< \frac{1+s}{2} = 1 - \frac{1-s}{2}$, then, by [86], the value must be $< s'$ by definition of s' . Hence we have the wanted reduction. ■

To finish this section, note that the CSP promise problems that can actually be verified by a multiprover interactive proof are ${}_g\text{CSP}_w^Q(A)_{c,s}$ and ${}_g\text{MIP}_w^Q(A)_{c,s}$, with $1 \geq c > s \geq 0$ and $Q \in \{d, q, co\}$ depending on the model of entanglement.

3.1.4 Some equivalences based on stability

We first need some general stability results that will be useful throughout.

Lemma 3.1.14. Let $\mathcal{M}$ be a finite von Neumann algebra with a faithful tracial state tr . Then,

- (i) For any $X \in \mathcal{M}$, there exists a unitary U such that $\|X - U\|_{\text{tr}} \leq \|X^*X - 1\|_{\text{tr}}$.

- (ii) For any unitary $U \in \mathcal{M}$, there exists an order- d unitary W such that $\|U - W\|_{\text{tr}} \leq \|U^d - 1\|_{\text{tr}}$.
- (iii) For any $X \in \mathcal{M}$, there exists an order- d unitary W such that $\|X - W\|_{\text{tr}} \leq d(\|X\|^d + 2)\|X^*X - 1\|_{\text{tr}} + \|X^d - 1\|_{\text{tr}}$.

Proof. (i) Let $X = UP$ be the polar decomposition, and let $P = \int_0^\infty \lambda dE(\lambda)$ be the spectral decomposition. Since $\mathcal{M}$ is finite, we may choose U to be unitary. Then, $\|X - U\|_{\text{tr}} = \|P - 1\|_{\text{tr}}$ by unitary invariance. For any $x \in \mathbb{R}_{\geq 0}$, we have that $(x^2 - 1)^2 = (x + 1)^2(x - 1)^2 \geq (x - 1)^2$. Noting that the spectral decomposition of $P - 1 = \int(\lambda - 1)dE(\lambda)$ and writing μ for the spectral measure $\text{tr} \circ E$, we get via the measurable functional calculus that

$$\|P - 1\|_{\text{tr}}^2 = \int (\lambda - 1)^2 d\mu(\lambda) \leq \int (\lambda^2 - 1)^2 d\mu(\lambda) = \|P^2 - 1\|_{\text{tr}}^2.$$

Using unitary invariance again gives $\|X - U\|_{\text{tr}}^2 \leq \|X^*X - 1\|_{\text{tr}}^2$.

- (ii) Let $U = \int_{S^1} \lambda dE(\lambda)$ be the spectral decomposition of U . We can write any $\lambda \in S^1$ uniquely as $\lambda = \omega^{a(\lambda)} e^{i\theta(\lambda)}$ where ω is a primitive d -th root of unity, $a(\lambda) \in \{0, \dots, d-1\}$, and $\theta(\lambda) \in (-\pi/d, \pi/d]$. Note that a and θ are piecewise continuous functions, so they are measurable. Let $W = \int_{S^1} \omega^{a(\lambda)} dE(\lambda)$. Also, note that for any $\theta \in (-\pi/d, \pi/d]$, $|e^{id\theta} - 1| = |e^{i\theta} - 1| |1 + e^{i\theta} + \dots + e^{i(d-1)\theta}| = |e^{i\theta} - 1| \left| \frac{\sin(d\theta/2)}{\sin(\theta/2)} \right| \geq |e^{i\theta} - 1|$. Therefore, writing again the spectral measure as $\mu = \text{tr} \circ E$,

$$\begin{aligned} \|U - W\|_{\text{tr}}^2 &= \int |\lambda - \omega^{a(\lambda)}|^2 d\mu(\lambda) = \int |e^{i\theta(\lambda)} - 1|^2 d\mu(\lambda) \\ &\leq \int |e^{id\theta(\lambda)} - 1|^2 d\mu(\lambda) = \|U^d - 1\|_{\text{tr}}^2. \end{aligned}$$

- (iii) Using (i) and (ii), let U be the nearest unitary to X , and let W be the nearest order- d unitary to U . Then,

$$\begin{aligned} \|X - W\|_{\text{tr}} &\leq \|X - U\|_{\text{tr}} + \|U - W\|_{\text{tr}} \\ &\leq \|X^*X - 1\|_{\text{tr}} + \|U^d - 1\|_{\text{tr}} \\ &\leq \|X^*X - 1\|_{\text{tr}} + \|X^d - 1\|_{\text{tr}} + \|U^d - X^d\|_{\text{tr}}. \end{aligned}$$

To finish, note that

$$\begin{aligned} \|U^d - X^d\|_{\text{tr}} &\leq \|U^d - U^{d-1}X\|_{\text{tr}} + \|U^{d-1}X - U^{d-2}X^2\|_{\text{tr}} + \dots + \|UX^{d-1} - X^d\|_{\text{tr}} \\ &\leq (\|X\|^{d-1} + \dots + \|X\| + 1)\|U - X\|_{\text{tr}}, \end{aligned}$$

and then that $\|U - X\|_{\text{tr}} \leq \|X^*X - 1\|_{\text{tr}}$ and $\|X\|^{d-1} + \dots + \|X\| + 1 \leq d \max\{\|X\|^{d-1}, 1\} \leq d(\|X\|^d + 1)$. $\blacksquare$

Lemma 3.1.15. Let $(\mathcal{A}, \mu)$ be a weighted algebra generated as the free product of finite-order unitaries $u_1, \dots, u_N$. Let Q be a set of finite von Neumann algebras, and write $\overline{Q}$ for the closure of Q under ultraproducts. Then, the infimum of the defects of tracial states over Q is $\leq \delta$ iff there exists a tracial state τ over $\overline{Q}$ such that $\text{def}(\tau) \leq \delta$.

Note that $\bar{d} = d$, $\bar{c} = c$, $\bar{q} = \overline{qa} = qa$, and $\bar{co} = co$.

Proof. Write d_i for the order of the generator u_i . First, suppose that the infimum of the defects of tracial states over Q is $\leq \delta$. Then, there exists a sequence of tracial states $(\tau_n)_n$ over Q such that $\lim_{n \rightarrow \infty} \text{def}(\tau_n) \leq \delta$. Now, let $\varphi_n : \mathcal{A} \rightarrow (\mathcal{M}_n, \text{tr}_n)$ be GNS representations of the τ_n , for some von Neumann algebras $\mathcal{M}_n \in Q$ and tracial states tr_n on $\mathcal{M}_n$. Then, for any free ultrafilter ω , let $(\mathcal{M}, \text{tr}) = \prod_{\omega} (\mathcal{M}_n, \text{tr}_n) \in \overline{Q}$, so we get a $*$ -homomorphism $\varphi = \prod_{\omega} \varphi_n : \mathcal{A} \rightarrow \mathcal{M}$. Let $\tau = \text{tr} \circ \varphi$. Then, τ is a tracial state over $\overline{Q}$ such that $\text{def}(\tau) = \lim_{n \rightarrow \omega} \text{def}(\tau_n) \leq \delta$.

For the converse, suppose that τ is a tracial state on $\mathcal{A}$ over $\overline{Q}$ such that $\text{def}(\tau) \leq \delta$. Let $\varphi : \mathcal{A} \rightarrow (\mathcal{M}, \text{tr})$ be GNS representation of the τ , for some von Neumann algebra $\mathcal{M} \in \overline{Q}$ and tracial state tr on $\mathcal{M}$. By construction, there exists an ultrafilter ω and a sequence of von Neumann algebras $\mathcal{M}_n \in Q$ and tracial states tr_n on $\mathcal{M}_n$ such that $(\mathcal{M}, \text{tr}) = \prod_{\omega} (\mathcal{M}_n, \text{tr}_n)$. Without loss of generality, we may assume the tracial states are faithful. Because φ is a $*$ -homomorphism, it can be realised as a sequence of maps $\varphi_n : \mathcal{A} \rightarrow \mathcal{M}_n$ satisfying $\lim_{n \rightarrow \omega} \|\varphi_n(a)\varphi_n(b) - \varphi_n(ab)\|_{\text{tr}_n}^2 = 0$ and $\lim_{n \rightarrow \omega} \|\varphi_n(a^*) - \varphi_n(a)^*\|_{\text{tr}_n}^2 = 0$, for all $a, b \in \mathcal{A}$. Let $U_{n,i}$ be the nearest order- d_i unitary to $\varphi_n(u_i)$. By Lemma 3.1.14, we know that $\|\varphi_n(u_i) - U_{n,i}\|_{\text{tr}_n}^2 \xrightarrow{\omega} 0$. Define the $*$ -representation $\varphi'_n : \mathcal{A} \rightarrow \mathcal{M}_n$ on the generators by $\varphi'_n(u_i) = U_{n,i}$. By construction, we have that for any $a \in \mathcal{A}$, $\|\varphi'_n(a) - \varphi_n(a)\|_{\text{tr}_n}^2 \xrightarrow{\omega} 0$. Let $\tau_n = \text{tr}_n \circ \varphi'_n$; this is a tracial state over Q for all n . We have immediately that $\lim_{n \rightarrow \omega} \text{def}(\tau_n) = \text{def}(\tau) \leq \delta$. Then, the infimum of the defects of tracial states over Q is $\leq \delta$. $\blacksquare$

Proposition 3.1.16. Let A be a relational structure, let $w \in \{\text{occ}, \text{ocv}, \text{no}\}$, let Q be a set of finite von Neumann algebras, and $\overline{Q}$ be the closure of Q under ultraproducts. Then, ${}^e\text{CSP}_{\overline{w}}^{\overline{Q}}(A)_{c,s} = \text{CSP}_w^Q(A)_{c,s}$ and ${}^e\text{MIP}_{\overline{w}}^{\overline{Q}}(A)_{c,s} = \text{MIP}_w^Q(A)_{c,s}$.

The proof follows immediately from Lemma 3.1.15.

3.1.5 Complexity of entangled CSPs

In this section, we show basic undecidability results for entangled CSPs. All our subsequent complexity statements will be based on reductions from these results.

Theorem 3.1.17. For $w \in \{ocv, no\}$, there exists a relational structure A and $s < 1$ such that

- (i) $\text{MIP}_w^q(A)_{1,s}$ and ${}^e\text{MIP}_w^q(A)_{1,s}$ are RE-complete via polynomial-time reductions from the halting problem.
- (ii) $\text{MIP}_w^{co}(A)_{1,s}$ is coRE-complete via a polynomial-time reduction from the cohalting problem.

Since the runtime of reductions is immaterial in RE and coRE, this gives analogous completeness results for the non-succinct versions of these problems, albeit without the nice polynomial-time reduction from the halting/cohalting problem. Note also that the theorem implies hardness for these classes in the gapless ($s = 1$), but not necessarily completeness in (i).

This theorem relies on the powerful undecidability results known for nonlocal games.

Theorem 3.1.18 ([60]). There exists a constant ℓ_a such that $\text{MIP}_{1,1/2}^*(\text{poly}, \ell_a) = \text{RE}$ via a polynomial-time reduction from the halting problem. Further, the instances are synchronous nonlocal games such that, for a yes instance, there exists an oracularisable and finite-dimensional perfect strategy.

This is [60, Theorem 6.11]. Note that this also implies the containment of RE in the gapless analogue $\text{MIP}_{1,1}^*(\text{poly}, \ell_a)$, however not equality. In fact, we know that $\text{MIP}_{1,1}^*(\text{poly}, \text{poly}) = \Pi_2^0$, strictly larger than RE [89]. However, it is not obvious how to decrease the answer length here to constant, as the answer reduction protocol from [60] relies on the nonzero gap.

Theorem 3.1.19 ([83]). There exists a constant ℓ_a such that $\text{MIP}_{1,1/2}^{co}(\text{poly}, \ell_a) = \text{coRE}$ via a polynomial-time reduction from the cohalting problem. Further, the instances are synchronous nonlocal games such that, for a yes instance, there exists an oracularisable perfect strategy.

This theorem follows by applying the answer reduction of [60] to the main result of [83]. Note that unlike in the quantum case, we preserve the equality $\text{MIP}_{1,1}^{co}(\text{poly}, \ell_a) = \text{coRE}$ when passing to the zero-gap setting [110].

Proof of Theorem 3.1.17. Let A be the relational structure with alphabet $[\ell_a]$ and binary relations $R \subseteq [\ell_a] \times [\ell_a]$ for every subset. Take the powerset $\sigma = \mathcal{P}([\ell_a]^2)$ for the signature of A , where $R^A = R$. Now, given a synchronous nonlocal game $\mathbf{G} = (X, X, [\ell_a], [\ell_a], \mu, V)$, define the σ -structure $Y(\mathbf{G})$ with domain $\mathfrak{D}(Y(\mathbf{G})) = X$ and for $R \subseteq [\ell_a]^2$, let $R^{Y(\mathbf{G})} = \{(x, y) \in X^2 \mid R = \{(a, b) \in A^2 \mid V(a, b|x, y) = 1\}\}$. Since every pair $(x, y) \in X^2$ belongs to at most one constraint $Y(\mathbf{G})$, we can see μ as a probability distribution on the constraints of $Y(\mathbf{G})$. Then, it is clear to see that $\mathbf{G}_{no}(Y(\mathbf{G}), A, \mu) = \mathbf{G}$ and $\mathbf{G}_{ocv}(Y(\mathbf{G}), A, \mu) = \mathbf{G}^{\text{orac}}$; and that there is a polynomial time mapping from a probabilistic Turing machine that samples question pairs of $\mathbf{G}$ to a probabilistic Turing machine that samples the constraints of $Y(\mathbf{G})$. Hence, composing the mapping $\mathbf{G} \mapsto Y(\mathbf{G})$ with Theorem 3.1.18 gives (i). Note that hardness of the exact CSP is also achieved since the perfect strategy is finite-dimensional. Identically, composing the mapping $\mathbf{G} \mapsto Y(\mathbf{G})$ with Theorem 3.1.19 gives (ii), by the oracularisability of the perfect strategies in the yes case. ■

3.2 Quantum relational structure homomorphisms

In this section, we show that perfect strategies for CSP nonlocal games give rise to generalised homomorphisms of relational structures. In Section 3.2.1, we introduce the algebras whose representations correspond to perfect strategies of CSP nonlocal games. In Section 3.2.2, we introduce a category-theoretic framework that captures these homomorphism algebras: the cocategory. In Section 3.2.3, we define the cocategory arising from relational structure homomorphisms.

3.2.1 The structure of perfect strategies for CSP nonlocal games

We can characterise perfect strategies for CSP nonlocal games via representations of an algebra with more structure than the weighted algebras discussed in the previous section.

Definition 3.2.1. Let A and B be σ -structures. The **homomorphism algebra** $A \rightarrow B$ is the C^* -algebra $\text{Mor}^+(A, B)$ generated by $p_{a,b}$ for $a \in A$ and $b \in B$, subject to the relations

$$\begin{aligned} p_{a,b} &= p_{a,b}^* = p_{a,b}^2 \quad \forall a \in A, b \in B \\ \sum_{b \in B} p_{a,b} &= 1 \quad \forall a \in A \\ p_{a_1, b_1} \cdots p_{a_{\text{ar}(R)}, b_{\text{ar}(R)}} &= 0 \quad \forall R \in \sigma, \mathbf{a} \in R^A, \mathbf{b} \notin R^B. \end{aligned}$$

The **oracular homomorphism algebra** $A \rightarrow B$ is the quotient C^* -algebra $\text{Mor}^{o+}(A, B)$ of $\text{Mor}^+(A, B)$ by the ideal generated by the relations

$$[p_{a_i, b}, p_{a_j, b'}] = 0 \quad \forall R \in \sigma, \mathbf{a} \in R^A, i, j \in [\text{ar}(R)], b, b' \in B.$$

For $\mathbf{a} \in A^n$ and $\mathbf{b} \in B^n$, we write $p_{\mathbf{a}, \mathbf{b}} = p_{a_1, b_1} \cdots p_{a_n, b_n}$. In $\text{Mor}^{o+}(A, B)$, for $\mathbf{a} \in R^A$, $\{p_{\mathbf{a}, \mathbf{b}}\}_{\mathbf{b} \in R^B}$ forms a PVM.

We can also see homomorphism algebras in terms of quantum spaces. The **quantum space of homomorphisms** $A \rightarrow B$ is the quantum space $\text{mor}^+(A, B)$ with algebra of functions $C(\text{mor}^+(A, B)) = \text{Mor}^+(A, B)$. Identically, the **oracular quantum space of homomorphisms** $A \rightarrow B$ is the quantum space $\text{mor}^{o+}(A, B)$ with algebra of functions $C(\text{mor}^{o+}(A, B)) = \text{Mor}^{o+}(A, B)$.

Lemma 3.2.2. Let A and B be σ -structures.

- (i) For a connected probability distribution π_0 on pairs of constraints of A ,

$$C_u^*(\text{Mor}^{occ}(A, B) / \langle\langle \text{supp } \mu_{occ, \pi_0} \rangle\rangle) \cong \text{Mor}^{o+}(A, B).$$

- (ii) For a fully-supported probability distribution π_1 on the constraints of A ,

$$C_u^*(\text{Mor}^{ocv}(A, B) / \langle\langle \text{supp } \mu_{ocv, \pi_1} \rangle\rangle) \cong \text{Mor}^{o+}(A, B).$$

- (iii) For a fully-supported probability distribution π_2 on the constraints of A ,

$$C_u^*(\text{Mor}^{no}(A, B) / \langle\langle \text{supp } \mu_{no, \pi_2} \rangle\rangle) \cong \text{Mor}^+(A, B).$$

Proof. For (i), consider the $*$ -homomorphism $\varphi : \text{Mor}^{occ}(A, B) \rightarrow \text{Mor}^{o+}(A, B)$ defined on the generators as $\varphi(\Phi_{R, \mathbf{a}, \mathbf{b}}) = p_{\mathbf{a}, \mathbf{b}}$. As noted above, $\{\varphi(\Phi_{R, \mathbf{a}, \mathbf{b}})\}_{\mathbf{b} \in R^B}$ is a PVM. Hence the $*$ -homomorphism is well-defined. Next, note that if $x \in \text{supp } \mu_{occ, \pi_0}$, then $x = \Phi_{R, \mathbf{a}, \mathbf{b}} \Phi_{S, \mathbf{a}', \mathbf{b}'}$ such that $\pi_0((R, \mathbf{a}), (S, \mathbf{a}')) > 0$ and there exist $i \in [\text{ar}(R)]$ and $j \in [\text{ar}(S)]$ such that $a_i = a'_j$ but $b_i \neq b'_j$. Therefore,

$$\varphi(x) = p_{\mathbf{a}, \mathbf{b}} p_{\mathbf{a}', \mathbf{b}'} = p_{\mathbf{a}, \mathbf{b}} p_{a_i, b_i} p_{a'_j, b'_j} p_{\mathbf{a}', \mathbf{b}'} = 0.$$

As such, φ factors through to a $*$ -homomorphism $\text{Mor}^{occ}(A, B) / \langle\langle \text{supp } \mu_{occ, \pi_0} \rangle\rangle \rightarrow \text{Mor}^{o+}(A, B)$. By the universal property of the universal C^* -algebra [97], this induces a $*$ -homomorphism $\bar{\varphi} : C_u^*(\text{Mor}^{occ}(A, B) / \langle\langle \text{supp } \mu_{occ, \pi_0} \rangle\rangle) \rightarrow \text{Mor}^{o+}(A, B)$. Now, we need to show that $\bar{\varphi}$ is invertible. To do so, note first that if $\pi_0((R, \mathbf{a}), (S, \mathbf{a}')) > 0$ and $a_i = a'_j$, then $\sum_{\mathbf{b} \in R^B: b_i = b} \Phi_{R, \mathbf{a}, \mathbf{b}} =$

$\sum_{b' \in S^B: b'_j = b} \Phi_{S, a', b'}$ in $C_u^*(\text{Mor}^{occ}(A, B) / \langle\langle \text{supp } \mu_{occ, \pi_0} \rangle\rangle)$, since $\Phi_{R, a, b} \Phi_{S, a', b'} = 0$ whenever $b_i \neq b'_j$. Since π_0 is connected, this holds for every pair $\mathbf{a}, \mathbf{a}'$ such that $a_i = a'_j$. Let $\tilde{p}_{a, b} = \sum_{b \in R^B: b_i = b} \Phi_{R, a, b}$ for some $\mathbf{a}$ such that $a_i = a$. By the discussion above, this does not depend on the choice of $\mathbf{a}$. Since we have that $\Phi_{R, a, b} = \tilde{p}_{a_1, b_1} \cdots \tilde{p}_{a_{\text{ar}(R)}, b_{\text{ar}(R)}}$, the map $p_{a, b} \mapsto \tilde{p}_{a, b}$ extends to a $*$ -homomorphism $\text{Mor}^{o+}(A, B) \rightarrow C_u^*(\text{Mor}^{occ}(A, B) / \langle\langle \text{supp } \mu_{occ, \pi_0} \rangle\rangle)$ and is the inverse of $\bar{\varphi}$.

(ii) and (iii) are proved identically based on the $*$ -homomorphisms $\text{Mor}^{ocv}(A, B) \rightarrow \text{Mor}^{o+}(A, B)$ defined by $\Phi_{R, a, b} \mapsto p_{a, b}$ and $p_{a, b} \mapsto p_{a, b}$; and $\text{Mor}^{no}(A, B) \rightarrow \text{Mor}^+(A, B)$ defined by $p_{a, b} \mapsto p_{a, b}$, respectively. $\blacksquare$

3.2.2 Cocategories

Categories admit a quantisation similar to groups, introduced in [20]. We develop this framework here, and apply it to the context of quantum homomorphisms of relational structures in the following section. The power of using this model arises from the fact that properties in the cocategory setting carry over to dual versions of those properties in representations, which are quantum homomorphisms. This allows properties to be shown for many classes of quantum strategy at once, by showing that a dual version holds for the generating cocategory.

Definition 3.2.3. A C^* -**cocategory** $\mathcal{C}$ consists of a class of objects $\text{ob}(\mathcal{C})$ and for all pairs $a, b \in \text{ob}(\mathcal{C})$, a quantum space $\text{mor}_{\mathcal{C}}^+(a, b)$ – where we denote the C^* -algebra of functions $C(\text{mor}_{\mathcal{C}}^+(a, b)) = \text{Mor}_{\mathcal{C}}^+(a, b)$ and drop the subscript $\mathcal{C}$ when clear from context – such that for all $a, b, c \in \text{ob}(\mathcal{C})$, there exists a $*$ -homomorphism $\Delta_{a, c}^b : \text{Mor}^+(a, c) \rightarrow \text{Mor}^+(a, b) \otimes \text{Mor}^+(b, c)$, called **cocomposition**, satisfying the coassociativity conditions

$$(\Delta_{a, c}^b \otimes \text{id}_{\text{Mor}^+(b, d)}) \Delta_{a, d}^c = (\text{id}_{\text{Mor}^+(a, b)} \otimes \Delta_{b, d}^c) \Delta_{a, d}^b;$$

and for all $a \in \text{ob}(\mathcal{C})$, there exists a $*$ -homomorphism $\epsilon_a : \text{Mor}^+(a, a) \rightarrow \mathbb{C}$, called **coidentity**, satisfying the counit conditions

$$(\epsilon_a \otimes \text{id}_{\text{Mor}^+(a, b)}) \Delta_{a, b}^a = (\text{id}_{\text{Mor}^+(a, b)} \otimes \epsilon_b) \Delta_{a, b}^b = \text{id}_{\text{Mor}^+(a, b)}.$$

Every category $\mathcal{C}$ where $\text{mor}_{\mathcal{C}}(a, b)$ is a compact Hausdorff space is a C^* -cocategory, where we simply take $\text{mor}^+(a, b) = \text{mor}(a, b)$. Then, $\text{Mor}^+(a, b)$ is the space of continuous functions $\text{mor}(a, b) \rightarrow \mathbb{C}$, the cocomposition $\Delta_{a, c}^b : C(\text{mor}(a, c)) \rightarrow C(\text{mor}(a, b) \times \text{mor}(b, c))$ is $\Delta_{a, c}^b(\phi)(f, g) = \phi(g \circ f)$, and the coidentity $\epsilon_a : C(\text{mor}(a, a)) \rightarrow \mathbb{C}$ is $\epsilon_a(\phi) = \phi(\text{id}_A)$.

By passing to the space of representations of the algebras of functions of the morphism spaces, a cocategory induces a category. Different categories arise from representations into different classes of von Neumann algebra, which are well-defined if the class has the following structure.

Definition 3.2.4. We call a class of von Neumann algebras containing $\mathbb{C}$ and closed under the tensor product a **tensor monoid**.

Identifying $\mathbb{C} \otimes \mathcal{M}$ and $\mathcal{M} \otimes \mathbb{C}$ with $\mathcal{M}$ as usual gives a tensor monoid a monoid structure, where the product is given by the tensor product and the unit by $\mathbb{C}$. All of the relevant classes of strategy for entangled CSPs correspond to tensor monoids: the sets d , c , q , qa , and co introduced in [Section 3.1.3](#). It is easy to check that these are closed under the tensor product of von Neumann algebras. Further, we consider the tensor monoid of all von Neumann algebras vN . This does not correspond to tracial strategies, but does give rise to a distinct collection of representations.

Definition 3.2.5. Let $\mathcal{C}$ be a C^* -cocategory and let Q be a tensor monoid. The Q -**morphism category of $\mathcal{C}$** is the category $\mathcal{C}^Q$ with objects $\text{ob}(\mathcal{C}^Q) = \text{ob}(\mathcal{C})$ and morphisms

$$\text{mor}_{\mathcal{C}^Q}(a, b) = \left\{ \pi : \text{Mor}^+(a, b) \rightarrow \mathcal{M} \mid \pi \text{ is a } *- \text{homomorphism and } \mathcal{M} \in Q \right\},$$

with composition defined for $\pi \in \text{mor}_{\mathcal{C}^Q}(a, b)$ and $\varphi \in \text{mor}_{\mathcal{C}^Q}(b, c)$ as

$$\varphi \circ \pi = (\pi \otimes \varphi) \Delta_{a,c}^b.$$

When $\mathcal{C}$ is clear from context, we write $\text{mor}^Q(a, b) = \text{mor}_{\mathcal{C}^Q}(a, b)$. We write $\pi : a \xrightarrow{Q} b$ to mean $\pi \in \text{mor}^Q(a, b)$.

Lemma 3.2.6. Let $\mathcal{C}$ be a C^* -cocategory and Q be a tensor monoid. Then $\mathcal{C}^Q$ is well-defined as a category.

Proof. First, note that composition is well-defined. Let $\pi \in \text{mor}^Q(a, b)$ and $\varphi \in \text{mor}^Q(b, c)$. Then, there exist $\mathcal{A}, \mathcal{B} \in Q$ such that $\pi : \text{Mor}^+(a, b) \rightarrow \mathcal{A}$ and $\varphi : \text{Mor}^+(b, c) \rightarrow \mathcal{B}$ are $*$ -homomorphisms. As such, $\varphi \circ \pi : \text{Mor}^Q(a, c) \rightarrow \mathcal{A} \otimes \mathcal{B}$ is an element of $\text{mor}^Q(a, c)$.

Next, we need to show that the composition is associative. Let $\pi \in \text{mor}^Q(a, b)$, $\varphi \in$

$\text{mor}^Q(b, c)$, and $\psi \in \text{mor}^Q(c, d)$. Then, by coassociativity,

$$\begin{aligned}
\psi \circ (\varphi \circ \pi) &= ((\varphi \circ \pi) \otimes \psi) \Delta_{a,d}^c \\
&= (\pi \otimes \varphi \otimes \psi) (\Delta_{a,c}^b \otimes \text{id}) \Delta_{a,d}^c \\
&= (\pi \otimes \varphi \otimes \psi) (\text{id} \otimes \Delta_{b,d}^c) \Delta_{a,d}^b \\
&= (\pi \otimes (\psi \circ \varphi)) \Delta_{a,d}^b \\
&= (\psi \circ \varphi) \circ \pi.
\end{aligned}$$

Finally, we need to show the existence of an identity for each $a \in \text{ob}(\mathcal{C}^Q)$. We claim $\epsilon_a : \text{Mor}^+(a, a) \rightarrow \mathbb{C}$ is this element. In fact, for any $\pi \in \text{mor}^Q(a, b)$, we get by the counit property that

$$\begin{aligned}
\pi \circ \epsilon_a &= (\epsilon_a \otimes \pi) \Delta_{a,b}^a = \pi \\
\epsilon_b \circ \pi &= (\pi \otimes \epsilon_b) \Delta_{a,b}^b = \pi.
\end{aligned}$$

■

3.2.3 The cocategory of relational structures

Quantum homomorphisms between relational structures are captured by a cocategory structure.

Definition 3.2.7. The C^* -cocategory of σ -structures is the C^* -cocategory $\mathcal{R}^+(\sigma)$ where the objects are σ -structures; for $A, B \in \text{ob}(\mathcal{R}^+(\sigma))$, $\text{mor}_{\mathcal{R}^+(\sigma)}^+(A, B) = \text{mor}^+(A, B)$ is the quantum space of homomorphisms $A \rightarrow B$; the cocomposition is defined via $\Delta_{A,C}^B(p_{a,c}) = \sum_{b \in B} p_{a,b} \otimes p_{b,c}$, and the coidentity is defined by $\epsilon_A(p_{a,a'}) = \delta_{a,a'}$.

The **oracular C^* -cocategory of σ -structures** is the C^* -cocategory $\mathcal{R}^{o+}(\sigma)$ defined identically except that $\text{mor}_{\mathcal{R}^{o+}(\sigma)}^+(A, B) = \text{mor}^{o+}(A, B)$.

Lemma 3.2.8. $\mathcal{R}^+(\sigma)$ and $\mathcal{R}^{o+}(\sigma)$ are well-defined as C^* -cocategories.

Proof. We begin with $\mathcal{R}^+(\sigma)$. To show this, we need that the cocomposition and coidentity extend to $*$ -homomorphisms that satisfy the coproduct and counit relations. First, the relations constraining $\{p_{a,c}\}_{c \in C} \subseteq \text{Mor}^+(A, C)$ to be PVMs are preserved under $\Delta_{A,C}^B$ as

$$\begin{aligned}
\Delta_{A,C}^B(p_{a,c})^* &= \left(\sum_{b \in B} p_{a,b} \otimes p_{b,c} \right)^* = \sum_{b \in B} p_{a,b}^* \otimes p_{b,c}^* = \sum_{b \in B} p_{a,b} \otimes p_{b,c} = \Delta_{A,C}^B(p_{a,c}), \\
\Delta_{A,C}^B(p_{a,c})^2 &= \sum_{b, b' \in B} p_{a,b} p_{a,b'} \otimes p_{b,c} p_{b',c} = \sum_{b, b' \in B} \delta_{b,b'} p_{a,b} \otimes p_{b,c} p_{b',c} = \sum_{b \in B} p_{a,b} \otimes p_{b,c} = \Delta_{A,C}^B(p_{a,c}), \\
\sum_{c \in C} \Delta_{A,C}^B(p_{a,c}) &= \sum_{\substack{b \in B \\ c \in C}} p_{a,b} \otimes p_{b,c} = \sum_{b \in B} p_{a,b} \otimes 1 = 1 \otimes 1 = \Delta_{A,C}^B(1).
\end{aligned}$$

Now, let $R \in \sigma$, and let $\mathbf{a} \in R^A$ and $\mathbf{c} \notin R^C$. Writing $n = \text{ar}(R)$, the relation $p_{a_1, c_1} \cdots p_{a_n, c_n} = 0$ is preserved under $\Delta_{A,C}^B$ as

$$\begin{aligned} \Delta_{A,C}^B(p_{a_1, c_1}) \cdots \Delta_{A,C}^B(p_{a_n, c_n}) &= \sum_{b_1, \dots, b_n \in B} p_{a_1, b_1} \cdots p_{a_n, b_n} \otimes p_{b_1, c_1} \cdots p_{b_n, c_n} \\ &= \sum_{b \in R^B} p_{a_1, b_1} \cdots p_{a_n, b_n} \otimes p_{b_1, c_1} \cdots p_{b_n, c_n} = 0, \end{aligned}$$

as $\mathbf{a} \in R^A$. Therefore, $\Delta_{A,C}^B$ extends to a $*$ -homomorphism. Next, we need coassociativity of the cocomposition. Since the $\Delta_{A,C}^B$ are $*$ -homomorphisms, we need only show this property on the generators. In fact,

$$\begin{aligned} (\Delta_{A,C}^B \otimes \text{id})\Delta_{A,D}^C(p_{a,d}) &= \sum_{c \in C} \Delta_{A,C}^B(p_{a,c}) \otimes p_{c,d} \\ &= \sum_{\substack{b \in B \\ c \in C}} p_{a,b} \otimes p_{b,c} \otimes p_{c,d} \\ &= \sum_{b \in B} p_{a,b} \otimes \Delta_{B,D}^C(p_{b,d}) \\ &= (\text{id} \otimes \Delta_{B,D}^C)\Delta_{A,D}^B(p_{a,d}). \end{aligned}$$

We also need to show that the coidentity is a $*$ -homomorphism. We approach this the same way:

$$\begin{aligned} \epsilon_A(p_{a,a'})^* &= \delta_{a,a'}^* = \delta_{a,a'} = \epsilon_A(p_{a,a'}), \\ \epsilon_A(p_{a,a'})^2 &= \delta_{a,a'}^2 = \epsilon_A(p_{a,a'}), \\ \sum_{a' \in A} \epsilon_A(p_{a,a'}) &= \sum_{a' \in A} \delta_{a,a'} = 1 = \epsilon_A(1). \end{aligned}$$

For the remaining relations, suppose $R \in \sigma$, $\mathbf{a} \in R^A$, and $\mathbf{a}' \notin R^A$. Writing $n = \text{ar}(R)$, we find that

$$\epsilon_A(p_{a_1, a'_1}) \cdots \epsilon_A(p_{a_n, a'_n}) = \delta_{a_1, a'_1} \cdots \delta_{a_n, a'_n} = \delta_{\mathbf{a}, \mathbf{a}'} = 0,$$

as $\mathbf{a}$ and $\mathbf{a}'$ cannot be equal. Hence ϵ_A is a $*$ -homomorphism. As such, to show it also satisfies the counit relations, it suffices to show it on the generators. In fact,

$$\begin{aligned} (\epsilon_A \otimes \text{id})\Delta_{A,B}^A(p_{a,b}) &= \sum_{a' \in A} \epsilon_A(p_{a,a'}) \otimes p_{a',b} = \sum_{a' \in A} \delta_{a,a'} p_{a',b} = p_{a,b}, \\ (\text{id} \otimes \epsilon_B)\Delta_{A,B}^B(p_{a,b}) &= \sum_{b' \in B} p_{a,b'} \otimes \epsilon_B(p_{b',b}) = \sum_{b' \in B} \delta_{b',b} p_{a,b'} = p_{a,b}. \end{aligned}$$

Hence $\mathcal{R}^+(\sigma)$ is well-defined as a C^* -cocategory.

To show that $\mathcal{R}^{o+}(\sigma)$ is also well-defined, it suffices to show that the cocomposition and the coidentity preserve the additional commutation relations $[p_{a_i,c}, p_{a_j,c'}] = 0$ for all $R \in \sigma$, $\mathbf{a} \in R^A$, $1 \leq i, j \leq \text{ar}(R)$, and $c, c' \in C$. In fact, writing $n = \text{ar}(R)$ and assuming without loss of generality that $i < j$, we have that

$$\begin{aligned} [\Delta_{A,C}^B(p_{a_i,c}), \Delta_{A,C}^B(p_{a_j,c'})] &= \sum_{b,b' \in B} [p_{a_i,b} \otimes p_{b,c}, p_{a_j,b'} \otimes p_{b',c'}] \\ &= \sum_{b,b' \in B} p_{a_i,b} p_{a_j,b'} \otimes [p_{b,c}, p_{b',c'}] \\ &= \sum_{b_1, \dots, b_n \in B} p_{a_1,b_1} \cdots p_{a_n,b_n} \otimes [p_{b_i,c}, p_{b_j,c'}] \\ &= \sum_{b \in R^B} p_{a_1,b_1} \cdots p_{a_n,b_n} \otimes [p_{b_i,c}, p_{b_j,c'}] = 0, \end{aligned}$$

by the commutation relations on $\text{Mor}^{o+}(B, C)$. It is clear that $[\epsilon_A(p_{a_i,a}), \epsilon_A(p_{a_j,a'})] = 0$, as the codomain of ϵ_A is a commutative algebra. Hence $\mathcal{R}^{o+}(\sigma)$ is well-defined as well. $\blacksquare$

To close this section, we see how this unified picture gives a precise way in which quantum homomorphisms quantise the classical homomorphisms of relational structures. First, for a tensor monoid Q , we write the associated Q -morphism categories $\mathcal{R}^Q(\sigma) = \mathcal{R}^+(\sigma)^Q$ and $\mathcal{R}^{oQ}(\sigma) = \mathcal{R}^{o+}(\sigma)^Q$. Now, we show that the d -morphism category is essentially the same as the classical category of relational structures.

Lemma 3.2.9. The categories $\mathcal{R}^d(\sigma)$, $\mathcal{R}^{od}(\sigma)$, and $\mathcal{R}(\sigma)$ are isomorphic.

Proof. First, note that for a $*$ -homomorphism $\pi : \text{Mor}^+(A, B) \rightarrow \mathbb{C}$, every commutator in $\text{Mor}^+(A, B)$ is in the kernel of π , and hence it factors through $\text{Mor}^{o+}(A, B)$. As such, we get a bijection between $\text{mor}^d(A, B)$ and $\text{mor}^{od}(A, B)$. In the following, we show the isomorphism $\mathcal{R}^d(\sigma) \cong \mathcal{R}(\sigma)$ — by the bijections above, an identical argument can be used to show the other isomorphism $\mathcal{R}^{od}(\sigma) \cong \mathcal{R}(\sigma)$.

Next, for such a $*$ -homomorphism, $\{\pi(p_{a,b})\}_{b \in B}$ is a PVM in $\mathbb{C}$. So $\pi(p_{a,b}) = 0$ or $\pi(p_{a,b}) = 1$, and for each a there is a unique b such that $\pi(p_{a,b}) = 1$. This induces a unique function $f : A \rightarrow B$ such that $\pi(p_{a,b}) = \delta_{b,f(a)}$. We define the covariant functors $F : \mathcal{R}^d(\sigma) \rightarrow \mathcal{R}(\sigma)$ and $G : \mathcal{R}(\sigma) \rightarrow \mathcal{R}^d(\sigma)$ on objects as $F(A) = G(A) = A$ and on morphisms as $F(\pi) = f$ and $G(f)$ as the extension of a map $p_{a,b} \mapsto \delta_{b,f(a)}$ to a $*$ -homomorphism. It is clear that these are inverses of each other if they are well-defined. It remains to show that they are in fact functors. To show that G is a functor, we need

to first show that for all $f \in \text{mor}(A, B)$, $G(f) \in \text{mor}^d(A, B)$. This follows because $G(f)$ preserves the relations of $\text{Mor}^+(A, B)$: let $a \in A$, $b, b' \in B$, $R \in \sigma$, $\mathbf{a} \in R^A$, and $\mathbf{b} \notin R^B$, and writing $n = \text{ar}(R)$,

$$\begin{aligned} G(f)(p_{a,b})^* &= \delta_{b,f(a)} = G(f)(p_{a,b}), \\ G(f)(p_{a,b})^2 &= \delta_{b,f(a)}^2 = G(f)(p_{a,b}), \\ \sum_{b \in B} G(f)(p_{a,b}) &= \sum_{b \in B} \delta_{b,f(a)} = 1, \\ G(f)(p_{a_1,b_1}) \cdots G(f)(p_{a_n,b_n}) &= \delta_{b_1,f(a_1)} \cdots \delta_{b_n,f(a_n)} = \delta_{\mathbf{b},f(\mathbf{a})} = 0, \end{aligned}$$

as $f(\mathbf{a}) \in R^B$ since f is a homomorphism. Next, we need that G preserves compositions. Let $f \in \text{mor}(A, B)$ and $g \in \text{mor}(B, C)$. Then, we have that

$$\begin{aligned} (G(g) \circ G(f))(p_{a,c}) &= (G(f) \otimes G(g))\Delta_{A,C}^B(p_{a,c}) \\ &= \sum_{b \in B} G(f)(p_{a,b}) \otimes G(g)(p_{b,c}) \\ &= \sum_{b \in B} \delta_{b,f(a)} \delta_{c,g(b)} \\ &= \delta_{c,g(f(a))} = G(g \circ f)(p_{a,c}). \end{aligned}$$

As the maps are $*$ -homomorphisms, this implies that $G(g) \circ G(f) = G(g \circ f)$ as wanted. Finally, G preserves the identity, as $G(\text{id}_A)(p_{a,a'}) = \delta_{a,a'} = \epsilon_A(p_{a,a'})$, so G is a functor.

We also need to show that F is a functor. Let $\pi \in \text{mor}^d(A, B)$ and suppose that $F(\pi) \notin \text{mor}(A, B)$. Then, there exists $R \in \sigma$ and $\mathbf{a} \in R^A$ such that $F(\pi)(\mathbf{a}) \notin R^B$. This means that

$$0 = \pi(p_{a_1,F(\pi)(a_1)} \cdots p_{a_n,F(\pi)(a_n)}) = \delta_{F(\pi)(a_1),F(\pi)(a_1)} \cdots \delta_{F(\pi)(a_n),F(\pi)(a_n)} = 1,$$

a contradiction. Hence, $F(\pi)$ is a homomorphism. Next, we need that F preserves compositions. Let $\pi \in \text{mor}^d(A, B)$ and $\varphi \in \text{mor}^d(B, C)$, so $F(\varphi \circ \pi)(a)$ is such that $(\varphi \circ \pi)(p_{a,c}) = \delta_{c,F(\varphi \circ \pi)(a)}$. But

$$(\varphi \circ \pi)(p_{a,c}) = (\pi \otimes \varphi)\Delta_{A,C}^B(p_{a,c}) = \sum_{b \in B} \pi(p_{a,b}) \otimes \varphi(p_{b,c}) = \sum_{b \in B} \delta_{b,F(\pi)(a)} \delta_{c,F(\varphi)(b)} = \delta_{c,F(\varphi) \circ F(\pi)(b)},$$

so $F(\varphi \circ \pi) = F(\varphi) \circ F(\pi)$. Finally, $F(\epsilon_A)(a) = a$, so it is the identity. Hence, F is a functor as well, finishing the proof. ■

3.3 Almost perfect strategies for projection games are approximately tracial

In [117], it was noted that the rounding result from general quantum correlations to synchronous correlations can be extended to apply to projection games. In this section, we will show a stronger version of this result directly, and extend it to commuting operator correlations, as was done in [86]. Recall that projection games, as defined in Section 2.12.3, are nonlocal games where for any one of Alice's answers, Bob has at most one correct answer. For a pair of σ -structures A and B , the oracular constraint-variable game $\mathsf{G}_{ocv}(A, B, \pi)$ is always a projection, and it will be the main example of game to which we apply the results of this section.

In Section 3.3.1, we study a useful property of the structure of projection games. In Sections 3.3.2 and 3.3.3, we show that perfect quantum and commuting operator strategies for projection games are both tracial. In Section 3.3.4, we round near-perfect strategies for projection games to near-perfect tracial strategies. In Section 3.3.5, we apply our results in the context of oracular constraint-variable CSP nonlocal games, and oracularisations of games.

3.3.1 Degenerate and non-degenerate projection games

In this section, we introduce a simple new definition which helps us control the structure of a projection game.

Definition 3.3.1. Let $\mathsf{G} = (X, Y, \{A_x\}_{x \in X}, \{B_y\}_{y \in Y}, \mu, V)$ be a projection game with projection functions $p_{x,y}$. We say G is **non-degenerate** if the function $p_x : A_x \rightarrow \prod_{\substack{y \in Y: \\ \mu(x,y) > 0}} B_y$ defined as $p_x(a) = (p_{x,y}(a))_y$ is injective for all $x \in X$.

The major examples of projection games, oracular constraint-variable games and oracularisations, are always non-degenerate.

If a game is degenerate, whenever Alice answers a , she has the freedom to choose any element of $p_x^{-1}(p_x(a))$ without affecting the winning probability. This can cause some problems in pinpointing the behaviour of the players. However, we can always reduce to the case of non-degenerate projection games, while preserving the important properties of the game.

Lemma 3.3.2. Let $\mathbf{G}$ be a projection game. Then, there exists a non-degenerate projection game $\tilde{\mathbf{G}}$ such that for every correlation p for $\mathbf{G}$, there exists a correlation $\tilde{p}$ for $\tilde{\mathbf{G}}$ such that $\mathfrak{w}(\mathbf{G}, p) = \mathfrak{w}(\tilde{\mathbf{G}}, \tilde{p})$, and vice versa. Further, the mappings from p to $\tilde{p}$ and from $\tilde{p}$ to p preserve the classes of quantum, commuting operator, and tracial correlations.

Though it is not relevant here, the mappings also preserve every other natural (closed under local classical post-processing) class of correlations by construction, notably the classical correlations, the non-signalling correlations, and the levels of the NPA hierarchy.

Proof. Let $\tilde{A}_x \subseteq A_x$ be a set of representatives for the partition of A_x induced by p_x — that is, for every $a \in A_x$, there exists a unique $\tilde{a} \in \tilde{A}_x$ such that $p_x(\tilde{a}) = p_x(a)$. Then, we can define $\tilde{V} = \{V(a, b|x, y)\}_{x \in X, y \in Y, a \in \tilde{A}_x, b \in B_y}$, and the nonlocal game $\tilde{\mathbf{G}} = (X, Y, \{\tilde{A}_x\}_{x \in X}, \{B_y\}_{y \in Y}, \mu, \tilde{V})$. By definition, this is a projection game via the projection functions $\tilde{p}_{x,y} = p_{x,y}|_{\tilde{A}_x}$, and it is non-degenerate. Now, let p be a correlation for $\mathbf{G}$. Define $\tilde{p}$ as the correlation $\tilde{p}(a, b|x, y) = \sum_{a' \in p_x^{-1}(p_x(a))} p(a', b|x, y)$. Since $\tilde{p}$ is defined from p via classical post-processing on Alice's system, the mapping preserves quantum, commuting operator, and tracial correlations. It also preserves the winning probability as

$$\begin{aligned} \mathfrak{w}(\tilde{\mathbf{G}}, \tilde{p}) &= \sum_{x,y} \mu(x, y) \sum_{a \in \tilde{A}_x, b \in B_y} \tilde{V}(a, b|x, y) \tilde{p}(a, b|x, y) \\ &= \sum_{x,y} \mu(x, y) \sum_{a \in \tilde{A}_x} \tilde{p}(a, p_{x,y}(a)|i, j) \\ &= \sum_{x,y} \mu(x, y) \sum_{a \in \tilde{A}_x} \sum_{a' \in p_x^{-1}(p_y(a))} p(a', p_{x,y}(a)|i, j) \\ &= \sum_{x,y} \mu(x, y) \sum_{a' \in A_x} p(a', p_{x,y}(a')|i, j) \\ &= \mathfrak{w}(\mathbf{G}, p). \end{aligned}$$

For the other direction, let $\tilde{p}$ be a correlation for $\tilde{\mathbf{G}}$. Then, define the correlation p for $\mathbf{G}$ as $p(a, b|x, y) = \tilde{p}(a, b|x, y)$ when $a \in \tilde{A}_x$ and $p(a, b|x, y) = 0$ otherwise. As before, it is easy to see that this mapping preserves the sets of quantum, commuting-operator, and tracial correlations. Finally it preserves the winning probability as

$$\begin{aligned} \mathfrak{w}(\mathbf{G}, p) &= \sum_{x,y} \mu(x, y) \sum_{a \in A_x, b \in B_y} V(a, b|x, y) p(a, b|x, y) \\ &= \sum_{x,y} \mu(x, y) \sum_{a \in \tilde{A}_x, b \in B_y} \tilde{V}(a, b|x, y) \tilde{p}(a, b|x, y) \\ &= \mathfrak{w}(\tilde{\mathbf{G}}, \tilde{p}). \end{aligned} \quad \blacksquare$$

For the majority of this section, we will assume that we are dealing with non-degenerate projection games, but using the above lemma, we can guarantee that the main result also holds for degenerate games.

3.3.2 Perfect quantum case

In this section, we extend the result of [43, 1] showing that perfect correlations for constraint-variable CS games are tracial to the setting of general projection games. We express the proofs using the von Neumann algebra language that we make use of throughout this work.

Lemma 3.3.3 (e.g. [1]). Let $p = \{p(a, b|x, y)\}_{x \in X, y \in Y, a \in A_x, b \in B_y}$ be a quantum correlation. Then, there exist a finite-dimensional Hilbert space $H = \mathbb{C}^d$, POVMs $\{P_a^x\}_{a \in A_x} \subseteq \mathcal{B}(H)$ and $\{Q_b^y\}_{b \in B_y} \subseteq \mathcal{B}(H)$, and a unit vector $|\phi\rangle = \sum_{n=1}^d \lambda_n |n\rangle |n\rangle \in H \otimes H$ with $\lambda_n > 0$ such that $p(a, b|x, y) = \langle \phi | P_a^x \otimes Q_b^y | \phi \rangle$.

Proof. Since p is a quantum correlation, there exist some finite-dimensional Hilbert spaces H_A, H_B , POVMs $\{A_a^x\}_{a \in A_x} \subseteq \mathcal{B}(H_A)$, $\{B_b^y\}_{b \in B_y} \subseteq \mathcal{B}(H_B)$, and unit vector $|\psi\rangle \in H_A \otimes H_B$ such that $p(a, b|x, y) = \langle \psi | A_a^x \otimes B_b^y | \psi \rangle$. Let $|\psi\rangle = \sum_{n=1}^d \lambda_n |\psi_n\rangle |\phi_n\rangle$ be the Schmidt decomposition of $|\psi\rangle$, for d the Schmidt rank of $|\psi\rangle$. Then, letting $H = \mathbb{C}^d$, define the isometries $U : H \rightarrow H_A$ and $V : H \rightarrow H_B$ as $U = \sum_n |\psi_n\rangle \langle n|$ and $V = \sum_n |\phi_n\rangle \langle n|$. Letting $P_a^x = U^* A_a^x U$ and $Q_b^y = V^* B_b^y V$, we get that

$$p(a, b|x, y) = \langle \psi | A_a^x \otimes B_b^y | \psi \rangle = \langle \phi | (U \otimes V)^* (A_a^x \otimes B_b^y) (U \otimes V) | \phi \rangle = \langle \phi | P_a^x \otimes Q_b^y | \phi \rangle. \blacksquare$$

In the main theorem of this section, we make use of the following result. We say POVMs $\{A_i^1\}_{i \in I_1}, \dots, \{A_i^n\}_{i \in I_n}$ are **jointly measurable** if there exists a POVM $\{A_{i_1, \dots, i_n}\}_{i_1, \dots, i_n}$, called the **joint measurement**, such that $A_i^j = \sum_{\mathbf{i}: i_j = i} A_{\mathbf{i}}$.

Proposition 3.3.4 ([73]). If PVMs $\{A_i^1\}, \dots, \{A_i^n\}$ are jointly measurable, then the joint measurement is a PVM and has the form $A_{i_1, \dots, i_n} = A_{i_1}^1 \cdots A_{i_n}^n$.

Note that the above proposition holds independently of dimension.

Theorem 3.3.5. Let $\mathbf{G} = (X, Y, \{A_x\}_{x \in X}, \{B_y\}_{y \in Y}, \mu, V)$ be a non-degenerate projection game, and let $p(a, b|x, y) = \langle \phi | P_a^x \otimes Q_b^y | \phi \rangle$ be a quantum correlation for $\mathbf{G}$ of the form guaranteed by Lemma 3.3.3. Suppose p is perfect for $\mathbf{G}$. Define $D = \sum_{n=1}^d \lambda_n |n\rangle \langle n|$. Then, we have that

- P_a^x and Q_b^y are projections

- D commutes with P_a^x and Q_b^y
- $P_a^x(Q_b^y)^T = 0$ whenever $p_{x,y}(a) \neq b$.

This theorem was shown in the case of binary oracular constraint-variable games in [43] and for oracular constraint-variable games over a general alphabet in [1].

Proof. For p to be perfect, we need that $p(a, b|x, y) = 0$ whenever $V(a, b|x, y) = 0$. Hence, $\langle \phi | P_a^x \otimes Q_b^y | \phi \rangle = 0$ if $p_{x,y}(a) \neq b$. In this case, we have $\sqrt{P_a^x \otimes Q_b^y} | \phi \rangle = 0$, which implies $(P_a^x \otimes Q_b^y) | \phi \rangle = 0$. Rewriting this as a matrix gives $P_a^x D(Q_b^y)^T = 0$. Write $P_b^{x,y} = \sum_{a \in p_{x,y}^{-1}(b)} P_a^x$. By summing over all $b \neq p_{x,y}(a)$, we find that $P_a^x D(1 - (Q_b^y)^T) = 0$ if $b = p_{x,y}(a)$. Then, summing over a such that $p_{x,y}(a) = b$, we find that $P_b^{x,y} D(Q_b^y)^T = P_b^{x,y} D$. On the other hand, summing over those a such that $p_{x,y}(a) \neq b$, we can get $(1 - P_b^{x,y}) D(Q_b^y)^T$. Together $P_b^{x,y} D(Q_b^y)^T = P_b^{x,y} D = D(Q_b^y)^T$.

Then, $(P_b^{x,y})^2 D = P_b^{x,y} D(Q_b^y)^T = P_b^{x,y} D$, so by invertibility of D , $(P_b^{x,y})^2 = P_b^{x,y}$. As such, $P_b^{x,y}$ is a projection. For each x , we define a measurement on $\prod_{y \in Y: \mu(x,y) > 0} B_y$ as $P_b^x = \sum_{a \in p_x^{-1}(b)} P_a^x$. By construction, this is the joint measurement of the $\{P_b^{x,y}\}$ over y . Using Proposition 3.3.4, P_b^x is a projection and $P_b^x = \prod_y P_b^{x,y}$. Since $\mathbf{G}$ is non-degenerate, P_b^x is the sum of at most 1 term, so we get that $P_a^x = \prod_y P_{p_{x,y}(a)}^{x,y}$ and is a projection. In the same way, $D((Q_b^y)^2)^T = P_b^{x,y} D(Q_b^y)^T = D(Q_b^y)^T$, so $(Q_b^y)^T$ and hence Q_b^y is a projection as well.

Next, $P_b^{x,y} D^2 = D(Q_b^y)^T D = D(D(Q_b^y)^T)^* = D^2 P_b^{x,y}$, so P_a^x commutes with D^2 . By the continuous functional calculus, P_a^x commutes with $(D^2)^{1/2} = D$. In the same way, Q_b^y commutes with D .

To finish, note that if $p_{x,y}(a) \neq b$, $0 = P_a^x D(Q_b^y)^T = P_a^x (Q_b^y)^T D$. By invertibility of D , $P_a^x (Q_b^y)^T = 0$. ■

Note that, if $\mathbf{G}$ were degenerate, we could not rebuild P_a^x from the $P_b^{x,y}$. Hence P_a^x might not be a projection and might not commute with D . Nevertheless the other results would hold.

Corollary 3.3.6. Let G be a non-degenerate projection game, and let p be a quantum correlation. Suppose that p is a perfect strategy for G . Then, p is a tracial correlation.

This corollary follows for the case of oracular constraint-variable games from [1, 98, 43].

Proof. By [Theorem 3.3.5](#), we know that $p(a, b|x, y) = \text{Tr}(P_a^x D (Q_b^y)^T D)$, where $D > 0$ is invertible and $\text{Tr}(D^2) = 1$, P_a^x and Q_b^y are projectors, D commutes with P_a^x and Q_b^y . First note that $p(a, b|x, y) = \text{Tr}(P_a^x (Q_b^y)^T D^2)$. Now, let $\mathcal{M} \subseteq \mathcal{B}(H)$ be the von Neumann algebra generated by the P_a^x and $(Q_b^y)^T$. By construction, we have that $D \in \mathcal{M}'$. Let τ be the state on $\mathcal{M}$ defined as $\tau(x) = \text{Tr}(x D^2)$. As $D^2 \in \mathcal{M}'$, τ is tracial. Let $A_a^x = P_a^x \in \mathcal{M}$ and $B_b^y = (Q_b^y)^T \in \mathcal{M}$. We have that $p(a, b|x, y) = \tau(A_a^x B_b^y)$, so it is a tracial correlation. $\blacksquare$

3.3.3 Perfect commuting operator case

Using ideas from [Section 2.4.4](#), it is possible to extend the results of the previous section from quantum correlations to commuting-operator correlations.

Theorem 3.3.7. Let $\mathbf{G} = (X, Y, \{A_x\}_{x \in X}, \{B_y\}_{y \in Y}, \mu, V)$ be a non-degenerate projection game, and let $p(a, b|x, y) = \langle \psi | P_a^x Q_b^y | \psi \rangle$ be a commuting operator correlation that is perfect for $\mathbf{G}$. Then, p is a tracial correlation.

Proof. Let $\mathcal{M} \subseteq \mathcal{B}(H)$ be the von Neumann algebra generated by the Q_b^y . Then $P_a^x \in \mathcal{M}'$. The map $\mathcal{M} \rightarrow \mathbb{C}$, $T \mapsto \langle \psi | T | \psi \rangle$ is an element of $\mathcal{M}_*$, so there exists a corresponding element $h_\psi \in L^1(\mathcal{M})$. Note that h_ψ is positive and $\text{tr}(h_\psi) = \psi(1) = 1$. By a standard argument, there exist POVMs $\{\bar{P}_a^x\}_{a \in A_i} \subseteq \mathcal{M}$ such that $\langle \psi | P_a^x Q_b^y | \psi \rangle = \text{tr}(\bar{P}_a^x h_\psi^{1/2} Q_b^y h_\psi^{1/2})$ (see the proof of Theorem 2.1 in [\[86\]](#)). As p is perfect, we know that $p(a, b|x, y) = 0$ whenever $p_{x,y}(a) \neq b$. By tracial property of tr , we have in this case that

$$0 = \text{tr}((\bar{P}_a^x)^{1/2} h_\psi^{1/2} Q_b^y h_\psi^{1/2} (\bar{P}_a^x)^{1/2}) = \|(\bar{P}_a^x)^{1/2} h_\psi^{1/2} (Q_b^y)^{1/2}\|_{\text{tr}}^2,$$

and hence $(\bar{P}_a^x)^{1/2} h_\psi^{1/2} (Q_b^y)^{1/2} = 0$ (in terms of the strong product). Therefore, $\bar{P}_a^x h_\psi^{1/2} Q_b^y = 0$. Let $\bar{P}_b^{x,y} = \sum_{a \in p_{x,y}^{-1}(b)} \bar{P}_a^x$. Fixing any $b \neq b' \in B_y$, we have $\bar{P}_b^{x,y} h_\psi^{1/2} Q_{b'}^y = \sum_{a \in p_{x,y}^{-1}(b)} \bar{P}_a^x h_\psi^{1/2} Q_{b'}^y = 0$ as $p_{x,y}(a) = b \neq b'$. Now, $\bar{P}_b^{x,y} h_\psi^{1/2} = \sum_{b' \in B_j} \bar{P}_b^{x,y} h_\psi^{1/2} Q_{b'}^y = \bar{P}_b^{x,y} h_\psi^{1/2} Q_b^y$ and $h_\psi^{1/2} Q_b^y = \sum_{b' \in B_y} \bar{P}_{b'}^{x,y} h_\psi^{1/2} Q_b^y = \bar{P}_b^{x,y} h_\psi^{1/2} Q_b^y$. Now, in the same way as in the finite-dimensional case except that we are using the strong product, $(\bar{P}_b^{x,y})^2 h_\psi^{1/2} = \bar{P}_b^{x,y} h_\psi^{1/2} Q_b^y = \bar{P}_b^{x,y} h_\psi^{1/2}$; $h_\psi^{1/2} (Q_b^y)^2 = \bar{P}_b^{x,y} h_\psi^{1/2} Q_b^y = h_\psi^{1/2} Q_b^y$; and $Q_b^y h_\psi = (h_\psi^{1/2} Q_b^y)^* h_\psi^{1/2} = h_\psi^{1/2} \bar{P}_b^{x,y} h_\psi^{1/2} = h_\psi^{1/2} Q_b^y$. As the Q_b^y generate $\mathcal{M}$, h_ψ is affiliated with $\mathcal{M}'$. Let q be the projection onto the support of h_ψ . Then, $q\mathcal{M} \subseteq \mathcal{B}(qc(H))$ is a von Neumann algebra, and $p_b^{x,y} = \bar{P}_b^{x,y} q$ and $q_b^y = Q_b^y q$ are projections in $q\mathcal{M}$. Since the $p_b^{x,y}$ are jointly measurable and G is non-degenerate, we have that $p_a^x = \prod_{y \in Y} p_{p_{x,y}(a)}^{x,y}$ is a PVM as in [Theorem 3.3.5](#); note that $p_a^x = \bar{P}_a^x q$. Now, let $\rho : q\mathcal{M} \rightarrow \mathbb{C}$ be defined as $\rho(T) = \text{tr}(T h_\psi)$. ρ is a tracial state: as $\rho(1_{q\mathcal{M}}) = \rho(q) = \text{tr}(h_\psi) = 1$, for $T \geq 0$,

$\rho(T) = \text{tr}(h_\psi^{1/2} T h_\psi^{1/2}) \geq 0$, and $\rho(TS) = \text{tr}(T S h_\psi) = \text{tr}(T h_\psi S) = \text{tr}(S T h_\psi) = \rho(ST)$. Putting it together,

$$p(a, b|x, y) = \text{tr}(p_a^x h_\psi^{1/2} q_b^y h_\psi^{1/2}) = \text{tr}(p_a^x q_b^y h_\psi) = \rho(p_a^x q_b^y),$$

which is a tracial correlation. ■

Corollary 3.3.8. In the notation of the previous proof, we have for a perfect commuting operator strategy that $p_a^x q_b^y = 0$ whenever $p_{x,y}(a) \neq b$.

3.3.4 Approximate case

In this section, we prove the following theorem, extending to the most general case of almost perfect commuting operator strategies.

Theorem 3.3.9. Let $G = (X, Y, \{A_x\}_{x \in X}, \{B_y\}_{y \in Y}, \mu, V)$ be a non-degenerate projection game, and let $p(a, b|x, y) = \langle \phi | P_a^x Q_b^y | \phi \rangle$ be a commuting-operator correlation for $\mathbf{G}$. Suppose that $\mathfrak{w}(\mathbf{G}, p) \geq 1 - \varepsilon$. Then, there exists a tracial correlation p' such that $\mathfrak{w}(\mathbf{G}, p') \geq 1 - 148(L\varepsilon)^{1/4}$, where $L = \max_{x,y:\mu(x,y)>0} \frac{\mu(x)}{\mu(x,y)}$ for $\mu(x) = \sum_{y \in Y} \mu(x, y)$.

Our main technical tool is the following lemma of [86], which that work uses to show the analogous result for synchronous games.

Lemma 3.3.10 ([86]). Let $\mathcal{M} \subseteq \mathcal{B}(H)$ be a von Neumann algebra, $h \in L^1(\mathcal{M})$ be positive such that $\text{tr}(h) = 1$, and $q = \chi_{(1,\infty)}(h)$. Let X and A be finite sets, and let π be a symmetric probability distribution on $X \times X$ with marginal ν on X . Suppose that $\{p_a^x\}_{a \in A}$ are PVMs such that $\sum_{x \in X} \nu(x) \sum_a \|[p_a^x, h^{1/2}]\|_2^2 \leq \delta$. Then, there exist PVMs $\{r_a^x\}_{a \in A} \subseteq qc(\mathcal{M})q$ such that

$$\sum_{x,y \in X} \pi(x, y) \sum_{a,b \in A} \left| \text{tr}(p_a^x h^{1/2} p_b^y h^{1/2}) - \text{tr}(r_a^x r_b^y) \right| \leq 38\delta^{1/4}.$$

Note that $\tau|_{qc(\mathcal{M})q}$ is a tracial state as $\tau(q) = \text{tr}(h) = 1$.

The following lemma shows that (with respect to KMS norm) if jointly measurable POVMs are approximately PVMs, then their joint measurement is also approximately a PVM. Note that this result is the only place where the dependence on the number of variables per constraint in Theorem 3.3.9 arises.

Lemma 3.3.11. Let $\mathcal{M} \subseteq \mathcal{B}(H)$ be a von Neumann algebra and let $h \in L^1(\mathcal{M})$ be a positive element such that $\text{tr}(h) = 1$. Let π be a strictly positive probability distribution on $[n]$. Suppose that $\{A_i\}_{i \in [k]^n} \subseteq \mathcal{M}$ is a POVM such that $\sum_{j=1}^n \pi(j) \sum_{i=1}^k \|h^{1/4} A_i^j h^{1/4}\|_{\text{tr}}^2 \geq 1 - \delta$. Then,

$$\sum_{i \in [k]^n} \|h^{1/4} A_i h^{1/4}\|_{\text{tr}}^2 \geq 1 - M\delta,$$

where $M = \max_i \frac{1}{\pi(i)}$.

Before passing to the proof, note that this bound is tight in the limit of small δ . Take $k = 2$, π uniform, $\mathcal{M} = \mathcal{B}(\mathbb{C}^2)^{\otimes n}$ with tracial state $\frac{1}{2^n} \text{Tr}$, and $A_i = \otimes_j \left((1 - \varepsilon) |i_j\rangle\langle i_j| + \frac{\varepsilon}{2} I \right)$. Then, for any j ,

$$\sum_{i=1,2} \frac{1}{2^n} \text{Tr}((A_i^j)^2) = \frac{1}{2} \sum_{i=1,2} \text{Tr}\left(\left((1 - \varepsilon) |i\rangle\langle i| + \frac{\varepsilon}{2} I\right)^2\right) = \frac{1}{2} \sum_{i=1,2} \text{Tr}\left(\left((1 - \varepsilon) |i\rangle\langle i| + \frac{\varepsilon^2}{4} I\right)\right) = 1 - \varepsilon + \frac{\varepsilon^2}{2},$$

giving $\delta = \varepsilon - \frac{\varepsilon^2}{2}$. In the same way, we have that

$$\sum_i \frac{1}{2^n} \text{Tr}(A_i^2) = \left(\frac{1}{2} \sum_{i=1,2} \text{Tr}\left(\left((1 - \varepsilon) |i\rangle\langle i| + \frac{\varepsilon}{2} I\right)^2\right) \right)^n = (1 - \delta)^n \leq 1 - n\delta + \binom{n}{2} \delta^2.$$

For small $\delta \ll \frac{1}{n}$, this is approximately $1 - n\delta$.

Proof. Write $\delta_j = 1 - \sum_{i=1}^k \|h^{1/4} A_i^j h^{1/4}\|_{\text{tr}}^2$. Then,

$$\sum_i \|h^{1/4} \sum_{i \in [k]^n: i_j = i} A_i h^{1/4}\|_{\text{tr}}^2 = \sum_i \|h^{1/4} A_i^j h^{1/4}\|_{\text{tr}}^2 \geq 1 - \delta_j.$$

Hence,

$$\begin{aligned} \delta_j &\geq \left\| \sum_i h^{1/4} A_i h^{1/4} \right\|_{\text{tr}}^2 - \sum_i \left\| \sum_{i: i_j = i} h^{1/4} A_i h^{1/4} \right\|_{\text{tr}}^2 \\ &= \sum_{i, i'} \text{tr}(A_i h^{1/2} A_{i'} h^{1/2}) - \sum_{i, i': i_j = i'_j} \text{tr}(A_i h^{1/2} A_{i'} h^{1/2}) \\ &= \sum_{i, i': i_j \neq i'_j} \text{tr}(A_i h^{1/2} A_{i'} h^{1/2}). \end{aligned}$$

Putting these together via a union bound,

$$\begin{aligned}
\sum_i \left\| h^{1/4} A_i h^{1/4} \right\|_{\text{tr}}^2 &= 1 - \sum_{i, i': i \neq i'} \text{tr}(A_i h^{1/2} A_{i'} h^{1/2}) \\
&\geq 1 - \sum_j \sum_{i, i': i_j \neq i'_j} \text{tr}(A_i h^{1/2} A_{i'} h^{1/2}) \\
&\geq 1 - \sum_j \delta_j \geq 1 - M \sum_j \pi(j) \delta_j \geq 1 - M\delta. \quad \blacksquare
\end{aligned}$$

Proof of Theorem 3.3.9. We may also assume that $\varepsilon \leq \frac{1}{L}$, as otherwise the result holds trivially.

As in the previous section, we take $\mathcal{M} \subseteq \mathcal{B}(H)$ to be the von Neumann algebra generated by the Q_b^y . Then, there exist a positive $h_\psi \in L^1(\mathcal{M})$ and POVMs $\{\bar{P}_a^x\}_{a \in A_x} \subseteq \mathcal{M}$ such that $\langle \psi | P_a^x Q_b^y | \psi \rangle = \text{tr}(\bar{P}_a^x h_\psi^{1/2} Q_b^y h_\psi^{1/2})$. Therefore, writing $\bar{P}_b^{x,y} = \sum_{a \in p_{x,y}^{-1}(b)} \bar{P}_a^x$,

$$\begin{aligned}
\mathfrak{w}(\mathbf{G}, p) &= \sum_{x,y} \mu(x, y) \sum_{a \in A_x} \langle \psi | P_a^x Q_{p_{x,y}(a)}^y | \psi \rangle \\
&= \sum_{x,y} \mu(x, y) \sum_{a \in A_x} \text{tr}(\bar{P}_a^x h_\psi^{1/2} Q_{p_{x,y}(a)}^y h_\psi^{1/2}) \\
&= \sum_{x,y,b} \mu(x, y) \langle h_\psi^{1/4} \bar{P}_b^{x,y} h_\psi^{1/4}, h_\psi^{1/4} Q_b^y h_\psi^{1/4} \rangle,
\end{aligned}$$

where $\langle \cdot, \cdot \rangle$ is the inner product on $L^2(\mathcal{M})$. From here, we often make use of two Hilbert spaces: $V_1 = L^2(\mathcal{M}) \bigcup_{y \in Y} X \times \{y\} \times B_y$ equipped with the inner product

$$\langle (u_{x,y,b})_{x,y,b}, (v_{x,y,b})_{x,y,b} \rangle_{V_1} = \sum_{x,y,b} \mu(x, y) \langle u_{x,y,b}, v_{x,y,b} \rangle,$$

and $V_2 = L^2(\mathcal{M}) \bigcup_{x \in X} \{x\} \times Y \times A_x$ equipped with the inner product

$$\langle (u_{x,y,a})_{x,y,a}, (v_{x,y,a})_{x,y,a} \rangle_{V_2} = \sum_{x,y,a} \mu(x, y) \langle u_{x,y,a}, v_{x,y,a} \rangle.$$

Consider $(h_\psi^{1/4} \bar{P}_b^{x,y} h_\psi^{1/4})_{x,y,b}$ and $(h_\psi^{1/4} Q_b^y h_\psi^{1/4})_{x,y,b}$ as vectors in V_1 . First, these vectors have norm at most 1. In fact,

$$\begin{aligned}
\| (h_\psi^{1/4} \bar{P}_b^{x,y} h_\psi^{1/4})_{x,y,b} \|_{V_1}^2 &= \sum_{x,y,b} \mu(x, y) \text{tr}(\bar{P}_b^{x,y} h_\psi^{1/2} \bar{P}_b^{x,y} h_\psi^{1/2}) \\
&\leq \sum_{x,y,b} \mu(x, y) \text{tr}(h_\psi^{1/2} \bar{P}_b^{x,y} h_\psi^{1/2}) \\
&\leq \sum_{x,y} \mu(x, y) \text{tr}(h_\psi) = 1,
\end{aligned}$$

and identically $\|(h_\psi^{1/4} Q_b^y h_\psi^{1/4})_{x,y,b}\|_{V_1} \leq 1$. By hypothesis, we know that

$$\langle (h_\psi^{1/4} \bar{P}_b^{x,y} h_\psi^{1/4})_{x,y,b}, (h_\psi^{1/4} Q_b^y h_\psi^{1/4})_{x,y,b} \rangle_{V_1} \geq 1 - \varepsilon,$$

so by Cauchy-Schwarz inequality,

$$\begin{aligned} 1 - \varepsilon &\leq \|(h_\psi^{1/4} \bar{P}_b^{x,y} h_\psi^{1/4})_{x,y,b}\|_{V_1} \|(h_\psi^{1/4} Q_b^y h_\psi^{1/4})_{x,y,b}\|_{V_1} \\ &\leq \min\left\{ \|(h_\psi^{1/4} \bar{P}_b^{x,y} h_\psi^{1/4})_{x,y,b}\|_{V_1}, \|(h_\psi^{1/4} Q_b^y h_\psi^{1/4})_{x,y,b}\|_{V_1} \right\}. \end{aligned}$$

Hence both vectors have norm between $1 - \varepsilon$ and 1. This gives that the commutator

$$\begin{aligned} \sum_{x,y,b} \mu(x,y) \left\| [Q_b^y, h_\psi^{1/2}] \right\|_2^2 &= \sum_{x,y,b} \mu(x,y) \operatorname{tr} \left((Q_b^y h_\psi^{1/2} - h_\psi^{1/2} Q_b^y) (h_\psi^{1/2} Q_b^y - Q_b^y h_\psi^{1/2}) \right) \\ &= 2 \sum_{x,y,b} \mu(x,y) \operatorname{tr} \left((Q_b^y)^2 h_\psi - h_\psi^{1/2} Q_b^y h_\psi^{1/2} Q_b^y \right) \\ &\leq 2 - 2 \|(h_\psi^{1/4} Q_b^y h_\psi^{1/4})_{x,y,b}\|_{V_1}^2 \leq 2 - 2(1 - \varepsilon)^2 \leq 4\varepsilon. \end{aligned}$$

Next, from the fact that $\sum_{x,y,b} \mu(x,y) \left\| h_\psi^{1/4} \bar{P}_b^{x,y} h_\psi^{1/4} \right\|_2^2 \geq (1 - \varepsilon)^2 \geq 1 - 2\varepsilon$, we can use [Lemma 3.3.11](#) to get that $\sum_{x \in X, y \in Y, a \in A_x} \mu(x,y) \left\| h_\psi^{1/4} \bar{P}_a^x h_\psi^{1/4} \right\|_2^2 \geq 1 - 2L\varepsilon$. Hence, it follows identically that

$$\sum_{x,y,a} \mu(x,y) \left\| [\bar{P}_a^x, h_\psi^{1/2}] \right\|_2^2 \leq 4L\varepsilon.$$

From here, it follows that

$$\begin{aligned} 1 - 2\varepsilon &\leq \sum_{x,y,b} \mu(x,y) \operatorname{tr} (Q_b^y h_\psi^{1/2} Q_b^y h_\psi^{1/2}) \\ &= \sum_{x,y,b} \mu(x,y) \langle h_\psi^{1/2} Q_b^y, Q_b^y h_\psi^{1/2} \rangle \\ &\leq \sum_{x,y,b} \mu(x,y) \|h_\psi^{1/2} Q_b^y\|_2 \|Q_b^y h_\psi^{1/2}\|_2 \\ &= \sum_{x,y,b} \mu(x,y) \operatorname{tr} ((Q_b^y)^2 h_\psi). \end{aligned}$$

In the same way,

$$\sum_{x,y,a} \mu(x,y) \operatorname{tr} ((\bar{P}_a^x)^2 h_\psi) \geq 1 - 2L\varepsilon.$$

Since for all $T \in \mathcal{M}$, $\text{tr}(Th_\psi) = \langle \psi | T | \psi \rangle$, the map $T \mapsto \text{tr}(Th_\psi)$ is a normal state on $\mathcal{M}$. Hence, by [55, Theorem 2], there exist PVMs $\{p_a^x\}_{a \in A_x}, \{q_b^y\}_{b \in B_y} \subseteq \mathcal{M}$ such that

$$\begin{aligned} \|((Q_b^y - q_b^y)h_\psi^{1/2})_{x,y,b}\|_{V_1}^2 &= \sum_{x,y,b} \mu(x,y) \text{tr}((Q_b^y - q_b^y)^2 h_\psi) \leq 18\varepsilon, \\ \|((\bar{P}_a^x - p_a^x)h_\psi^{1/2})_{x,y,a}\|_{V_2}^2 &= \sum_{x,y,a} \mu(x,y) \text{tr}((\bar{P}_a^x - p_a^x)^2 h_\psi) \leq 18L\varepsilon. \end{aligned}$$

Now, using again the norm on V_1 , we get the approximate commutation

$$\begin{aligned} \sum_{x,y,b} \mu(x,y) \|[q_b^y, h_\psi^{1/2}]\|_2^2 &= \|([q_b^y, h_\psi^{1/2}])_{x,y,b}\|_{V_1}^2 \\ &= \|([Q_b^y, h_\psi^{1/2}])_{x,y,b} - ((Q_b^y - q_b^y)h_\psi^{1/2})_{x,y,b} + (h_\psi^{1/2}(Q_b^y - q_b^y))_{x,y,b}\|_{V_1}^2 \\ &\leq \left(\|([Q_b^y, h_\psi^{1/2}])_{x,y,b}\|_{V_1} + 2\|((Q_b^y - q_b^y)h_\psi^{1/2})_{x,y,b}\|_{V_1} \right)^2 \\ &\leq (\sqrt{4\varepsilon} + 2\sqrt{18\varepsilon})^2 \leq 110\varepsilon, \end{aligned}$$

and in the same way using V_2 , $\sum_{x,y,a} \mu(x,y) \|[p_a^x, h_\psi^{1/2}]\|_2^2 \leq 110L\varepsilon$. To finish the proof, we want to make use of Lemma 3.3.10. Define $\tilde{X} = X \sqcup Y$ and $\tilde{A} = \sqcup_{x,y} A_x \sqcup B_y$. Take π to be the probability distribution on $\tilde{X} \times \tilde{X}$ defined as

$$\pi(x, x') = \begin{cases} \frac{1}{2}\mu(x, x') & x \in X \wedge x' \in Y \\ \frac{1}{2}\mu(x', x) & x' \in X \wedge x \in Y \\ 0 & \text{else} \end{cases}$$

Note that π is symmetric and the marginal of π on one copy of $\tilde{X}$ is

$$\nu(x) = \begin{cases} \frac{1}{2} \sum_{y \in Y} \mu(x, y) & x \in X \\ \frac{1}{2} \sum_{y \in X} \mu(y, x) & x \in Y \end{cases}.$$

Let $\tilde{p}_a^x = p_a^x$ if $x \in X$ and $\tilde{p}_a^x = q_a^x$ if $x \in Y$. By the above, we have that

$$\sum_{x \in \tilde{X}} \nu(x) \sum_a \|\tilde{p}_a^x, h_\psi^{1/2}\|_2^2 \leq 110L\varepsilon.$$

Then, using Lemma 3.3.10, we get that there exist PVMs $\{r_a^x\}_{a \in A} \subseteq qc(\mathcal{M})q$ such that

$$\sum_{x,y \in \tilde{X}} \pi(x,y) \sum_{a,b \in A} |\text{tr}(\tilde{p}_a^x h_\psi^{1/2} \tilde{p}_b^y h_\psi^{1/2}) - \tau(r_a^x r_b^y)| \leq 38(110L\varepsilon)^{1/4} \leq 124(L\varepsilon)^{1/4},$$

which implies

$$\sum_{x,y,a} \mu(x,y) \left| \text{tr}(p_a^x h_\psi^{1/2} q_{p_{i,j}(a)}^y h_\psi^{1/2}) - \tau(r_a^x r_{p_{x,y}(a)}^y) \right| \leq 124(L\varepsilon)^{1/4}.$$

Now, note that

$$\begin{aligned} \sum_{x,y,a} \mu(x,y) \text{tr}\left(\left(\bar{P}_a^x - (\bar{P}_a^y)^{1/2}\right)^2 h_\psi\right) &= \sum_{x,y,a} \mu(x,y) \text{tr}\left(\left((\bar{P}_a^x)^2 + \bar{P}_a^x - 2(\bar{P}_a^x)^{3/2}\right) h_\psi\right) \\ &\leq \sum_{x,y,a} \mu(x,y) \text{tr}\left(\left(\bar{P}_a^x - (\bar{P}_a^x)^2\right) h_\psi\right) \\ &\leq 1 - (1 - 2L\varepsilon) = 2L\varepsilon, \end{aligned}$$

and in the same way, $\sum_{x,y,b} \mu(x,y) \text{tr}\left((Q_b^y - (Q_b^y)^{1/2})^2 h_\psi\right) \leq 2\varepsilon$. Then, we can bound

$$\begin{aligned} &\left| \sum_{x,y,a} \mu(x,y) \text{tr}(\bar{P}_a^x h_\psi^{1/2} Q_{p_{x,y}(a)}^y h_\psi^{1/2}) - \sum_{x,y,a} \mu(x,y) \text{tr}(p_a^x h_\psi^{1/2} q_{p_{x,y}(a)}^y h_\psi^{1/2}) \right| \\ &= \left| \left\| ((\bar{P}_a^x)^{1/2} h_\psi^{1/2} (Q_{p_{x,y}(a)}^y)^{1/2})_{x,y,a} \right\|_{V_2}^2 - \left\| (p_a^x h_\psi^{1/2} q_{p_{x,y}(a)}^y)_{x,y,a} \right\|_{V_2}^2 \right| \\ &\leq 2 \left| \left\| ((\bar{P}_a^x)^{1/2} h_\psi^{1/2} (Q_{p_{x,y}(a)}^y)^{1/2})_{x,y,a} \right\|_{V_2} - \left\| (p_a^x h_\psi^{1/2} q_{p_{x,y}(a)}^y)_{x,y,a} \right\|_{V_2} \right| \\ &\leq 2 \left\| ((\bar{P}_a^x)^{1/2} h_\psi^{1/2} (Q_{p_{x,y}(a)}^y)^{1/2})_{x,y,a} - (p_a^x h_\psi^{1/2} q_{p_{x,y}(a)}^y)_{x,y,a} \right\|_{V_2} \\ &\leq 2 \left\| ((p_a^x - (\bar{P}_a^x)^{1/2}) h_\psi^{1/2} (Q_{p_{x,y}(a)}^y)^{1/2})_{x,y,a} \right\|_{V_2} + 2 \left\| (p_a^x h_\psi^{1/2} (q_{p_{x,y}(a)}^y - (Q_{p_{x,y}(a)}^y)^{1/2}))_{x,y,a} \right\|_{V_2} \\ &= 2 \sqrt{\sum_{x,y,a} \mu(x,y) \text{tr}((p_a^x - (\bar{P}_a^x)^{1/2})^2 h_\psi^{1/2} Q_{p_{x,y}(a)}^y h_\psi^{1/2})} + 2 \sqrt{\sum_{x,y,b} \mu(x,y) \text{tr}(p_b^{x,y} h_\psi^{1/2} (q_b^y - (Q_b^y)^{1/2})^2 h_\psi^{1/2})} \\ &\leq 2 \sqrt{\sum_{x,y,a} \mu(x,y) \text{tr}((p_a^x - (\bar{P}_a^x)^{1/2})^2 h_\psi)} + 2 \sqrt{\sum_{x,y,b} \mu(x,y) \text{tr}((q_b^y - (Q_b^y)^{1/2})^2 h_\psi)} \\ &= 2 \left\| ((p_a^x - (\bar{P}_a^x)^{1/2}) h_\psi^{1/2})_{x,y,a} \right\|_{V_2} + 2 \left\| ((q_b^y - (Q_b^y)^{1/2}) h_\psi^{1/2})_{x,y,b} \right\|_{V_1} \\ &\leq 2 \left\| ((p_a^x - \bar{P}_a^x) h_\psi^{1/2})_{x,y,a} \right\|_{V_2} + \left\| ((\bar{P}_a^x - (\bar{P}_a^x)^{1/2}) h_\psi^{1/2})_{x,y,a} \right\|_{V_2} + 2 \left\| ((q_b^y - Q_b^y) h_\psi^{1/2})_{x,y,b} \right\|_{V_1} \\ &\quad + 2 \left\| ((Q_b^y - (Q_b^y)^{1/2}) h_\psi^{1/2})_{x,y,b} \right\|_{V_1} \\ &\leq 2\sqrt{18L\varepsilon} + 2\sqrt{2L\varepsilon} + 2\sqrt{18\varepsilon} + 2\sqrt{2\varepsilon} \\ &\leq 23(L\varepsilon)^{1/2}. \end{aligned}$$

Via the triangle inequality, we get that

$$\left| \sum_{x,y,a} \mu(x,y) \left(\langle \psi | P_a^x Q_{p_{x,y}(a)}^y | \psi \rangle - \tau(r_a^x r_{p_{x,y}(a)}^y) \right) \right| \leq 23(L\varepsilon)^{1/2} + 124(L\varepsilon)^{1/4} \leq 147(L\varepsilon)^{1/4}.$$

To finish, take $\{\hat{p}_a^x\}_{a \in A_x}$ to be the completion of $\{r_a^x\}_{a \in A_x}$ to a PVM by adding $1 - \sum_{a \in A_x} r_a^x$ to an arbitrary element; and construct $\{\hat{q}_b^y\}_{b \in B_y}$ from $\{r_b^y\}_{b \in B_y}$ in the same way. Then $p'(a, b|x, y) = \tau(\hat{p}_a^x \hat{q}_b^y)$ is a tracial correlation and we find that

$$\begin{aligned} \mathfrak{w}(\mathbf{G}, p') &= \sum_{x,y} \mu(x, y) \sum_{a \in A_x} \tau(\hat{p}_a^x \hat{q}_{p_{x,y}(a)}^y) \\ &\geq \sum_{x,y} \mu(x, y) \sum_{a \in A_x} \tau(r_a^x r_{p_{x,y}(a)}^y) \\ &\geq \mathfrak{w}(\mathbf{G}, p) - \left| \sum_{x,y,a} \mu(x, y) \left(\langle \psi | P_a^x Q_{p_{x,y}(a)}^y | \psi \rangle - \tau(r_a^x r_{p_{x,y}(a)}^y) \right) \right| \\ &\geq 1 - 148(L\varepsilon)^{1/4}. \end{aligned}$$

Corollary 3.3.12. Let $\mathbf{G}$ be as in [Theorem 3.3.9](#). Suppose $p(a, b|x, y) = \langle \psi | P_a^x \otimes Q_b^y | \psi \rangle$ is a quantum correlation such that $\mathfrak{w}(\mathbf{G}, p) \geq 1 - \varepsilon$. Then, there exists a finite-dimensional tracial correlation p' such that $\mathfrak{w}(\mathbf{G}, p') \geq 1 - 148(L\varepsilon)^{1/4}$.

This provides an tightening of [[117](#), Theorem 4.6] in the context of projection games where Bob has few answers, as that is the case in which where L is small. Oracularisations and oracular constraint-variable games both fulfill this criterion.

Proof. Take $\mathcal{M}$ to be the von Neumann algebra generated by the $1 \otimes Q_b^y$; this is a finite-dimensional von Neumann algebra contained in $\mathcal{B}(H_A \otimes H_B)$. Then, following [Example 2.4.2](#), $c(\mathcal{M}) = L^\infty(\mathbb{R}) \otimes \mathcal{M}$. From the theorem, there exists a correlation $p'(a, b|x, y) = \tau(p_a^x q_b^y)$ such that $\mathfrak{w}(\mathbf{G}, p') \geq 1 - 148(L\varepsilon)^{1/4}$, where τ is the semifinite weight $\tau(T) = \int_{-\infty}^{\infty} \text{Tr}(T(t)) e^{-t} dt$ on $c(\mathcal{M})$, and $p_a^x, q_b^y \in qc(\mathcal{M})q$ for some projection $q \in c(\mathcal{M})$ satisfying $\tau(q) = 1$. By construction, $q(t), p_a^x(t), q_b^y(t)$ are projections for every $t \in \mathbb{R}$, so

$$\tau(p_a^x q_b^y) = \int_{-\infty}^{\infty} \text{Tr}(p_a^x(t) q_b^y(t)) e^{-t} dt = \int_{\{t | \text{Tr}(q(t)) \neq 0\}} \frac{1}{\text{Tr}(q(t))} \text{Tr}(p_a^x(t) q_b^y(t)) \text{Tr}(q(t)) e^{-t} dt$$

is a convex combination of finite-dimensional tracial correlations. This implies that at least one of the correlations has winning probability $\geq 1 - 148(L\varepsilon)^{1/4}$, completing the proof. ■

Now, we show that the same result can be recovered for degenerate games.

Corollary 3.3.13. Suppose $\mathbf{G}$ is an arbitrary projection game. Suppose there exists a commuting operator correlation p for $\mathbf{G}$ such that $\mathfrak{w}(\mathbf{G}, p) \geq 1 - \varepsilon$. Then, there exists a tracial correlation p' for $\mathbf{G}$ such that $\mathfrak{w}(\mathbf{G}, p') \geq 1 - 148(L\varepsilon)^{1/4}$. If p is quantum, then p' can be chosen to be finite-dimensional.

Proof. If G is non-degenerate, this is exactly the result of [Theorem 3.3.9](#) and [Corollary 3.3.12](#). Now, suppose G is degenerate. Let $\tilde{\mathsf{G}}$ be the non-degenerate game constructed in [Lemma 3.3.2](#). By that lemma, there exists a commuting operator correlation $\tilde{p}$ such that $\mathfrak{w}(\tilde{\mathsf{G}}, \tilde{p}) = \mathfrak{w}(\mathsf{G}, p) \geq 1 - \varepsilon$. Then, by [Theorem 3.3.9](#), there exists a tracial correlation $\tilde{p}'$ for $\tilde{\mathsf{G}}$ such that $\mathfrak{w}(\tilde{\mathsf{G}}, \tilde{p}') \geq 1 - 148(L\varepsilon)^{1/4}$. Using again [Lemma 3.3.2](#), there exists a tracial correlation p' for G such that $\mathfrak{w}(\mathsf{G}, p') = \mathfrak{w}(\tilde{\mathsf{G}}, \tilde{p}') \geq 1 - 148(L\varepsilon)^{1/4}$. The result for quantum correlations holds using [Corollary 3.3.12](#). $\blacksquare$

3.3.5 Application to CSPs and oracularisation

In this section, we study the implications of our main result for oracular constraint-variable games and oracularisations. First, we consider oracularisations.

Recall that the oracularisation of a nonlocal game $\mathsf{G} = (X, Y, \{A_x\}_x, \{B_y\}_y, \mu, V)$, as defined in [Section 2.12.3](#), is the nonlocal game G^{orac} where Alice receives both questions and Bob receives one of them, and they win if Alice satisfies the predicate and Bob's answer is consistent. The oracularisation of a game is always a projection game via the projection functions $p_{(x,y),x}((a,b)) = a$ and $p_{(x,y),y}((a,b)) = b$. Prior work [\[76, 87\]](#) has made use of a more involved synchronous version of oracularisation to reduce to tracial strategies, but we show that this is not necessary.

Corollary 3.3.14. Let G be a nonlocal game and let p be a commuting operator correlation for G^{orac} . If $\mathfrak{w}(\mathsf{G}^{\text{orac}}, p) \geq 1 - \varepsilon$, then there exists a tracial correlation p' for G^{orac} such that $\mathfrak{w}(\mathsf{G}^{\text{orac}}, p') \geq 1 - 177\varepsilon^{1/4}$. If p is quantum, then p' can be chosen to be finite-dimensional.

Proof. This follows directly from [Theorem 3.3.9](#) and [Corollary 3.3.12](#) by noting that

$$L = \max_{x,y,z: \mu^{\text{orac}}((x,y),z) \neq 0} \frac{\mu^{\text{orac}}((x,y))}{\mu^{\text{orac}}((x,y),z)} = \max_{x,y: \mu(x,y) \neq 0} \frac{\mu(x,y)}{\frac{1}{2}\mu(x,y)} = 2$$

and then that $148(2)^{1/4} \leq 177$. $\blacksquare$

Now, we turn to oracular constraint-variable games, as defined in [Definition 3.1.2](#).

Corollary 3.3.15. Let X and A be σ -structures, and let π be probability distribution over the constraints of X . If p is a commuting operator correlation for $\mathsf{G}_{\text{ocv}}(X, A, \pi)$ such that $\mathfrak{w}(\mathsf{G}_{\text{ocv}}(X, A, \pi), p) \geq 1 - \varepsilon$, then there exists a tracial correlation p' for $\mathsf{G}_{\text{ocv}}(X, A, \pi)$ such that $\mathfrak{w}(\mathsf{G}_{\text{ocv}}(X, A, \pi), p') \geq 1 - 148(L\varepsilon)^{1/4}$, where $L = \max_{R \in \sigma} \text{ar}(R)$. If p is quantum, p' can be chosen to be finite-dimensional.

Proof. This follows directly from [Theorem 3.3.9](#) and [Corollary 3.3.12](#) by noting that

$$L = \max_{R \in \sigma, \mathbf{x} \in R^X, i \in \text{ar}(R): \mu((R, \mathbf{x}), x_i) \neq 0} \frac{\mu((R, \mathbf{x}))}{\mu((R, \mathbf{x}), x_i)} = \max_{R \in \sigma, \mathbf{x} \in R^X: \pi(R, \mathbf{x}) > 0} \frac{\pi(R, \mathbf{x})}{\frac{\pi(R, \mathbf{x})}{\text{ar}(R)}} = \max_{R \in \sigma} \text{ar}(R). \blacksquare$$

This also implies that results in the weighted algebra setting, which considers only tracial states, can now be translated directly to the game setting, without having to work with a synchronised version of the game with consistency checks.

Corollary 3.3.16. Let X , A , π , and L be as in the previous corollary.

1. If there exists a commuting operator correlation p such that $\mathfrak{w}(\mathbb{G}_{ocv}(X, A, \pi), p) \geq 1 - \varepsilon$, then there exists a tracial state τ on $\text{Mor}_\pi^{ocv}(X, A)$ such that $\text{def}(\tau) \leq 148(L\varepsilon)^{1/4}$.
2. If there exists a quantum correlation p such that $\mathfrak{w}(\mathbb{G}_{ocv}(X, A, \pi), p) \geq 1 - \varepsilon$, then there exists a finite-dimensional tracial state τ on $\text{Mor}_\pi^{ocv}(X, A)$ such that $\text{def}(\tau) \leq 148(L\varepsilon)^{1/4}$.

Hence, we find a version of [Lemma 3.1.13](#) for the oracular constraint-variable framework.

Lemma 3.3.17. Let A be a σ -structure (where σ is finite), let $0 \leq s \leq 1$, and let $Q \in \{q, co\}$. Then, there exists $0 \leq s' \leq 1$ such that $s' = 1 \iff s = 1$ for which there are polynomial time reductions from $\text{CSP}_{ocv}^Q(A)_{1,s}$ to ${}_g\text{CSP}_{ocv}^Q(A)_{1,s'}$, ${}^e\text{CSP}_{ocv}^Q(A)_{1,s}$ to ${}^e\text{CSP}_{ocv}^Q(A)_{1,s'}$, $\text{MIP}_{ocv}^Q(A)_{1,s}$ to ${}_g\text{MIP}_{ocv}^Q(A)_{1,s'}$, and ${}^e\text{MIP}_{ocv}^Q(A)_{1,s}$ to ${}^e\text{MIP}_{ocv}^Q(A)_{1,s'}$.

The proof is identical to [Lemma 3.1.13](#), using [Theorem 3.3.9](#) and [Corollary 3.3.12](#) instead of [\[86\]](#).

Chapter 4

Commutativity gadgets

In this chapter, we develop the theory of commutativity gadgets. In [Section 4.1](#), we define commutativity gadgets and show how they can be used to show reductions. In [Section 4.2](#), we show that graph colouring admits an oracular commutativity gadget. In [Section 4.3](#), we relate commutativity gadgets of a relational structure to the behaviour of the polymorphism algebra. In [Section 4.4](#), we fully characterise the polymorphism algebras of the complete graphs. In [Section 4.5](#), we show separations between different types of commutativity gadget. In [Section 4.6](#), we study how hard it is to decide if a commutativity gadget exists for a given entangled CSP.

4.1 Reductions based on commutativity gadgets

In this section, we study the properties afforded by the existence of a commutativity gadget. In [Section 4.1.1](#), we define the different notions of commutativity gadget. In [Section 4.1.2](#), we study the complexity-theoretic implications of the existence of a commutativity gadget. In [Section 4.1.3](#), we study a stronger C^* -algebraic notion of gadget which will be useful in existence proofs. In [Section 4.1.4](#), we provide some of the examples of commutativity gadget known from prior work.

4.1.1 Definitions

A commutativity gadget for an entangled CSP is a gadget that allows classical relations on the variables of the CSP to be expressed.

Definition 4.1.1. Let Q be a tensor monoid, and let A be a σ -structure. A **(non-oracular) Q -commutativity gadget** for A is a σ -structure G along with two elements $x, y \in G$, called the **distinguished variables**, such that

- (i) For every $a, b \in A$, there exists a $*$ -homomorphism $\pi_{a,b} : \text{Mor}^+(G, A) \rightarrow \mathcal{A}$ such that $\mathcal{A} \in Q$ and $\pi_{a,b}(p_{x,a}) = \pi_{a,b}(p_{y,b}) = 1$.
- (ii) For every $*$ -homomorphism $\pi : \text{Mor}^+(G, A) \rightarrow \mathcal{A}$ such that $\mathcal{A} \in Q$, $[\pi(p_{x,a}), \pi(p_{y,b})] = 0$ for all $a, b \in A$.

An **oracular Q -commutativity gadget** for A is a relational structure G over σ along with two distinguished variables $x, y \in G$ such that

- (i) For every $a, b \in A$, there exists a $*$ -homomorphism $\pi_{a,b} : \text{Mor}^{o+}(G, A) \rightarrow \mathcal{A}$ such that $\mathcal{A} \in Q$ and $\pi_{a,b}(p_{x,a}) = \pi_{a,b}(p_{y,b}) = 1$.
- (ii) For every $*$ -homomorphism $\pi : \text{Mor}^{o+}(G, A) \rightarrow \mathcal{A}$ such that $\mathcal{A} \in Q$, $[\pi(p_{x,a}), \pi(p_{y,b})] = 0$ for all $a, b \in A$.

Robust commutativity gadgets extend the expressive power of commutativity gadgets to the gapped setting.

Definition 4.1.2. A **robust Q -commutativity gadget** for A is a Q -commutativity gadget (G, x, y) such that for all $\varepsilon > 0$, there exists $\delta > 0$ such that if $\text{def}(\tau) \leq \delta$ for a tracial state τ on $\text{Mor}_{\mathfrak{u}}^{no}(G, A)$ over Q , then

$$\sum_{a,b \in A} \|[p_{x,a}, p_{y,b}]\|_{\tau}^2 \leq \varepsilon.$$

A **robust oracular Q -commutativity gadget** for A is an oracular Q -commutativity gadget (G, x, y) such that for all $\varepsilon > 0$, there exists $\delta > 0$ such that if $\text{def}(\tau) \leq \delta$ for a tracial state τ on $\text{Mor}_{\mathfrak{u}}^{ocv}(G, A)$ over Q , then

$$\sum_{a,b \in A} \|[p_{x,a}, p_{y,b}]\|_{\tau}^2 \leq \varepsilon.$$

When dealing with robust commutativity gadgets, we want to restrict to those tensor monoids where the von Neumann algebras admit tracial states.

Definition 4.1.3. We say a tensor monoid is **finite** if it consists of finite von Neumann algebras.

The tensor monoids d , c , q , qa , and co are all finite. A tensor monoid is finite if and only if it is a subset of co . Note also that the tensor monoid vN is not finite. Given a finite tensor monoid, each element may be associated a tracial state, but this choice is not unique; for the moment, this non-uniqueness does not matter.

4.1.2 Hardness from the existence of a commutativity gadget

The expressivity provided by commutativity gadgets immediately translates to hardness results. First, we show that any entangled CSP may be reduced to the CSP whose constraints are instances of an NP-complete CSP A . Then, we will use a commutativity gadget to split these instances up into separate constraints of A .

Lemma 4.1.4. Let A be a σ -structure, B be a τ -structure, and let Q be a finite tensor monoid. Suppose that $\text{CSP}(A)$ is NP-complete. Let $\bar{A}$ be the relational structure with the same domain as A and relations labelled by σ -structures X with arity $|X|$, such that $X^{\bar{A}} = \text{mor}(X, A)$. For all $s \leq 1$ there exists $s' \leq 1$ with $s' = 1 \iff s = 1$ such that there are polynomial-time reductions from ${}^e\text{MIP}_{ocv}^Q(B)_{1,s}$ to ${}^e\text{MIP}_{ocv}^Q(\bar{A})_{1,s'}$ and from $\text{MIP}_{ocv}^Q(B)_{1,s}$ to $\text{MIP}_{ocv}^Q(\bar{A})_{1,s'}$.

Proof. Since $\text{CSP}(A)$ is NP-complete, there exists a minion homomorphism $\xi : \text{pol}(A) \rightarrow \text{pol}(B)$. The signature of $\bar{A}$ is the set of σ -structures, which we call $\bar{\sigma}$; a relation of a $\bar{\sigma}$ -structure labelled by a σ -structure X has arity given by the cardinality of X , so the tuples will be indexed by the elements of $\mathfrak{D}(X)$. Let M be an instance of ${}^e\text{MIP}_{ocv}^Q(B)_{1,s}$. For each constraint $(R, \mathbf{y})$ of $X(M)$, construct the σ -structure $X_{R,\mathbf{y}}$ from the τ -structure $Y_{R,\mathbf{y}}$ with domain $\{y_1, \dots, y_{\text{ar}(R)}\}$ and single nontrivial relation $R^{Y_{R,\mathbf{y}}} = \{\mathbf{y}\}$ according to the prescription of Theorem 2.10.2. We define a $\bar{\sigma}$ -structure $\bar{X}_{R,\mathbf{y}}$ with the same domain as $X_{R,\mathbf{y}}$ and a single nontrivial relation $X_{R,\mathbf{y}}^{\bar{X}_{R,\mathbf{y}}} = \{\mathbf{x}^{R,\mathbf{y}}\}$ where $x_z^{R,\mathbf{y}} = z$ for $z \in X_{R,\mathbf{y}}$. By Theorem 2.10.2, there exists a map $\Gamma_{R,\mathbf{y}}$ that sends satisfying assignments to $Y_{R,\mathbf{y}}$ to satisfying assignments to $\bar{X}_{R,\mathbf{y}}$, and a map $\Delta_{R,\mathbf{y}}$ that sends satisfying assignments to $\bar{X}_{R,\mathbf{y}}$ to satisfying assignments to $Y_{R,\mathbf{y}}$. Now, let $X_0 = \bigsqcup_{R \in \tau, \mathbf{y} \in R^Y} \bar{X}_{R,\mathbf{y}}$, and let $X = X_0 / \sim$, where the equivalence relation identifies variables of the form $[(y, \mathbf{a})]$ with the same label between the terms of the disjoint union. Define M' as the probabilistic Turing machine that samples a constraint $(R, \mathbf{y})$ using M , and then outputs $(X_{R,\mathbf{y}}, \mathbf{x}^{R,\mathbf{y}})$. The mapping $M \mapsto M'$ is clearly polynomial-time.

Let $L = \max_{R \in \tau, \mathbf{y} \in R^Y} |X_{R,\mathbf{y}}|$, and $s' = 1 - \left(\frac{1-s}{267|A|^{|B|}|A|^{\frac{3}{2}|A||B|}L} \right)^2$.

Next, suppose M is a yes instance of ${}^e\text{MIP}_{ocv}^Q(B)_{1,s}$. Then, there exists a $*$ -homomorphism $\varphi : \text{Mor}^{o+}(X(M), B) \rightarrow \mathcal{M} \in Q$. Using [Lemma 3.2.2](#), define a $*$ -homomorphism $\varphi' : \text{Mor}^{o+}(X(M'), A) \rightarrow \mathcal{M}$ via $\varphi'(\Phi_{X_{R,\mathbf{y}}, \mathbf{x}^{R,\mathbf{y}}, \mathbf{a}}) = \sum_{\mathbf{b} \in \Gamma_{R,\mathbf{y}}^{-1}(\mathbf{a})} \varphi(\Phi_{R,\mathbf{y},\mathbf{b}})$. Clearly, $\{\varphi'(\Phi_{X_{R,\mathbf{y}}, \mathbf{x}^{R,\mathbf{y}}, \mathbf{a}})\}_{\mathbf{a}}$ is a PVM. Now, suppose $R, R' \in \tau$, $\mathbf{y} \in R^Y$, $\mathbf{y}' \in R'^Y$, $\mathbf{a} \in X_{R,\mathbf{y}}^A$, $\mathbf{a}' \in X_{R',\mathbf{y}'}^A$, $z \in X_{R,\mathbf{y}}$, and $z' \in X_{R',\mathbf{y}'}$ such that $x_z^{R,\mathbf{y}} = z = x_{z'}^{R',\mathbf{y}'} = z'$ but $a_z \neq a'_{z'}$. Then, we must have that there exists $y \in Y$, $\mathbf{a}_0 \in A^{\mathfrak{D}(B)}$, $k \in \text{ar}(R)$, and $l \in \text{ar}(R')$ such that $z = z' = [(y, \mathbf{a}_0)]$ and $y_k = y'_l = y$. We can expand

$$\varphi'(\Phi_{X_{R,\mathbf{y}}, \mathbf{x}^{R,\mathbf{y}}, \mathbf{a}}) \varphi'(\Phi_{X_{R',\mathbf{y}'}', \mathbf{x}^{R',\mathbf{y}'}, \mathbf{a}'}) = \sum_{\mathbf{b} \in \Gamma_{R,\mathbf{y}}^{-1}(\mathbf{a}), \mathbf{b}' \in \Gamma_{R',\mathbf{y}'}^{-1}(\mathbf{a}')} \varphi(\Phi_{R,\mathbf{y},\mathbf{b}} \Phi_{R',\mathbf{y}',\mathbf{b}'})$$

If $b_k \neq b'_l$, then $\varphi(\Phi_{R,\mathbf{y},\mathbf{b}} \Phi_{R',\mathbf{y}',\mathbf{b}'}) = 0$. Now, suppose that for one of these terms, $b_k = b'_l$. Using [Theorem 2.10.2](#), $\mathbf{a} = \Gamma_{R,\mathbf{y}}(\mathbf{b})$ implies that $a_z = \mathbf{a}_0(b_k)$ and $\mathbf{a}' = \Gamma_{R',\mathbf{y}'}(\mathbf{b}')$ implies that $a'_{z'} = \mathbf{a}_0(b'_l)$. But $a_z \neq a'_{z'}$, so this is a contradiction. Hence, $\varphi'(\Phi_{X_{R,\mathbf{y}}, \mathbf{x}^{R,\mathbf{y}}, \mathbf{a}}) \varphi'(\Phi_{X_{R',\mathbf{y}'}', \mathbf{x}^{R',\mathbf{y}'}, \mathbf{a}'}) = 0$, and therefore φ' is a $*$ -homomorphism, and M' is a yes instance of ${}^e\text{MIP}_{ocv}^Q(\bar{A})_{1,s'}$.

Now, suppose M is a no instance of ${}^e\text{MIP}_{ocv}^Q(B)_{1,s}$. Suppose, for a contradiction, that there exists a tracial state τ on $\text{Mor}_{\pi_{M'}}^{ocv}(X(M'), \bar{A})$ over Q such that $\text{def } \tau \leq 1 - s'$. Let $\varphi : \text{Mor}^{ocv}(X(M'), \bar{A}) \rightarrow (\mathcal{M}, \text{tr}) \in Q$ be a GNS representation of τ . For every $y \in Y$, we want to construct a PVM in $\mathcal{M}$ valued in the polymorphisms of A at level $|B|$. For every constraint $(R, \mathbf{y})$ such that $y_i = y$ for some i , $\mathbf{a} \in A^{\mathfrak{D}(B)}$, and $\phi \in \text{mor}(A^{\mathfrak{D}(B)}, A)$, let $\Phi_{X_{R,\mathbf{y}}, \mathbf{x}^{R,\mathbf{y}}, y, \phi} = \prod_{\mathbf{a} \in A^{\mathfrak{D}(B)}} \Phi_{X_{R,\mathbf{y}}, \mathbf{x}^{R,\mathbf{y}}, [(y, \mathbf{a})], \phi(\mathbf{a})}$, where we recall that for $z \in X_{R,\mathbf{y}}$ and $\mathbf{a} \in \bar{A}$, $\Phi_{X_{R,\mathbf{y}}, \mathbf{x}^{R,\mathbf{y}}, z, \mathbf{a}} = \sum_{\mathbf{a} \in X_{R,\mathbf{y}}^{\bar{A}} : a_z = \mathbf{a}} \Phi_{X_{R,\mathbf{y}}, \mathbf{x}^{R,\mathbf{y}}, \mathbf{a}}$. For $y \in Y$, write $\pi(y) = \sum_{R \in \tau, \mathbf{y} \in R^Y, i \in \text{ar}(R) : y_i = y} \frac{\pi(R, \mathbf{y})}{\text{ar}(R)}$, and write $\pi(R, \mathbf{y}|y) = \frac{\pi(R, \mathbf{y})}{\text{ar}(R)\pi(y)}$ if there exists i such that $y_i = y$ and $\pi(R, \mathbf{y}|y) = 0$ otherwise. With this notation, we have that

$$\begin{aligned} & \sum_{y \in Y} \pi(y) \sum_{R \in \tau, \mathbf{y} \in R^Y} \pi(R, \mathbf{y}|y) \sum_{\phi \in \text{mor}(A^{\mathfrak{D}(B)}, A)} \left\| \Phi_{X_{R,\mathbf{y}}, \mathbf{x}^{R,\mathbf{y}}, y, \phi} - \prod_{\mathbf{a} \in A^{\mathfrak{D}(B)}} p_{[(y, \mathbf{a})], \phi(\mathbf{a})} \right\|_{\tau}^2 \\ & \leq |A|^{|B|} \sum_{y \in Y} \pi(y) \sum_{R \in \tau, \mathbf{y} \in R^Y} \pi(R, \mathbf{y}|y) \sum_{\mathbf{a} \in A^{\mathfrak{D}(B)}, a \in A} \left\| \Phi_{X_{R,\mathbf{y}}, \mathbf{x}^{R,\mathbf{y}}, [(y, \mathbf{a})], a} - p_{[(y, \mathbf{a})], a} \right\|_{\tau}^2 \\ & \leq 2|A|^{|B|} L \sum_{R \in \tau, \mathbf{y} \in R^Y, i \in \text{ar}(R)} \frac{\pi(R, \mathbf{y})}{\text{ar}(X_{R,\mathbf{y}})} \sum_{\mathbf{a} \in A^{\mathfrak{D}(B)}, a \in A} \tau(\Phi_{X_{R,\mathbf{y}}, \mathbf{x}^{R,\mathbf{y}}, [(y_i, \mathbf{a})], a} (1 - p_{[(y_i, \mathbf{a})], a})) \\ & \leq 2|A|^{|B|} L \text{def}(\tau). \end{aligned}$$

Let $q_{y,\phi} = \prod_{a \in A^{\mathfrak{D}(B)}} \varphi(p[(y,a)], \phi(a))$, and let $Q_{y,\phi} = q_{y,\phi}^* q_{y,\phi} + \Delta$, where $\Delta = \frac{1 - \sum_{\phi \in \text{mor}(A^{\mathfrak{D}(B)}, A)} q_{y,\phi}^* q_{y,\phi}}{|\text{mor}(A^{\mathfrak{D}(B)}, A)|}$. Since $\Delta \geq 0$, $\{Q_{y,\phi}\}_{\phi \in \text{mor}(A^{\mathfrak{D}(B)}, A)}$ is a POVM. We also have that

$$\begin{aligned} \sum_{\phi \in \text{mor}(A^{\mathfrak{D}(B)}, A)} \left\| \varphi(\Phi_{X_{R,y}, \mathbf{x}^{R,y}, y, \phi}) - Q_{y,\phi} \right\|_{\text{tr}}^2 &\leq 4 \sum_{\phi \in \text{mor}(A^{\mathfrak{D}(B)}, A)} \left\| \varphi(\Phi_{X_{R,y}, \mathbf{x}^{R,y}, y, \phi}) - q_{y,\phi}^* q_{y,\phi} \right\|_{\text{tr}}^2 \\ &\leq 16 \sum_{\phi \in \text{mor}(A^{\mathfrak{D}(B)}, A)} \left\| \varphi(\Phi_{X_{R,y}, \mathbf{x}^{R,y}, y, \phi}) - q_{y,\phi} \right\|_{\text{tr}}^2, \end{aligned}$$

so

$$\begin{aligned} &\sum_{y \in Y} \pi(y) \left(1 - \sum_{\phi \in \text{mor}(A^{\mathfrak{D}(B)}, A)} \text{tr}(Q_{y,\phi}^2) \right) \\ &= \sum_{y \in Y} \pi(y) \sum_{R \in \tau, \mathbf{y} \in R^Y} \pi(R, \mathbf{y} | y) \sum_{\phi \in \text{mor}(A^{\mathfrak{D}(B)}, A)} \left(\left\| \varphi(\Phi_{X_{R,y}, \mathbf{x}^{R,y}, y, \phi}) \right\|_{\text{tr}}^2 - \|Q_{y,\phi}\|_{\text{tr}}^2 \right) \\ &\leq 2 \sum_{y \in Y} \pi(y) \sum_{R \in \tau, \mathbf{y} \in R^Y} \pi(R, \mathbf{y} | y) \sum_{\phi \in \text{mor}(A^{\mathfrak{D}(B)}, A)} \left\| \varphi(\Phi_{X_{R,y}, \mathbf{x}^{R,y}, y, \phi}) - Q_{y,\phi} \right\|_{\text{tr}} \\ &\leq 2 \sqrt{32 |A|^{|B|} |A|^{|A|^{|B|}} L \text{def}(\tau)} \end{aligned}$$

Therefore, by [55], there exist PVMs $\{P_{y,\phi}\}_{\phi \in \text{mor}(A^{\mathfrak{D}(B)}, A)} \subseteq \mathcal{M}$ for all $y \in Y$ such that $\sum_{y \in Y} \pi(y) \|Q_{y,\phi} - P_{y,\phi}\|_{\text{tr}}^2 \leq 18 \sqrt{32 |A|^{|B|} |A|^{|A|^{|B|}} L \text{def}(\tau)}$.

Now, we define the $*$ -homomorphism $\varphi' : \text{Mor}^{ocv}(X(M), B) \rightarrow \mathcal{M}$ on the generators as $\varphi'(\Phi_{R,y}, b) = \sum_{a \in \Delta_{R,y}^{-1}(b)} \varphi(\Phi_{X_{R,y}, \mathbf{x}^{R,y}, a})$ and $\varphi'(p_{y,b}) = \sum_{\phi \in \text{mor}(A^{\mathfrak{D}(B)}, A) : \xi(\phi)(\text{id}_B) = b} P_{y,\phi}$. Note that if $\Delta_{R,y}(\mathbf{a}) = \mathbf{b}$, then for any $i \in \text{ar}(R)$, the polymorphism $A^{\mathfrak{D}(B)} \rightarrow A$, $\mathbf{a}' \mapsto a[(y, \mathbf{a}')]]$ maps under ξ to a polymorphism $B^{\mathfrak{D}(B)} \rightarrow B$ under which $\text{id} \mapsto b_i$. As such, $\varphi'(\Phi_{R,y,i}, b) = \sum_{a \in X_{R,A}^{-1} : \Delta_{R,y}(a)_i = b} \varphi(\Phi_{X_{R,y}, \mathbf{x}^{R,y}, a}) = \sum_{\phi \in \text{mor}(A^{\mathfrak{D}(B)}, A) : \xi(\phi)(\text{id}_B) = b} \varphi(\Phi_{X_{R,y}, \mathbf{x}^{R,y}, y_i, \phi})$. Let $\tau' = \text{tr} \circ \varphi$. τ' is a tracial state on $\text{Mor}^{ocv}(Y, B)$ and it has defect

$$\begin{aligned} \text{def}(\tau) &= \frac{1}{2} \sum_{\substack{R \in \tau, \mathbf{y} \in R^Y \\ i \in \text{ar}(R), b \in B}} \frac{\pi(R, \mathbf{y})}{\text{ar}(R)} \|\Phi_{R,y,i}, b - p_{y,b}\|_{\tau'}^2 \\ &\leq \frac{|A|^{|A|^{|B|}}}{2} \sum_{\substack{R \in \tau, \mathbf{y} \in R^Y \\ y \in Y, \phi \in \text{mor}(A^{\mathfrak{D}(B)}, A)}} \pi(y) \pi(R, \mathbf{y} | y) \left\| \varphi(\Phi_{X_{R,y}, \mathbf{x}^{R,y}, y, \phi}) - P_{y,\phi} \right\|_{\text{tr}}^2 \\ &\leq 2 |A|^{|A|^{|B|}} \sum_{\substack{R \in \tau, \mathbf{y} \in R^Y \\ y \in Y, \phi \in \text{mor}(A^{\mathfrak{D}(B)}, A)}} \pi(y) \pi(R, \mathbf{y} | y) \left(\left\| \varphi(\Phi_{X_{R,y}, \mathbf{x}^{R,y}, y, \phi}) - Q_{y,\phi} \right\|_{\text{tr}}^2 + \|Q_{y,\phi} - P_{y,\phi}\|_{\text{tr}}^2 \right) \end{aligned}$$

$$\begin{aligned}
&\leq 2|A|^{|A|^{|B|}} \left(32|A|^{|B|} L \text{def}(\tau) + 18\sqrt{32|A|^{|B|}|A|^{|A|^{|B|}}} L \text{def}(\tau) \right) \\
&\leq 267|A|^{|B|}|A|^{\frac{3}{2}|A|^{|B|}} L\sqrt{1-s'} = 1-s.
\end{aligned}$$

Hence, M is not a no instance of ${}^e\text{MIP}_{ocv}^Q(B)_{1,s}$, which is a contradiction, finishing the proof. The non-exact case is proved identically. $\blacksquare$

Lemma 4.1.5. Let A and B be relational structures and let Q be a finite tensor monoid. Suppose that $\text{CSP}(A)$ is NP-complete.

- (i) If A admits an oracular Q -commutativity gadget, then there is a polynomial-time reduction from ${}^e\text{MIP}_{ocv}^Q(B)_{1,1}$ to ${}^e\text{MIP}_{ocv}^Q(A)_{1,1}$.
- (ii) If A admits a Q -commutativity gadget, then there is a polynomial-time reduction from ${}^e\text{MIP}_{ocv}^Q(B)_{1,1}$ to ${}^e\text{MIP}_{no}^Q(A)_{1,1}$.
- (iii) If A admits a robust oracular Q -commutativity gadget, then for all $s < 1$ there exists $s' < 1$ such that there is a polynomial-time reduction from ${}^e\text{MIP}_{ocv}^Q(B)_{1,s}$ to ${}^e\text{MIP}_{ocv}^Q(A)_{1,s'}$ and a polynomial-time reduction from $\text{MIP}_{ocv}^Q(B)_{1,s}$ to $\text{MIP}_{ocv}^Q(A)_{1,s'}$.
- (iv) If A admits a robust Q -commutativity gadget, then for all $s < 1$ there exists $s' < 1$ such that there is a polynomial-time reduction from ${}^e\text{MIP}_{ocv}^Q(B)_{1,s}$ to ${}^e\text{MIP}_{no}^Q(A)_{1,s'}$ and a polynomial-time reduction from $\text{MIP}_{ocv}^Q(B)_{1,s}$ to $\text{MIP}_{no}^Q(A)_{1,s'}$.

In the proof below, we make use of [Lemma 4.1.4](#) to reduce each constraint of B to a collection of constraints of A , and a commutativity gadget to separate these collections into single constraints. It is possible, however, to replace [Lemma 4.1.4](#) with any local classical reduction from B to A , which may be simpler than the one coming from the minion homomorphism in some cases.

Proof. Using [Lemma 4.1.4](#), we may assume that $B = \overline{A}$, the relational structure where the relations are σ -structures; and using the reduction in the proof, we may assume that these σ -structures have finite size at most N depending only on A and B .

Now, for all four cases, we implement the reduction in the same way. Let Y be a $\overline{\sigma}$ -structure. Then, every constraint $(Z, \mathbf{y})$ of Y corresponds to a σ -structure $Z_{\mathbf{y}}$ with relations $R^{Z_{\mathbf{y}}} = \{(y_{z_1}, \dots, y_{z_{\text{ar}(R)}}) \mid \mathbf{z} \in R^Z\}$ for $R \in \sigma$. For $y, y' \in Y$ that belong to the same relation of $\overline{A}$, let $G_{y,y'}$ be a copy of the Q -commutativity gadget with distinguished elements x_G and y_G . Let $X_0 = \sqcup_{Z \in \tau, \mathbf{y} \in Z^Y} (Z_{\mathbf{y}} \sqcup \sqcup_{y,y' \in Z_{\mathbf{y}}} G_{y,y'})$, and let $X_0 = X / \sim$ subject

to the equivalence relation $(Z, \mathbf{y}, y) \sim (y, y', x_G) \sim (y', y, y_G)$ and $(Z, \mathbf{y}, y) \sim (Z', \mathbf{y}', y)$. That is, the variables are exactly those of Y (so we take y as a representative of the class $[(Z, \mathbf{y}, y)]$), with the variables apart from the distinguished variables from the gadgets added. Also, given a Turing machine M sampling the constraints of $Y = X(M)$, we can construct a Turing machine M' sampling the constraints of $X = X(M')$ in polynomial time, by first sampling a constraint $(Z, \mathbf{y})$ using M , and then, with probability $\frac{1}{2}$, outputting a uniform random constraint of $Z_{\mathbf{y}}$, and, with probability $\frac{1}{2}$, sampling a uniform random pair of variables y, y' from $Z_{\mathbf{y}}$ and outputting a uniform random constraint of $G_{y, y'}$.

- (i) To show completeness, suppose that there is a $*$ -homomorphism $\varphi : \text{Mor}^{o+}(X(M), \overline{A}) \rightarrow \mathcal{M} \in \mathcal{Q}$. Now, since G is an oracular Q -commutativity gadget, for $a, b \in A$, let $\pi_{a,b} : \text{Mor}^{o+}(G, A) \rightarrow \mathcal{N} \in \mathcal{Q}$ be such that $\pi_{a,b}(p_{x_G, a}) = \pi_{a,b}(p_{y_G, b}) = 1$ — we may assume that that codomains are the same for all the representations $\pi_{a,b}$ by taking the tensor product of the codomains. Now, define a $*$ -homomorphism $\varphi' : \text{Mor}^{o+}(X(M'), A) \rightarrow \mathcal{M} \otimes \mathcal{N}$ on the generators as $\varphi'(p_{y, a}) = \varphi(p_{y, a}) \otimes 1$ for $y \in Y$, and $\varphi'(p_{[(y, y', g)], a}) = \sum_{a_1, a_2 \in A} \varphi(p_{y, a_1} p_{y', a_2}) \otimes \pi_{a_1, a_2}(p_{g, a})$. First, note that φ' is well-defined on the generators in the sense that it respects the equivalence relation $\sim$. In fact, for $y, y' \in Z_{\mathbf{y}}$,

$$\varphi'(p_{[(y, y', x_G)], a}) = \sum_{a_1, a_2 \in A} \varphi(p_{y, a_1} p_{y', a_2}) \otimes \pi_{a_1, a_2}(p_{x_G, a}) = \sum_{a_2 \in A} \varphi(p_{y, a} p_{y', a_2}) \otimes 1 = \varphi'(p_{y, a});$$

and identically $\varphi'(p_{[(y, y', y_G)], a}) = \varphi'(p_{y', a})$. Next, it is clear to see that by construction, $\{\varphi'(p_{y, a})\}_{a \in A}$ and $\{\varphi'(p_{[(y, y', g)], a})\}_{a \in A}$ are PVMs. Now, let $R \in \sigma$, $\mathbf{x} \in R^X$, and $\mathbf{a} \notin R^A$; write $n = \text{ar}(R)$. There are two cases: $\mathbf{x}$ comes from $Z_{\mathbf{y}}$ or from $G_{y, y'}$. In the first case, there exists $\mathbf{z} \in R^Z$ such that $x_i = y_{z_i}$, and since the relation is satisfied in Z , we have $\varphi'(p_{x_1, a_1}) \cdots \varphi'(p_{x_n, a_n}) = \varphi(p_{y_{z_1}, a_1}) \cdots \varphi(p_{y_{z_n}, a_n}) \otimes 1 = 0$. In the second case, there exists $\mathbf{g} \in R^G$ such that $x_i = [(y, y', g_i)]$. Hence,

$$\begin{aligned} & \varphi'(p_{x_1, a_1}) \cdots \varphi'(p_{x_n, a_n}) \\ &= \sum_{a_1^1, a_1^2, \dots, a_n^1, a_n^2 \in A} \varphi(p_{y, a_1^1} p_{y', a_1^2} \cdots p_{y, a_n^1} p_{y', a_n^2}) \otimes \pi_{a_1^1, a_1^2}(p_{g_1, a_1}) \cdots \pi_{a_n^1, a_n^2}(p_{g_n, a_n}) \\ &= \sum_{a^1, a^2 \in A} \varphi(p_{y, a^1} p_{y', a^2}) \otimes \pi_{a^1, a^2}(p_{g_1, a_1} \cdots p_{g_n, a_n}) \\ &= 0. \end{aligned}$$

Finally, for any $a, a' \in A$, we find in the same way as above that in the first case $[\varphi'(p_{x_i, a}), \varphi'(p_{x_j, a'})] = [\varphi(p_{y_{z_i}, a}), \varphi(p_{y_{z_j}, a'})] \otimes 1 = 0$; and in the second case $[\varphi'(p_{x_i, a}), \varphi'(p_{x_j, a'})] = \sum_{a_1, a_2 \in A} \varphi(p_{y, a_1} p_{y', a_2}) \otimes \pi_{a_1, a_2}([p_{g_i, a}, p_{g_j, a'}]) = 0$.

To show soundness, suppose that there is a $*$ -homomorphism $\varphi : \text{Mor}^{o+}(X(M'), A) \rightarrow \mathcal{M} \in Q$. Define a $*$ -homomorphism $\varphi' : \text{Mor}^{o+}(X(M), \overline{A}) \rightarrow \mathcal{M} \in Q$ on the generators by $\varphi'(p_{y,a}) = \varphi(p_{y,a})$. For any $y, y' \in Z_{\mathbf{y}}$, we have that $\varphi'(p_{y,a}) = \varphi(p_{[(y,y',x_G)],a})$ and $\varphi'(p_{y',a'}) = \varphi(p_{[(y,y',y_G)],a'})$. Since G is an oracular Q -commutativity gadget, these must commute. As such, all the PVMs corresponding to variables in $Z_{\mathbf{y}}$ commute and they satisfy the relations of $Z_{\mathbf{y}}$, so the constraint $(Z, \mathbf{y})$ in $X(M)$ is satisfied. Hence, φ' is a $*$ -homomorphism.

- (ii) To show completeness, suppose that there is a $*$ -homomorphism $\varphi : \text{Mor}^{o+}(X(M), \overline{A}) \rightarrow \mathcal{M} \in Q$. We construct $\varphi' : \text{Mor}^+(X(M'), A) \rightarrow \mathcal{M} \otimes \mathcal{N} \in Q$ using a Q -commutativity gadget G in the same way as in (i), and it satisfies the relations of $\text{Mor}^+(X(M'), A)$ in the same way.

To show soundness, suppose that there is a $*$ -homomorphism $\varphi : \text{Mor}^+(X(M'), A) \rightarrow \mathcal{M} \in Q$. In the same way as (i), we define the $*$ -homomorphism $\varphi' : \text{Mor}^{o+}(X(M), A) \rightarrow \mathcal{M} \in Q$ on the generators as $\varphi'(p_{y,a}) = \varphi(p_{y,a})$. For all $y, y' \in Z_{\mathbf{y}}$, $\varphi'(p_{y,a})$ and $\varphi'(p_{y',a'})$ commute as G is a Q -commutativity gadget. Thus, φ' is a $*$ -homomorphism in the same way.

- (iii) Since G is a robust oracular Q -commutativity gadget, there exists a convex function $\delta : [0, 1] \rightarrow [0, 1]$ such that $\delta(\varepsilon) \leq \varepsilon$ and if $\text{def}(\tau) < \delta(\varepsilon)$ for a tracial state on $\text{Mor}_{\mathbf{u}}^{ocv}(G, A)$ over Q , then $\sum_{a,b \in A} \|p_{x,a}, p_{y,b}\|_{\tau} < \varepsilon$ (we know such a modulus function δ exists by taking any modulus function δ_0 and letting δ be the function whose epigraph is the convex hull of the epigraph of $\varepsilon \mapsto \min\{\delta_0(\varepsilon), \varepsilon\}$). Let $m = \sum_{R \in \sigma} |R^G|$, $M = \max_{Z \in \bar{\sigma}: \exists \mathbf{y}. \pi_M(Z, \mathbf{y}) > 0} \sum_{R \in \sigma} |R^Z|$, $L = \max_{R \in \sigma} \text{ar}(R)$, and let $s' = 1 - \delta\left(\left(\frac{1-s}{84MN^2L^2|A|^{3N}}\right)^2\right)$.

Completeness holds in the same way as (i). For soundness, suppose τ is a tracial state on $\text{Mor}_{\pi_{M'}}^{ocv}(X(M'), A)$ over Q such that $\text{def}(\tau) \leq 1 - s'$. Let $\varphi : \text{Mor}^{ocv}(X(M'), A) \rightarrow (\mathcal{M}, \text{tr})$ be the GNS representation of τ . Then, for $y, y' \in Z_{\mathbf{y}}$, φ induces a $*$ -homomorphism $\varphi_{y,y'} : \text{Mor}^{ocv}(G, A) \rightarrow \mathcal{M}$ defined on the generators by $\varphi_{y,y'}(p_{g,a}) = \varphi(p_{[(y,y',g)],a})$. Hence $\tau_{y,y'} = \text{tr} \circ \varphi_{y,y'}$ is a tracial state on $\text{Mor}^{ocv}(G, A)$. By the robust oracular Q -commutativity gadget property, $\sum_{a,b \in A} \|p_{y,a}, p_{y',b}\|_{\tau}^2 \leq \delta^{-1}(\text{def}(\tau_{y,y'}))$. Now, we want to patch all the PVMs for the variables of $Z_{\mathbf{y}}$ into a single PVM. For $\mathbf{a} \in A^Z$, let $q_{\mathbf{y},\mathbf{a}} = \prod_{z \in Z} \varphi(p_{y_z, a_z})$ and $Q_{\mathbf{y},\mathbf{a}} = q_{\mathbf{y},\mathbf{a}}^* q_{\mathbf{y},\mathbf{a}}$. Writing $Z = \{z_1, \dots, z_{|Z|}\}$, we have that

$$\sum_{\mathbf{a} \in A^Z} \|Q_{\mathbf{y},\mathbf{a}} - q_{\mathbf{y},\mathbf{a}}\|_{\text{tr}}^2 = \sum_{\mathbf{a} \in A^Z} \|p_{y_{z_1}, a_{z_1}} \cdots p_{y_{z_{|Z|}}, a_{z_{|Z|}}} \cdots p_{y_{z_1}, a_{z_1}} - p_{y_{z_1}, a_{z_1}} \cdots p_{y_{z_{|Z|}}, a_{z_{|Z|}}}\|_{\tau}^2$$

$$\begin{aligned}
&\leq |Z| \sum_{\mathbf{a} \in A^Z} \sum_{i=1}^{|Z|-1} \|p_{y_{z_1}, a_{z_1}} \cdots p_{y_{z_i}, a_{z_i}} [p_{y_{z_{i+1}}, a_{z_{i+1}}} \cdots p_{y_{z_{|Z|}}, a_{z_{|Z|}}}, p_{y_{z_i}, a_{z_i}}] p_{y_{z_{i-1}}, a_{z_{i-1}}} \cdots p_{y_{z_1}, a_{z_1}}\|_{\tau}^2 \\
&\leq |Z| \sum_{\mathbf{a} \in A^Z} \sum_{i=1}^{|Z|-1} \|[p_{y_{z_{i+1}}, a_{z_{i+1}}} \cdots p_{y_{z_{|Z|}}, a_{z_{|Z|}}}, p_{y_{z_i}, a_{z_i}}]\|_{\tau}^2 \\
&\leq |Z|^2 \sum_{\mathbf{a} \in A^Z} \sum_{i=1}^{|Z|-1} \sum_{j=i+1}^{|Z|} \| [p_{y_{z_j}, a_{z_j}}, p_{y_{z_i}, a_{z_i}}] \|_{\tau}^2 \\
&\leq |Z|^2 |A|^{|Z|} \sum_{y, y' \in Z_y} \delta^{-1}(\text{def}(\tau_{y, y'})).
\end{aligned}$$

Write $\varepsilon_{Z, \mathbf{y}} = |Z|^2 |A|^{|Z|} \sum_{y, y' \in Z_y} \delta^{-1}(\text{def}(\tau_{y, y'}))$. This implies that

$$\begin{aligned}
\varepsilon_{Z, \mathbf{y}} &\geq \left(\sqrt{\sum_{\mathbf{a} \in A^Z} \|Q_{\mathbf{y}, \mathbf{a}}\|_{\text{tr}}^2} - \sqrt{\sum_{\mathbf{a} \in A^Z} \|q_{\mathbf{y}, \mathbf{a}}\|_{\text{tr}}^2} \right)^2 \\
&= \left(1 - \sqrt{\sum_{\mathbf{a} \in A^Z} \text{tr}(Q_{\mathbf{y}, \mathbf{a}}^2)} \right)^2,
\end{aligned}$$

giving that $\sum_{\mathbf{a} \in A^Z} \text{tr}(Q_{\mathbf{y}, \mathbf{a}}^2) \geq (1 - \sqrt{\varepsilon_{Z, \mathbf{y}}})^2 \geq 1 - 2\sqrt{\varepsilon_{Z, \mathbf{y}}}$. Next, if $\mathbf{a} \notin Z^{\bar{A}}$, that implies that there exists $R \in \sigma$ and $\mathbf{z} \in R^Z$ such that $(a_{z_1}, \dots, a_{z_{\text{ar}(R)}}) \notin R^A$. Let $\tau_{Z, \mathbf{y}}$ be the tracial state on $\text{Mor}_u^{\text{ocu}}(Z_{\mathbf{y}}, A)$ defined by the restriction of τ to this algebra. Then,

$$\begin{aligned}
\sum_{\mathbf{a} \notin Z^{\bar{A}}} \text{tr}(Q_{\mathbf{y}, \mathbf{a}}) &\leq \sum_{R \in \sigma, \mathbf{z} \in R^Z, \mathbf{a} \notin R^A} \text{tr} \left(\sum_{\mathbf{a}' \in A^Z: \mathbf{a} = (a_{z_1}, \dots, a_{z_{\text{ar}(R)}})} Q_{\mathbf{y}, \mathbf{a}'} \right) \\
&\leq \sum_{R \in \sigma, \mathbf{z} \in R^Z, \mathbf{a} \notin R^A} \left\| \sum_{\mathbf{a}' \in A^Z: \mathbf{a} = (a_{z_1}, \dots, a_{z_{\text{ar}(R)}})} Q_{\mathbf{y}, \mathbf{a}'} \right\|_{\text{tr}} \\
&\leq \sum_{R \in \sigma, \mathbf{z} \in R^Z, \mathbf{a} \notin R^A} \left(\|p_{y_{z_1}, a_1} \cdots p_{y_{z_{\text{ar}(R)}}, a_{\text{ar}(R)}}\|_{\tau} + \sum_{\mathbf{a}' \in A^Z: \mathbf{a} = (a_{z_1}, \dots, a_{z_{\text{ar}(R)}})} \|Q_{\mathbf{y}, \mathbf{a}'} - q_{\mathbf{y}, \mathbf{a}'}\|_{\text{tr}} \right) \\
&\leq \sum_{R \in \sigma, \mathbf{z} \in R^Z, \mathbf{a} \notin R^A} \|p_{y_{z_1}, a_1} \cdots p_{y_{z_{\text{ar}(R)}}, a_{\text{ar}(R)}} - \Phi_{R, (y_{z_1}, \dots, y_{z_{\text{ar}(R)}}), \mathbf{a}}\|_{\tau} + \sum_{R \in \sigma} |R^Z| \sum_{\mathbf{a}' \in A^Z} \|Q_{\mathbf{y}, \mathbf{a}'} - q_{\mathbf{y}, \mathbf{a}'}\|_{\text{tr}} \\
&\leq \sum_{\substack{R \in \sigma, \mathbf{z} \in R^Z \\ \mathbf{a} \in A, i \in \text{ar}(R)}} \|p_{y_{z_i}, a} - \Phi_{R, (y_{z_1}, \dots, y_{z_{\text{ar}(R)}}), i, a}\|_{\tau} + \sum_{R \in \sigma} |R^Z| \sum_{\mathbf{a}' \in A^Z} \|Q_{\mathbf{y}, \mathbf{a}'} - q_{\mathbf{y}, \mathbf{a}'}\|_{\text{tr}}
\end{aligned}$$

$$\begin{aligned}
&\leq \sqrt{\sum_{R \in \sigma} |R^Z| \ar(R) |A| \sum_{\substack{R \in \sigma, z \in R^Z \\ a \in A, i \in \ar(R)}} \|p_{y_{z_i}, a} - \Phi_{R, (y_{z_1}, \dots, y_{z_{\ar(R)}}), i, a}\|_\tau^2} \\
&\quad + \sum_{R \in \sigma} |R^Z| \sqrt{|A|^{|Z|} \sum_{a' \in A^Z} \|Q_{y, a'} - q_{y, a'}\|_{\text{tr}}^2} \\
&\leq \sqrt{\frac{|A|}{2}} M L^2 \sqrt{\text{def}(\tau_{Z, y})} + M \sqrt{|A|^{|Z|} \varepsilon_{Z, y}}.
\end{aligned}$$

Let $\varepsilon'_{Z, y} = \sqrt{\frac{|A|}{2}} M L^2 \sqrt{\text{def}(\tau_{Z, y})} + M \sqrt{|A|^{|Z|} \varepsilon_{Z, y}}$ and define the POVM $\{\hat{Q}_{y, a}\}_{a \in Z^{\bar{A}}}$ by $\hat{Q}_{y, a} = Q_{y, a} + \Delta$ where $\Delta = \frac{1}{|Z^{\bar{A}}|} \left(1 - \sum_{a \in Z^{\bar{A}}} Q_{y, a}\right)$. Then,

$$\sum_{a \in Z^{\bar{A}}} \text{tr}(\hat{Q}_{y, a}^2) \geq \sum_{a \in Z^{\bar{A}}} \text{tr}(Q_{y, a}^2) \geq 1 - 2\sqrt{\varepsilon_{Z, y}} - \varepsilon'_{Z, y}.$$

By [55], there exists a PVM $\{P_{Z, y, a}\}_{a \in Z^{\bar{A}}} \subseteq \mathcal{M}$ such that $\sum_{a \in Z^{\bar{A}}} \|P_{Z, y, a} - \hat{Q}_{y, a}\|_{\text{tr}}^2 \leq 18\sqrt{\varepsilon_{Z, y}} + 9\varepsilon'_{Z, y}$. This implies

$$\begin{aligned}
\sum_{a \in Z^{\bar{A}}} \|P_{Z, y, a} - q_{y, a}\|_{\text{tr}}^2 &\leq 3 \sum_{a \in Z^{\bar{A}}} \left(\|P_{Z, y, a} - \hat{Q}_{y, a}\|_{\text{tr}}^2 + \|\hat{Q}_{y, a} - Q_{y, a}\|_{\text{tr}}^2 + \|Q_{y, a} - q_{y, a}\|_{\text{tr}}^2 \right) \\
&\leq 54\sqrt{\varepsilon_{Z, y}} + 27\varepsilon'_{Z, y} + 3|A|^{|Z|} \varepsilon'_{Z, y} + 3\varepsilon_{Z, y} \\
&\leq 27|A|^{|Z|+1} M L^2 \sqrt{\text{def}(\tau_{Z, y})} + 84|A|^{\frac{3}{2}|Z|} M \sqrt{\varepsilon_{Z, y}} =: \varepsilon''_{Z, y}.
\end{aligned}$$

Define the $*$ -homomorphism $\varphi' : \text{Mor}^{ocv}(X(M), \bar{A}) \rightarrow \mathcal{M}$ on the generators as $\varphi'(\Phi_{Z, y, a}) = P_{Z, y, a}$ and $\varphi'(p_{y, a}) = \varphi(p_{y, a})$. By construction of the $P_{Z, y, a}$, this is well-defined. Let $\tau' = \tau' \circ \varphi'$. Next, note that for any $z \in Z$,

$$\begin{aligned}
\sum_{a \in A} \|\Phi_{Z, y, z, a} - p_{y, z, a}\|_{\tau'}^2 &= \sum_{a \in A} \left\| \sum_{a \in A^Z : a_z = a} (P_{Z, y, a} - q_{y, a}) \right\|_{\text{tr}}^2 \\
&\leq |A|^{|Z|} \sum_{a \in A^Z} \|P_{Z, y, a} - q_{y, a}\|_{\text{tr}}^2 \\
&\leq |A|^{|Z|} \varepsilon''_{Z, y}.
\end{aligned}$$

This implies that the defect

$$\begin{aligned}
\text{def}(\tau') &= \sum_{\substack{Z \in \bar{\sigma}, \mathbf{y} \in Z^{X(M)} \\ a \in A, z \in Z}} \frac{\pi_M(Z, \mathbf{y})}{2|Z|} \|\Phi_{Z, \mathbf{y}, z, a} - p_{yz, a}\|_{\tau'}^2 \\
&\leq \sum_{Z \in \bar{\sigma}, \mathbf{y} \in Z^{X(M)}} \frac{\pi_M(Z, \mathbf{y})}{2} |A|^{|Z|} \varepsilon''_{Z, \mathbf{y}} \\
&\leq 84MN^2L^2|A|^{3N} \sum_{Z \in \bar{\sigma}, \mathbf{y} \in Z^{X(M)}} \frac{\pi_M(Z, \mathbf{y})}{2} \left(\sqrt{\text{def}(\tau_{Z, \mathbf{y}})} + \sqrt{\frac{1}{|Z|^2} \sum_{y, y' \in Z_{\mathbf{y}}} \delta^{-1}(\text{def}(\tau_{y, y'}))} \right) \\
&\leq 84MN^2L^2|A|^{3N} \sqrt{\delta^{-1}(\text{def}(\tau))} \leq 1 - s,
\end{aligned}$$

finishing the proof of soundness.

- (iv) Completeness is shown in the same way as (ii), and soundness is shown in the same way as (iii). ■

Theorem 4.1.6. Let A be a relational structure such that $\text{CSP}(A)$ is NP-complete.

- (i) If A admits a q -commutativity gadget, then there is a polynomial-time reduction from the halting problem to ${}^e\text{MIP}_{no}^q(A)_{1,1}$.
- (ii) If A admits an oracular q -commutativity gadget, then there is a polynomial-time reduction from the halting problem to ${}^e\text{MIP}_{ocv}^q(A)_{1,1}$.
- (iii) If A admits a robust q -commutativity gadget, then there exists $\varepsilon > 0$ such that there are polynomial-time reductions from the halting problem to $\text{MIP}_{no}^q(A)_{1,1-\varepsilon}$ and ${}^e\text{MIP}_{no}^q(A)_{1,1-\varepsilon}$, making these RE-complete.
- (iv) If A admits a robust oracular q -commutativity gadget, then there exists $\varepsilon > 0$ such that there are polynomial-time reductions from the halting problem to $\text{MIP}_{ocv}^q(A)_{1,1-\varepsilon}$ and ${}^e\text{MIP}_{ocv}^q(A)_{1,1-\varepsilon}$, making these RE-complete.
- (v) If A admits a qa -commutativity gadget, then there is a polynomial-time reduction from the halting problem to $\text{MIP}_a^{qa}(A)_{1,1}$.
- (vi) If A admits an oracular qa -commutativity gadget, then there is a polynomial-time reduction from the halting problem to $\text{MIP}_{ocv}^{qa}(A)_{1,1}$.

- (vii) If A admits a robust qa -commutativity gadget, then there exists $\varepsilon > 0$ such that there is a polynomial-time reduction from the halting problem to $\text{MIP}_{no}^{qa}(A)_{1,1-\varepsilon}$, making this RE-complete.
- (viii) If A admits a robust oracular qa -commutativity gadget, then there exists $\varepsilon > 0$ such that there is a polynomial-time reduction from the halting problem to $\text{MIP}_{ocv}^{qa}(A)_{1,1-\varepsilon}$, making this RE-complete.
- (ix) If A admits a co -commutativity gadget, then there is a polynomial-time reduction from the cohalting problem to $\text{MIP}_{no}^{co}(A)_{1,1}$, making this coRE-complete.
- (x) If A admits an oracular co -commutativity gadget, then there is a polynomial-time reduction from the cohalting problem to $\text{MIP}_{ocv}^{co}(A)_{1,1}$, making this coRE-complete.
- (xi) If A admits a robust co -commutativity gadget, then there exists $\varepsilon > 0$ such that there is a polynomial-time reduction from the cohalting problem to $\text{MIP}_{no}^{co}(A)_{1,1-\varepsilon}$, making this coRE-complete.
- (xii) If A admits a robust oracular co -commutativity gadget, then there exists $\varepsilon > 0$ such that there is a polynomial-time reduction from the cohalting problem to $\text{MIP}_{ocv}^{co}(A)_{1,1-\varepsilon}$, making this coRE-complete.

The classes RE and coRE do not depend on runtime, just computability, so completeness (or hardness) of any succinct problem for this class also implies completeness (or hardness) of the non-succinct version.

Proof. The proof follows immediately by applying [Lemma 4.1.5](#) to [Theorem 3.1.17](#). ■

4.1.3 Algebraic commutativity gadgets

As seen in the previous section, working with robust commutativity gadgets can be quite unwieldy. However, it often suffices to work with the following stronger notion.

Definition 4.1.7. Let A be a σ -structure. An **algebraic commutativity gadget** for A is a q -commutativity gadget (G, x, y) such that for all $a, b \in A$, $[p_{x,a}, p_{y,b}] = 0$ in $\text{Mor}^+(G, A)$. An **oracular algebraic commutativity gadget** for A is an oracular q -commutativity gadget (G, x, y) such that the same commutation relations hold in $\text{Mor}^{o+}(G, A)$.

Since every C^* -algebra admits a faithful representation into the bounded operators on some Hilbert space, every algebraic commutativity gadgets is a vN -commutativity gadget.

Lemma 4.1.8. Suppose (G, x, y) is an algebraic commutativity gadget for A . Then it is a robust Q -commutativity gadget for any finite tensor monoid Q .

Proof. Since $[p_{x,a}, p_{y,b}] = 0$ in $\text{Mor}^+(G, A)$ for each $a, b \in A$, we know that $[p_{x,a}, p_{y,b}]$ is an infinitesimal element of the quotient of $\text{Mor}_{\mathfrak{u}}^{no}(G, A)$ by $\langle \text{supp } \mu_{no, \mathfrak{u}} \rangle$ by Lemma 3.2.2. Let $\varepsilon > 0$. Then, $\| [p_{x,a}, p_{y,b}] \|^2 \leq \frac{\varepsilon}{2|A|^2}$ in $\text{Mor}_{\mathfrak{u}}^{no}(G, A) / \langle \text{supp } \mu_{no, \mathfrak{u}} \rangle$, so $\sum_{a,b \in A} \| [p_{x,a}, p_{y,b}] \|^2 \leq \frac{\varepsilon}{2}$. Therefore, there exists a positive element $p \in \text{Mor}_{\mathfrak{u}}^{no}(G, A)$ such that $\frac{\varepsilon}{2} - \sum_{a,b \in A} \| [p_{x,a}, p_{y,b}] \|^2 - p \in \langle \text{supp } \mu_{no, \mathfrak{u}} \rangle$. Then, we have that there exist $N \in \mathbb{N}$, $r_i \in \text{supp } \mu_{no, \mathfrak{u}}$, and $v_i, w_i \in \text{Mor}_{\mathfrak{u}}^{no}(G, A)$ for $i \in [N]$ such that $\sum_{a,b \in A} \| [p_{x,a}, p_{y,b}] \|^2 + p - \frac{\varepsilon}{2} = \sum_{i=1}^N v_i r_i w_i$. Since $\text{Mor}_{\mathfrak{u}}^{no}(G, A)$ is archimedean with respect to the $*$ -positive cone of sums-of-squares, there exist $\alpha_i \in \mathbb{R}_{\geq 0}$ such that $\| w_i v_i \|^2 \leq \alpha_i$. Now, let τ be a tracial state on $\text{Mor}_{\mathfrak{u}}^{no}(G, A)$ over Q . Then,

$$\begin{aligned} \sum_{a,b \in A} \| [p_{x,a}, p_{y,b}] \|^2_{\tau} &= \frac{\varepsilon}{2} - \tau(p) + \sum_{i=1}^N \tau(v_i r_i w_i) \\ &\leq \frac{\varepsilon}{2} + \sum_{i=1}^N \sqrt{\alpha_i} \| r_i \|_{\tau} \\ &\leq \frac{\varepsilon}{2} + \sum_{i=1}^N \sqrt{m \alpha_i \text{def}(\tau)}, \end{aligned}$$

where $m = \sum_{R \in \sigma} |R^G|$. Write $C = \sum_{i=1}^N \sqrt{m \alpha_i}$. Now, let $\delta = \left(\frac{\varepsilon}{2C} \right)^2$. By the above, if τ is a tracial state such that $\text{def}(\tau) \leq \delta$, then $\sum_{a,b} \| [p_{x,a}, p_{y,b}] \|^2_{\tau} \leq \varepsilon$. $\blacksquare$

Lemma 4.1.9. Suppose (G, x, y) is an oracular algebraic commutativity gadget for A . Then it is a robust oracular Q -commutativity gadget for every finite tensor monoid Q .

We omit the proof as it follows along exactly the same lines as the previous lemma.

4.1.4 Examples of commutativity gadget

In this section, we note that the notion of a commutativity gadget is not vacuous, and there are important CSPs that admit them. The first example of a commutativity gadget is for 3-colouring, due to Ji.

Theorem 4.1.10 ([75]). The complete graph on 3 vertices K_3 admits an algebraic commutativity gadget and an oracular algebraic commutativity gadget.

The same gadget is valid in both the oracular and non-oracular framework, and takes the form of a triangular prism, *i.e.* two copies of K_3 with the corresponding vertices connected (see Fig. 4.1). [75] only considered the gadget for non-oracular quantum strategies, but in hindsight the proof holds at the level of the C^* -algebra, which implies the possibility of gapped reductions. Further, the gadget holds in the oracular framework, as quantum homomorphisms to K_3 are oracularisable, in the sense that $\text{Mor}^{o+}(X, K_3) \cong \text{Mor}^+(X, K_3)$.

This example has been recently extended to all odd cycles (as $K_3 = C_3$).

Theorem 4.1.11 ([38]). For $n \in \mathbb{N}$, the cycle graph C_{2n+1} admits a q -commutativity gadget and an oracular q -commutativity gadget.

Again, these graphs have oracularisable quantum homomorphisms, so the oracular and non-oracular commutativity gadgets are identical.

The other important class of examples arises for boolean CSPs.

Theorem 4.1.12 ([49]). Let A be a boolean relational structure (*i.e.* $\mathfrak{D}(A) = \{0, 1\}$), and suppose that $\text{CSP}(A)$ is NP-complete. Then, A admits an oracular algebraic commutativity gadget.

This result also extends in some cases to non-oracular gadgets.

Since all of the relational structures discussed in this section have NP-complete CSP, Theorem 4.1.6 immediately implies the associated hardness of the entangled analogues.

4.2 Commutativity gadget for oracular entangled graph colouring

In this section, we show that the complete graph K_n admits an oracular algebraic commutativity gadget, which implies hardness results via Lemma 4.1.8 and Theorem 4.1.6

Theorem 4.2.1. For all $n \in \mathbb{N}$, the complement of the cycle graph $(\overline{C}_{2n}, 0, 1)$ is an oracular algebraic commutativity gadget for K_n .

Note that for $n = 3$, $(\overline{C}_6, 0, 1)$ is Ji's triangular prism gadget [75]. See the discussion in Section 4.1.4, or Fig. 4.1 for a visualisation.

Lemma 4.2.2. $\text{End}^{o+}(K_n)$ is commutative and $\sum_{a \in K_n} p_{a,b} = 1$ for all $b \in K_n$.




Figure 4.1: The graphs $\overline{C}_6$ and $\overline{C}_8$ presented as prisms

Proof. Every pair of distinct vertices in K_n is connected by an edge, hence by construction every pair of generators in $\text{End}^{o+}(K_n)$ commutes. So $\text{End}^{o+}(K_n)$ is a commutative C^* -algebra, proving the first part of the lemma.

Next, consider the following:

$$\begin{aligned}
 \sum_{b \in K_n} \left(1 - \sum_{a \in K_n} p_{a,b} \right)^2 &= \sum_{b \in K_n} \left(1 - 2 \sum_{a \in K_n} p_{a,b} + \sum_{a, a' \in K_n} p_{a,b} p_{a',b} \right) \\
 &= n + \sum_{b \in K_n} \left(- \sum_{a \in K_n} p_{a,b} + \sum_{a \sim_{K_n} a'} p_{a,b} p_{a',b} \right) \\
 &= n - \sum_{a \in K_n} \sum_{b \in K_n} p_{a,b} = n - n = 0.
 \end{aligned}$$

In a C^* -algebra, a sum of positive elements is 0 if and only if each term in the sum is 0, so it follows that $\sum_{a \in K_n} p_{a,b} = 1$ for all b . $\blacksquare$

Lemma 4.2.3. In $\text{Mor}^{o+}(\overline{C}_{2n}, K_n)$, $[p_{2a,b}, p_{2a+1,c}] = [p_{0,b}, p_{1,c}]$ for all $a, b, c \in \mathbb{Z}_n = V(K_n)$.

Proof. For all $i \neq j$, $2i+1 \sim_{\overline{C}_{2n}} 2j+1$. As such, there is a $*$ -homomorphism $\text{End}^{o+}(K_n) \rightarrow \text{Mor}^{o+}(\overline{C}_{2n}, K_n)$ defined on the generators as $p_{a,b} \mapsto p_{2a+1,b}$. Then, by Lemma 4.2.2, the $p_{2a+1,b}$ commute and $\sum_{a \in \mathbb{Z}_n} p_{2a+1,b} = 1$. This implies that $\sum_{a \in \mathbb{Z}_n} [p_{0,b}, p_{2a+1,c}] = [p_{0,b}, 1] = 0$. For all $a \neq 0$ or -1 , we have $0 \sim_{\overline{C}_{2n}} 2a+1$, giving $[p_{0,b}, p_{2a+1,c}] = 0$. Hence, $[p_{0,b}, p_{1,c}] = -[p_{0,b}, p_{-1,c}]$.

In the same way, for all $i \neq j$, $2i \sim_{\overline{C}_{2n}} 2j$, so $\sum_{a \in \mathbb{Z}_n} p_{2a,b} = 1$. Taking the commutator with $p_{-1,c}$ gives $\sum_{a \in \mathbb{Z}_k} [p_{2a,b}, p_{-1,c}] = 0$. As $2a \sim_{\overline{C}_{2n}} -1$ for all $a \neq 0$ or -1 , this implies that $[p_{0,b}, p_{-1,c}] = -[p_{-2,b}, p_{-1,c}]$. Putting these together gives $[p_{0,b}, p_{1,c}] = [p_{-2,b}, p_{-1,c}]$. Proceeding inductively gives the result. $\blacksquare$

Proof of Theorem 4.2.1. We need to show that $(\overline{C}_{2n}, 0, 1)$ satisfies the two conditions of an oracular algebraic commutativity gadget. For the first condition, consider the maps $f, g : \overline{C}_{2n} \rightarrow K_n$ defined as $f(2a) = f(2a + 1) = a$ and $g(2a) = g(2a - 1) = a$. These are graph homomorphisms as the vertices of $\overline{C}_{2n}$ that map to the same vertex of K_n are not connected by an edge. Now, let σ_a be a permutation of $\mathbb{Z}_n$ that maps 0 to a , and for $a \neq b$ let $\sigma_{a,b}$ be a permutation of $\mathbb{Z}_n$ that maps 0 to a and 1 to b . Now, for any $a, b \in K_n$, we take $\pi_{a,b}$ as the classical homomorphism $\sigma_a \circ f$ if $a = b$, and $\sigma_{a,b} \circ g$ if $a \neq b$.

To show the second condition, we need that $[p_{0,b}, p_{1,c}] = 0$ for all $b, c \in K_n$. In fact,

$$\begin{aligned} [p_{0,b}, p_{1,c}] &= \sum_{a \in \mathbb{Z}_n} p_{2a,b} [p_{0,b}, p_{1,c}] \\ &= \sum_{a \in \mathbb{Z}_n} p_{2a,b} [p_{2a+2,b}, p_{2a+3,c}] \\ &= \sum_{a \in \mathbb{Z}_n} [p_{2a,b} p_{2a+2,b}, p_{2a+3,c}] \\ &= \sum_{a \in \mathbb{Z}_n} [0, p_{2a+3,c}] = 0, \end{aligned}$$

where the first line follows from $\sum_{a \in \mathbb{Z}_k} p_{2a,b} = 1$, the second line follows from Lemma 4.2.3, the third line follows from the fact that $2a \sim_{\overline{C}_{2n}} 2a + 3$ and hence $p_{2a,b}$ commutes with $p_{2a+3,c}$, and the last line follows from the n -colouring condition. ■

Corollary 4.2.4. For every $k \geq 3$, there exists $s < 1$ such that there are polynomial-time reductions from the halting problem to $\text{MIP}_{ocv}^q(K_n)_{1,s}$, ${}^e\text{MIP}_{ocv}^q(K_n)_{1,s}$, and $\text{MIP}_{ocv}^{qa}(K_n)_{1,s}$; and from the cohalting problem to $\text{MIP}_{ocv}^{co}(K_n)_{1,s}$.

The corollary follows by using Theorem 4.2.1 in Lemma 4.1.8 and Theorem 4.1.6.

4.3 Polymorphism characterisation

In this section, we show that the existence of commutativity gadgets for a relational structure may be fully characterised by the commutativity of its corresponding polymorphism algebra. In Section 4.3.1, we give a proof that this holds in the exact case. In Section 4.3.2, we study a stable notion of commutation which we use, in Section 4.3.3, to extend the characterisation to the robust case. In Section 4.3.4, we study the relationships between different notions of commutativity gadget implied by the polymorphism characterisation.

4.3.1 Exact commutativity gadgets

In recent work, Ciardo, Joubert, and Mottet [38] prove an equivalence between the existence of q -commutativity gadgets and commutativity of the quantum polymorphisms. We extend this characterisation to all tensor monoids. The proofs work in exactly the same way, but we include them here, as they inform the proofs in the robust cases.

Theorem 4.3.1 ([38]). Let A be a relational structure and let Q be a tensor monoid. The following are equivalent.

- (i) $\text{mor}^Q(A^k, A) \subseteq \text{mor}^c(A^k, A)$ for all $k \in \mathbb{N}$.
- (ii) $\text{mor}^Q(A^{|A|^2}, A) \subseteq \text{mor}^c(A^{|A|^2}, A)$.
- (iii) A admits a Q -commutativity gadget.

It is clear that (i) $\Rightarrow$ (ii). Next, we show that (iii) $\Rightarrow$ (i) in contrapositive.

Lemma 4.3.2. Suppose that there exists $k \in \mathbb{N}$ such that $\text{mor}^Q(A^k, A) \not\subseteq \text{mor}^c(A^k, A)$. Then A does not admit a Q -commutativity gadget.

Proof. Let π be a non-classical $*$ -representation of $\text{Mor}^+(A^k, A)$ over Q . Then, there exist $\mathbf{a}, \mathbf{b} \in A^k$ and $a, b \in A$ such that $\pi(p_{\mathbf{a}, \mathbf{a}})$ and $\pi(p_{\mathbf{b}, \mathbf{b}})$ do not commute. Now, suppose (G, x, y) is a Q -commutativity gadget for A . By definition of a Q -commutativity gadget, for each $a, b \in A$, there exists a $*$ -representation $\pi_{a,b} : \text{Mor}^+(G, A) \rightarrow \mathcal{M}_{a,b}$ such that $\mathcal{M}_{a,b} \in Q$ and $\pi_{a,b}(p_{x,a}) = \pi_{a,b}(p_{y,b}) = 1$. Now consider the map $\varphi : \text{Mor}^+(G, A^k) \rightarrow \mathcal{M}_{a_1, b_1} \otimes \cdots \otimes \mathcal{M}_{a_k, b_k}$ defined on the generators as $\varphi(p_{z, \mathbf{c}}) = \pi_{a_1, b_1}(p_{z, c_1}) \otimes \cdots \otimes \pi_{a_k, b_k}(p_{z, c_k})$. I claim this extends to a $*$ -representation of $\text{Mor}^+(G, A^k)$ and therefore is an element of $\text{mor}^Q(G, A^k)$. To see this, simply note that φ satisfies the defining relations of $\text{Mor}^+(G, A^k)$. In fact, the $\varphi(p_{z, \mathbf{c}})$ are projections as they are tensor products of projections. Next,

$$\sum_{\mathbf{c} \in A^k} \varphi(p_{z, \mathbf{c}}) = \sum_{c_1 \in A} \pi_{a_1, b_1}(p_{z, c_1}) \otimes \cdots \otimes \sum_{c_k \in A} \pi_{a_k, b_k}(p_{z, c_k}) = 1.$$

Finally, suppose $R \in \sigma$, $\mathbf{z} \in R^G$, and $\mathbf{c} \notin R^{A^k}$; this implies that $(c_{1i}, \dots, c_{ni}) \notin R^A$ for some i . Writing $n = \text{ar}(R)$, we have that

$$\varphi(p_{z_1, \mathbf{c}_1}) \cdots \varphi(p_{z_n, \mathbf{c}_n}) = \bigotimes_{i=1}^k \pi_{a_i, b_i}(p_{z_1, c_{1i}} \cdots p_{z_n, c_{ni}}) = 0.$$

Thus, $\varphi \in \text{mor}^Q(G, A^k)$. Note also that by construction $\varphi(p_{x,a}) = \varphi(p_{y,b}) = 1$. Composing φ with π as Q -morphisms, we have that $\pi \circ \varphi \in \text{mor}^Q(G, A)$. But we have that

$$\begin{aligned} (\pi \circ \varphi)(p_{x,a}) &= \sum_{c \in A^k} \varphi(p_{x,c}) \otimes \pi(p_{c,a}) = 1 \otimes \pi(p_{a,a}) \\ (\pi \circ \varphi)(p_{y,b}) &= \sum_{c \in A^k} \varphi(p_{y,c}) \otimes \pi(p_{c,b}) = 1 \otimes \pi(p_{b,a}), \end{aligned}$$

which do not commute by hypothesis. Contradiction, so A does not admit a commutativity gadget. $\blacksquare$

We finish the proof of [Theorem 4.3.1](#) by showing (ii) $\Rightarrow$ (iii), which relies on [Lemma 2.10.1](#).

Lemma 4.3.3. Suppose A is a relational structure such that $\text{mor}^Q(A^{|A|^2}, A) \subseteq \text{mor}^c(A^{|A|^2}, A)$. Then, A admits a Q -commutativity gadget.

Proof. Let $k = |A^2| = |A|^2$. By [Lemma 2.10.1](#), A pp-defines the relation $\mathfrak{D}(A)^2$, as the set of polymorphisms that preserve $\mathfrak{D}(A)^2$, which are all maps, contains $\text{pol}(A)$. Let $\mathbf{x}, \mathbf{y} \in A^k$ be such that $x_{(a,b)} = a$ and $y_{(a,b)} = b$ for all indices $(a,b) \in A^2$. Now, we claim $(A^k, \mathbf{x}, \mathbf{y})$ is a commutativity gadget for A . To show property (i), let $a, b \in A$. By the construction of [Lemma 2.10.1](#), there exists a polymorphism $f : A^k \rightarrow A$ such that $f(\mathbf{x}) = a$ and $f(\mathbf{y}) = b$, given by a projection map. So, we can take $\pi_{a,b}(p_{a,c}) = \delta_{c,f(a)}$. Property (ii) follows directly from the hypothesis that $\text{mor}^Q(A^k, A) \subseteq \text{mor}^c(A^k, A)$. $\blacksquare$

This proof extends identically to the oracular case.

Theorem 4.3.4 ([38]). Let A be a relational structure and let Q be a tensor monoid. The following are equivalent.

1. $\text{mor}^{oQ}(A^k, A) \subseteq \text{mor}^c(A^k, A)$ for all $k \in \mathbb{N}$.
2. $\text{mor}^{oQ}(A^{|A|^2}, A) \subseteq \text{mor}^c(A^{|A|^2}, A)$.
3. A admits an oracular Q -commutativity gadget.

Proof.

(i) $\Rightarrow$ (ii) Immediate.

(ii) $\Rightarrow$ (iii) Suppose $\text{mor}^{oQ}(A^{|A|^2}, A) \subseteq \text{mor}^c(A^{|A|^2}, A)$. Following the notation of [Lemma 4.3.3](#), we claim $(A^{|A|^2}, \mathbf{x}, \mathbf{y})$ is an oracular commutativity gadget for A . In fact, property (i) holds identically; and property (ii) holds since every oracular Q -morphism is classical.

(iii) $\Rightarrow$ (i) To extend the proof of [Lemma 4.3.2](#), it suffices to show that the constructed map φ also satisfies the additional commutation relations of $\text{Mor}^{o+}(G, A^k)$. Let $\sigma \in R$, $\mathbf{z} \in R^G$, $i \neq j \in [\text{ar}(R)]$, and $\mathbf{c}, \mathbf{d} \in A^k$. Then,

$$\begin{aligned} \varphi(p_{z_i, \mathbf{c}}) \varphi(p_{z_j, \mathbf{d}}) &= \pi_{a_1, b_1}(p_{z_i, c_1}) \pi_{a_1, b_1}(p_{z_j, d_1}) \otimes \cdots \otimes \pi_{a_n, b_n}(p_{z_i, c_n}) \pi_{a_n, b_n}(p_{z_j, d_n}) \\ &= \pi_{a_1, b_1}(p_{z_j, d_1}) \pi_{a_1, b_1}(p_{z_i, c_1}) \otimes \cdots \otimes \pi_{a_n, b_n}(p_{z_j, d_n}) \pi_{a_n, b_n}(p_{z_i, c_n}) \\ &= \varphi(p_{z_j, \mathbf{d}}) \varphi(p_{z_i, \mathbf{c}}), \end{aligned}$$

as wanted. The rest of the proof proceeds identically. ■

4.3.2 Stable commutation

To extend the characterisation of commutativity gadgets in terms of polymorphisms to the robust setting, we make use of an approximate notion of commutativity.

Definition 4.3.5. Let $(\mathcal{A}, \mu)$ be a weighted algebra, and let Q be a tensor monoid. We say $\mathcal{A}$ is **Q -stably commutative** with respect to a finite generating set $X \subseteq \mathcal{A}$ if for all $\varepsilon > 0$, there exists $\delta > 0$ such that for any tracial state $\tau : \mathcal{A} \rightarrow \mathbb{C}$ over Q , $\text{def}(\tau) \leq \delta$ implies that $\| [x, y] \|_\tau^2 \leq \varepsilon$ for all $x, y \in X$.

It is clear that if $\mathcal{A}$ is Q -stably commutative, then $\mathcal{A}/\langle\langle \text{supp } \mu \rangle\rangle$ is commutative with respect to any representation to any finite von Neumann algebra in Q . Also, for $Q \subseteq Q'$, if $\mathcal{A}$ is Q' -stably commutative, then $\mathcal{A}$ is Q -stably commutative.

Next we show that stable commutativity is independent of the choice of finite generating set.

Lemma 4.3.6. Let $(\mathcal{A}, \mu)$ be a weighted algebra that is archimedean with respect to the $*$ -positive cone of sums of squares, and let X and Y be finite generating sets of $\mathcal{A}$ such that $X = X^*$ and $Y = Y^*$. Then $\mathcal{A}$ is Q -stably commutative with respect to X if and only if it is Q -stably commutative with respect to Y .

Proof. By symmetry, we need only show one direction. Suppose $\mathcal{A}$ is Q -stably commutative with respect to X . Write $X = \{x_1, \dots, x_n\}$ and $Y = \{y_1, \dots, y_m\}$. For each

$i \in [m]$, there exists a noncommutative $*$ -polynomial p_i such that $y_i = p_i(x_1, \dots, x_n) = \sum_M c_{i,M} M(x_1, \dots, x_n)$ where the sum is over monomials M . Fix a tracial state τ on $\mathcal{A}$ over Q . First note that for any $a, b, c \in \mathcal{A}$, the commutator

$$\begin{aligned} \|[ab, c]\|_\tau &= \|abc - acb + acb - cab\|_\tau = \|a[b, c] + [a, c]b\|_\tau \\ &\leq \|a\| \| [b, c] \|_\tau + \|b\| \| [a, c] \|_\tau, \end{aligned}$$

where $\|\cdot\|$ is the seminorm induced by the order on $\mathcal{A}$. Hence, by induction, for any $a_1, \dots, a_N, b \in \mathcal{A}$, we have that

$$\begin{aligned} \|[a_1 \cdots a_N, b]\|_\tau &\leq \|a_1\| \cdots \|a_{\lfloor N/2 \rfloor}\| \| [a_{\lfloor N/2 \rfloor + 1} \cdots a_N, b] \|_\tau + \|a_{\lfloor N/2 \rfloor + 1}\| \cdots \|a_N\| \| [a_1 \cdots a_{\lfloor N/2 \rfloor}, b] \|_\tau \\ &\leq \sum_{i=1}^N \prod_{j \neq i} \|a_j\| \cdot \| [a_i, b] \|_\tau. \end{aligned}$$

As such, given two monomials $M(x_1, \dots, x_n) = x_{i_1} \cdots x_{i_K}$ and $M'(x_1, \dots, x_n) = x_{j_1} \cdots x_{j_{K'}}$,

$$\begin{aligned} \|[M(x_1, \dots, x_n), M'(x_1, \dots, x_n)]\|_\tau &\leq \sum_{k=1}^K \prod_{l \neq k} \|x_{i_l}\| \| [x_{i_k}, x_{j_1} \cdots x_{j_{K'}}] \|_\tau \\ &\leq \sum_{k=1}^K \sum_{k'=1}^{K'} \prod_{l \neq k} \|x_{i_l}\| \prod_{l' \neq k'} \|x_{j_{l'}}\| \| [x_{i_k}, x_{j_{k'}}] \|_\tau. \end{aligned}$$

Let $C(M, M') = \sum_{k=1}^K \sum_{k'=1}^{K'} \prod_{l \neq k} \|x_{i_l}\| \prod_{l' \neq k'} \|x_{j_{l'}}\|$ for monomials M, M' as above. Set $C = \max_{i, i' \in [m]} \sum_{M, M'} |c_{i,M} c_{i',M'}| C(M, M')$.

Now, let $\varepsilon > 0$. As $\mathcal{A}$ is Q -stably commutative with respect to X , then there exists $\delta > 0$ such that for any tracial state τ on $\mathcal{A}$ over Q and $\text{def}(\tau) < \delta$, the norms $\|[x_j, x_{j'}]\|_\tau^2 < \frac{\varepsilon}{C^2}$. Suppose τ is a tracial state on $\mathcal{A}$ over Q such that $\text{def}(\tau) < \delta$. We have therefore that

$$\begin{aligned} \|[y_i, y_{i'}]\|_\tau &\leq \sum_{M, M'} |c_{i,M} c_{i',M'}| \| [M(x_1, \dots, x_n), M'(x_1, \dots, x_n)] \|_\tau \\ &\leq \sum_{M, M'} |c_{i,M} c_{i',M'}| C(M, M') \frac{\sqrt{\varepsilon}}{C} = \sqrt{\varepsilon}, \end{aligned}$$

giving $\|[y_i, y_{i'}]\|_\tau^2 \leq \varepsilon$. As such, $\mathcal{A}$ is Q -stably commutative with respect to Y . ■

Next, we show that stable commutativity is a weaker condition than commutativity of the quotient algebra.

Lemma 4.3.7. Suppose $(\mathcal{A}, \mu)$ is a weighted algebra with finite generating set X , that is archimedean with respect to the $*$ -positive cone of sums of squares. If $C_u^*(\mathcal{A}/\langle\langle \text{supp } \mu \rangle\rangle)$ is commutative, then $\mathcal{A}$ is Q -stably commutative for any tensor monoid Q .

Proof. Since $C_u^*(\mathcal{A}/\langle\langle \text{supp } \mu \rangle\rangle)$ is commutative, $[x, y] = 0$ in $C_u^*(\mathcal{A}/\langle\langle \text{supp } \mu \rangle\rangle)$ for all $x, y \in X$. This implies that $[x, y]$ is in the infinitesimal ideal of $\mathcal{A}/\langle\langle \text{supp } \mu \rangle\rangle$. Let $\varepsilon > 0$. By the above, we know $|[x, y]|^2 \leq \frac{\varepsilon}{2}$ in $\mathcal{A}/\langle\langle \text{supp } \mu \rangle\rangle$. By definition, there exists $p_{x,y} \in \mathcal{A}/\langle\langle \text{supp } \mu \rangle\rangle$ positive such that $\frac{\varepsilon}{2} - |[x, y]|^2 = p_{x,y}$. In $\mathcal{A}$, this gives that $|[x, y]|^2 - \frac{\varepsilon}{2} + p_{x,y} \in \langle\langle \text{supp } \mu \rangle\rangle$, so there exist $N_{x,y} \in \mathbb{N}$, $r_{x,y,i} \in \text{supp } \mu$, and $a_{x,y,i}, b_{x,y,i} \in \mathcal{A}$ such that

$$|[x, y]|^2 - \frac{\varepsilon}{2} + p_{x,y} = \sum_{i=1}^{N_{x,y}} a_{x,y,i} r_{x,y,i} b_{x,y,i}.$$

Now, let $M = \max_{x,y} \sum_{i=1}^{N_{x,y}} \frac{\|b_{x,y,i} a_{x,y,i}\|}{\sqrt{\mu(r_{x,y,i})}}$ and set $\delta = \left(\frac{\varepsilon}{2M}\right)^2$.

Suppose τ is a tracial state on $\mathcal{A}$ over Q such that $\text{def}(\tau) < \delta$. Then, for all $x, y \in X$,

$$\begin{aligned} \|[x, y]\|_\tau^2 &= \tau(|[x, y]|^2) = \tau\left(\frac{\varepsilon}{2} - p_{x,y} + \sum_{i=1}^{N_{x,y}} a_{x,y,i} r_{x,y,i} b_{x,y,i}\right) \\ &= \frac{\varepsilon}{2} - \tau(p_{x,y}) + \sum_{i=1}^{N_{x,y}} \tau(b_{x,y,i} a_{x,y,i} r_{x,y,i}) \\ &\leq \frac{\varepsilon}{2} + \sum_{i=1}^{N_{x,y}} \|b_{x,y,i} a_{x,y,i}\| \|r_{x,y,i}\|_\tau \\ &\leq \frac{\varepsilon}{2} + \sum_{i=1}^{N_{x,y}} \|b_{x,y,i} a_{x,y,i}\| \sqrt{\frac{\text{def}(\tau)}{\mu(r_{x,y,i})}} \\ &< \frac{\varepsilon}{2} + M\sqrt{\delta} = \varepsilon. \end{aligned} \quad \blacksquare$$

To finish this section, we find some equivalences for stable commutativity between different min-tensor monoids. Here, we follow the notation of [36] for ultrafilters.

Lemma 4.3.8. Let $(\mathcal{A}, \mu)$ be a weighted algebra where $\mathcal{A} = \mathbb{C}\mathbb{Z}_{d_1} * \cdots * \mathbb{C}\mathbb{Z}_{d_k}$. If $\mathcal{A}$ is q -stably commutative, then $\mathcal{A}$ is qa -stably commutative.

Proof. Let $\varepsilon > 0$ and fix the generating set $X = \{u_1, u_1^*, \dots, u_k, u_k^*\}$ of $\mathcal{A}$ such that $u_i^{d_i} = u_i^* u_i = 1$ with no other relations. Since $\mathcal{A}$ is q -stably commutative, there exists $\delta > 0$

such that if $\text{def}(\tau) < \delta$ for a tracial state $\tau : \mathcal{A} \rightarrow \mathbb{C}$ over q , then $\|[x, y]\|_\tau^2 < \varepsilon$ for all $x, y \in X$. Now, suppose τ is a tracial state on $\mathcal{A}$ over qa such that $\text{def}(\tau) < \frac{\delta}{2}$. Now consider the GNS representation of τ ; we may assume that it is minimal, giving a $*$ -representation $\pi : \mathcal{A} \rightarrow \mathcal{M}$, where $(\mathcal{M}, \text{tr})$ is a finite von Neumann algebra that embeds into $(\mathcal{R}^\omega, \tau_\omega)$ and $\tau = \text{tr} \circ \pi$. In particular tr is faithful. Further, by Lemma 2.4.1, $\mathcal{R}^\omega$ embeds into $(\prod_{\omega'} \mathcal{M}_n, \tau_{\omega'})$, where ω' is some free ultrafilter and $\mathcal{M}_n$ are finite-dimensional von Neumann algebras. By composing π with the embedding $\mathcal{M} \hookrightarrow \prod_{\omega'} \mathcal{M}_n$, we get a $*$ -representation $\pi' : \mathcal{A} \rightarrow \prod_{\omega'} \mathcal{M}_n$ such that $\tau = \tau_{\omega'} \circ \pi'$. This gives that π' can be realised as a sequence of maps (π_n) such that for all $a, b \in \mathcal{A}$, $\lim_{n \rightarrow \omega'} \|\pi_n(a)\pi_n(b) - \pi_n(ab)\|_{\text{tr}_n}^2 = 0$ and $\lim_{n \rightarrow \omega'} \|\pi_n(a^*) - \pi_n(a)^*\|_{\text{tr}_n}^2 = 0$. Now, we need to round the π_n to $*$ -representations of $\mathcal{A}$. Let $U_{n,i}$ be the nearest order- d_i unitary to $\pi_n(u_i)$. By Lemma 3.1.14, we know that $\|\pi_n(u_i) - U_{n,i}\|_{\text{tr}_n} \xrightarrow{\omega'} 0$. Then, we define the $*$ -representation $\pi'_n : \mathcal{A} \rightarrow \mathcal{M}_n$ by $\pi'_n(u_i) = U_{n,i}$. We have that $\|\pi_n(a) - \pi'_n(a)\|_{\text{tr}_n} \xrightarrow{\omega'} 0$ for all $a \in \mathcal{A}$, so the sequence (π'_n) also induces the representation π' on $\mathcal{R}^\omega$. Since $\text{def}(\tau) < \frac{\delta}{2}$, we know that $\text{def}(\text{tr}_n \circ \pi_n) \xrightarrow{\omega} \text{def}(\tau) < \frac{\delta}{2}$. Let $N = \{n \in \mathbb{N} \mid \text{def}(\text{tr}_n \circ \pi_n) < \delta\}$. By convergence, $N \in \omega'$ and by q -stable commutativity, for all $n \in N$, $\|[x, y]\|_{\text{tr}_n \circ \pi_n}^2 < \varepsilon$ for all $x, y \in X$. Now, suppose there exist $x, y \in X$ such that $\|[x, y]\|_\tau^2 \geq \varepsilon$. Then, there exists $M \in \omega'$ such that for all $n \in M$, $\|[x, y]\|_{\text{tr}_n \circ \pi_n}^2 \geq \varepsilon$. However, we get that M and N are disjoint, so $\emptyset = M \cap N \in \omega'$. Contradiction, so $\|[x, y]\|_\tau^2 < \varepsilon$ for all $x, y \in X$, giving that $\mathcal{A}$ is qa -stably commutative. $\blacksquare$

Lemma 4.3.9. Let $(\mathcal{A}, \mu)$ be a weighted algebra. If every representation of $C_u^*(\mathcal{A}/\langle\langle \text{supp } \mu \rangle\rangle)$ into $\mathcal{R}^\omega$ (for some non-principal ultrafilter ω) is commutative, then $\mathcal{A}$ is q -stably commutative.

Proof. Fix a finite generating set X of $\mathcal{A}$, and for a tracial state τ write $\text{comm-def}(\tau) = \sum_{x, y \in X} \|[x, y]\|_\tau^2$. We proceed by contrapositive. Suppose $\mathcal{A}$ is not q -stably commutative. Then, there exists $\varepsilon > 0$ such that for all $\delta > 0$, there exists a finite-dimensional tracial state τ on $\mathcal{A}$ such that $\text{def}(\tau) < \delta$ but $\text{comm-def}(\tau) \geq \varepsilon$. In particular, for each $n \in \mathbb{N}$, there exists a finite-dimensional tracial state τ_n such that $\text{def}(\tau_n) < \frac{1}{2^n}$ but $\text{comm-def}(\tau_n) \geq \varepsilon$. Now, consider the GNS representation $\pi_n : \mathcal{A} \rightarrow \mathcal{M}_n$, where $\mathcal{M}_n$ is a finite-dimensional von Neumann algebra with tracial state tr_n ; we may without loss of generality assume that $\dim \mathcal{M}_n \rightarrow \infty$ in n . Fix a non-principal ultrafilter ω , and let $\mathcal{M} = \prod_\omega \mathcal{M}_n$ the ultraproduct. $\mathcal{M}$ has a tracial state tr defined as $\text{tr}(x) = \lim_{n \rightarrow \omega} \text{tr}_n(x_n)$, and embeds in $\mathcal{R}^\omega$ by Lemma 2.4.1. Let $\pi : \mathcal{A} \rightarrow \mathcal{M}$ be defined as $\pi(a) = (\pi_n(a))$. This is a $*$ -homomorphism and induces a tracial state $\tau = \text{tr} \circ \pi$ on $\mathcal{A}$. We also have that

$$\text{def}(\tau) = \sum_{a \in \mathcal{A}} \mu(a) \tau(a^*a) = \sum_{a \in \mathcal{A}} \mu(a) \lim_{n \rightarrow \omega} \text{tr}_n(\pi_n(a^*a)) = \lim_{n \rightarrow \omega} \sum_{a \in \mathcal{A}} \mu(a) \|a\|_{\tau_n}^2 = \lim_{n \rightarrow \omega} \frac{1}{2^n} = 0.$$

As such, π factors through $\mathcal{A}/\langle\langle \text{supp } \mu \rangle\rangle$ to give a representation $\bar{\pi} : \mathcal{A}/\langle\langle \text{supp } \mu \rangle\rangle \rightarrow \mathcal{M}$. By universal property, this induces a representation $\bar{\bar{\pi}} : C_u^*(\mathcal{A}/\langle\langle \text{supp } \mu \rangle\rangle) \rightarrow \mathcal{M}$. To conclude, it remains to see that $\text{im } \bar{\bar{\pi}} = \overline{\text{im } \pi}^{\|\cdot\|}$ is not commutative. Let $a = \sum_{x,y \in X} |[x, y]|^2$. Since tr is faithful, if $\tau(a) = \text{comm-def}(\tau) \neq 0$, then $\text{im } \bar{\bar{\pi}}$ is not commutative. We have that

$$\tau(a) = \lim_{n \rightarrow \omega} \tau_n(a) = \lim_{n \rightarrow \omega} \text{comm-def}(\tau_n).$$

However, $\text{comm-def}(\tau_n) \geq \varepsilon$ for all n , and therefore $\{n \in \mathbb{N} \mid |\text{comm-def}(\tau_n)| < \varepsilon\} = \emptyset \notin \omega$. Thus, $(\text{comm-def}(\tau_n))$ does not tend to 0, so $\bar{\bar{\pi}}$ is a noncommutative representation of $C_u^*(\mathcal{A}/\langle\langle \text{supp } \mu \rangle\rangle)$ into $\mathcal{M} \hookrightarrow \mathcal{R}^\omega$. $\blacksquare$

Lemma 4.3.10. Let $(\mathcal{A}, \mu)$ be a weighted algebra. If every representation of $C_u^*(\mathcal{A}/\langle\langle \text{supp } \mu \rangle\rangle)$ into $\mathcal{B} \in co$ is commutative, then $\mathcal{A}$ is co -stably commutative.

The proof follows in the same way as [Lemma 4.3.9](#), by noting that co is closed under ultraproducts.

4.3.3 Robust commutativity gadgets

Using the stability notions of [Section 4.3.2](#), we can fully characterise when robust commutativity gadgets exist in terms of properties of the quantum polymorphisms.

Theorem 4.3.11. Let A be a relational structure and let Q be a finite tensor monoid. The following are equivalent.

- (i) $\text{Mor}_{\mathbf{u}}^{no}(A^k, A)$ is Q -stably commutative for all $k \in \mathbb{N}$.
- (ii) $\text{Mor}_{\mathbf{u}}^{no}(A^{|A|^2}, A)$ is Q -stably commutative.
- (iii) A admits a robust Q -commutativity gadget.

Proof.

(i) $\Rightarrow$ (ii) Immediate.

(ii) $\Rightarrow$ (iii) Suppose $\text{Mor}_{\mathbf{u}}^{no}(A^{|A|^2}, A)$ is Q -stably commutative. Following the notation of [Lemma 4.3.3](#), we claim $(A^{|A|^2}, \mathbf{x}, \mathbf{y})$ is a robust Q -commutativity gadget for A . Since $\text{Mor}_{\mathbf{u}}^{no}(A^{|A|^2}, A)$ is Q -stably commutative, we know that every $*$ -representation of the quantum homomorphism algebra $\text{Mor}^+(A^{|A|^2}, A) = C_u^*(\text{Mor}_{\mathbf{u}}^{no}(A^{|A|^2}, A)/\langle\langle \text{supp } \mu_{no, \mathbf{u}} \rangle\rangle)$ over Q

is commutative, and hence $\text{mor}^Q(A^{|A|^2}, A) \subseteq \text{mor}^c(A^{|A|^2}, A)$. By Lemma 4.3.3, $(A^{|A|^2}, \mathbf{x}, \mathbf{y})$ is a commutativity gadget for A . For the additional robustness property, it follows from the stable commutativity property that for any $\varepsilon > 0$, there exists $\delta > 0$ such that for any tracial state τ on $\text{Mor}_{\mathfrak{u}}^{no}(A^{|A|^2}, A)$ over Q , if $\text{def}(\tau) < \delta$, then $\| [p_{\mathbf{a}, \mathbf{a}}, p_{\mathbf{b}, \mathbf{b}}] \|_{\tau}^2 \leq \frac{\varepsilon}{|A|^2}$. So, this implies that $\sum_{\mathbf{a}, \mathbf{b} \in A} \| [p_{\mathbf{x}, \mathbf{a}}, p_{\mathbf{y}, \mathbf{b}}] \|_{\tau}^2 < \varepsilon$, and hence $(A^{|A|^2}, \mathbf{x}, \mathbf{y})$ is a robust commutativity gadget.

(iii) $\Rightarrow$ (i) As in Lemma 4.3.2, we show the contrapositive by contradiction. Suppose (G, x, y) is a robust commutativity gadget. Also, there exists $k \in \mathbb{N}$ such that $\text{Mor}_{\mathfrak{u}}^{no}(A^k, A)$ is not Q -stably commutative. Write $m = \sum_{R \in \sigma} |R^G|$ and $M = \sum_{R \in \sigma} |R^{A^k}|$. Then, there exists $\varepsilon > 0$ such that for any $\delta > 0$ there is a tracial state τ on $\text{Mor}_{\mathfrak{u}}^{no}(A^k, A)$ over Q such that $\text{def}(\tau) < \frac{\delta}{M^2}$ but $\| [p_{\mathbf{a}, \mathbf{a}}, p_{\mathbf{b}, \mathbf{b}}] \|_{\tau}^2 \geq \varepsilon$ for some $\mathbf{a}, \mathbf{b} \in A^k$ and $a, b \in A$. Fix $\delta > 0$, and the corresponding $\tau, \mathbf{a}, \mathbf{b}, a, b$. Let $\pi : \text{Mor}_{\mathfrak{u}}^{no}(A^k, A) \rightarrow \mathcal{M}$ be the GNS representation of τ , with a tracial state ρ on $\mathcal{M}$ such that $\tau = \rho \circ \pi$. Let φ be the representation of $\text{Mor}^+(G, A^k)$ defined in Lemma 4.3.2. Write ν for the tracial state on $\text{im } \varphi$. Now, define the $*$ -representation π' of $\text{Mor}_{\mathfrak{u}}^{no}(G, A)$ on the generators by

$$\pi'(p_{z, c}) = \sum_{c \in A^k} \varphi(p_{z, c}) \otimes \pi(p_{c, c}).$$

This definition clearly extends to a $*$ -representation as the $\pi'(p_{z, c})$ are projections and

$$\sum_{c \in A} \pi'(p_{z, c}) = \sum_{c \in A, c \in A^k} \varphi(p_{z, c}) \otimes \pi(p_{c, c}) = 1.$$

Note also that $(\nu \otimes \rho)$ is a tracial state on $\text{im } \pi'$, which induces a tracial state $\tau' = (\nu \otimes \rho) \circ \pi'$ on $\text{Mor}_{\mathfrak{u}}^{no}(G, A)$. The defect of this state is upper-bounded as

$$\begin{aligned} \text{def}(\tau') &= \frac{1}{m} \sum_{R \in \sigma, z \in R^G, c \notin R^A} \| p_{z_1, c_1} \cdots p_{z_{\text{ar}(R)}, c_{\text{ar}(R)}} \|_{\tau'}^2 \\ &= \frac{1}{m} \sum_{R \in \sigma, z \in R^G, c \notin R^A} \left\| \sum_{c_1, \dots, c_k \in A^k} \varphi(p_{z_1, c_1} \cdots p_{z_{\text{ar}(R)}, c_{\text{ar}(R)}}) \otimes \pi(p_{c_1, c_1} \cdots p_{c_{\text{ar}(R)}, c_{\text{ar}(R)}}) \right\|_{\nu \otimes \rho}^2 \\ &= \frac{1}{m} \sum_{R \in \sigma, z \in R^G, c \notin R^A} \left\| \sum_{c \in R^{A^k}} \varphi(p_{z_1, c_1} \cdots p_{z_{\text{ar}(R)}, c_{\text{ar}(R)}}) \otimes \pi(p_{c_1, c_1} \cdots p_{c_{\text{ar}(R)}, c_{\text{ar}(R)}}) \right\|_{\nu \otimes \rho}^2 \\ &\leq \frac{1}{m} \sum_{R \in \sigma, z \in R^G, c \notin R^A, c \in R^{A^k}} |R^{A^k}| \cdot \| p_{z_1, c_1} \cdots p_{z_{\text{ar}(R)}, c_{\text{ar}(R)}} \|_{\nu \circ \varphi}^2 \cdot \| p_{c_1, c_1} \cdots p_{c_{\text{ar}(R)}, c_{\text{ar}(R)}} \|_{\tau}^2 \\ &\leq \frac{1}{m} \sum_{R \in \sigma, c \notin R^A, c \in R^{A^k}} |R^G| \cdot |R^{A^k}| \cdot \| p_{c_1, c_1} \cdots p_{c_{\text{ar}(R)}, c_{\text{ar}(R)}} \|_{\tau}^2 \\ &\leq M^2 \text{def}(\tau) < \delta. \end{aligned}$$

On the other hand, since $\pi'(p_{x,a}) = \sum_{c \in A^k} \varphi(p_{x,c}) \otimes \pi(p_{c,a}) = 1 \otimes \pi(p_{a,a})$ and similarly $\pi'(p_{y,b}) = 1 \otimes \pi(p_{b,b})$, so

$$\sum_{c,d \in A} \|[p_{x,c}, p_{y,d}]\|_{\tau'}^2 \geq \|[p_{x,a}, p_{y,b}]\|_{\tau'}^2 = \|[p_{a,a}, p_{b,b}]\|_{\tau}^2 \geq \varepsilon.$$

Hence (G, x, y) is not a robust Q -commutativity gadget, giving the wanted contradiction. $\blacksquare$

Theorem 4.3.12. Let A be a relational structure and Q be a finite tensor monoid. The following are equivalent.

- (i) $\text{Mor}_{\mathbf{u}}^{\text{ocv}}(A^k, A)$ is Q -stably commutative for all $k \in \mathbb{N}$.
- (ii) $\text{Mor}_{\mathbf{u}}^{\text{ocv}}(A^{|A|^2}, A)$ is Q -stably commutative.
- (iii) A admits a robust oracular Q -commutativity gadget.

Proof.

$\boxed{\text{(i)} \Rightarrow \text{(ii)}}$ Immediate.

$\boxed{\text{(ii)} \Rightarrow \text{(iii)}}$ Suppose $\text{Mor}_{\mathbf{u}}^{\text{ocv}}(A^{|A|^2}, A)$ is Q -stably commutative. Following the notation of [Lemma 4.3.3](#), we claim $(A^{|A|^2}, \mathbf{x}, \mathbf{y})$ is a robust oracular Q -commutativity gadget for A . Since $\text{Mor}_{\mathbf{u}}^{\text{ocv}}(A^{|A|^2}, A)$ is Q -stably commutative, we know that every representation of the oracular quantum homomorphism algebra $\text{Mor}^{o+}(A^{|A|^2}, A) = C_u^*(\text{Mor}_{\mathbf{u}}^{\text{ocv}}(A^{|A|^2}, A) / \langle\langle \text{supp } \mu_{\text{ocv}, \mathbf{u}} \rangle\rangle\rangle$ over Q is commutative, and hence $\text{mor}^Q(A^{|A|^2}, A) \subseteq \text{mor}^c(A^{|A|^2}, A)$. By [Theorem 4.3.4](#), $(A^{|A|^2}, \mathbf{x}, \mathbf{y})$ is an oracular Q -commutativity gadget for A . For the additional robustness property, it follows from the stable commutativity property that for any $\varepsilon > 0$, there exists $\delta > 0$ such that for any tracial state τ on $\text{Mor}_{\mathbf{u}}^{\text{ocv}}(A^{|A|^2}, A)$ over Q , if $\text{def}(\tau) < \delta$, then $\|[p_{a,a}, p_{b,b}]\|_{\tau}^2 \leq \frac{\varepsilon}{|A|^2}$. So, this implies that $\sum_{a,b \in A} \|[p_{x,a}, p_{y,b}]\|_{\tau}^2 < \varepsilon$, and hence $(A^{|A|^2}, \mathbf{x}, \mathbf{y})$ is an robust oracular Q -commutativity gadget.

$\boxed{\text{(iii)} \Rightarrow \text{(i)}}$ First, for $R \in \sigma$, $\mathfrak{c} \in R^{A^k}$, write $\Phi_{R, \mathfrak{c}, i, c} = \sum_{c \in R^A: c_i = c} \Phi_{R, \mathfrak{c}, c}$. We have that

$$\begin{aligned} \|\Phi_{R, \mathfrak{c}, i, c} - p_{c_i, c}\|_{\tau}^2 &= \tau(\Phi_{R, \mathfrak{c}, i, c} - p_{c_i, c}) + 2\tau(\Phi_{R, \mathfrak{c}, i, c}(1 - p_{c_i, c})) \\ &\leq 2 \sum_{c \in A} \tau(\Phi_{R, \mathfrak{c}, i, c}(1 - p_{c_i, c})) \\ &= 2 \sum_{c \in R^A} \|\Phi_{R, \mathfrak{c}, c}(1 - p_{c_i, c_i})\|_{\tau}^2 \\ &\leq 2ML \text{def}(\tau), \end{aligned}$$

where as before $M = \sum_{R \in \sigma} |R^{A^k}|$ and $L = \max_{R \in \sigma} \text{ar}(R)$. This implies that

$$\begin{aligned} \|\Phi_{R, \mathfrak{c}, i, c}, p_{\mathfrak{d}, d}\|_\tau^2 &\leq 3\|p_{\mathfrak{c}, i, c}, p_{\mathfrak{d}, d}\|_\tau^2 + 12ML \text{def}(\tau) \\ \|\Phi_{R, \mathfrak{c}, i, c}, \Phi_{S, \mathfrak{d}, j, d}\|_\tau^2 &\leq 5\|p_{\mathfrak{c}, i, c}, p_{\mathfrak{d}, j, d}\|_\tau^2 + 40ML \text{def}(\tau) \end{aligned}$$

As in [Lemma 4.3.2](#), we show the contrapositive by contradiction. Suppose (G, x, y) is a robust oracular Q -commutativity gadget. Also, there exists $k \in \mathbb{N}$ such that $\text{Mor}_{\mathfrak{u}}^{\text{ocv}}(A^k, A)$ is not Q -stably commutative. In particular that means that it is not Q -stably commutative with respect to the generating set $\{\Phi_{R, \mathfrak{c}, i, c}\}_{R \in \sigma, \mathfrak{c} \in R^{A^k}, i \in [\text{ar}(R)], c \in A} \cup \{p_{\mathfrak{c}, c}\}_{\mathfrak{c} \in A^k, c \in A}$ via [Lemma 4.3.6](#). Write $m = \sum_{R \in \sigma} |R^G|$. Then, there exists $\varepsilon > 0$ such that for any $\delta > 0$ there is a trace τ on $\text{Mor}_{\mathfrak{u}}^{\text{ocv}}(A^k, A)$ over Q such that $\text{def}(\tau) \leq \frac{\delta}{M}$ but $\|\Phi_{R, \mathfrak{a}, i, a}, \Phi_{S, \mathfrak{b}, j, b}\|_\tau^2 > 10\varepsilon$ for some $R, S \in \sigma$, $\mathfrak{a} \in R^{A^k}$, $\mathfrak{b} \in S^{A^k}$, $i \in [\text{ar}(R)]$, $j \in [\text{ar}(S)]$, and $a, b \in A$; or $\|\Phi_{R, \mathfrak{a}, i, a}, p_{\mathfrak{b}, b}\|_\tau^2 > 10\varepsilon$ for some $R \in \sigma$, $\mathfrak{a} \in R^{A^k}$, $i \in [\text{ar}(R)]$, $\mathfrak{b} \in A^k$, and $a, b \in A$; or $\|p_{\mathfrak{a}, a}, p_{\mathfrak{b}, b}\|_\tau^2 > 10\varepsilon$ for some $\mathfrak{a}, \mathfrak{b} \in A^k$ and $a, b \in A$. By the argument at the beginning of the proof, this implies that in either of the three cases, $\|p_{\mathfrak{a}, a}, p_{\mathfrak{b}, b}\|_\tau^2 > 2\varepsilon - 8ML \text{def}(\tau)$ for some $\mathfrak{a}, \mathfrak{b} \in A^k$ and $a, b \in A$. If $\frac{\delta}{M} \leq \frac{\varepsilon}{8ML}$, we get that $\|p_{\mathfrak{a}, a}, p_{\mathfrak{b}, b}\|_\tau^2 > \varepsilon$

Fix $\delta > 0$ satisfying $\frac{\delta}{M} \leq \frac{\varepsilon}{8ML}$, and the corresponding $\tau, \mathfrak{a}, \mathfrak{b}, a, b$ as above. Let $\pi : \text{Mor}_{\mathfrak{u}}^{\text{ocv}}(A^k, A) \rightarrow \mathcal{M}$ be the GNS representation of τ , with a tracial state ρ on $\mathcal{M}$ such that $\tau = \rho \circ \pi$. Let φ be the representation of $\text{Mor}^+(G, A^k)$ defined in [Lemma 4.3.2](#). Write ν for the tracial state on $\text{im } \varphi$. Now, define the $*$ -representation π' of $\text{Mor}_{\mathfrak{u}}^{\text{ocv}}(G, A)$ on the generators by

$$\begin{aligned} \pi'(\Phi_{R, \mathfrak{z}, c}) &= \sum_{\mathfrak{c} \in R^{A^k}} \varphi(p_{z_1, c_1} \cdots p_{z_{\text{ar}(R)}, c_{\text{ar}(R)}}) \otimes \pi(\Phi_{R, \mathfrak{c}, c}), \\ \pi'(p_{z, c}) &= \sum_{\mathfrak{c} \in A^k} \varphi(p_{z, c}) \otimes \pi(p_{\mathfrak{c}, c}). \end{aligned}$$

Below we write as usual $p_{\mathfrak{z}, \mathfrak{c}} = p_{z_1, c_1} \cdots p_{z_{\text{ar}(R)}, c_{\text{ar}(R)}}$. To see that π' is in fact a $*$ -representation, it suffices to note that the relations of $\text{Mor}_{\mathfrak{u}}^{\text{ocv}}(A^{|A|^2}, A)$ are satisfied. First, $\pi'(\Phi_{R, \mathfrak{z}, c})$ and $\pi'(p_{z, c})$ are projections as they are sums of orthogonal projections. Also,

$$\sum_{\mathfrak{c} \in R^A} \pi'(\Phi_{R, \mathfrak{z}, c}) = \sum_{\mathfrak{c} \in R^A} \sum_{\mathfrak{c} \in R^{A^k}} \varphi(p_{\mathfrak{z}, \mathfrak{c}}) \otimes \pi(\Phi_{R, \mathfrak{c}, c}) = \sum_{\mathfrak{c} \in R^{A^k}} \varphi(p_{\mathfrak{z}, \mathfrak{c}}) \otimes 1 = 1,$$

and

$$\sum_{c \in A} \pi'(p_{z, c}) = \sum_{c \in A} \sum_{\mathfrak{c} \in A^k} \varphi(p_{z, c}) \otimes \pi(p_{\mathfrak{c}, c}) = \sum_{\mathfrak{c} \in A^k} \varphi(p_{z, c}) \otimes 1 = 1.$$

Now, since $(\nu \otimes \rho)$ is a tracial state on $\text{im } \pi'$, $\tau' = (\nu \otimes \rho) \circ \pi'$ is a tracial state on $\text{Mor}_{\mathbf{u}}^{\text{ocv}}(G, A)$. The defect of this tracial state is

$$\begin{aligned}
\text{def}(\tau') &= \frac{1}{m} \sum_{\substack{R \in \sigma, z \in R^G, \\ i \in [\text{ar}(R)], c \in R^A}} \frac{1}{\text{ar}(R)} \tau'(\Phi_{R,z,c}(1 - p_{z_i, c_i})) \\
&= \frac{1}{m} \sum_{\substack{R \in \sigma, z \in R^G, \\ i \in [\text{ar}(R)], c \neq d \in A}} \frac{1}{\text{ar}(R)} \tau'(\Phi_{R,z,i,c} p_{z_i, d}) \\
&= \frac{1}{m} \sum_{\substack{R \in \sigma, z \in R^G, \\ i \in [\text{ar}(R)], c \neq d \in A}} \frac{1}{\text{ar}(R)} \sum_{\mathbf{c} \in R^{A^k}, \mathbf{d} \in A^k} (\nu \circ \varphi)(p_{z_1, \mathbf{c}_1} \cdots p_{z_{\text{ar}(R)}, \mathbf{c}_{\text{ar}(R)}} p_{z_i, \mathbf{d}}) \tau(\Phi_{R, \mathbf{c}, i, c} p_{\mathbf{d}, d}) \\
&= \frac{1}{m} \sum_{\substack{R \in \sigma, z \in R^G, \\ i \in [\text{ar}(R)], c \neq d \in A}} \frac{1}{\text{ar}(R)} \sum_{\mathbf{c} \in R^{A^k}} (\nu \circ \varphi)(p_{z_1, \mathbf{c}_1} \cdots p_{z_{\text{ar}(R)}, \mathbf{c}_{\text{ar}(R)}}) \tau(\Phi_{R, \mathbf{c}, i, c} p_{\mathbf{c}_i, d}) \\
&\leq \frac{1}{m} \sum_{\substack{R \in \sigma, z \in R^G, \\ i \in [\text{ar}(R)], c \neq d \in A}} \frac{1}{\text{ar}(R)} \sum_{\mathbf{c} \in R^{A^k}} \tau(\Phi_{R, \mathbf{c}, i, c} p_{\mathbf{c}_i, d}) \\
&= M \text{def}(\tau) \leq M \frac{\delta}{M} = \delta
\end{aligned}$$

On the other hand, since $\pi'(p_{x,a}) = \sum_{\mathbf{c} \in A^k} \varphi(p_{x,\mathbf{c}}) \otimes \pi(p_{\mathbf{c},a}) = 1 \otimes \pi(p_{\mathbf{a},a})$ and similarly $\pi'(p_{y,b}) = 1 \otimes \pi(p_{\mathbf{b},b})$, so

$$\sum_{c,d \in A} \|[p_{x,c}, p_{y,d}]\|_{\tau'}^2 \geq \|[p_{x,a}, p_{y,b}]\|_{\tau'}^2 = \|[p_{\mathbf{a},a}, p_{\mathbf{b},b}]\|_{\tau}^2 > \varepsilon.$$

Hence (G, x, y) is not a robust oracular Q -commutativity gadget, giving the wanted contradiction. $\blacksquare$

4.3.4 Implications and equivalences

We use the characterisation of robust commutativity gadgets, as well the properties of stable commutativity from [Section 4.3.2](#) to show some implications and equivalences between different types of commutativity gadget.

First, we show a lemma relating oracular and non-oracular commutativity gadgets.

Lemma 4.3.13. Suppose (G, x, y) is a Q -commutativity gadget (*resp.* robust Q -commutativity gadget, algebraic commutativity gadget) for A such that $\pi_{a,b} \in \text{mor}^c(G, A)$ for all $a, b \in A$.

Then, (G, x, y) is an oracular commutativity gadget (*resp.* robust oracular Q -commutativity gadget, oracular algebraic commutativity gadget).

Note that the hypothesis of this lemma is always satisfied for commutativity gadgets constructed from the polymorphism spaces, but not for all gadgets.

Proof. First, if (G, x, y) is a commutativity gadget with $\pi_{a,b}$ classical, then $\pi_{a,b}$ factors through $\text{Mor}^{o+}(G, A)$. Hence, (G, x, y) satisfies property (i) of an oracular commutativity gadget. Next, to show property (ii), suppose that $\pi \in \text{mor}^{oQ}(G, A)$. Then, by precomposing with the quotient map $\text{Mor}^+(G, A) \rightarrow \text{Mor}^{o+}(G, A)$, this induces a representation of $\text{Mor}^+(G, A)$ and by the commutativity gadget property, we have that $[\pi(p_{x,a}), \pi(p_{y,b})] = 0$ for all $a, b \in A$. The same proof with π replaced by the identity also shows that if (G, x, y) is an algebraic commutativity gadget under the same assumption, then it is also an oracular algebraic commutativity gadget.

Now, suppose (G, x, y) is an robust commutativity gadget. Let τ be a tracial state on $\text{Mor}_{\mathbf{u}}^{ocv}(G, A)$. By precomposing with the inclusion map $\text{Mor}_{\mathbf{u}}^{no}(G, A) \rightarrow \text{Mor}_{\mathbf{u}}^{ocv}(G, A)$, τ induces a tracial state on $\text{Mor}_{\mathbf{u}}^{no}(G, A)$, which we refer to as τ' . Let $m = \sum_{R \in \sigma} |R^G|$. We have that

$$\begin{aligned}
\text{def}(\tau') &= \frac{1}{m} \sum_{R \in \sigma, \mathbf{a} \in R^A, \mathbf{b} \notin R^B} \|p_{a_1, b_1} \cdots p_{a_{\text{ar}(R)}, b_{\text{ar}(R)}}\|_{\tau}^2 \\
&= \frac{1}{m} \sum_{R \in \sigma, \mathbf{a} \in R^A, \mathbf{b} \notin R^B} \left\| (p_{a_1, b_1} - \Phi_{R, \mathbf{a}, 1, b_1}) p_{a_2, b_2} \cdots p_{a_{\text{ar}(R)}, b_{\text{ar}(R)}} \right. \\
&\quad \left. + \Phi_{R, \mathbf{a}, 1, b_1} (p_{a_2, b_2} - \Phi_{R, \mathbf{a}, 2, b_2}) p_{a_3, b_3} \cdots p_{a_{\text{ar}(R)}, b_{\text{ar}(R)}} + \cdots \right. \\
&\quad \left. + \Phi_{R, \mathbf{a}, 1, b_1} \cdots \Phi_{R, \mathbf{a}, \text{ar}(R)-1, b_{\text{ar}(R)-1}} (p_{a_{\text{ar}(R)}, b_{\text{ar}(R)}} - \Phi_{R, \mathbf{a}, \text{ar}(R), b_{\text{ar}(R)}}) + \Phi_{R, \mathbf{a}, \mathbf{b}} \right\|_{\tau}^2 \\
&\leq \frac{1}{m} \sum_{R \in \sigma, \mathbf{a} \in R^A, \mathbf{b} \in B} \text{ar}(R) \sum_{i=1}^{\text{ar}(R)} \|p_{a_i, b} - \Phi_{R, \mathbf{a}, i, b}\|_{\tau}^2 \\
&= \frac{2}{m} \sum_{R \in \sigma, \mathbf{a} \in R^A, \mathbf{b} \in R^B, i \in [\text{ar}(R)]} \text{ar}(R) \|\Phi_{R, \mathbf{a}, \mathbf{b}}(1 - p_{a_i, b_i})\|_{\tau}^2 \leq 2L^2 \text{def}(\tau),
\end{aligned}$$

where $L = \max_{R \in \sigma} \text{ar}(R)$. Now, let $\varepsilon > 0$. As (G, x, y) is an robust commutativity gadget, there exists $\delta > 0$ such that if φ is a tracial state on $\text{Mor}_{\mathbf{u}}^{no}(G, A)$ such that $\text{def}(\varphi) < \delta$, then $\sum_{a, b \in A} \| [p_{x, a}, p_{y, b}] \|_{\varphi}^2 < \varepsilon$. Suppose that τ is a tracial state on $\text{Mor}_{\mathbf{u}}^{ocv}(G, A)$ such that $\text{def}(\tau) < \frac{\delta}{2L^2}$. By the above, τ' is a tracial state on $\text{Mor}_{\mathbf{u}}^{no}(G, A)$ such that $\text{def}(\tau') < \delta$, and hence $\sum_{a, b \in A} \| [p_{x, a}, p_{y, b}] \|_{\tau'}^2 = \sum_{a, b \in A} \| [p_{x, a}, p_{y, b}] \|_{\tau}^2 < \varepsilon$. So (G, x, y) is a robust oracular commutativity gadget. $\blacksquare$




Figure 4.2: Relationships between commutativity gadget classes shown by [Theorem 4.3.14](#). Trivial implications are denoted by solid grey arrows, the implications (i) – (iv) are denoted by solid black arrows, the implications (v) and (vi) are denoted by dashed black arrows, the implications (vii) and (viii) are denoted by dotted black arrows, and the equivalences (ix) – (xii) are denoted by blue boxes.

Theorem 4.3.14. Let A be a relational structure. Let Q and Q' be finite tensor monoids.

- (i) If $Q \subseteq Q'$ and A admits a Q' -commutativity gadget, then A admits a Q -commutativity gadget.
- (ii) If $Q \subseteq Q'$ and A admits an oracular Q' -commutativity gadget, then A admits an oracular Q -commutativity gadget.
- (iii) If $Q \subseteq Q'$ and A admits a robust Q' -commutativity gadget, then A admits a robust Q -commutativity gadget.
- (iv) If $Q \subseteq Q'$ and A admits a robust oracular Q' -commutativity gadget, then A admits a robust oracular Q -commutativity gadget.
- (v) If A admits Q -commutativity gadget, then A admits an oracular Q -commutativity gadget.
- (vi) If A admits a robust Q -commutativity gadget, then A admits a robust oracular Q -commutativity gadget.
- (vii) If A admits an algebraic commutativity gadget, then A admits a robust Q -commutativity gadget for all finite Q .

- (viii) If A admits an oracular algebraic commutativity gadget, then A admits a robust oracular Q -commutativity gadget for all finite Q .
- (ix) A admits a qa -commutativity gadget iff A admits a robust qa -commutativity gadget iff A admits a robust q -commutativity gadget.
- (x) A admits an oracular qa -commutativity gadget iff A admits a robust oracular qa -commutativity gadget iff A admits a robust oracular q -commutativity gadget.
- (xi) A admits a co -commutativity gadget iff A admits a robust co -commutativity gadget.
- (xii) A admits an oracular co -commutativity gadget iff A admits a robust oracular co -commutativity gadget.

The results of this theorem are summarised in [Fig. 4.2](#)

Proof. (i) to (iv) follow as the commutativity gadget constructed in [Lemma 4.3.3](#) has deterministic morphisms $\pi_{a,b}$, and every Q -morphism is a Q' -morphism. (v) and (vi) follow as the commutativity gadgets constructed in [Theorems 4.3.1](#) and [4.3.11](#) satisfy the conditions of [Lemma 4.3.13](#). (vii) and (viii) follow using [Lemma 4.3.7](#) to note that a commutative quantum polymorphism algebra implies Q -stable commutativity. (ix) and (x) follow by using [Lemma 4.3.8](#) to equate q -stable commutativity and qa -stable commutativity, and using [Lemma 4.3.9](#) to equate q -stable commutativity and commutativity under Connes-embeddable representations. Finally, (xi) and (xii) follow by [Lemma 4.3.10](#). ■

4.4 Polymorphism algebras of the complete graphs

In this section, we study the entangled polymorphisms of the complete graphs K_n for $n \geq 3$. This gives an example of how quantum polymorphisms relate to the classical ones, and will be useful in a few applications, beginning in the following section.

Before studying the quantum case, we mention the classical case, which will motivate our analysis. It is known that the endomorphisms of a complete graph are the permutations, and that the polymorphisms take the form of a projection down to one of the components, followed by a permutation. That is, the polymorphisms $f : K_n^k \rightarrow K_n$ have the form $f = \sigma \circ \pi_i$, where $\pi_i : K_n^k \rightarrow K_n$ is the canonical projection $\pi_i(\mathbf{x}) = x_i$ and $\sigma : K_n \rightarrow K_n$ is a permutation. As such, it follows that the commutative algebra $\text{Mor}(K_n^k, K_n)$ is generated by the projections $p_{\mathbf{x},a}(\sigma \circ \pi_i) = \delta_{\sigma(x_i),a}$. To realise the bijection $\text{mor}(K_n^k, K_n) = S_n^{\sqcup k}$ at

the level of the algebra, let $\mathbf{x}, \mathbf{x}' \in K_n^k$ such that $x_i = x'_i$ but $x_j \neq x'_j$ for all $j \neq i$, and consider $\Pi_i = \sum_a p_{x,a} p_{x',a}$. Then, we have that $\Pi_i(\sigma \circ \pi_j) = \delta_{i,j}$, which gives that Π_i is independent of the choice of $\mathbf{x}, \mathbf{x}'$ satisfying $x_j = x'_j \iff j = i$ for all j , and that $\{\Pi_i\}_{i \in [k]}$ forms a resolution of identity. By construction, $\Pi_i \text{Mor}(K_n^k, K_n)$ admits the coproduct $\Delta(f)(\sigma \circ \pi_j, \tau \circ \pi_{j'}) = \delta_{i,j,j'} f(\sigma \circ \tau \circ \pi_i)$, giving that $\Pi_i \text{Mor}(K_n^k, K_n) \cong C(S_n)$.

Theorem 4.4.1. For $n \geq 3$, $\text{mor}^+(K_n^k, K_n)$ is isomorphic to $(S_n^+)^{\sqcup k}$ in the following way: there exists a central PVM $\{\Pi_i\}_{i=1}^k$ in $\text{Mor}^+(K_n^k, K_n)$ such that for each i , the C^* -algebra

$$\Pi_i \text{Mor}^+(K_n^k, K_n) \cong \text{End}^+(K_n) = S_n^+$$

via the $*$ -homomorphism extending $\Pi_i p_{x,a} \mapsto p_{x_i,a}$.

In the following we always assume $n \geq 3$.

Lemma 4.4.2. Let $\mathbf{x}_1, \dots, \mathbf{x}_n \in K_n^k$ be such that $x_{ij} \neq x_{i'j}$ for all $i \neq i'$. Then, for all $a \in K_n$, $p_{x_1,a} + \dots + p_{x_n,a} = 1$ in $\text{Mor}^+(K_n^k, K_n)$.

Proof. Since $\mathbf{x}_i \sim_{K_n^k} \mathbf{x}_j$ for all $i \neq j$ by hypothesis, we have that $p_{x_i,a} p_{x_j,a} = 0$. As such, the $p_{x_i,a}$ are orthogonal projections, giving $p_{x_1,a} + \dots + p_{x_n,a} \leq 1$. Summing over a gives

$$n = \sum_i 1 = \sum_{i,a} p_{x_i,a} \leq \sum_a 1 = n,$$

and hence the inequalities must all be equalities. ■

Lemma 4.4.3. Let $\mathbf{x}, \mathbf{x}' \in K_n^k$ such that $x_i \neq x'_i$ for all i . Fix $S \subseteq [k]$ and define $\mathbf{y}, \mathbf{y}' \in K_n^k$ as $y_i = x_i$ and $y'_i = x'_i$ if $i \in S$, and $y_i = x'_i$ and $y'_i = x_i$ otherwise. Then, for all $a \in K_n$, $p_{x,a} + p_{x',a} = p_{y,a} + p_{y',a}$ in $\text{Mor}^+(K_n^k, K_n)$.

Proof. There exist $\mathbf{x}_3, \dots, \mathbf{x}_n \in K_n^k$ such that $x_{ij} \neq x_{i'j}$ for all $i \neq i'$, and $x_{ij} \neq x_j, x'_j$. Then, by Lemma 4.4.2, we have that $p_{x,a} + p_{x',a} + p_{x_3,a} + \dots + p_{x_n,a} = 1$ and $p_{y,a} + p_{y',a} + p_{x_3,a} + \dots + p_{x_n,a} = 1$. Putting these together gives the wanted result. ■

Lemma 4.4.4. For all $\mathbf{x}, \mathbf{y} \in K_n^k$ and $a \in K_n$, $p_{x,a}$ and $p_{y,a}$ commute in $\text{Mor}^+(K_n^k, K_n)$.

Proof. Consider first the case that $x_i \neq y_i$ for all i . Then, as $\mathbf{x} \sim_{K_n^k} \mathbf{y}$, $p_{x,a} p_{y,a} = 0$, and hence they commute by the defining relations of $\text{Mor}^+(K_n^k, K_n)$. In the other case, let $S = \{i \mid x_i = y_i\} \neq \emptyset$. Now, let $\mathbf{x}', \mathbf{y}' \in K_n^k$ such that $x'_i = y'_i \neq x_i$ for $i \in S$, and $x'_i = y_i$

and $y'_i = x_i$ for $i \notin S$. Hence $\mathbf{x}, \mathbf{x}', \mathbf{y}, \mathbf{y}'$ satisfy the conditions of [Lemma 4.4.3](#), and hence $p_{\mathbf{x},a} + p_{\mathbf{x}',a} = p_{\mathbf{y},a} + p_{\mathbf{y}',a}$. Hence, we find that

$$[p_{\mathbf{x},a}, p_{\mathbf{y},a}] = [p_{\mathbf{x},a}, p_{\mathbf{x},a} + p_{\mathbf{x}',a} - p_{\mathbf{y}',a}] = -[p_{\mathbf{x},a}, p_{\mathbf{y}',a}].$$

In particular, this tells us that $[p_{\mathbf{x},a}, p_{\mathbf{y},a}]$ does not depend on y_i for $i \notin S$. Let $\mathbf{z} \in K_n^k$ be such that $z_i = x_i = y_i$ for all $i \in S$, and $z_i \neq x_i$ and $z_i \neq y_i$ for all $i \notin S$. Therefore,

$$[p_{\mathbf{x},a}, p_{\mathbf{y},a}] = [p_{\mathbf{x},a}, p_{\mathbf{z},a}] = -[p_{\mathbf{z},a}, p_{\mathbf{x},a}] = -[p_{\mathbf{z},a}, p_{\mathbf{y},a}] = -[p_{\mathbf{x},a}, p_{\mathbf{y},a}],$$

giving that $[p_{\mathbf{x},a}, p_{\mathbf{y},a}] = 0$. ■

Lemma 4.4.5. Let $S \subseteq [k]$ and define $\Pi_S = \sum_{a \in K_n} p_{\mathbf{x},a} p_{\mathbf{y},a}$ for some $\mathbf{x}, \mathbf{y} \in K_n^k$ such that $x_i = y_i \iff i \in S$. Then,

- (i) Π_S does not depend on the choice of $\mathbf{x}, \mathbf{y}$ satisfying $x_i = y_i \iff i \in S$;
- (ii) Π_S is a projection;
- (iii) Π_S is central;
- (iv) $\Pi_S p_{\mathbf{x},a}$ does not depend on x_i for $i \notin S$; and
- (v) For $T \subseteq S$, $\Pi_S = \Pi_T + \Pi_{S \setminus T}$.

Proof. (i) Let $\tilde{\mathbf{x}}, \tilde{\mathbf{y}} \in K_n^k$ such that $\tilde{x}_i = \tilde{y}_i \iff i \in S$. Define $\mathbf{x}', \mathbf{y}' \in K_n^k$ such that $x'_i = y'_i \neq x_i, \tilde{x}_i$ for all $i \in S$, and $x'_i = y_i$ and $y'_i = x_i$ for all $i \notin S$. Then,

$$\Pi_S = \sum_{a \in K_n} p_{\mathbf{x},a} p_{\mathbf{y},a} = \sum_{a \in K_n} p_{\mathbf{x},a} (p_{\mathbf{x},a} + p_{\mathbf{x}',a} - p_{\mathbf{y}',a}) = 1 - \sum_{a \in K_n} p_{\mathbf{x},a} p_{\mathbf{y}',a}$$

Next, let $\mathbf{x}'', \mathbf{y}'' \in K_n^k$ such that $x''_i = y''_i$ and $y''_i = x_i$ for all $i \in S$, and $x''_i = y''_i \neq x_i, \tilde{x}_i$ for all $i \notin S$. Then,

$$\Pi_S = 1 - \sum_{a \in K_n} (p_{\mathbf{y}',a} + p_{\mathbf{y}'',a} - p_{\mathbf{x}'',a}) p_{\mathbf{y}',a} = \sum_{a \in K_n} p_{\mathbf{x}'',a} p_{\mathbf{y}',a}.$$

Then, let $\mathbf{x}''', \mathbf{y}''' \in K_n^k$ such that $x'''_i = y'''_i \neq y_i, \tilde{y}_i$ for all $i \in S$, and $x'''_i = y'_i$ and $y'''_i = x'_i$ for all $i \notin S$. Then,

$$\Pi_S = \sum_{a \in K_n} p_{\mathbf{x}'',a} (p_{\mathbf{x}'',a} + p_{\mathbf{x}''',a} - p_{\mathbf{y}''',a}) = 1 - \sum_{a \in K_n} p_{\mathbf{x}'',a} p_{\mathbf{y}''',a}.$$

By construction, $\mathbf{x}''$ is arbitrary such that $x_i'' \neq x_i, \tilde{x}_i$ for all i , and $\mathbf{y}'''$ is arbitrary such that $y_i''' \neq y_i, \tilde{y}_i$ for all i . As such, we can reverse the process to get

$$\Pi_S = 1 - \sum_{a \in K_n} p_{\mathbf{x}'',a} p_{\mathbf{y}''',a} = \sum_{a \in K_n} p_{\tilde{\mathbf{x}},a} p_{\tilde{\mathbf{y}},a},$$

giving the wanted independence.

(ii) As $p_{\mathbf{x},a}$ and $p_{\mathbf{y},a}$ commute by [Lemma 4.4.4](#), $p_{\mathbf{x},a} p_{\mathbf{y},a}$ is a projection. Further, if $a \neq a'$,

$$(p_{\mathbf{x},a} p_{\mathbf{y},a})(p_{\mathbf{x},a'} p_{\mathbf{y},a'}) = p_{\mathbf{y},a} p_{\mathbf{x},a} p_{\mathbf{x},a'} p_{\mathbf{y},a'} = 0,$$

so Π_S is a sum of orthogonal projections, and hence it is a projection.

(iii) We have that $p_{\mathbf{x},a} \Pi_S = p_{\mathbf{x},a} p_{\mathbf{y},a} = p_{\mathbf{y},a} p_{\mathbf{x},a} = \Pi_S p_{\mathbf{x},a}$.

(iv) Let $\mathbf{x}, \mathbf{x}' \in K_n^k$ such that $x_i = x'_i$ for all $i \in S$. Then, let $\mathbf{y} \in K_n^k$ such that $y_i = x_i$ for all $i \in S$ and $y_i \neq x_i, x'_i$ for all $i \notin S$. Then,

$$\Pi_S p_{\mathbf{x},a} = p_{\mathbf{x},a} p_{\mathbf{y},a} = \Pi_S p_{\mathbf{y},a} = \Pi_S p_{\mathbf{x}',a}.$$

(v) Let $\mathbf{x}, \mathbf{y} \in K_n^k$ such that $x_i = y_i \iff i \in T$. Define $\mathbf{y}', \mathbf{z}, \mathbf{z}' \in K_n^k$ such that $y'_i = z'_i \neq x_i$ and $z_i = x_i$ for $i \in T$, $y'_i = z_i = x_i$ and $z'_i = y_i$ for $i \in S \setminus T$, and $y'_i = z'_i \neq x_i, y_i$ and $z_i = y_i$ for $i \notin S$. By construction, [Lemma 4.4.3](#) implies that $p_{\mathbf{y},a} + p_{\mathbf{y}',a} = p_{\mathbf{z},a} + p_{\mathbf{z}',a}$, and

$$\Pi_T = \sum_{a \in K_n} p_{\mathbf{x},a} p_{\mathbf{y},a} = \sum_{a \in K_n} p_{\mathbf{x},a} (p_{\mathbf{z},a} + p_{\mathbf{z}',a} - p_{\mathbf{y}',a}) = \Pi_S - \Pi_{S \setminus T}. \quad \blacksquare$$

Proof of Theorem 4.4.1. Let $\Pi_i = \Pi_{\{i\}}$, as defined in [Lemma 4.4.5](#). Via [Lemma 4.4.5](#) (ii) and (iii) these are central projections; and via [Lemma 4.4.5](#) (v), $1 = \Pi_{[k]} = \Pi_1 + \Pi_{[k] \setminus \{1\}} = \dots = \Pi_1 + \dots + \Pi_k$, so we get a PVM as wanted. Now, fix $i \in [k]$, and let $\varphi : \text{End}^+(K_n) \rightarrow \Pi_i \text{Mor}^+(K_n^k, K_n)$ be the $*$ -homomorphism defined on the generators $\varphi(p_{\mathbf{x},a}) = \Pi_i p_{\mathbf{x},a}$ for $\mathbf{x}$ with $x_i = x$. First, due to [Lemma 4.4.5](#) (iv), this is well-defined on the generators. To see that it extends to a $*$ -homomorphism, it suffices to see that it satisfies the relations. First, since Π_i is a central projection, $\Pi_i p_{\mathbf{x},a}$ is a projection. Next,

$$\sum_a \varphi(p_{\mathbf{x},a}) = \sum_a \Pi_i p_{\mathbf{x},a} = \Pi_i = \varphi(1).$$

Finally, for all $x \neq y$, there exists some $\mathbf{x}, \mathbf{y}$ such that $x_i = x$, $y_i = y$, and $x_j \neq y_j$ for all j . Then

$$\varphi(p_{\mathbf{x},a}) \varphi(p_{\mathbf{y},a}) = \Pi_i p_{\mathbf{x},a} p_{\mathbf{y},a} = 0.$$

So φ is a $*$ -homomorphism.

For the other direction, let $\psi : \text{Mor}^+(K_n^k, K_n) \rightarrow \text{End}^+(K_n)$ be the $*$ -homomorphism defined on the generators as $\psi(p_{\mathbf{x},a}) = p_{x_i,a}$. To see that it is in fact a $*$ -homomorphism, we verify the relations: $p_{x_i,a}$ is a projection;

$$\sum_a \psi(p_{\mathbf{x},a}) = \sum_a p_{x_i,a} = 1;$$

and if $x_j \neq y_j$ for all j , then

$$\psi(p_{\mathbf{x},a})\psi(p_{\mathbf{y},a}) = p_{x_i,a}p_{y_i,a} = 0.$$

Also, since $\psi(\Pi_i) = \sum_a p_{x_i,a}p_{x_i,a} = 1$, ψ restricts to a $*$ -homomorphism $\bar{\psi} : \Pi_i \text{Mor}^+(K_n^k, K_n) \rightarrow \text{End}^+(K_n)$. It is clear that $\bar{\psi} = \varphi^{-1}$, giving the wanted isomorphism. ■

Definition 4.4.6. Let $f \in \text{mor}(A, B)$. Define the character $\hat{f} : \text{Mor}^+(A, B) \rightarrow \mathbb{C}$ on the generators $\hat{f}(p_{a,b}) = \delta_{f(a),b}$.

The map $f \mapsto \hat{f}$ induces the isomorphism $\mathcal{R}(\sigma) \cong \mathcal{R}^d(\sigma)$ of [Lemma 3.2.9](#), so it is well-defined. Recall that $\pi_i : A^k \rightarrow A$ is the projection onto the i -th component.

Corollary 4.4.7. Let π be an irreducible representation of $\text{Mor}^+(K_n^k, K_n)$. Then, there exists $i \in [k]$ and a representation π' of $\text{Aut}^+(K_n)$ such that $\pi = \pi' \circ \hat{\pi}_i$.

Proof. Since π is irreducible, there exists a unique $i \in [k]$ such that $\pi(\Pi_i) \neq 0$ — else π is reducible into representations on the $\text{supp } \pi(\Pi_i)$. Now, let $\pi' = \pi\varphi$, where φ is the embedding into the i -th component from the proof of [Theorem 4.4.1](#). To finish, we have that

$$\begin{aligned} (\pi' \circ \hat{\pi}_i)(p_{\mathbf{x},a}) &= (\hat{\pi}_i \otimes \pi')\Delta_{K_n^k, K_n}^{K_n}(p_{\mathbf{x},a}) = \sum_{b \in K_n} \hat{\pi}_i(p_{\mathbf{x},b}) \otimes \pi'(p_{b,a}) = \sum_{b \in K_n} \delta_{x_i,b} \pi\varphi(p_{ba}) \\ &= \pi\varphi(p_{x_i,a}) = \pi(\Pi_i p_{\mathbf{x},a}) = \pi(p_{\mathbf{x},a}). \end{aligned} \quad \blacksquare$$

4.5 Separations

In this section, we use the polymorphism characterisation of commutativity gadgets to show separations between some of the different types of gadget. In [Section 4.5.1](#), we show that there exist relational structures admitting no commutativity gadget. In [Section 4.5.2](#),

we show that complete graphs, which admit an oracular algebraic commutativity gadget, admit no non-oracular commutativity gadget. In [Section 4.5.3](#), we use constructions based on solution groups and complete graphs to separate q - and qa -commutativity gadgets, and conditionally separate qa - and co -commutativity gadgets. The results of [Sections 4.5.2](#) and [4.5.3](#) are summarised in [Fig. 4.4](#).

4.5.1 Relational structures with no commutativity gadget

In this section, we work out a technique for showing that graphs have nonclassical quantum endomorphisms, and use this to construct small examples of graphs admitting no commutativity gadget. To do so, we generalise a result of Schmidt [\[106\]](#), used to show that quantum automorphism groups of graphs are nonclassical.

Theorem 4.5.1 (Schmidt criterion [\[106\]](#)). Let G be a graph and let $\sigma, \tau \in \text{aut}(G) \setminus \{\text{id}\}$.

1. If σ and τ have disjoint supports, then $\text{aut}^+(G) \neq \text{aut}(G)$.
2. If σ and τ have disconnected supports, then $\text{aut}^{o+}(G) \neq \text{aut}(G)$.

We extend this by considering graph endomorphisms that are not permutations, which will give rise to more general quantum endomorphisms.

Definition 4.5.2. Let X be a finite set and let $f : X \rightarrow X$ be a function. The **support** of f is the set $\text{supp}(f) = \{x \in X \mid f(x) \neq x\}$.

This agrees with the notion of support for permutations, but not the notion for real- or complex-valued functions, where the support is the set of points that are not mapped to 0. The results of [\[106\]](#) hold generally for permutations, but extending them to non-invertible functions requires the use of an additional assumption.

Definition 4.5.3. Let $f, g \in \text{end}(G)$. We say f and g are **weakly adjacency congruent (WAC)** if, for all $x \in \text{supp}(f)$ and $y \in \text{supp}(g)$ such that $x \sim_G y$, it holds that $f(x) \sim_G g(y)$. This induces a relation on $\text{end}(G)$, denoted $f \approx_G g$.

The WAC relation is a weaker form of the adjacency relation on the exponential graph G^G , where $f \sim_{G^G} g$ if $x \sim_G y$ implies $f(x) \sim_G g(y)$. We work out some properties of the WAC relation.

Lemma 4.5.4. Let G be a graph.

- (i) The WAC relation is reflexive and symmetric.
- (ii) For all $f \in \text{end}(G)$, $\text{id}_G \approx_G f$.
- (iii) Let $f, g \in \text{end}(G)$ with disjoint supports. Suppose that for all $x \sim_G y$ such that $x, g(y) \in \text{supp}(f)$ and $y, f(x) \in \text{supp}(g)$, it holds that $f(x) \sim_G g(y)$. Then, $f \approx_G g$.
- (iv) If $f \approx_G g$, then $f^i \approx_G g^j$ for all $i, j \in \mathbb{N}$.

Proof. (i) Let $x, y \in \text{supp}(f)$. If $x \sim_G y$, then $f(x) \sim_G f(y)$, since f is an endomorphism. Hence $f \approx_G f$. Also, the definition of $f \approx_G g$ is symmetric, so it implies $g \approx_G f$.

(ii) $\text{id}_G \approx_G f$ vacuously, as the support of id_G is empty.

(iii) Suppose that f and g satisfy the condition in the statement. Let $x \sim_G y$ such that $x \in \text{supp}(f)$ and $y \in \text{supp}(g)$. We need to show that $f(x) \sim_G g(y)$. Consider first the case that $g(y)$ is a fixed point of f . As the supports of f and g are disjoint, we must have that x is a fixed point of g . As g is an endomorphism, this implies that $g(y) \sim_G g(x) = x$. Next, as f is an endomorphism, $g(y) = f(g(y)) \sim_G f(x)$. The case that $f(x)$ is a fixed point of g follows by symmetry. Finally, the case that $g(y) \in \text{supp}(f)$ and $f(x) \in \text{supp}(g)$ follows by hypothesis. Hence $f \approx_G g$.

(iv) Suppose, for a contradiction, that $f \approx_G g$ but $f^i \not\approx_G g^j$ for some (i, j) . Then, there exist $x \sim_G y$ such that $x \in \text{supp}(f)$, $y \in \text{supp}(g)$, and $f^i(x) \not\sim_G g^j(y)$. We may assume that (i, j) is minimal in the sense that $f^{i'} \approx_G g^{j'}$ for all (i', j') such that $(i' \leq i \wedge j' < j) \vee (i' < i \wedge j' \leq j)$. We may also assume that $f^{i'}(x) \in \text{supp}(f)$ for all $i' < i$, since if $f^{i'}(x)$ is a fixed point, then $f^{i'}(x) = f^{i'+1}(x) = \dots = f^i(x)$, so we may replace i with i' . In the same way, we may assume $g^{j'}(y) \in \text{supp}(g)$ for all $j' < j$. With this, since $f^{i-1} \approx_G g^{j-1}$, we have $f^{i-1}(x) \sim_G g^{j-1}(y)$. Since we also have $f^{i-1}(x) \in \text{supp}(f)$ and $g^{j-1}(y) \in \text{supp}(g)$, this implies that $f^i(x) \sim_G g^j(y)$, since f and g are WAC, giving the wanted contradiction. ■

Lemma 4.5.5. Suppose that $f, g \in \text{end}(G)$ with disjoint supports.

- (i) If f and g commute, then $f \approx_G g$.
- (ii) If f and g have disconnected supports, then $f \approx_G g$.

In particular, (i) implies that a pair of permutations with disjoint supports are WAC.

Proof. (i) Let $x \sim_G y$ such that $x \in \text{supp}(f)$ and $y \in \text{supp}(g)$. Then, as the supports of f and g are disjoint, x is a fixed point of g and y is a fixed point of f , giving $g(x) = x$ and $f(y) = y$. Also, $g \circ f = f \circ g$ is an endomorphism, so $g(f(x)) \sim_G g(f(y))$. Putting these together with the commutation assumption gives

$$f(x) = f(g(x)) = g(f(x)) \sim_G g(f(y)) = g(y).$$

(ii) In this case, if $x \sim_G y$, we cannot simultaneously have $x \in \text{supp}(f)$ and $y \in \text{supp}(g)$. Thus we vacuously get $f \approx_G g$. $\blacksquare$

Proposition 4.5.6 (Schmidt criterion for endomorphisms). Suppose that $f, g \in \text{end}(G)$ with nonempty supports.

- (i) If f and g have disjoint supports and are WAC, then $\text{end}^q(G) \not\subseteq \text{end}^c(G)$.
- (ii) If f and g have disconnected supports, then $\text{end}^{oq}(G) \not\subseteq \text{end}^c(G)$.

Proof. For (i), consider the C^* -algebra

$$C^*(\mathbb{Z}_2 * \mathbb{Z}_2) = \langle p_0, p_1, q_0, q_1 : p_i^2 = p_i^* = p_i, q_i^2 = q_i^* = q_i, p_0 + p_1 = q_0 + q_1 = 1 \rangle.$$

Let $\varphi : \text{End}^+(G) \rightarrow C^*(\mathbb{Z}_2 * \mathbb{Z}_2)$ be the $*$ -homomorphism defined on the generators as $\varphi(p_{x,y}) = \delta_{x,y}(p_0 + q_0 - 1) + \delta_{f(x),y}p_1 + \delta_{g(x),y}q_1$. We want to show that this is a well-defined $*$ -homomorphism by showing that the $\varphi(p_{x,y})$ satisfy the relations of $\text{End}^+(G)$. If x is a fixed point of f , the $\varphi(p_{x,y}) = \delta_{x,y}q_0 + \delta_{g(x),y}q_1$, which is a projection. On the other hand, if x is not a fixed point of f , then it must be a fixed point of g , so $\varphi(p_{x,y}) = \delta_{x,y}p_0 + \delta_{f(x),y}p_1$. Also,

$$\sum_{y \in G} \varphi(p_{x,y}) = (p_0 + q_0 - 1) + p_1 + q_1 = 1.$$

So $\{\varphi(p_{x,y})\}_{y \in G}$ is a PVM for each x . For the remaining relations, let $x \sim_G x'$ and $y \approx_G y'$. We need to show that $\varphi(p_{x,y})\varphi(p_{x',y'}) = 0$. Since f and g are endomorphisms, $f(x) \sim_G f(x')$ and $g(x) \sim_G g(x')$. First, suppose that x and x' are fixed points of f . Then,

$$\varphi(p_{x,y})\varphi(p_{x',y'}) = (\delta_{x,y}q_0 + \delta_{g(x),y}q_1)(\delta_{x',y'}q_0 + \delta_{g(x'),y'}q_1) = \delta_{x,y}\delta_{x',y'}q_0 + \delta_{g(x),y}\delta_{g(x'),y'}q_1 = 0,$$

as $x \sim_G x'$ and $g(x) \sim_G g(x')$ but $y \approx_G y'$. Next, if $x, x' \in \text{supp}(f)$, then they are fixed points of g , so we have by symmetry that $\varphi(p_{x,y})\varphi(p_{x',y'}) = 0$. Lastly, suppose that $x \in \text{supp}(f)$ and x' is a fixed point of f (the case that x is a fixed point of f and $x' \in \text{supp}(f)$ will follow by symmetry). Then,

$$\varphi(p_{x,y})\varphi(p_{x',y'}) = \delta_{x,y}\delta_{x',y'}p_0q_0 + \delta_{x,y}\delta_{g(x'),y'}p_0q_1 + \delta_{f(x),y}\delta_{x',y'}p_1q_0 + \delta_{f(x),y}\delta_{g(x'),y'}p_1q_1.$$

The first term is 0 as $x \sim_G x'$ but $y \not\sim_G y'$. The second term is 0 as $x = g(x) \sim_G g(x')$, and similarly for the third term as $f(x) \sim_G f(x') = x'$. Finally, for the final term, if x' is also a fixed point of g , then it is 0 for the same reason as the third term; else $x' \in \text{supp}(g)$, so by WAC property, $f(x) \sim_G g(x')$. Hence $\varphi(p_{x,y})\varphi(p_{x',y'}) = 0$, so φ is a $*$ -homomorphism.

Composing φ with a $*$ -homomorphism to a finite-dimensional C^* -algebra gives an element of $\text{end}^q(G)$. Since the supports of f and g are disjoint and nonempty, there exist distinct $x \in \text{supp}(f)$ and $y \in \text{supp}(g)$. Then, $\varphi(p_{x,f(x)}) = p_0$ and $\varphi(p_{y,g(y)}) = q_0$. By composing with the representation extending $p_0 \mapsto |0\rangle\langle 0|$ and $q_0 \mapsto |+\rangle\langle +|$, we get a non-classical quantum endomorphism.

For (ii), we want to show that φ defined in the same way on $\text{End}^{o+}(G)$ also extends to a $*$ -homomorphism. Due to Lemma 4.5.5(ii), f and g with disconnected supports are WAC, so all the relations but the commutativity relations hold by the proof of (i). To show these, let $x \sim_G x'$ and let $y, y' \in G$. First, if $x, x' \in \text{supp}(f)$, $\varphi(p_{x,y}), \varphi(p_{x',y'}) \in \text{span}\{p_0, p_1\}$ so they commute. Similarly, if x and x' are fixed points of f , the $\varphi(p_{x,y}), \varphi(p_{x',y'}) \in \text{span}\{q_0, q_1\}$, so they commute. Finally, in the case that $x \in \text{supp}(f)$ and x' is a fixed point of f , we know that x' must also be a fixed point of g , as it is adjacent to an element of $\text{supp}(f)$. As such, x' is a fixed point of both f and g , so $\varphi(p_{x',y'}) = \delta_{x',y'}$, which is central. Hence, $\varphi : \text{End}^{o+}(G) \rightarrow C^*(\mathbb{Z}_2 * \mathbb{Z}_2)$ is a $*$ -homomorphism, and we can find that $\text{end}^{oq}(G) \neq \text{end}^c(G)$ by the same argument as for (i). ■

Example 4.5.7. The **diamond graph** is the graph D defined by $V(D) = [4]$ and

$$E(D) = \{\{1, 2\}, \{2, 3\}, \{3, 4\}, \{4, 1\}, \{2, 4\}\}.$$

See also Fig. 4.3. Consider the maps $f, g : [4] \rightarrow [4]$ defined as

$$f = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 1 & 2 & 1 & 4 \end{pmatrix} \quad g = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 3 & 2 & 3 & 4 \end{pmatrix}.$$

It is easy to check that these are endomorphisms of D with disconnected supports. Hence by Proposition 4.5.6 and Theorem 4.3.1, D has no oracular or non-oracular commutativity gadget. This gives an example of a relational structure with alphabet size 4 with no commutativity gadgets.

Example 4.5.7 shows that there is a relational structure with NP-complete CSP on an alphabet of size 4 that does not admit an oracular q -commutativity gadget. Is this a minimal example? Due to [49], every NP-complete relational structure with alphabet size 2 has an oracular commutativity gadget. However, the situation in the intermediate case of alphabet size 3 is not clear. Since the only non-bipartite graph on 3 vertices is the




Figure 4.3: The diamond graph D

triangle K_3 , every graph with NP-complete CSP on 3 vertices has both a commutativity gadget and an oracular commutativity gadget — the other graphs on three vertices are bipartite, for which the classical and quantum CSPs are equivalent and in P [33]. Hence, Example 4.5.7 is a minimal example among graphs, but not necessarily among all relational structures.

4.5.2 Oracular and non-oracular gadgets

In this section, we show a strong separation between oracular and non-oracular commutativity gadgets. More precisely, we exhibit a relational structure with an oracular algebraic commutativity gadget but not even a q -commutativity gadget.

Proposition 4.5.8. For $n \geq 4$, the complete graph K_n does not admit a q -commutativity gadget.

By Theorem 4.3.14, this implies that K_n does not admit a Q -commutativity gadget or robust Q -commutativity gadget for any $Q \supseteq q$, or an algebraic commutativity gadget. However, due to Theorem 4.2.1, we know that K_n admits an oracular algebraic commutativity gadget.

Proof. By definition, $\text{aut}^+(K_n) = S_n^+$. It is a standard result in the theory of quantum groups that $C(S_n^+)$ is noncommutative for $n \geq 4$ [119]. Further, we can explicitly show that in fact $\text{End}^+(K_n)$ has a finite-dimensional noncommutative representation. It is clear that the following map extends to a noncommutative representation $\pi : \text{End}^+(K_n) \rightarrow B(\mathbb{C}^2)$: $\pi(p_{0,0}) = \pi(p_{1,1}) = |0\rangle\langle 0|$, $\pi(p_{0,1}) = \pi(p_{1,0}) = |1\rangle\langle 1|$, $\pi(p_{2,2}) = \pi(p_{3,3}) = |+\rangle\langle +|$, $\pi(p_{2,3}) = \pi(p_{3,2}) = |-\rangle\langle -|$, and $\pi(p_{i,j}) = \delta_{i,j}$ otherwise. As such $\text{end}^q(K_n) \not\subseteq \text{end}^c(K_n)$, so by Theorem 4.3.1, there is no q -commutativity gadget for K_n . ■




Figure 4.4: Separations between commutativity gadget classes; robust commutativity gadgets are not shown in this figure, due to the equivalences in [Theorem 4.3.14](#). The separation between oracular and non-oracular commutativity gadgets is due to [Proposition 4.5.8](#), and the remaining separations are due to [Theorem 4.5.9](#). The dashed line indicates a separation relying the existence of a non-hyperlinear group.

4.5.3 Separations based on solution groups

In this section, we show the following result, which relies on the full characterisation of the quantum polymorphisms of the complete graph studied in [Section 4.4](#).

Theorem 4.5.9. There exists a relational structure that admits a q -commutativity gadget but no qa -commutativity gadget.

If there exists a non-hyperlinear group, then there exists a relational structure that admits a qa -commutativity gadget but no co -commutativity gadget.

Here we make use of the solution group presentation of the entangled homomorphism algebras corresponding to linear relational structures, introduced in [\[42\]](#). Recall that LIN is the boolean structure over $\sigma_{\text{LIN}} = \{L_{b,n} \mid n \in \mathbb{N}, b \in \{0, 1\}\}$ with relations $L_{b,n}^{\text{LIN}} = \{\mathbf{a} \in \mathbb{Z}_2^n \mid a_1 + \dots + a_n = b\}$. We call a relational structure A **linear** if it is over σ_{LIN} .

The **solution group** of A is the finitely-presented group $\Gamma(A)$ with generators x_a for $a \in A$ and J subject to the relations $x_a^2 = 1$ for all $a \in A$, $J^2 = 1$, $x_a J = J x_a$ for all $a \in A$, and $x_{a_1} \dots x_{a_n} = J^b$ and $x_{a_i} x_{a_j} = x_{a_j} x_{a_i}$ if $\mathbf{a} \in L_{b,n}^A$ and $i, j \in [n]$. For linear relational structures $\text{Mor}^{o+}(A, \text{LIN}) \cong C^*\Gamma(A)/\langle\langle J+1 \rangle\rangle$, the full group C^* -algebra modulo the relation $J = -1$, via the $*$ -homomorphism $p_{ab} \mapsto \frac{1}{2}(1 + (-1)^b x_a)$.

A linear relational structure is **homogeneous** if $L_{1,n}^A = \emptyset$ for all n , and for a given linear relational structure A , its **homogenisation** is the homogeneous linear relational structure A_H with domain $\mathfrak{D}(A_H) = \mathfrak{D}(A)$ and relations $L_{0,n}^{A_H} = L_{0,n}^A \cup L_{1,n}^A$. The **homogeneous solution group** of A is the quotient group $\Gamma_H(A) = \Gamma(A)/\langle\langle J \rangle\rangle$. Then, $\Gamma(A_H) = \Gamma_H(A) \times \mathbb{Z}_2$, so if A is homogeneous $\text{Mor}^{o+}(A, \text{LIN}) \cong C^*\Gamma_H(A)$.

The separation relies on prior work of Slofstra that constructed solution groups with no noncommutative finite-dimensional or Connes-embeddable representations (*i.e.* representations with image embedding into $\mathcal{R}^\omega$).

Theorem 4.5.10 ([109, 110]). There exists a linear relational structure A such that $\Gamma(A)$ is noncommutative and hyperlinear, and $J \mapsto 1$ in every finite-dimensional representation.

If there exists a non-hyperlinear group, then there exists a linear relational structure A such that $\Gamma(A)$ is noncommutative and $J \mapsto 1$ in every Connes-embeddable representation.

We begin by extending these results to get solution groups with no non-trivial finite-dimensional or Connes-embeddable representations. A similar construction has been independently developed in [59].

Lemma 4.5.11. There exists a homogeneous linear relational structure A with two distinguished elements $a_1, a_2 \in A$ such that $\Gamma_H(A)$ is hyperlinear, $[x_{a_1}, x_{a_2}] \neq 1$, $[x_{a_1}, x_{a_2}]^2 = 1$, and $\Gamma_H(A)/\langle\langle [x_{a_1}, x_{a_2}] \rangle\rangle$ is trivial.

In particular, we have that the subgroup generated by x_{a_1} and x_{a_2} is isomorphic to the dihedral group D_4 . Note that when referring to group elements, the commutator $[\cdot, \cdot]$ represents the group commutator $[x, y] = xyx^{-1}y^{-1}$.

Proof. Let A be the linear relational structure defined as follows: $\mathfrak{D}(A)$ is the set of order-two elements of the alternating group A_7 , $L_{0,3}^A$ is the set of tuples (a, b, c) such that $abc = 1$ in A_7 ,¹ and $R^A = \emptyset$ for all the remaining relations. Now, due to [58, Remark 6.9], $\Gamma_H(A)$ is isomorphic to the unique perfect central extension of A_7 by $\mathbb{Z}_3$; thus it is finite and therefore hyperlinear. Also, the map $\phi : \Gamma_H(A) \rightarrow A_7$, $x_a \mapsto a$ is a surjective group homomorphism.

Next, let $a_1 = (12)(34)$ and $a_2 = (23)(56)$. In A_7 , $[a_1, a_2] = (14)(23)$ which is a nontrivial order-two element. By construction $x_{a_1}x_{a_2}x_{a_1}$ is an order-two element, whose image under the homomorphism $\phi : \Gamma_H(A) \rightarrow A_7$ is $a_1a_2a_1 = (13)(24)$. As such, there exists $z \in \ker \phi$ such that $x_{a_1}x_{a_2}x_{a_1} = x_{a_1a_2a_1}z$. But, squaring both sides gives that $1 = (x_{a_1}x_{a_2}x_{a_1})^2 = x_{a_1a_2a_1}^2z^2 = z^2$, as the kernel is central. But since $\ker \phi \cong \mathbb{Z}_3$, we must have $z = 1$ and thus $x_{a_1}x_{a_2}x_{a_1} = x_{a_1a_2a_1}$. Since $a = a_2$, $b = a_1a_2a_1$, and $c = a_1a_2a_1a_2$ are all order-two elements of A_7 satisfying $abc = 1$, we have that $x_ax_bx_c = 1$ in $\Gamma_H(A)$. This implies that

$$x_c = x_bx_a = x_{a_1}x_{a_2}x_{a_1}x_{a_2} = [x_{a_1}, x_{a_2}],$$

¹Note that the two conditions together imply that the subgroup $\langle a, b, c \rangle \subseteq A_7$ is abelian: if $abc = 1$ and a, b, c have order 2, then $ab = c^{-1} = c = (ab)^{-1} = b^{-1}a^{-1} = ba$.

which is therefore a non-identity order-two element of $\Gamma_H(A)$.

It remains to show that the normal subgroup generated by x_c is the whole of $\Gamma_H(A)$. Consider the subgroup $H = \langle\langle x_c \rangle\rangle \ker \phi$. Since $A_7 \cong \Gamma_H(A)/\ker \phi$ is simple, $H/\ker \phi$ is either trivial or isomorphic to A_7 . The first case would imply that $x_c \in \ker \phi$, which is impossible as it has order two. As such, the second case must hold, giving that $H = \Gamma_H(A)$. If $\langle\langle x_c \rangle\rangle \cap \ker \phi = \{1\}$, then $\Gamma_H(A)/\langle\langle x_c \rangle\rangle \cong \ker \phi$. Since $\ker \phi$ is abelian and $\Gamma_H(A)$ is perfect, this is a contradiction. This implies that $\ker \phi \subseteq \langle\langle x_c \rangle\rangle$, and hence $\langle\langle x_c \rangle\rangle = H = \Gamma_H(A)$. ■

Lemma 4.5.12. Let A be a linear relational structure. Then, there exists a homogeneous linear relational structure B such that $\Gamma(A) \cong \Gamma_H(B)$.

Proof. Define B as follows. Let $\mathfrak{D}(B) = \mathfrak{D}(A) \cup \{j\}$, for every $\mathbf{a} \in L_{0,n}^A$ put $\mathbf{a}$ in $L_{0,n}^B$, for every $\mathbf{a} \in L_{1,n}^A$ put $(a_1, \dots, a_n, j)$ in $L_{0,n+1}^B$, and for every $a \in A$ put (a, j, a, j) in $L_{0,4}^B$. Then it is clear that $x_a \mapsto x_a$ and $J \mapsto x_j$ induces an isomorphism $\Gamma(A) \rightarrow \Gamma_H(B)$. ■

Lemma 4.5.13. There exists a homogeneous linear relational structure A with distinguished elements $a \neq j \in A$ such that $\Gamma_H(A)$ is hyperlinear, $x_a \neq 1$ in $\Gamma_H(A)$, and $\Gamma_H(A)/\langle\langle x_j \rangle\rangle$ is trivial.

Proof. Let B be the linear relational structure corresponding to the Mermin-Peres magic square [88, 101], i.e. $\mathfrak{D}(B) = [9]$, and relations

$$L_{0,3}^B = \{(1, 2, 3), (4, 5, 6), (7, 8, 9), (1, 4, 7), (2, 5, 8)\}$$

and $L_{1,3}^B = \{(3, 6, 9)\}$. We know $\Gamma(B)$ is finite, and therefore hyperlinear, and that $[x_1, x_5] = [x_2, x_4] = J$. Let B' be the homogeneous linear relational structure such that $\Gamma(B) \cong \Gamma_H(B')$ constructed in Lemma 4.5.12. So, in $\Gamma_H(B')$, $[x_1, x_5] = [x_2, x_4] = x_j$. Now, let A_1 and A_2 be two distinct copies of the relational structure from Lemma 4.5.11 with distinguished elements $a_{11}, a_{12} \in A_1$ and $a_{21}, a_{22} \in A_2$. Take A to be the homogeneous linear relational structure with domain $\mathfrak{D}(A) = \mathfrak{D}(B') \sqcup \mathfrak{D}(A_1) \sqcup \mathfrak{D}(A_2)$ and relations $R^A = R^{B'} \cup R^{A_1} \cup R^{A_2}$, except for $R = L_{0,2}$ where $L_{0,2}^A = L_{0,2}^{B'} \cup L_{0,2}^{A_1} \cup L_{0,2}^{A_2} \cup \{(a_{11}, 1), (a_{12}, 5), (a_{21}, 2), (a_{22}, 4)\}$. Let the distinguished element $a = 1$.

By construction, $\Gamma_H(A)$ is the amalgamated free product of $\Gamma_H(B')$, $\Gamma_H(A_1)$, and $\Gamma_H(A_2)$ over the dihedral groups $D_4 \cong \langle x_1, x_5 \rangle \cong \langle x_{a_{11}}, x_{a_{12}} \rangle$ and $D_4 \cong \langle x_2, x_4 \rangle \cong \langle x_{a_{21}}, x_{a_{22}} \rangle$. Since we know that $\Gamma_H(B')$, $\Gamma_H(A_1)$, and $\Gamma_H(A_2)$ are hyperlinear, we have that $\Gamma_H(A)$ is hyperlinear. This also implies that $x_a = x_1 \neq 1$, and that $[x_{a_{11}}, x_{a_{12}}] = [x_1, x_5] = x_j$ and $[x_{a_{21}}, x_{a_{22}}] = [x_2, x_4] = x_j$ in $\Gamma_H(A)$. Now, we take the quotient by the relation $x_j = 1$. Since $[x_{a_{11}}, x_{a_{12}}] = x_j$, we take the quotient of $\Gamma_H(A_1)$ by $[x_{a_{11}}, x_{a_{12}}]$. By Lemma 4.5.11,

this implies that the image of $\Gamma_H(A_1)$ in $\Gamma_H(A)/\langle\langle x_j \rangle\rangle$ is trivial, and hence $x_{a_{11}} = x_{a_{12}} = 1$, which gives $x_1 = x_5 = 1$. Proceeding identically with $\Gamma_H(A_2)$, we also get $x_2 = x_4 = 1$. Now, the relations of $\Gamma(B)$ imply that $x_3 = x_1x_2 = 1$, $x_6 = x_4x_5 = 1$, $x_7 = x_1x_4 = 1$, $x_8 = x_2x_5 = 1$, and $x_9 = x_3x_6x_j = 1$. Therefore $\Gamma_H(A)/\langle\langle x_j \rangle\rangle$ is trivial. ■

Lemma 4.5.14. Let A be a linear relational structure. Then, there exists a linear relational structure A' such that there is a subset $S \subseteq \mathfrak{D}(A)$ such that $\{x_a \mid a \in S\}$ generates $\Gamma(A')$, for all $a \in S$, $x_a \neq 1$ or J in $\Gamma(A')$, and there is an embedding $\Gamma(A) \hookrightarrow \Gamma(A')$ such that $J \mapsto J$. Further, if $\Gamma(A)$ is hyperlinear, so is $\Gamma(A')$.

Proof. Let A' be the linear relational structure with domain $\mathfrak{D}(A') = \mathfrak{D}(A) \cup \mathfrak{D}(A) \times [2]$ and relations $R^{A'} = R^A$, except for $R = L_{0,3}$, where we take $L_{0,3}^{A'} = L_{0,3}^A \cup \{(a, (a, 1), (a, 2)) \mid a \in A\}$. This means that $\Gamma(A')$ is constructed from the $\Gamma(A)$ by splitting each x_a into the product $x_{(a,1)}x_{(a,2)}$, where we know $x_{(a,1)}$ and $x_{(a,2)}$ commute. Since $\Gamma(A) \subseteq \Gamma(A')$, this immediately gives a natural embedding. Next, we take $S = \mathfrak{D}(A) \times [2]$. It is clear that $\{x_a \mid a \in S\}$ generates $\Gamma(A')$; and for each a , $x_{(a,1)}$ is free except that it has order 2 and commutes with x_a , and hence we must have that neither $x_{(a,1)}$ nor $x_{(a,2)} = x_{(a,1)}x_a$ are trivial.

To finish, we need to show that if $\Gamma(A)$ is hyperlinear, so is $\Gamma(A')$. For any $a \in A$, adding $x_{(a,1)}$ and $x_{(a,2)}$ can be seen as taking the amalgamated free product of $\Gamma(A)$ with a power of $\mathbb{Z}_2$ corresponding to $\langle J, x_{(a,1)}, x_{(a,2)} \rangle$. If $x_a = 1$ or J in $\Gamma(A)$, then it corresponds to amalgamating with $\mathbb{Z}_2^2$ over $\mathbb{Z}_2 \cong \langle J \rangle$; else, it corresponds to amalgamating with $\mathbb{Z}_2^3$ over $\mathbb{Z}_2^2 \cong \langle J, x_a \rangle \cong \langle J, x_{(a,1)}x_{(a,2)} \rangle$. As amalgamated free products over finite abelian groups preserve hyperlinearity, we get the wanted result. ■

Lemma 4.5.15. Let A be a linear relational structure such that $J \neq 1$ in $\Gamma(A)$. Then, there exists a homogeneous linear relational structure B with a distinguished element $j \in B$ such that there exists an embedding $\Gamma(A) \hookrightarrow \Gamma_H(B)$ such that $J \mapsto x_j$, and $\Gamma_H(B)/\langle\langle x_j \rangle\rangle$ is trivial. Further, if $\Gamma(A)$ is hyperlinear, so is $\Gamma_H(B)$.

Proof. Let A' be the linear relational structure and let $S \subseteq \mathfrak{D}(A')$ be as constructed from A in Lemma 4.5.14. Let A'' be the homogeneous linear relational structure such that $\Gamma(A') \cong \Gamma_H(A'')$ constructed in Lemma 4.5.12. For each $a \in S$, let A_a be a copy of the linear relational structure from Lemma 4.5.13 with distinguished elements $\alpha_a \neq j_a \in A_a$. Now, define B as the homogeneous linear relational structure with domain $\mathfrak{D}(B) = \mathfrak{D}(A'') \sqcup \bigsqcup_{a \in S} \mathfrak{D}(A_a)$, and relations $R^B = R^{A''} \cup \bigcup_{a \in S} R^{A_a}$ except for $R = L_{0,2}$ where $L_{0,2}^B = L_{0,2}^{A''} \cup \bigcup_{a \in S} (L_{0,2}^{A_a} \cup \{(a, \alpha_a), (j, j_a)\})$.

Now, $\Gamma_H(B)$ is the amalgamated free product of $\Gamma_H(A'')$ and the $\Gamma_H(A_a)$ over the subgroups $\mathbb{Z}_2^2 \cong \langle x_a, x_j \rangle \cong \langle x_{\alpha_a}, x_{j_a} \rangle$, since $x_{\alpha_a} \neq 1$ and $x_a \neq 1$ for all $a \in S$ by

[Lemma 4.5.13](#) and [Lemma 4.5.14](#), respectively. By construction, if $\Gamma(A)$ is hyperlinear, so is $\Gamma_H(B)$. Next, we can take ι to be the homomorphism $\Gamma(A) \rightarrow \Gamma_H(B)$ extending $x_a \mapsto x_a$ for all $a \in A$, and $J \mapsto x_j$, which is an injection [\[77\]](#).

Next, consider the quotient of $\Gamma_H(B)$ by the relation $x_j = 1$. Since $x_{j_a} = 1$ in the image of $\Gamma_H(A_a)$, we have that that image is trivial by [Lemma 4.5.13](#). In particular, $x_{\alpha_a} = x_a = 1$ for all $a \in S$, so $\Gamma_H(B)$ is trivial. ■

The following two theorems construct the solution groups with no nontrivial finite dimensional or Connes-embeddable representations.

Theorem 4.5.16. There exists a homogeneous linear relational structure A such that $\Gamma_H(A)$ is noncommutative and hyperlinear but $\Gamma_H(A)$ is trivial under any finite-dimensional representation.

Proof. Let A_0 be the linear relational structure such that $\Gamma(A_0)$ is noncommutative and hyperlinear, and $J \mapsto 1$ in every finite-dimensional representation from [Theorem 4.5.10](#). Next, let A be the relational structure constructed from A_0 in [Lemma 4.5.15](#). Suppose π is a finite-dimensional representation of $\Gamma_H(A)$. Then, $\pi \circ \iota$ is a finite-dimensional representation of $\Gamma(A_0)$. Therefore, $\pi(x_j) = 1$. As such, π factors through to a representation of $\Gamma_H(A)/\langle\langle x_j \rangle\rangle$. But we know, this quotient is trivial, so π is trivial. ■

Theorem 4.5.17. Suppose there exists a non-hyperlinear group. Then, there exists a homogeneous linear relational structure A such that $\Gamma_H(A)$ is noncommutative but $\Gamma_H(A)$ is trivial under any Connes-embeddable representation.

Proof. We proceed in the same way as in [Theorem 4.5.16](#) starting from the relational structure A_0 such that $\Gamma(A_0)$ is noncommutative and $J \mapsto 1$ in every Connes-embeddable representation from [Theorem 4.5.10](#). ■

Before passing to the proof of the main theorem of this section, we establish the connection between the solution groups studied above and polymorphisms.

Definition 4.5.18. Let A be a relational structure. We define the **completion** of A to be the relational structure A_K with domain $\mathfrak{D}(A_K) = \mathfrak{D}(A)$ and relations $\mathfrak{R}(A_K) = \mathfrak{R}(A) \cup \{E\}$, where $E = \{(a, b) \in A^2 \mid a \neq b\}$ is the (arity-2) edge relation of the complete graph.

Lemma 4.5.19. Let A be a relational structure with $|A| \geq 3$. Then, the quantum polymorphism space of the completion is $\text{mor}^+(A_K^k, A_K) \cong \text{end}^+(A_K)^{\sqcup k}$.

Proof. Let $n = |A|$. We can see $\text{Mor}^+(A_K^k, A_K)$ as the quotient of $\text{Mor}^+(K_n^k, K_n)$ by the additional relations coming from A ; and similarly for $\text{End}^+(A_K)$ with regards to $\text{End}^+(K_n)$. Recall the isomorphisms implied by [Theorem 4.4.1](#): $\Psi : \text{Mor}^+(K_n^k, K_n) \rightarrow \text{End}^+(K_n)^{\oplus k}$ and $\Phi : \text{End}^+(K_n)^{\oplus k} \rightarrow \text{Mor}^+(K_n^k, K_n)$ defined on the generators as $\Psi(p_{x,a}) = p_{x_1,a} \oplus \dots \oplus p_{x_k,a}$ and $\Phi((p_{x,a})_i) = \Pi_i p_{x,a}$ where $(p_{x,a})_i = 0 \oplus \dots \oplus 0 \oplus p_{x,a} \oplus 0 \oplus \dots \oplus 0$, for $p_{x,a}$ in the i -th component, and $\mathbf{x}$ is such that $x_i = x$. Also, let $\iota_1 : \text{Mor}^+(K_n^k, K_n) \rightarrow \text{Mor}^+(A_K^k, A_K)$ and $\iota_2 : \text{End}^+(K_n)^{\oplus k} \rightarrow \text{End}^+(A_K)^{\oplus k}$ be the quotient maps. First, we show that $\iota_2 \Psi$ satisfies the additional relations of $\text{Mor}^+(A_K^k, A_K)$. In fact, let $R \in \mathfrak{R}(A)$, $\mathbf{x}_1, \dots, \mathbf{x}_k \in R$, and $\mathbf{a} \notin R$. Then,

$$\begin{aligned} & \iota_2 \Psi(p_{(x_{11}, \dots, x_{k1}), a_1}) \cdots \iota_2 \Psi(p_{(x_{1 \text{ ar}(R)}, \dots, x_{k \text{ ar}(R)}), a_{\text{ar}(R)}}) \\ &= (p_{x_{11}, a_1} \oplus \dots \oplus p_{x_{k1}, a_1}) \cdots (p_{x_{1 \text{ ar}(R)}, a_{\text{ar}(R)}} \oplus \dots \oplus p_{x_{k \text{ ar}(R)}, a_{\text{ar}(R)}}) \\ &= (p_{x_{11}, a_1} \cdots p_{x_{1 \text{ ar}(R)}, a_{\text{ar}(R)}}) \oplus \dots \oplus (p_{x_{k1}, a_1} \cdots p_{x_{k \text{ ar}(R)}, a_{\text{ar}(R)}}) = 0, \end{aligned}$$

as $p_{x_{i1}, a_1} \cdots p_{x_{i \text{ ar}(R)}, a_{\text{ar}(R)}} = 0$ in $\text{End}^+(A_K)$. Therefore, Ψ induces a $*$ -homomorphism $\bar{\Psi} : \text{Mor}^+(A_K^k, A_K) \rightarrow \text{End}^+(A_K)^{\otimes k}$.

For the other direction, we want to show that $\iota_1 \Phi$ respects the additional relations of $\text{End}^+(A_K)^{\oplus k}$. In fact, let $i \in [k]$, $R \in \mathfrak{R}(A)$, $\mathbf{x} \in R$, and $\mathbf{a} \notin R$. Then, letting $x_{ij} = x_i$ for all $i \in [k]$,

$$\begin{aligned} \iota_1 \Phi((p_{x_1, a_1})_i) \cdots \iota_1 \Phi((p_{x_{\text{ar}(R)}, a_{\text{ar}(R)}})_i) &= \Pi_i p_{x_1, a_1} \cdots \Pi_i p_{x_{\text{ar}(R)}, a_{\text{ar}(R)}} \\ &= \Pi_i (p_{x_1, a_1} \cdots p_{x_{\text{ar}(R)}, a_{\text{ar}(R)}}) = 0. \end{aligned}$$

Therefore, Φ induces a $*$ -homomorphism $\bar{\Phi} : \text{End}^+(A_K)^{\oplus k} \rightarrow \text{Mor}^+(A_K^k, A_K)$. To finish, note that $\bar{\Phi} = \bar{\Psi}^{-1}$, which follows since $\Phi = \Psi^{-1}$. $\blacksquare$

Next, we make use of a graph encoding of the homomorphism structure of a pair of relational structures, from [\[58\]](#).

Definition 4.5.20. Let A and B be σ -structures. $G(A, B)$ is the bipartite coloured graph with vertex set

$$\{(a, b) \mid a \in A, b \in B\} \cup \{(R, \mathbf{a}, \mathbf{b}) \mid R \in \sigma, \mathbf{a} \in R^A, \mathbf{b} \in R^B : a_i = a_j \Rightarrow b_i = b_j\},$$

edges $(a, b) \sim_{G(A, B)} (R, \mathbf{a}, \mathbf{b})$ iff there exists there exists $i \in [\text{ar}(R)]$ such that $a_i = a$ and $b_i = b$, and unary colour relations $C_a = \{(a, b) \mid b \in B\}$ for all $a \in A$ and $C_{(R, \mathbf{a})} = \{(R, \mathbf{a}, \mathbf{b}) \mid \mathbf{b} \in R^B\}$ for all $R \in \sigma$ and $\mathbf{a} \in R^A$.

For linear constraint systems, the completion of the graph encoding has quantum transformations given by the solution group.

Theorem 4.5.21. Let A be a linear relational structure. Then $\text{end}^+(G(A, \text{LIN})_K) \cong \text{mor}^{o+}(A_H, \text{LIN})$.

Recall that A_H is the homogenisation of A , where all relations are mapped to homogeneous relations.

Proof. First, note that, in $\text{End}^+(G(A, \text{LIN})_K)$, due to the unary relations on $G(A, \text{LIN})$, $p_{(a,b),(a',b')} = 0$ if $a \neq a'$, $p_{(R,a,b),(R',a',b')} = 0$ if $(R, \mathbf{a}) \neq (R', \mathbf{a}')$, and $p_{(a,b),(R,a,b)} = p_{(R,a,b),(a,b)} = 0$ always. Then, define $p_{(a,b),b'} := p_{(a,b),(a,b')}$ and $p_{(R,a,b),b'} := p_{(R,a,b),(R,a,b')}$. These generate $\text{End}^+(G(A, \text{LIN})_K)$ subject to the relations $\sum_{b' \in \mathbb{Z}_2} p_{(a,b),b'} = \sum_{b \in \mathbb{Z}_2} p_{(a,b),b'} = \sum_{b' \in R^{\text{LIN}}} p_{(R,a,b),b'} = \sum_{b \in R^{\text{LIN}}} p_{(R,a,b),b'} = 1$, and $p_{(a,b),b'} p_{(R,a,b),b'} = 0$ if there exists $i \in [\text{ar}(R)]$ such that $a_i = a$ and for all such i , $b_i = b$ but $b'_i \neq b'$. First, note that the first two relations are $p_{(a,0),0} + p_{(a,0),1} = p_{(a,1),0} + p_{(a,1),1} = p_{(a,0),0} + p_{(a,1),0} = p_{(a,0),1} + p_{(a,1),1} = 1$. This implies that $p_{(a,0),0} = p_{(a,1),1}$ and $p_{(a,0),1} = p_{(a,1),0}$. Next, we have that

$$p_{(R,a,b),b'} = p_{(a_i,b_i),b'_i} p_{(a_j,b_j),b'_j} p_{(R,a,b),b'} = p_{(a_j,b_j),b'_j} p_{(a_i,b_i),b'_i} p_{(R,a,b),b'},$$

so $[p_{(a_i,b_i),b'_i}, p_{(a_j,b_j),b'_j}] p_{(R,a,b),b'} = 0$. But, since LIN is boolean, this implies that for all $c, c' \in \mathbb{Z}_2$, $[p_{(a_i,b_i),c}, p_{(a_j,b_j),c'}] p_{(R,a,b),b'} = 0$, so we can sum over b' and get $[p_{(a_i,b_i),c}, p_{(a_j,b_j),c'}] = 0$. By the relations on the $p_{(a,b),b'}$, this gives in fact that $[p_{(a_i,b),c}, p_{(a_j,b'),c'}] = 0$ for all $b, b', c, c' \in \mathbb{Z}_2$. Then, we have that

$$p_{(R,a,b),b'} = p_{(a_1,b_1),b'_1} \cdots p_{(a_{\text{ar}(R)},b_{\text{ar}(R)}),b'_{\text{ar}(R)}} p_{(R,a,b),b'} = p_{(a_1,b_1),b'_1} \cdots p_{(a_{\text{ar}(R)},b_{\text{ar}(R)}),b'_{\text{ar}(R)}}.$$

In particular, this implies that $p_{(R,a,b),b'} = p_{(R,a,\mathbf{0}),b+b'}$.

Now, define the map $\varphi : \text{End}^+(G(A, \text{LIN})_K) \rightarrow \text{Mor}^{o+}(A_H, \text{LIN})$ on the generators by $\varphi(p_{(a,b),b'}) = p_{a,b+b'}$ and $\varphi(p_{(R,a,b),b'}) = p_{a_1,b_1+b'_1} \cdots p_{a_{\text{ar}(R)},b_{\text{ar}(R)}+b'_{\text{ar}(R)}}$. To see that this extends to a $*$ -homomorphism, it remains to see that the relations of $\text{Mor}^+(G(A, \text{LIN})_K)$ remain

satisfied. In fact, the $\varphi(p_{(a,b),b'})$ and $\varphi(p_{(R,a,b),b'})$ are projections;

$$\begin{aligned}\sum_{b' \in \mathbb{Z}_2} \varphi(p_{(a,b),b'}) &= \sum_{b' \in \mathbb{Z}_2} p_{a,b+b'} = 1, \\ \sum_{b \in \mathbb{Z}_2} \varphi(p_{(a,b),b'}) &= \sum_{b \in \mathbb{Z}_2} p_{a,b+b'} = 1, \\ \sum_{b' \in R^{\text{LIN}}} \varphi(p_{(R,a,b),b'}) &= \sum_{b' \in R^{\text{LIN}}} p_{a_1,b_1+b'_1} \cdots p_{a_{\text{ar}(R)},b_{\text{ar}(R)}+b'_{\text{ar}(R)}} = \sum_{b \in \mathbb{Z}_2^{\text{ar}(R)}} p_{a_1,b_1} \cdots p_{a_{\text{ar}(R)},b_{\text{ar}(R)}} = 1, \\ \sum_{b \in R^{\text{LIN}}} \varphi(p_{(R,a,b),b'}) &= \sum_{b \in R^{\text{LIN}}} p_{a_1,b_1+b'_1} \cdots p_{a_{\text{ar}(R)},b_{\text{ar}(R)}+b'_{\text{ar}(R)}} = \sum_{b \in \mathbb{Z}_2^{\text{ar}(R)}} p_{a_1,b_1} \cdots p_{a_{\text{ar}(R)},b_{\text{ar}(R)}} = 1;\end{aligned}$$

and for $\mathbf{a} \in R^A$, $\mathbf{b}, \mathbf{b}' \in R^{\text{LIN}}$, and $b' \neq b'_i$,

$$\begin{aligned}\varphi(p_{(a_i,b_i),b'})\varphi(p_{(R,a,b),b'}) &= p_{a_i,b_i+b'}p_{a_1,b_1+b'_1} \cdots p_{a_{\text{ar}(R)},b_{\text{ar}(R)}+b'_{\text{ar}(R)'}} \\ &= p_{a_i,b_i+b'}p_{a_i,b_i+b'_i}p_{a_1,b_1+b'_1} \cdots p_{a_{\text{ar}(R)},b_{\text{ar}(R)}+b'_{\text{ar}(R)'}} = 0.\end{aligned}$$

Next, define the map $\psi : \text{Mor}^+(A_H, \text{LIN}) \rightarrow \text{End}^+(G(A, \text{LIN})_K)$ on the generators by $\psi(p_{a,b}) = p_{(a,0),b}$. Again, to see that this extends to a $*$ -homomorphism, it suffices to see that the relations of $\text{Mor}^+(A_H, \text{LIN})$ remain satisfied. In fact, the $\psi(p_{a,b})$ are projectors;

$$\sum_{b \in \mathbb{Z}_2} \psi(p_{a,b}) = \sum_{b \in \mathbb{Z}_2} p_{(a,0),b} = 1;$$

and for $\mathbf{a} \in R^{A_H}$ and $\mathbf{b} \notin R^{\text{LIN}}$,

$$\psi(p_{a_1,b_1}) \cdots \psi(p_{a_{\text{ar}(R)},b_{\text{ar}(R)}}) = p_{(a_1,0),b_1} \cdots p_{(a_{\text{ar}(R)},0),b_{\text{ar}(R)}} = 0,$$

as $\sum_{b \in R^{\text{LIN}}} p_{(a_1,0),b_1} \cdots p_{(a_{\text{ar}(R)},0),b_{\text{ar}(R)}} = \sum_{b \in R^{\text{LIN}}} p_{(R,a,0),b} = 1$; and if $\mathbf{a} \in R^{A_H}$, we have that

$$[\psi(p_{a_i,b}), \psi(p_{a_j,b'})] = [p_{(a_i,0),b}, p_{(a_j,0),b'}] = 0.$$

Finally, to finish the proof, note that $\varphi = \psi^{-1}$ since they are inverses on the generators. ■

Now we can pass the the proof of the main theorem of this section.

Proof of Theorem 4.5.9. Let A be the homogeneous linear relational structure constructed in Theorem 4.5.16. Let $G = G(A, \text{LIN})_K$. Then, by Lemma 4.5.19 and Theorem 4.5.21 $\text{Mor}^+(G^k, G) \cong \text{End}^+(G)^{\oplus k} \cong \text{Mor}^+(A, \text{LIN})^{\oplus k} \cong C^*(\Gamma_H(A)^k)$. Now due to Theorem 4.5.16, every $*$ -representation of $\text{Mor}^+(G^k, G)$ into an element of q is trivial, but there is a non-commutative $*$ -representation into an element of qa .

For the second part, we proceed the same way using the homogeneous linear relational structure constructed in Theorem 4.5.17. ■

4.6 Complexity of deciding the existence of a commutativity gadget

In this section, we lower-bound the hardness of deciding if a relational structure admits certain types of commutativity gadget.

- Theorem 4.6.1.** (i) The problem of deciding if a given relational structure admits a q -commutativity gadget is undecidable.
- (ii) The problem of deciding if a given relational structure admits a qa -commutativity gadget is undecidable.
- (iii) The problem of deciding if a given relational structure admits a co -commutativity gadget is **coRE**-hard.
- (iv) The problem of deciding if a given relational structure admits an algebraic commutativity gadget is **coRE**-hard.

To show this, we make use of the hardness results for linear system games due to Slofstra [109, 110], as well as the reductions of Section 4.5.3.

Theorem 4.6.2 ([109, 110]). Given a linear relational structure A as input it is

- (i) Undecidable to decide if there exists a finite-dimensional representation of $\Gamma(A)$ where J is non-trivial.
- (ii) Undecidable to decide if there exists a Connes-embeddable representation of $\Gamma(A)$ where J is non-trivial.
- (iii) **coRE**-complete to decide if J is non-trivial in $\Gamma(A)$.

Proof of Theorem 4.6.1. For (i), we make use of Theorem 4.6.2(i). Let A be a linear relational structure. Then, in the same way as Theorem 4.5.16, there exists a homogeneous linear relational structure B such that $\Gamma(A)$ embeds into $\Gamma_H(B)$ and every finite-dimensional representation of $\Gamma_H(B)$ is trivial iff $J \mapsto 1$ in every finite-dimensional representation of $\Gamma(A)$. The forward implication follows from the fact that the construction of Lemma 4.5.15 allows any finite-dimensional representation of $\Gamma(A)$ where $J \not\mapsto 1$ to be lifted to a nontrivial finite-dimensional representation of $\Gamma_H(B)$, as it is based on solution groups with faithful finite-dimensional representations. Next, as in the proof of Theorem 4.5.9, there exists a

relational structure $G = G(B, \text{LIN})_K$ such that $\text{Mor}^+(G^k, G) \cong C^*(\Gamma_H(B)^k)$. Therefore, every finite-dimensional representation of $\text{Mor}^+(G^k, G)$ is abelian iff $J \mapsto 1$ in every finite-dimensional representation of $\Gamma(A)$. Due to [Theorem 4.3.1](#), G admits a q -commutativity gadget iff $J \mapsto 1$ in every finite-dimensional representation of $\Gamma(A)$. Due to [Theorem 4.6.2](#), it is undecidable to decide the latter, so it is undecidable to decide if a relational structure admits a q -commutativity gadget.

(ii) follows in the same way from [Theorem 4.6.2\(ii\)](#); and (iv) follows in the same way from [Theorem 4.6.2\(iii\)](#). Since relations in group algebras hold if and only if they hold for every *co*-representation, (iii) is equivalent to (iv). ■

Chapter 5

Reductions beyond commutativity gadgets

The reduction between CSPs based on minion homomorphisms of [Theorem 2.10.2](#) has been extended to gapless entangled CSPs in [\[38\]](#). In this chapter, we extend this further to capture the gapped case, and use this reduction to show hardness of entangled CSPs that do not admit commutativity gadgets. In [Section 5.1](#), we introduce minions, and define a dual notion called a cominion, for entangled polymorphisms. In [Section 5.2](#), we give a proof of the gapless reduction between CSPs implied by a minion homomorphism. In [Section 5.3](#), we prove a gapped reduction implied by a stable analogue of a minion homomorphism. In [Section 5.4](#), we show that cominion homomorphisms imply stable minion homomorphisms. In [Section 5.5](#), we use minion homomorphism-based reductions to show hardness for some CSPs which do not admit commutativity gadgets. Finally, in [Section 5.6](#), we use minion homomorphisms to show the equivalence of entangled CSPs in the oracular constraint-constraint and constraint-variable frameworks.

5.1 Entangled polymorphism minions

In this section, we show that entangled polymorphisms admit a minion structure.

Definition 5.1.1. Let A be a σ -structure, and let Q be a tensor monoid. The Q -polymorphism minion of A is the minion $\text{pol}^Q(A)$ with levels $\text{pol}^Q(A)_n = \text{mor}^Q(A^n, A)$ and minor operations defined via $\varphi_{/\pi}(p_{\mathbf{a},a}) = \varphi(p_{\mathbf{a} \circ \pi}, a)$ for $\pi : [n] \rightarrow [m]$ and $\varphi \in \text{pol}^Q(A)_n$.

The **oracular Q -polymorphism minion** of A is the minion $\text{pol}^Q(A)$ defined in the same way based on the morphism spaces $\text{mor}^{oQ}(A^n, A)$.

Lemma 5.1.2. $\text{pol}^Q(A)$ and $\text{pol}^{oQ}(A)$ are well-defined as minions.

The proof of this lemma requires that the minor operations send representations to representations. To do this in a clean way, we introduce a dual notion of a minion, in the spirit of [Section 3.2.2](#), and defer the proof to the end of this section.

Definition 5.1.3. A C^* -**cominion** $\mathcal{M}$ is a collection of C^* -algebras $\mathcal{M}_n$, called the **levels**, for $n \in \mathbb{N}$ such that for every function $\pi : [n] \rightarrow [m]$ there exists a $*$ -homomorphism $\Phi_\pi : \mathcal{M}_m \rightarrow \mathcal{M}_n$, called a **minor homomorphism**, satisfying $\Phi_{\text{id}} = \text{id}$, and for $\pi : [n] \rightarrow [m]$ and $\pi' : [m] \rightarrow [k]$, $\Phi_\pi \circ \Phi_{\pi'} = \Phi_{\pi' \circ \pi}$.

A C^* -cominion is a contravariant functor from a skeleton category of sets (with sets in bijection identified) to the category of C^* -algebras with morphisms given by the $*$ -homomorphisms. Relaxing the C^* -algebra condition to a semi-pre- C^* -algebra condition gives a **semi-pre- C^* -cominion**.

Definition 5.1.4. Given a C^* -cominion $\mathcal{M}$ and a tensor monoid Q , the Q -**minion** of $\mathcal{M}$ is the minion $\mathcal{M}^Q$ with levels

$$\mathcal{M}_n^Q = \{ \varphi : \mathcal{M}_n \rightarrow \mathcal{M} \mid \pi \text{ is a } * \text{-homomorphism and } \mathcal{M} \in Q \},$$

and minor operations $\varphi/\pi = \varphi \circ \Phi_\pi$.

Lemma 5.1.5. $\mathcal{M}^Q$ is well-defined as a minion.

Proof. First, for $\pi : [n] \rightarrow [m]$, $\Phi_\pi : \mathcal{M}_m \rightarrow \mathcal{M}_n$, and thus for $\varphi \in \mathcal{M}_n^Q$, $\varphi \circ \Phi_\pi \in \mathcal{M}_m^Q$. Next, since $\Phi_{\text{id}} = \text{id}$, $\varphi/\text{id} = \varphi \circ \text{id} = \varphi$. Finally, for $\pi : [n] \rightarrow [m]$ and $\pi' : [m] \rightarrow [k]$,

$$(\varphi/\pi)/\pi' = \varphi/\pi \circ \Phi_{\pi'} = \varphi \circ \Phi_\pi \circ \Phi_{\pi'} = \varphi \circ \Phi_{\pi' \circ \pi} = \varphi/\pi' \circ \pi. \quad \blacksquare$$

Before passing to the particular case of entangled polymorphisms, we show the following useful result about minion homomorphisms.

Lemma 5.1.6. Let $\mathcal{M}$ and $\mathcal{N}$ be C^* -cominions, and let Q be a tensor monoid. Suppose there are $*$ -homomorphisms $\Xi_n : \mathcal{N}_n \rightarrow \mathcal{M}_n$ for all $n \in \mathbb{N}$ that intertwine the Φ_π . Then, there exists a minion homomorphism $\xi : \mathcal{M}^Q \rightarrow \mathcal{N}^Q$.

We call the sequence of maps above a **cominion homomorphism** $\Xi : \mathcal{M} \rightarrow \mathcal{N}$ (note the change in the direction of the arrow).

Proof. For $\varphi \in \mathcal{M}_n^Q$, define $\xi_n(\varphi) = \varphi \circ \Xi_n$. Then, this induces a minion homomorphism since, for $\pi : [n] \rightarrow [m]$,

$$\xi_m(\varphi/\pi) = \varphi \circ \Phi_\pi \circ \Xi_m = \varphi \circ \Xi_n \circ \Phi_\pi = \xi_n(\varphi)/\pi. \quad \blacksquare$$

Definition 5.1.7. Let A be a σ -structure. The **polymorphism cominion** of A is the C^* -cominion $\text{Pol}^+(A)$ with levels $\text{Pol}^+(A)_n = \text{Mor}^+(A^n, A)$ and minor homomorphisms defined via $\Phi_\pi(p_{\mathbf{a},a}) = p_{\mathbf{a} \circ \pi, a}$ for $\pi : [n] \rightarrow [m]$, $\mathbf{a} \in A^m$, and $a \in A$.

The **oracular polymorphism cominion** of A is the C^* -cominion $\text{Pol}^{o+}(A)$ defined identically based on the C^* -algebras $\text{Mor}^{o+}(A^n, A)$.

Lemma 5.1.8. $\text{Pol}^+(A)$ and $\text{Pol}^{o+}(A)$ are well-defined as C^* -cominions.

Proof. First, we show the non-oracular case. For $\pi : [n] \rightarrow [m]$, we need to see that Φ_π is a $*$ -homomorphism, that is it preserves the relations of $\text{Mor}^+(A^m, A)$. First, it is direct to see that for any $\mathbf{a} \in A^m$, $\{\Phi_\pi(p_{\mathbf{a},a})\}_{a \in A} = \{p_{\mathbf{a} \circ \pi, a}\}_{a \in A}$ is a PVM in $\text{Mor}^+(A^n, A)$. Next, let $R \in \sigma$, and writing $k = \text{ar}(R)$, let $(\mathbf{a}_1, \dots, \mathbf{a}_k) \in R^{A^m}$ and $(a_1, \dots, a_k) \notin R^A$. Since inclusion in the relations of cartesian products of relational structures is decided componentwise, we have that $(\mathbf{a}_1 \circ \pi, \dots, \mathbf{a}_k \circ \pi) \in R^{A^n}$. As such, we get that $\Phi_\pi(p_{\mathbf{a}_1, a_1}) \cdots \Phi_\pi(p_{\mathbf{a}_k, a_k}) = p_{\mathbf{a}_1 \circ \pi, a_1} \cdots p_{\mathbf{a}_k \circ \pi, a_k} = 0$, so Φ_π is a $*$ -homomorphism. Next, it is easy to see that $\Phi_{\text{id}} = \text{id}$. To finish this case, note that

$$(\Phi_\pi \circ \Phi_{\pi'})(p_{\mathbf{a},a}) = \Phi_\pi(p_{\mathbf{a} \circ \pi', a}) = p_{\mathbf{a} \circ \pi' \circ \pi, a} = \Phi_{\pi' \circ \pi}(p_{\mathbf{a},a}),$$

which extends to the whole of $\text{Mor}^+(A^m, A)$ by the $*$ -homomorphism property.

To show the oracular case, it suffices to show that Φ_π also preserves the additional commutation relation of $\text{Mor}^{o+}(A^m, A)$. In fact, if $\mathbf{a}$ and $\mathbf{a}'$ appear in the same relation of A^m , then $\mathbf{a} \circ \pi$ and $\mathbf{a}' \circ \pi'$ appear in the same relation of A^n , so for any $a, a' \in A$,

$$[\Phi_\pi(p_{\mathbf{a},a}), \Phi_\pi(p_{\mathbf{a}',a'})] = [p_{\mathbf{a} \circ \pi, a}, p_{\mathbf{a}' \circ \pi, a'}] = 0. \quad \blacksquare$$

Proof of Lemma 5.1.2. It suffices to note that $\text{pol}^Q(A) = \text{Pol}^+(A)^Q$ and $\text{pol}^{oQ}(A) = \text{Pol}^{o+}(A)^Q$, and use Lemma 5.1.5. \blacksquare

5.2 Gapless reduction

In this section, we give a proof that minion homomorphisms between the associate polymorphism minions imply reductions between gapless CSPs.

Lemma 5.2.1. Let A be a σ -structure, and B and Y be τ -structures. Let X be the σ -structure constructed in [Theorem 2.10.2](#). Then, there exists a $*$ -homomorphism $\Phi : \text{Mor}^{o+}(X, A) \rightarrow \text{Mor}^{o+}(Y, B)$.

Proof. Recall that $X = X_0 / \sim$ where $X_0 = \bigsqcup_{y \in Y} A^{\mathfrak{D}(B)} \sqcup \bigsqcup_{R \in \tau, \mathbf{y} \in R^Y} A^{R^B}$ and the equivalence relation is given by $(y_i, \mathbf{a}) \sim (R, \mathbf{y}, \mathbf{a} \circ \pi_i)$. Now, define Φ via $\Phi(p_{[(y, \mathbf{a})], a}) = \sum_{b \in B: \mathbf{a}(b)=a} p_{y, b}$ and $\Phi(p_{[(R, \mathbf{y}, \mathbf{a})], a}) = \sum_{b \in R^B: \mathbf{a}(b)=a} p_{\mathbf{y}, b}$, where we recall that $p_{\mathbf{y}, b} = p_{y_1, b_1} \cdots p_{y_{\text{ar}(R)}, b_{\text{ar}(R)}}$. First, note that Φ is well-defined on the generators, since for $R \in \tau$, $\mathbf{y} \in R^Y$, and $\mathbf{a} \in A^{\mathfrak{D}(B)}$,

$$\Phi(p_{[(R, \mathbf{y}, \mathbf{a} \circ \pi_i)], a}) = \sum_{b \in R^B: \mathbf{a}(b_i)=a} p_{\mathbf{y}, b} = \sum_{b \in B^{\text{ar}(R)}: \mathbf{a}(b_i)=a} p_{\mathbf{y}, b} = \sum_{b \in B: \mathbf{a}(b)=a} p_{y_i, b} = \Phi(p_{[(y_i, \mathbf{a})], a}).$$

To extend Φ to a $*$ -homomorphism, it remains to show that it satisfies the relations of $\text{Mor}^{o+}(X, A)$. Since they are sums of orthogonal projectors, $\Phi(p_{[(y, \mathbf{a})], a})$ and $\Phi(p_{[(R, \mathbf{y}, \mathbf{a})], a})$ are projectors — note that for $\mathbf{y} \in R^Y$, $p_{\mathbf{y}, b}$ is a projector in $\text{Mor}^{o+}(Y, B)$ (but not necessarily in $\text{Mor}^+(Y, B)$). Next,

$$\begin{aligned} \sum_{a \in A} \Phi(p_{[(y, \mathbf{a})], a}) &= \sum_{b \in B} p_{y, b} = 1 \\ \sum_{a \in A} \Phi(p_{[(R, \mathbf{y}, \mathbf{a})], a}) &= \sum_{b \in R^B} p_{\mathbf{y}, b} = 1. \end{aligned}$$

Now, let $S \in \sigma$ and $\mathbf{x} \in S^X$, and let $\mathbf{a} \notin S^A$. Write $n = \text{ar}(S)$. Either $x_i = [(y, \mathbf{a}_i)]$ for $(\mathbf{a}_1, \dots, \mathbf{a}_n) \in S^{A^{\mathfrak{D}(B)}}$ or $x_i = [(R, \mathbf{y}, \mathbf{a}_i)]$ for $(\mathbf{a}_1, \dots, \mathbf{a}_n) \in S^{A^{R^B}}$, for all i . In the first case,

$$\Phi(p_{x_1, a_1}) \cdots \Phi(p_{x_n, a_n}) = \sum_{\substack{b_1, \dots, b_n \in B \\ \mathbf{a}_i(b_i)=a_i}} p_{y, b_1} \cdots p_{y, b_n} = \sum_{\substack{b \in B \\ \mathbf{a}_i(b)=a_i}} p_{y, b} = 0,$$

as the sum is empty, given that $\mathbf{a}_i(b) = a_i$ for all i would imply that $\mathbf{a} \in S^A$. In the second case,

$$\Phi(p_{x_1, a_1}) \cdots \Phi(p_{x_n, a_n}) = \sum_{\substack{b_1, \dots, b_n \in R^B \\ \mathbf{a}_i(b_i)=a_i}} p_{\mathbf{y}, b_1} \cdots p_{\mathbf{y}, b_n} = \sum_{\substack{b \in R^B \\ \mathbf{a}_i(b)=a_i}} p_{\mathbf{y}, b} = 0,$$

as the sum is empty as above, given that $\mathbf{a}_i(b) = a_i$ for all i would imply that $\mathbf{a} \in S^A$. For the final relations, let $S \in \sigma$, $\mathbf{x} \in S^X$, and $a, a' \in A$. Since $\Phi(p_{x_i, a})$ and $\Phi(p_{x_j, a'})$ are in the span of the same PVM in both of the cases (as considered above), they must commute. ■

Corollary 5.2.2 ([38]). Let A, B, X, Y, Φ be as above and let Q be a min-tensor monoid. If there exists a Q -morphism $\varphi \in \text{mor}^{oQ}(Y, B)$, then $\text{mor}^{oQ}(X, A)$ and $\text{mor}^Q(X, A)$ are nonempty.

Proof. We have that $\varphi \circ \Phi \in \text{mor}^{oQ}(X, A)$ and, taking $q : \text{Mor}^+(X, A) \rightarrow \text{Mor}^{o+}(X, A)$ to be the quotient map, we have that $\varphi \circ \Phi \circ q \in \text{mor}^Q(X, A)$. $\blacksquare$

Lemma 5.2.3 ([38]). Let A, B, X, Y, Q be as above. Supposing that there exists a minion homomorphism $\xi : \text{pol}^{oQ}(A) \rightarrow \text{pol}^{oQ}(B)$, if $\text{mor}^{oQ}(X, A)$ is nonempty, then $\text{mor}^{oQ}(Y, B)$ is nonempty.

Identically, supposing that there exists a minion homomorphism $\xi : \text{pol}^Q(A) \rightarrow \text{pol}^Q(B)$, if $\text{mor}^Q(X, A)$ is nonempty, then $\text{mor}^{oQ}(Y, B)$ is nonempty.

Proof. We show the oracular case. The non-oracular case is identical, as the proof does not depend significantly on the internal structure of the minions.

First, note that since X is constructed as a disjoint union of powers of A , we have for each $y \in Y$ a $*$ -homomorphism $\Phi_y : \text{Mor}^{o+}(A^{\mathfrak{D}(B)}, A) \rightarrow \text{Mor}^{o+}(X, A)$ given by $\Phi_y(p_{a,a}) = p_{[(y,a)],a}$; and for all $R \in \tau$, $\mathbf{y} \in R^Y$ a $*$ -homomorphism $\Phi_{R,\mathbf{y}} : \text{Mor}^{o+}(A^{R^B}, A) \rightarrow \text{Mor}^{o+}(X, A)$ given by $\Phi_{R,\mathbf{y}}(p_{a,a}) = p_{[(R,\mathbf{y},a)],a}$. Taking $\pi_i : R^B \rightarrow \mathfrak{D}(B)$ to be the projection onto the i -th component, the cominon structure of $\text{Pol}^{o+}(A)$ gives rise to a $*$ -homomorphism $\Phi_{\pi_i} : \text{Mor}^{o+}(A^{\mathfrak{D}(B)}, A) \rightarrow \text{Mor}^{o+}(A^{R^B}, A)$. It is easy to see that $\Phi_{R,\mathbf{y}} \circ \Phi_{\pi_i} = \Phi_{y_i}$, since

$$(\Phi_{R,\mathbf{y}} \circ \Phi_{\pi_i})(p_{a,a}) = p_{[(R,\mathbf{y},a \circ \pi_i)],a} = p_{[(y_i,a)],a} = \Phi_{y_i}(p_{a,a}).$$

Now, let $\varphi : \text{Mor}^{o+}(X, A) \rightarrow \mathcal{M} \in Q$ be a $*$ -homomorphism. We have Q -polymorphisms $\varphi \circ \Phi_y \in \text{mor}^{oQ}(A^{\mathfrak{D}(B)}, A)$ and $\varphi \circ \Phi_{R,\mathbf{y}} \in \text{mor}^{oQ}(A^{R^B}, A)$. The minion homomorphism can act on these polymorphisms to give $\xi(\varphi \circ \Phi_y) \in \text{mor}^{oQ}(B^{\mathfrak{D}(B)}, B)$ and $\xi(\varphi \circ \Phi_{R,\mathbf{y}}) \in \text{mor}^{oQ}(B^{R^B}, B)$. Note that these must all be homomorphisms into the same space. By taking $N \geq |R^B|, |\mathfrak{D}(B)|$, there exist surjective maps $\pi : [N] \rightarrow \mathfrak{D}(B)$ and $\pi_R : [N] \rightarrow R^B$, and thus using the cominon structure $\xi|_{\mathfrak{D}(B)} = \xi_N \circ \Phi_\pi$ and $\xi|_{R^B} = \xi_N \circ \Phi_{\pi_R}$. Hence, they have the same codomain, which we denote $\mathcal{N} \in Q$. Now, we define the $*$ -homomorphism $\varphi' : \text{Mor}^{o+}(Y, B) \rightarrow \mathcal{N}$ on the generators as $\varphi'(p_{y,b}) = \xi(\varphi \circ \Phi_y)(p_{\text{id}_B,b})$, where $\text{id}_B \in B^{\mathfrak{D}(B)}$ is the identity map. To see that this gives an element of $\text{mor}^{oQ}(Y, B)$, it suffices to see that it satisfies the relations of $\text{Mor}^{o+}(Y, B)$. First, since $\xi(\varphi \circ \Phi_y)$ is a $*$ -homomorphism and $\{p_{\text{id}_B,b}\}_{b \in B}$ is a PVM in $\text{Mor}^{o+}(B^{\mathfrak{D}(B)}, B)$, $\{\varphi'(p_{y,b})\}_{b \in B}$ is a PVM in $\mathcal{N}$. Next, let $R \in \tau$

and $\mathbf{y} \in R^Y$, and write $n = \text{ar}(R)$. If $\mathbf{b} \notin R^B$, then

$$\begin{aligned} \varphi'(p_{y_1, b_1}) \cdots \varphi'(p_{y_n, b_n}) &= \xi(\varphi \circ \Phi_{y_1})(p_{\text{id}_B, b_1}) \cdots \xi(\varphi \circ \Phi_{y_n})(p_{\text{id}_B, b_n}) \\ &= \xi(\varphi \circ \Phi_{R, \mathbf{y}})(\Phi_{\pi_1}(p_{\text{id}_B, b_1})) \cdots \xi(\varphi \circ \Phi_{R, \mathbf{y}})(\Phi_{\pi_n}(p_{\text{id}_B, b_n})) \\ &= \xi(\varphi \circ \Phi_{R, \mathbf{y}})(p_{\pi_1, b_1}) \cdots \xi(\varphi \circ \Phi_{R, \mathbf{y}})(p_{\pi_n, b_n}) = 0 \end{aligned}$$

as $\Phi_{R, \mathbf{y}}$ is a $*$ -homomorphism and $(\pi_1, \dots, \pi_n) \in R^{B^{R^B}}$. Finally, for $i, j \in [n]$ and $b, b' \in B$, we find in the same way that

$$\begin{aligned} [\varphi'(p_{y_i, b}), \varphi'(p_{y_j, b'})] &= [\xi(\varphi \circ \Phi_{y_i})(p_{\text{id}_B, b}), \xi(\varphi \circ \Phi_{y_j})(p_{\text{id}_B, b'})] \\ &= [\xi(\varphi \circ \Phi_{R, \mathbf{y}})(p_{\pi_i, b}), \xi(\varphi \circ \Phi_{R, \mathbf{y}})(p_{\pi_j, b'})] = 0. \end{aligned} \quad \blacksquare$$

Theorem 5.2.4. Let A be a σ -structure and B be a τ -structure, and let Q be a tensor monoid. If there exists a minion homomorphism $\text{pol}^{oQ}(A) \rightarrow \text{pol}^{oQ}(B)$, then there exist polynomial-time reductions from $\text{CSP}_{ocv}^Q(B)_{1,1}$ to $\text{CSP}_{ocv}^Q(A)_{1,1}$ and from ${}^e\text{CSP}_{ocv}^Q(B)_{1,1}$ to ${}^e\text{CSP}_{ocv}^Q(A)_{1,1}$.

If there exists a minion homomorphism $\text{pol}^Q(A) \rightarrow \text{pol}^Q(B)$, then there exist polynomial-time reductions from $\text{CSP}_{ocv}^Q(B)_{1,1}$ to $\text{CSP}_{no}^Q(A)_{1,1}$ and from ${}^e\text{CSP}_{ocv}^Q(B)_{1,1}$ to ${}^e\text{CSP}_{no}^Q(A)_{1,1}$.

Proof. We show the reduction from ${}^e\text{CSP}_{ocv}^Q(B)_{1,1}$ to ${}^e\text{CSP}_{ocv}^Q(A)_{1,1}$. The other cases are analogous. Let X be an instance of ${}^e\text{CSP}_{ocv}^Q(B)_{1,1}$. Let Y be the σ -structure constructed from X in [Theorem 2.10.2](#). It is clear that this instance can be constructed in polynomial time (in fact, in logspace [\[38\]](#)). Then, using [Corollary 5.2.2](#) and [Lemma 5.2.3](#) there exists $\varphi \in \text{mor}^{oQ}(X, B)$ if and only if there exists $\varphi' \in \text{mor}^{oQ}(Y, A)$. $\blacksquare$

With an additional technical condition, we can extend the proof to succinctly-presented entangled CSPs.

Definition 5.2.5. Let $R \subseteq D^n$ be a relation and $i, j \in [n]$ be distinct. Write $\Sigma_{R, i, j}$ for the graph with vertex set D and edges $a \sim b$ if there exist $\mathbf{a}, \mathbf{b} \in R$ such that $a_i = a$, $b_i = b$, and $a_j = b_j$.

A relational structure A is **well-connected** if, for every relation $R \in \mathfrak{R}(A)$ and $i, j \in [\text{ar}(R)]$ distinct, either $\Sigma_{R, i, j}$ or $\Sigma_{R, j, i}$ is connected.

It is easy to see that 3SAT, 3LIN, and K_n for $n \geq 3$ are well-connected.

Theorem 5.2.6. Let A be a σ -structure and B be a τ -structure that is well-connected, and let Q be a tensor monoid. If there exists a minion homomorphism $\text{pol}^{oQ}(A) \rightarrow \text{pol}^{oQ}(B)$, then there exist polynomial-time reductions from $\text{MIP}_{ocv}^Q(B)_{1,1}$ to $\text{MIP}_{ocv}^Q(A)_{1,1}$ and from ${}^e\text{MIP}_{ocv}^Q(B)_{1,1}$ to ${}^e\text{MIP}_{ocv}^Q(A)_{1,1}$.

If there exists a minion homomorphism $\text{pol}^Q(A) \rightarrow \text{pol}^{oQ}(B)$, then there exist polynomial-time reductions from $\text{MIP}_{ocv}^Q(B)_{1,1}$ to $\text{MIP}_{no}^Q(A)_{1,1}$ and from ${}^e\text{MIP}_{ocv}^Q(B)_{1,1}$ to ${}^e\text{MIP}_{no}^Q(A)_{1,1}$.

As noted above, 3SAT is well-connected, so a minion homomorphism to 3SAT is sufficient to show RE-hardness of the gapless problem. Also, note that [Theorem 5.2.4](#) gives exponential-time reductions directly, without the assumption of being well-connected.

Proof. We show the reduction from ${}^e\text{MIP}_{ocv}^Q(B)_{1,1}$ to ${}^e\text{MIP}_{ocv}^Q(A)_{1,1}$. The other cases are analogous. Let M be an instance of ${}^e\text{MIP}_{ocv}^Q(B)_{1,1}$. Let M' be a Turing machine that samples the constraints of the σ -structure constructed from $X(M)$ in [Theorem 2.10.2](#). This Turing machine can be made efficient by first sampling a constraint $(R, \mathbf{y})$ using M , and then, with probability $\frac{1}{2}$, sampling a uniformly random constraint $(S, \mathbf{x})$ from A^{R^B} and outputting $(S, [(R, \mathbf{y}, x_1)], \dots, [(R, \mathbf{y}, x_{\text{ar}(S)})])$; and, with probability $\frac{1}{2}$, sampling a uniformly random $i \in \text{ar}(R)$ and a uniformly random constraint $(S, \mathbf{x})$ from $A^{\mathfrak{D}(B)}$, and outputting $(S, [(y_i, x_1)], \dots, [(y_i, x_{\text{ar}(S)})])$. To not give the players information about the sampling process, we need to choose the right representative of the equivalence classes to send: if $(R, \mathbf{y}, \mathbf{a}) = (R, \mathbf{y}, \mathbf{a}' \circ \pi'_i)$, we send $(y_i, \mathbf{a}')$ instead. Due to the hypothesis that B is well-connected, this choice is unique unless $\mathbf{a}$ is constant, in which case we just send one of the entries (without loss of generality, we may assume the connectivity graph / Gaifman graph of $X(M)$ is connected, which implies $(y, \mathbf{a}') \sim (y', \mathbf{a}')$ for all y, y'). Then, using [Corollary 5.2.2](#) and [Lemma 5.2.3](#) there exists $\varphi \in \text{mor}^{oQ}(X(M), B)$ if and only if there exists $\varphi' \in \text{mor}^{oQ}(X(M'), A)$. ■

5.3 Gap-preserving reduction

In this section, we extend the results of the previous section to gapped CSPs. In [Section 5.3.1](#), we define a weighted-algebra version of a cominon and study the example of this structure coming from polymorphisms. In [Section 5.3.2](#), we use weighted cominions to study a stable notion of minion homomorphism. In [Section 5.3.3](#), we show that stable minion homomorphisms of the entangled polymorphism minions imply reductions between gapped CSPs.

5.3.1 Weighted cominions

Definition 5.3.1. A **weighted cominion** $\mathcal{A}$ is a sequence of weighted algebras $(\mathcal{A}_n, \mu_n)$ such that $\{\mathcal{A}_n\}_n$ is a semi-pre- C^* -cominion (with respect to the sum-of-squares cone on the semi-pre- C^* -algebras).

The examples of weighted cominions we consider arise from the weighted algebras of entangled polymorphisms.

Definition 5.3.2. Let A be a σ -structure. The **weighted polymorphism cominion** of A is $\text{Pol}^{no}(A) = \{\text{Mor}_{\mathbf{u}}^{no}(A^n, A)\}_n$ with minor homomorphisms $\Phi_\pi : \text{Mor}^{no}(A^m, A) \rightarrow \text{Mor}^{no}(A^n, A)$ for $\pi : [n] \rightarrow [m]$ given on the generators as $\Phi_\pi(p_{a,a}) = p_{a \circ \pi, a}$.

The **oracular weighted polymorphism cominion** of A is $\text{Pol}^{ocv}(A) = \{\text{Mor}_{\mathbf{u}}^{ocv}(A^n, A)\}_n$ with minor homomorphisms $\Phi_\pi : \text{Mor}^{ocv}(A^m, A) \rightarrow \text{Mor}^{ocv}(A^n, A)$ for $\pi : [n] \rightarrow [m]$ given on the generators as $\Phi_\pi(p_{a,a}) = p_{a \circ \pi, a}$ and $\Phi_\pi(\Phi_{R, \mathfrak{a}, a}) = \Phi_{R, (a_1 \circ \pi, \dots, a_{\text{ar}(R)} \circ \pi), a}$.

Lemma 5.3.3. $\text{Pol}^{ocv}(A)$ and $\text{Pol}^{no}(A)$ are well-defined as weighted cominions.

Proof. We show the non-oracular case; the oracular case is essentially identical. Let $\pi : [n] \rightarrow [m]$. Then, for every $\mathbf{a} \in A^m$, $\{\Phi_\pi(p_{a,a})\}_{a \in A}$ is a PVM, so Φ_π extends to a $*$ -homomorphism. In an identical way to [Lemma 5.1.8](#), we see that $\text{Pol}^{no}(A)$ is a cominion. $\blacksquare$

Weighted cominions can be used to study approximate representations of an associated C^* -cominion.

Definition 5.3.4. We say that a weighted cominion $\mathcal{A} = \{(\mathcal{A}_n, \mu_n)\}_n$ **approximates** a C^* -cominion $\mathcal{M}$ if $C_u^*(\mathcal{A}_n / \langle\langle \text{supp } \mu_n \rangle\rangle) \cong \mathcal{M}_n$ and the quotient maps $q_n : \mathcal{A}_n \rightarrow \mathcal{M}_n$ induced by the isomorphisms induce a cominion homomorphism $\mathcal{M} \rightarrow \{\mathcal{A}_n\}_n$.

We apply this definition in the context of polymorphism minions.

Lemma 5.3.5. $\text{Pol}^{ocv}(A)$ and $\text{Pol}^{no}(A)$ approximate $\text{Pol}^{o+}(A)$ and $\text{Pol}^+(A)$, respectively.

Proof. First, we show the non-oracular case. Due to [Lemma 3.2.2](#), $C_u^*(\text{Mor}^{no}(A^n, A) / \langle\langle \text{supp } \mu_{no, \mathbf{u}} \rangle\rangle) \cong \text{Mor}^+(A^n, A)$ for all n . Also, we see that $(q_n \circ \Phi_\pi)(p_{a,a}) = p_{a \circ \pi, a} = (\Phi_\pi \circ q_n)(p_{a,a})$, so the quotient maps induce a cominion homomorphism.

Now, we show the oracular case. Again due to [Lemma 3.2.2](#), $C_u^*(\text{Mor}^{ocv}(A^n, A) / \langle\langle \text{supp } \mu_{ocv, \mathbf{u}} \rangle\rangle) \cong \text{Mor}^{o+}(A^n, A)$ for all n . To finish, the quotient maps induce a cominion homomorphism as $(q_n \circ \Phi_\pi)(p_{a,a}) = p_{a \circ \pi, a} = (\Phi_\pi \circ q_n)(p_{a,a})$ and $(q_n \circ \Phi_\pi)(\Phi_{R, \mathfrak{a}, a}) = p_{a_1 \circ \pi, a_1} \cdots p_{a_{\text{ar}(R)} \circ \pi, a_{\text{ar}(R)}} = (\Phi_\pi \circ q_n)(\Phi_{R, \mathfrak{a}, a})$. $\blacksquare$

5.3.2 Stable minion homomorphisms

Now, we study morphisms of weighted cominions. Recall that a tensor monoid is finite if it consists of finite von Neumann algebras.

Definition 5.3.6. We say a finite tensor monoid is **measured** if every element is associated with a finite-dimensional tracial state, which are compatible in the sense the state of $\mathcal{M} \otimes \mathcal{N}$ is the tensor of the states of $\mathcal{M}$ and $\mathcal{N}$.

A finite tensor monoid consisting of factors is always measured. The tensor monoids d , c , q , qa , and co are all finite, and may be treated as measured. Note that for these tensor monoids, the choice of tracial state on the factors is unique, but there are multiple choices of tracial state on the remaining algebras — however, for our purposes, the choice is immaterial.

Definition 5.3.7. Let $\mathcal{A} = \{(\mathcal{A}_n, \mu_n)\}_n$ and $\mathcal{B} = \{(\mathcal{B}_n, \nu_n)\}_n$ be weighted cominions, and let Q be a finite measured tensor monoid. Fix a generating set S_n for $\mathcal{B}_n$. A **stable Q -minion homomorphism** $\xi : \mathcal{A} \rightarrow \mathcal{B}$ is a sequence of maps ξ_n on the $*$ -representations on $\mathcal{A}_n$ over Q to the $*$ -representations of $\mathcal{B}_n$ over Q such that

- (i) for all $\varepsilon > 0$ and $\pi : [n] \rightarrow [m]$, there exists $\delta > 0$ such that, for any $*$ -representation $\varphi : \mathcal{A}_n \rightarrow (\mathcal{M}, \text{tr}) \in Q$, if $\text{def}(\text{tr} \circ \varphi) \leq \delta$ and $\text{def}(\text{tr} \circ \varphi \circ \Phi_\pi) \leq \delta$, then $\|\xi_m(\varphi \circ \Phi_\pi)(x) - (\xi_n(\varphi) \circ \Phi_\pi)(x)\|_{\text{tr}'} \leq \varepsilon$ for all $x \in S_m$, where $\xi_n(\varphi) : \mathcal{B}_n \rightarrow (\mathcal{N}, \text{tr}')$; and
- (ii) for all $\varepsilon > 0$ and $n \in \mathbb{N}$, there exists $\delta > 0$ such that, for any $*$ -representation $\varphi : \mathcal{A}_n \rightarrow (\mathcal{M}, \text{tr}) \in Q$, if $\text{def}(\text{tr} \circ \varphi) \leq \delta$, then $\text{def}(\text{tr}' \circ \xi_n(\varphi)) \leq \varepsilon$, where $\xi_n(\varphi) : \mathcal{B}_n \rightarrow (\mathcal{N}, \text{tr}')$.

It is implicit in (i) that $\xi_m(\varphi \circ \Phi_\pi)$ and $\xi_n(\varphi)$ have the same codomain.

Lemma 5.3.8. If $\mathcal{B}_n$ is finitely-generated for all n , the stable Q -minion homomorphism property of ξ is independent of the choice of finite generating set of $\mathcal{B}_n$.

Proof. Suppose $S_m = \{x_1, \dots, x_K\}$ and $S'_m = \{y_1, \dots, y_L\}$ are two finite generating sets of $\mathcal{B}_m$. We aim to show that, if ξ is a stable Q -minion homomorphism with respect to S_m , it is also a stable Q -minion homomorphism with respect to S'_m . Let $\varepsilon > 0$, $\pi : [n] \rightarrow [m]$, and $\delta > 0$ to be determined later. Fix a $*$ -homomorphism $\varphi : \mathcal{A}_n \rightarrow (\mathcal{M}, \text{tr}) \in Q$ such that $\text{def}(\text{tr} \circ \varphi) \leq \delta$. Note that there exist noncommutative $*$ -polynomials f_i such that $y_i = f_i(x_1, \dots, x_K)$. We expand each f_i as a linear combination of monomials $f_i = \sum_M c_{i,M} M$

for some $c_{i,M} \in \mathbb{C}$. Let $N_i = \sum_M |c_{i,M}|$ and $N = \max_i N_i$, let D be the maximum degree of the f_i , and let $C = \max\{1, \max_{x \in S_m} \|x\|\}$. Then, writing $\lambda = \xi_m(\varphi \circ \Phi_\pi) - \xi_n(\varphi) \circ \Phi_\pi$

$$\|\xi_m(\varphi \circ \Phi_\pi)(y_i) - (\xi_n(\varphi) \circ \Phi_\pi)(y_i)\|_{\text{tr}'} \leq \sum_i |c_{i,M}| \|\lambda(M(x_1, \dots, x_K))\|_{\text{tr}'}$$

Now, consider a monomial $M(x_1, \dots, x_K) = x_{i_1}^{e_1} \cdots x_{i_d}^{e_d}$ where $e_j \in \{1, *\}$. We have that

$$\|\lambda(M(x_1, \dots, x_K))\|_{\text{tr}'} \leq \sum_{j=1}^d C^{d-1} \|\lambda(x_{i_j})\|_{\text{tr}'} \leq dC^{d-1} \max_i \|\lambda(x_i)\|_{\text{tr}'}$$

Therefore,

$$\|\xi_m(\varphi \circ \Phi_\pi)(y_i) - (\xi_n(\varphi) \circ \Phi_\pi)(y_i)\|_{\text{tr}'} \leq DC^{D-1} N \max_{x \in S_m} \|\xi_m(\varphi \circ \Phi_\pi)(x) - (\xi_n(\varphi) \circ \Phi_\pi)(x)\|_{\text{tr}'}$$

Choosing δ such that $\text{def}(\text{tr} \circ \varphi) \leq \delta$ implies that $\|\xi_m(\varphi \circ \Phi_\pi)(x) - (\xi_n(\varphi) \circ \Phi_\pi)(x)\|_{\text{tr}'} \leq \frac{\varepsilon}{DC^{D-1}N}$ gives the wanted result. $\blacksquare$

To finish this section, we show that stable minion homomorphisms are in fact a generalisation of minion homomorphisms in the usual sense.

Lemma 5.3.9. Let $\mathcal{M}$ and $\mathcal{N}$ C^* -cominions approximated by weighted cominions $\mathcal{A}$ and $\mathcal{B}$, respectively. Let Q be a finite measured tensor monoid. If there exists a stable Q -minion homomorphism $\xi : \mathcal{A} \rightarrow \mathcal{B}$, then there exists a minion homomorphism $\mathcal{M}^Q \rightarrow \mathcal{N}^Q$.

Proof. Define a minion homomorphism $\xi' : \mathcal{M}^Q \rightarrow \mathcal{N}^Q$ as follows. Let $\varphi : \mathcal{M}_n \rightarrow (\mathcal{M}, \text{tr}) \in Q$. Then, $\varphi \circ q_n$ is a $*$ -representation of $\mathcal{A}_n$, giving $\xi_n(\varphi \circ q_n) : \mathcal{B}_n \rightarrow (\mathcal{N}, \text{tr}')$. Since $\text{tr} \circ \varphi \circ q_n$ is a tracial state on $\mathcal{A}_n$ with defect 0, $\text{tr}' \circ \xi_n(\varphi \circ q_n)$ is a tracial state with defect 0 on $\mathcal{B}_n$. Using the universal property of the universal C^* -algebra, $\xi_n(\varphi \circ q_n)$ induces a $*$ -homomorphism $\mathcal{N}_n \rightarrow \mathcal{N} \in Q$, which we take to be $\xi'_n(\varphi)$.

Let $\pi : [n] \rightarrow [m]$. Then, as the defect is 0,

$$\xi_m(\varphi \circ \Phi_\pi \circ q_m) = \xi_m(\varphi \circ q_n \circ \Phi_\pi) = \xi_n(\text{tr} \circ \varphi \circ q_n) \circ \Phi_\pi.$$

Thus, using again the universal property, $\xi'_m(\varphi \circ \Phi_\pi) = \xi'_n(\varphi) \circ \Phi_\pi$. $\blacksquare$

5.3.3 Proof of the gap-preserving reduction

Theorem 5.3.10. Let Q be a finite measured tensor monoid. Let A be a σ -structure and B, Y be τ -structures. Let π be a probability distribution on the constraints of Y . Let X be the σ -structure constructed in Theorem 2.10.2. Suppose that there exists a Q -stable minion homomorphism $\xi : \text{Pol}^{ocv}(A) \rightarrow \text{Pol}^{ocv}(B)$. Then, there exists a probability distribution π' on the constraints of X such that, for any $\varepsilon > 0$, there exists $\delta > 0$ (depending only on A, B, Ξ , and ε) such that, for any tracial state τ on $\text{Mor}_{\pi'}^{ocv}(X, A)$ over Q , if $\text{def}(\tau) \leq \delta$, then there exists a tracial state τ' on $\text{Mor}_{\pi}^{ocv}(Y, B)$ over Q such that $\text{def}(\tau') \leq \varepsilon$.

Proof. Define the probability distribution π' as follows. Let $C = \sum_{S \in \sigma} |S^{A^{\mathfrak{D}(B)}}|$ be the number of constraints in $A^{\mathfrak{D}(B)}$, and for each $R \in \tau$, let $C_R = \sum_{S \in \sigma} |S^{A^{R^B}}|$ be the number of constraints in A^{R^B} . For $y \in Y$, let $\pi(y) = \sum_{\substack{R \in \tau, \mathbf{y} \in R^Y \\ i \in [\text{ar}(R)] : y_i = y}} \frac{\pi(R, \mathbf{y})}{\text{ar}(R)}$, which is the probability of sending Bob y . For $S \in \sigma$ and $\mathbf{x} \in S^X$, let

$$\pi'(S, \mathbf{x}) = \begin{cases} \frac{\pi(R, \mathbf{y})}{2C_R} & x_i = [(R, \mathbf{y}, \mathbf{a}_i)] \\ \frac{\pi(y)}{2C} & x_i = [(y, \mathbf{a}_i)]. \end{cases}$$

Now, let $\varepsilon > 0$, and fix $\delta > 0$ to be determined later. Let τ be a tracial state on $\text{Mor}_{\pi'}^{ocv}(X, A)$ over Q such that $\text{def}(\tau) \leq \delta$. Let $\varphi : \text{Mor}_{\pi'}^{ocv}(X, A) \rightarrow (\mathcal{M}, \text{tr}) \in Q$ be the GNS representation of τ . Analogously to the gapless case, let the embedding $*$ -homomorphisms $\Phi_y : \text{Mor}_{\mathbf{u}}^{ocv}(A^{\mathfrak{D}(B)}, A) \rightarrow \text{Mor}_{\pi'}^{ocv}(X, A)$ and $\Phi_{R, \mathbf{y}} : \text{Mor}_{\mathbf{u}}^{ocv}(A^{R^B}, A) \rightarrow \text{Mor}_{\pi'}^{ocv}(X, A)$, and we have that $\Phi_{R, \mathbf{y}} \circ \Phi_{\pi_i} = \Phi_{y_i}$. It is also direct to see that

$$\text{def}(\tau) = \sum_{R \in \tau, \mathbf{y} \in R^Y} \frac{\pi(R, \mathbf{y})}{2} \text{def}(\tau \circ \Phi_{R, \mathbf{y}}) + \sum_{y \in Y} \frac{\pi(y)}{2} \text{def}(\tau \circ \Phi_y).$$

Now, applying the minion homomorphism, $\xi(\varphi \circ \Phi_y) : \text{Mor}_{\mathbf{u}}^{ocv}(B^{\mathfrak{D}(B)}, B) \rightarrow (\mathcal{N}, \text{tr}') \in Q$ and $\xi(\varphi \circ \Phi_{R, \mathbf{y}}) : \text{Mor}_{\mathbf{u}}^{ocv}(B^{R^B}, B) \rightarrow (\mathcal{N}, \text{tr}') \in Q$ — we may take the codomain to be the same for all $*$ -homomorphisms in the same way as for the gapless case. Let δ_n be the maximal convex modulus function such that for a $*$ -representation η of $\text{Mor}_{\mathbf{u}}^{ocv}(A^n, A)$ over Q , $\text{def}(\text{tr}' \circ \xi(\eta)) \leq \delta_n^{-1}(\text{def}(\text{tr} \circ \eta))$. Note that $\delta_1 \geq \delta_2 \geq \dots$. Let also δ_{π} be the maximum convex modulus function such for a $*$ -representation η of $\text{Mor}_{\mathbf{u}}^{ocv}(A^n, A)$ over Q , $\|\xi(\eta \circ \Phi_{\pi})(p_{b,b}) - (\xi(\eta) \circ \Phi_{\pi})(p_{b,b})\|_{\text{tr}'}^2 \leq \delta_{\pi}^{-1}(\max\{\text{def}(\text{tr} \circ \eta), \text{def}(\text{tr} \circ \eta \circ \Phi_{\pi})\})$. Let Δ the pointwise minimum of $\delta_{|B|}$, $\delta_{|R^B|}$ for all $R \in \tau$, and δ_{π_i} for the projections $\pi_i : R^B \rightarrow \mathfrak{D}(B)$. Therefore $\text{def}(\text{tr}' \circ \xi(\varphi \circ \Phi_y)) \leq \Delta^{-1}(\text{def}(\text{tr} \circ \varphi \circ \Phi_y))$, $\text{def}(\text{tr}' \circ \xi(\varphi \circ \Phi_{R, \mathbf{y}})) \leq \Delta^{-1}(\text{def}(\text{tr} \circ \varphi \circ \Phi_{R, \mathbf{y}}))$, and $\|\xi(\varphi \circ \Phi_{y_i})(p_{b,b}) - (\xi(\varphi \circ \Phi_{R, \mathbf{y}}) \circ \Phi_{\pi_i})(p_{b,b})\|_{\text{tr}'}^2 \leq \Delta^{-1}(\max\{\text{def}(\text{tr} \circ \varphi \circ \Phi_{R, \mathbf{y}}), \text{def}(\text{tr} \circ \varphi \circ \Phi_{y_i})\})$.

Now, we define the $*$ -homomorphism $\varphi' : \text{Mor}_\pi^{\text{ocv}}(Y, B) \rightarrow \mathcal{N}$ on the generators as $\varphi'(p_{y,b}) = \xi(\varphi \circ \Phi_y)(p_{\text{id}_B, b})$ and $\varphi'(\Phi_{R,y,b}) = \xi(\varphi \circ \Phi_{R,y})(\Phi_{R,(\pi_1, \dots, \pi_{\text{ar}(R)}), b})$. This is well-defined as $\text{id}_B \in B^{\mathfrak{D}(B)}$ and $(\pi_1, \dots, \pi_{\text{ar}(R)}) \in R^{B^{R^B}}$; and this is a $*$ -homomorphism as the images of the generating PVMs are themselves PVMs. It remains to bound the defect of $\tau' = \text{tr}' \circ \varphi'$. We have that

$$\text{def}(\tau') = \sum_{\substack{R \in \tau, \mathbf{y} \in R^Y \\ b \in R^B, i \in \text{ar}(R)}} \frac{\pi(R, \mathbf{y})}{\text{ar}(R)} \|\xi(\varphi \circ \Phi_{R,y})(\Phi_{R,(\pi_1, \dots, \pi_{\text{ar}(R)}), b})(1 - \xi(\varphi \circ \Phi_{y_i})(p_{\text{id}_B, b_i}))\|_{\text{tr}'}^2.$$

As in the gapless case, we want to replace $\xi(\varphi \circ \Phi_{y_i})(p_{\text{id}_B, b_i})$ by $\xi(\varphi \circ \Phi_{R,y})(p_{\pi_i, b_i})$. Doing so, we find that

$$\begin{aligned} \text{def}(\tau') &\leq 2 \sum_{\substack{R \in \tau, \mathbf{y} \in R^Y \\ b \in R^B, i \in \text{ar}(R)}} \frac{\pi(R, \mathbf{y})}{\text{ar}(R)} \left(\|\Phi_{R,(\pi_1, \dots, \pi_{\text{ar}(R)}), b}(1 - p_{\pi_i, b_i})\|_{\text{tr}' \circ \xi(\varphi \circ \Phi_{R,y})}^2 \right. \\ &\quad \left. + \|\xi(\varphi \circ \Phi_{y_i})(p_{\text{id}_B, b}) - (\xi(\varphi \circ \Phi_{R,y}) \circ \Phi_{\pi_i})(p_{\text{id}_B, b})\|_{\text{tr}'}^2 \right) \\ &\leq 2 \sum_{R \in \tau, \mathbf{y} \in R^Y} \pi(R, \mathbf{y}) C_R \text{def}(\text{tr}' \circ \xi(\varphi \circ \Phi_{R,y})) \\ &\quad + 2 \sum_{\substack{R \in \tau, \mathbf{y} \in R^Y \\ i \in \text{ar}(R)}} \frac{\pi(R, \mathbf{y}) |R^B|}{\text{ar}(R)} \|\xi(\varphi \circ \Phi_{y_i})(p_{\text{id}_B, b}) - (\xi(\varphi \circ \Phi_{R,y}) \circ \Phi_{\pi_i})(p_{\text{id}_B, b})\|_{\text{tr}'}^2 \\ &\leq 2 \sum_{R \in \tau, \mathbf{y} \in R^Y} \pi(R, \mathbf{y}) C_R \Delta^{-1} (\text{def}(\text{tr} \circ \varphi \circ \Phi_{R,y})) \\ &\quad + 2 \sum_{\substack{R \in \tau, \mathbf{y} \in R^Y \\ i \in \text{ar}(R)}} \frac{\pi(R, \mathbf{y}) |R^B|}{\text{ar}(R)} \Delta^{-1} (\max\{\text{def}(\text{tr} \circ \varphi \circ \Phi_{R,y}), \text{def}(\text{tr} \circ \varphi \circ \Phi_{y_i})\}) \\ &\leq 2 \max_R (C_R) \Delta^{-1} \left(\sum_{R \in \tau, \mathbf{y} \in R^Y} \pi(R, \mathbf{y}) \text{def}(\text{tr} \circ \varphi \circ \Phi_{R,y}) \right) \\ &\quad + 2 \max_R |R^B| \Delta^{-1} \left(\sum_{\substack{R \in \tau, \mathbf{y} \in R^Y \\ i \in \text{ar}(R)}} \frac{\pi(R, \mathbf{y})}{\text{ar}(R)} (\text{def}(\text{tr} \circ \varphi \circ \Phi_{R,y}) + \text{def}(\text{tr} \circ \varphi \circ \Phi_{y_i})) \right) \\ &\leq 2(\max_R (C_R) + \max_R |R^B|) \Delta^{-1} (2 \text{def}(\tau)). \end{aligned}$$

Taking $\delta = \frac{1}{2} \Delta \left(\frac{\varepsilon}{2(\max_R (C_R) + \max_R |R^B|)} \right)$ gives the wanted result. $\blacksquare$

Theorem 5.3.11. Let Q be a finite measured tensor monoid. Let A be a σ -structure and B, Y be τ -structures, Let π be a probability distribution on the constraints of Y . Let X be

the σ -structure constructed in [Theorem 2.10.2](#). Suppose that there exists a Q -stable minion homomorphism $\xi : \text{Pol}^{no}(A) \rightarrow \text{Pol}^{ocv}(B)$. Then, there exists a probability distribution π' on the constraints of X such that, for any $\varepsilon > 0$, there exists $\delta > 0$ (depending only on A , B , ξ , and ε) such that, for any tracial state τ on $\text{Mor}_{\pi'}^{no}(X, A)$ over Q , if $\text{def}(\tau) \leq \delta$, then there exists a tracial state τ' on $\text{Mor}_{\pi}^{ocv}(Y, B)$ over Q such that $\text{def}(\tau') \leq \varepsilon$.

Proof. Noting that the proof of the previous theorem does not actually make use of the structure of the defect of $\text{Mor}_{\pi'}^{ocv}(X, A)$, the exact same proof holds here. ■

Theorem 5.3.12. Let A be a σ -structure and B be a τ -structure that is well-connected, let Q be a finite measured tensor monoid, and let $s < 1$. If there exists a Q -stable minion homomorphism $\text{Pol}^{ocv}(A) \rightarrow \text{Pol}^{ocv}(B)$, then there exists $s' < 1$ such that there are polynomial-time reductions from $\text{MIP}_{ocv}^Q(B)_{1,s}$ to $\text{MIP}_{ocv}^Q(A)_{1,s'}$ and from ${}^e\text{MIP}_{ocv}^Q(B)_{1,s}$ to ${}^e\text{MIP}_{ocv}^Q(A)_{1,s'}$.

If there exists a Q -stable minion homomorphism $\text{Pol}^{no}(A) \rightarrow \text{Pol}^{ocv}(B)$, then there exists $s' < 1$ such that there are polynomial-time reductions from $\text{MIP}_{ocv}^Q(B)_{1,s}$ to $\text{MIP}_{no}^Q(A)_{1,s'}$ and from ${}^e\text{MIP}_{ocv}^Q(B)_{1,s}$ to ${}^e\text{MIP}_{no}^Q(A)_{1,s'}$.

We also have the gapped analogue of the reduction in the non-succinct case, [Theorem 5.2.4](#), where the property that B is well-connected can be dropped.

Proof. As in [Theorem 5.2.6](#), we show the reduction from ${}^e\text{MIP}_{ocv}^Q(B)_{1,s}$ to ${}^e\text{MIP}_{ocv}^Q(A)_{1,s'}$. The other cases are analogous. Let M be an instance of ${}^e\text{MIP}_{ocv}^Q(B)_{1,s}$. Construct M' in the same way as in [Theorem 5.2.6](#). Completeness follows in the same way using [Corollary 5.2.2](#). For soundness, we take $\varepsilon = 1 - s$, and let $s' = 1 - \delta$, where δ is from [Theorem 5.3.10](#). Then soundness follows immediately from that theorem. ■

5.4 Stable minion homomorphisms from cominon homomorphisms

In this section, we show that a cominon homomorphism can be lifted to a stable minion homomorphism. This will be useful in showing gapped reductions based on [Theorem 5.3.12](#).

Lemma 5.4.1. Let $\mathcal{A} = \{(\mathcal{A}_n, \mu_n)\}_n$ and $\mathcal{B} = \{(\mathcal{B}_n, \nu_n)\}_n$ be weighted cominions that satisfy the following properties:

- (i) $\mathcal{A}_n$ and $\mathcal{B}_n$ are freely generated by finite-order unitaries $u_{n,1}, \dots, u_{n,K_n}$ and $v_{n,1}, \dots, v_{n,L_n}$, respectively;
- (ii) for $\pi : [n] \rightarrow [m]$, the minor homomorphisms satisfy $\Phi_\pi(u_{m,k}) = u_{n,f_\pi(k)}$ and $\Phi_\pi(v_{m,k}) = v_{n,g_\pi(k)}$ for some functions $f_\pi : [K_m] \rightarrow [K_n]$ and $g_\pi : [L_m] \rightarrow [L_n]$.

Let $\mathcal{M}$ and $\mathcal{N}$ be the C^* -cominions approximated by $\mathcal{A}$ and $\mathcal{B}$, respectively. Let $V_{n,k} \in \mathcal{A}_n$ for all $k \in [L_n]$, and suppose that the map $v_{n,k} \mapsto V_{n,k}$ induces a cominion homomorphism $\Xi : \mathcal{M} \rightarrow \mathcal{N}$. Then there is a Q -stable minion homomorphism $\mathcal{A} \rightarrow \mathcal{B}$ for all finite measured tensor monoids Q .

Note that the weighted polymorphism cominions satisfy properties (i) and (ii).

Proof. Write $e_{n,k}$ for the order of $v_{n,k}$. Let $\varphi : \mathcal{A}_n \rightarrow (\mathcal{M}, \text{tr}) \in Q$. First, we will define $\xi_n(\varphi)$, by taking $\xi_n(\varphi)(v_{n,k})$ to be the nearest order- $e_{n,k}$ unitary to $\varphi(V_{n,k})$. Clearly, $\xi_n(\varphi) : \mathcal{B}_n \rightarrow (\mathcal{M}, \text{tr})$ is a $*$ -homomorphism.

First we show the following claim, which we will use a few times: there exist $C_{n,k} \geq 0$ and for all $\delta_0 > 0$ there exists $C'_{n,k,\delta_0} > 0$ such that for any $*$ -homomorphism $\varphi : \mathcal{A}_n \rightarrow (\mathcal{M}, \text{tr}) \in Q$, $\|\varphi(V_{n,k}) - \xi_n(\varphi)(v_{n,k})\|_{\text{tr}}^2 \leq C_{n,k}\delta_0 + C'_{n,k,\delta_0}\sqrt{\text{def}(\text{tr} \circ \varphi)}$. Since $\mathcal{M}_n = C_u^*(\mathcal{A}_n / \langle\langle \text{supp } \mu_n \rangle\rangle)$, $V_{n,k}$ is an order- $e_{n,k}$ unitary in $\mathcal{M}_n$. Hence, $R = (V_{n,k}^* V_{n,k} - 1)^2 + |V_{n,k}^{e_{n,k}} - 1|^2$ is in the infinitesimal ideal of $\mathcal{A}_n / \langle\langle \text{supp } \mu_n \rangle\rangle$. Then, there exists $p \in \mathcal{A}_n / \langle\langle \text{supp } \mu_n \rangle\rangle$ positive such that $\delta_0 - R = p$. To get the relation in $\mathcal{A}_n$, we know there exist $r_1, \dots, r_N \in \text{supp } \mu_n$ and $a_i, b_i \in \mathcal{A}_n$ such that $\delta_0 - R = p + \sum_{i=1}^N a_i r_i b_i$. Now, using [Lemma 3.1.14](#), we know that

$$\begin{aligned}
\|\varphi(V_{n,k}) - \xi_n(\varphi)(v_{n,k})\|_{\text{tr}}^2 &\leq \left(e_{n,k}(\|V_{n,k}\|^{e_{n,k}} + 2) \|V_{n,k}^* V_{n,k} - 1\|_{\text{tr} \circ \varphi} + \|V_{n,k}^{e_{n,k}} - 1\|_{\text{tr} \circ \varphi} \right)^2 \\
&\leq 2e_{n,k}^2(\|V_{n,k}\|^{e_{n,k}} + 2)^2 \left(\|V_{n,k}^* V_{n,k} - 1\|_{\text{tr} \circ \varphi}^2 + \|V_{n,k}^{e_{n,k}} - 1\|_{\text{tr} \circ \varphi}^2 \right) \\
&= 2e_{n,k}^2(\|V_{n,k}\|^{e_{n,k}} + 2)^2 (\text{tr} \circ \varphi)(R) \\
&= 2e_{n,k}^2(\|V_{n,k}\|^{e_{n,k}} + 2)^2 \left(\delta_0 - (\text{tr} \circ \varphi)(p) - \sum_{i=1}^N (\text{tr} \circ \varphi)(a_i r_i b_i) \right) \\
&\leq 2e_{n,k}^2(\|V_{n,k}\|^{e_{n,k}} + 2)^2 \delta_0 + 2e_{n,k}^2(\|V_{n,k}\|^{e_{n,k}} + 2)^2 \sum_{i=1}^N \|b_i a_i\| \|r_i\|_{\text{tr} \circ \varphi} \\
&\leq 2e_{n,k}^2(\|V_{n,k}\|^{e_{n,k}} + 2)^2 \delta_0 + 2e_{n,k}^2(\|V_{n,k}\|^{e_{n,k}} + 2)^2 \sum_{i=1}^N \frac{\|b_i a_i\|}{\sqrt{\mu_n(r_i)}} \sqrt{\text{def}(\text{tr} \circ \varphi)}.
\end{aligned}$$

Taking $C_{n,k} = 2e_{n,k}^2(\|V_{n,k}\|^{e_{n,k}} + 2)^2$ and $C'_{n,k,\delta_0} = 2e_{n,k}^2(\|V_{n,k}\|^{e_{n,k}} + 2)^2 \sum_{i=1}^N \frac{\|b_i a_i\|}{\sqrt{\mu_n(r_i)}}$ proves the claim.

Next, we show property (i). Let $\varepsilon > 0$ and $\pi : [n] \rightarrow [m]$. Let $\varepsilon > 0$. Suppose that $\text{def}(\text{tr} \circ \varphi) \leq \delta$ and $\text{def}(\text{tr} \circ \varphi \circ \Phi_\pi) \leq \delta$, to which we will assign a value later. We want to bound $\|\xi_m(\varphi \circ \Phi_\pi)(v_{m,k}) - (\xi_n(\varphi) \circ \Phi_\pi)(v_{m,k})\|_{\text{tr}}^2$ by ε^2 . Let $\delta_0 > 0$, to be determined later. From the claim above, we know that $\|\xi_m(\varphi \circ \Phi_\pi)(v_{m,k}) - \varphi(\Phi_\pi(V_{m,k}))\|_{\text{tr}}^2 \leq C_{m,k}\delta_0 + C'_{m,k,\delta_0}\sqrt{\text{def}(\text{tr} \circ \varphi \circ \Phi_\pi)}$, and that $\|\xi_n(\varphi)(v_{n,g_\pi(k)}) - \varphi(V_{n,g_\pi(k)})\|_{\text{tr}}^2 \leq C_{n,g_\pi(k)}\delta_0 + C'_{n,g_\pi(k),\delta_0}\sqrt{\delta}$. Hence, since $\Phi_\pi(v_{m,k}) = v_{n,g_\pi(k)}$,

$$\begin{aligned} & \|\xi_m(\varphi \circ \Phi_\pi)(v_{m,k}) - (\xi_n(\varphi) \circ \Phi_\pi)(v_{m,k})\|_{\text{tr}}^2 \\ & \leq 3\|\Phi_\pi(V_{m,k}) - V_{n,g_\pi(k)}\|_{\text{tr} \circ \varphi}^2 + 3(C_{m,k} + C_{n,g_\pi(k)})\delta_0 + 3(C'_{m,k,\delta_0} + C'_{n,g_\pi(k),\delta_0})\sqrt{\text{def}(\text{tr} \circ \varphi)}. \end{aligned}$$

Since the map $v_{n,k} \mapsto V_{n,k}$ extends to a cominor homomorphism and $\mathcal{M}_n = C_u^*(\mathcal{A}_n / \langle\langle \text{supp } \mu_n \rangle\rangle)$, $\Phi_\pi(V_{m,k}) = V_{n,g_\pi(k)}$ in $\mathcal{M}_n$. This implies that $\Phi_\pi(V_{m,k}) - V_{n,g_\pi(k)}$ is in the infinitesimal ideal of $\mathcal{A}_n / \langle\langle \text{supp } \mu_n \rangle\rangle$. Fix $\delta_1 > 0$, to which we will assign a value later. There exists $p \in \mathcal{A}_n / \langle\langle \text{supp } \mu_n \rangle\rangle$ positive such that $\delta_1 - |\Phi_\pi(V_{m,k}) - V_{n,g_\pi(k)}|^2 = p$. In $\mathcal{A}_n$, this means that there are $r_1, \dots, r_N \in \text{supp } \mu_n$ and $a_i, b_i \in \mathcal{A}_n$ such that $\delta_1 - |\Phi_\pi(V_{m,k}) - V_{n,g_\pi(k)}|^2 - p = \sum_{i=1}^N a_i r_i b_i$. This implies as above

$$\begin{aligned} \|\Phi_\pi(V_{m,k}) - V_{n,g_\pi(k)}\|_{\text{tr} \circ \varphi}^2 &= \delta_1 - (\text{tr} \circ \varphi)(p) - \sum_{i=1}^N (\text{tr} \circ \varphi)(a_i r_i b_i) \\ &\leq \delta_1 + \sum_{i=1}^N \|b_i a_i\| \|r_i\|_{\text{tr} \circ \varphi} \\ &\leq \delta_1 + \sum_{i=1}^N \frac{\|b_i a_i\|}{\sqrt{\mu_n(r_i)}} \sqrt{\text{def}(\text{tr} \circ \varphi)}. \end{aligned}$$

Let $D_{k,\delta_1} = \sum_{i=1}^N \frac{\|b_i a_i\|}{\sqrt{\mu_n(r_i)}}$. Putting everything together, $\|\xi_m(\varphi \circ \Phi_\pi)(v_{m,k}) - (\xi_n(\varphi) \circ \Phi_\pi)(v_{m,k})\|_{\text{tr}}^2 \leq 3\delta_1 + 3(C_{m,k} + C_{n,g_\pi(k)})\delta_0 + (3C'_{m,k,\delta_0} + 3C'_{n,g_\pi(k),\delta_0} + D_{k,\delta_1})\sqrt{\delta}$. To finish, we take $\delta_0 = \min_k \frac{\varepsilon^2}{9(C_{m,k} + C_{n,g_\pi(k)})}$, $\delta_1 = \frac{\varepsilon^2}{9}$, and $\delta = \min_k \left(\frac{\varepsilon^2}{3(3C'_{n,g_\pi(k),\delta_0} + D_{k,\delta_1})} \right)^2$.

Now, we show property (ii). Let $\varepsilon > 0$. Suppose that $\text{def}(\text{tr} \circ \varphi) \leq \delta$, to which we will assign a value later. Fix also $\delta_0 > 0$, to be determined later. Using the first claim, we have that $\|\varphi(V_{n,k}) - \xi_n(\varphi)(v_{n,k})\|_{\text{tr}}^2 \leq C_{n,k}\delta_0 + C'_{n,k,\delta_0}\sqrt{\text{def}(\text{tr} \circ \varphi)}$. Let the upper bound be ε_k , and let $\varepsilon_0 = \max_k \varepsilon_k$. Next, note that every element of $\mathcal{B}_n$ can be seen as a noncommutative

polynomial in the generators $v_{n,k}$. This gives the expression

$$\text{def}(\text{tr} \circ \xi_n(\varphi)) = \sum_{r \in \mathcal{B}_n} \nu_n(r) \|r(\xi_n(\varphi)(v_{n,1}), \dots, \xi_n(\varphi)(v_{n,L_n}))\|_{\text{tr}}^2.$$

Now, fix some $r \in \text{supp } \nu_n$. We can expand $r = \sum_M c_M M$, where the sum is over the monomials and $c_M \in \mathbb{C}$. Let $N_r = \sum_M |c_M|$ and D_r the maximal degree of the monomials with nonzero coefficients. Then

$$\begin{aligned} & \left| \|r(\xi_n(\varphi)(v_{n,1}), \dots, \xi_n(\varphi)(v_{n,L_n}))\|_{\text{tr}} - \|r(\varphi(V_{n,1}), \dots, \varphi(V_{n,L_n}))\|_{\text{tr}} \right| \\ & \leq \|r(\xi_n(\varphi)(v_{n,1}), \dots, \xi_n(\varphi)(v_{n,L_n})) - r(\varphi(V_{n,1}), \dots, \varphi(V_{n,L_n}))\|_{\text{tr}} \\ & \leq \sum_M |c_M| \|M(\xi_n(\varphi)(v_{n,1}), \dots, \xi_n(\varphi)(v_{n,L_n})) - M(\varphi(V_{n,1}), \dots, \varphi(V_{n,L_n}))\|_{\text{tr}}. \end{aligned}$$

For a monomial $M(x_1, \dots, x_{L_n}) = x_{k_1} \cdots x_{k_d}$,

$$\begin{aligned} & \|M(\xi_n(\varphi)(v_{n,1}), \dots, \xi_n(\varphi)(v_{n,L_n})) - M(\varphi(V_{n,1}), \dots, \varphi(V_{n,L_n}))\|_{\text{tr}} \\ & \leq \sum_{i=1}^d \|\xi_n(\varphi)(v_{n,k_i}) - \varphi(V_{n,k_i})\|_{\text{tr}} \leq \sum_{i=1}^d \sqrt{\varepsilon_k} \leq D_r \sqrt{\varepsilon_0}, \end{aligned}$$

giving $\|r(\xi_n(\varphi)(v_{n,1}), \dots, \xi_n(\varphi)(v_{n,L_n}))\|_{\text{tr}} \leq \|r(\varphi(V_{n,1}), \dots, \varphi(V_{n,L_n}))\|_{\text{tr}} + N_r D_r \sqrt{\varepsilon_0}$. Taking the square, $\|r(\xi_n(\varphi)(v_{n,1}), \dots, \xi_n(\varphi)(v_{n,L_n}))\|_{\text{tr}}^2 \leq 2\|r(\varphi(V_{n,1}), \dots, \varphi(V_{n,L_n}))\|_{\text{tr}}^2 + 2N_r^2 D_r^2 \varepsilon_0$, so

$$\text{def}(\text{tr} \circ \xi_n(\varphi)) \leq 2 \sum_{r \in \mathcal{B}_n} \nu_n(r) \|r(V_{n,1}, \dots, V_{n,L_n})\|_{\text{tr} \circ \varphi}^2 + 2 \sum_{r \in \mathcal{B}_n} \nu_n(r) N_r^2 D_r^2 \varepsilon_0.$$

To finish, since $v_{n,k} \mapsto V_{n,k}$ extends to a $*$ -homomorphism $\mathcal{N}_n \rightarrow \mathcal{M}_n$, we have that $r(V_{n,1}, \dots, V_{n,L_n})$ is 0 in $\mathcal{M}_n = C_u^*(\mathcal{A}_n / \langle\langle \text{supp } \mu_n \rangle\rangle)$. Hence, it is in the infinitesimal ideal of $\mathcal{A}_n / \langle\langle \text{supp } \mu_n \rangle\rangle$. Fix $\delta_1 > 0$, to be determined later. There exists $p_r \in \mathcal{A}_n / \langle\langle \text{supp } \mu_n \rangle\rangle$ positive such that $\delta_1 - |r(V_{n,1}, \dots, V_{n,L_n})|^2 = p_r$. In $\mathcal{A}_n$, this means that there exist $s_{r,1}, \dots, s_{r,M_r} \in \text{supp } \mu_n$ and $a_{r,i}, b_{r,i} \in \mathcal{A}_n$ such that $|r(V_{n,1}, \dots, V_{n,L_n})|^2 = \delta_1 - p_r + \sum_{i=1}^{M_r} a_{r,i} s_{r,i} b_{r,i}$. Identically as above,

$$\begin{aligned} \|r(V_{n,1}, \dots, V_{n,L_n})\|_{\text{tr} \circ \varphi}^2 & \leq \delta_1 + \sum_{i=1}^{M_r} (\text{tr} \circ \varphi)(a_{r,i} s_{r,i} b_{r,i}) \leq \delta_1 + \sum_{i=1}^{M_r} \|b_{r,i} a_{r,i}\| \|s_{r,i}\|_{\text{tr} \circ \varphi} \\ & \leq \delta_1 + \sum_{i=1}^{M_r} \frac{\|b_{r,i} a_{r,i}\|}{\sqrt{\mu_n(s_{r,i})}} \sqrt{\text{def}(\text{tr} \circ \varphi)}. \end{aligned}$$

Now, we can put it all together. Let $C_0 = \max_k C_{n,k}$ and $C_1 = \max_k C''_{n,k,\delta_0}$, so that $\varepsilon_0 \leq C_0\delta_0 + C_1\sqrt{\delta}$. Let $C_2 = 2\sum_{r \in \mathcal{B}_n} \nu_n(r)$, $C_3 = 2\sum_{r \in \mathcal{B}_n} \nu_n(r) \sum_{i=1}^{M_r} \frac{\|b_{r,i}a_{r,i}\|}{\sqrt{\mu_n(s_{r,i})}}$, and $C_4 = 2\sum_{r \in \mathcal{B}_n} \nu_n(r)N_r^2D_r^2$. Then, we have that $\text{def}(\text{tr} \circ \xi_n(\varphi)) \leq C_2\delta_1 + C_3\sqrt{\delta} + C_4(C_0\delta_0 + C_1\sqrt{\delta})$. To get the wanted result, it suffices to take $\delta_0 = \frac{\varepsilon}{3C_0C_4}$, $\delta_1 = \frac{\varepsilon}{3C_2}$, and $\delta = \left(\frac{\varepsilon}{C_3+C_1C_4}\right)^2$. Note that δ_0 and δ_1 only depend on the properties of ν_n and the $V_{n,k}$, but δ may depend on the choice of δ_0, δ_1 . $\blacksquare$

5.5 Examples of reductions without commutativity gadgets

In this section, we give some examples of entangled CSPs whose hardness can be ascertained via reductions from minion homomorphisms, but not using commutativity gadgets. In [Section 5.5.1](#), we show the hardness of (non-oracular) entangled graph colouring; and in [Section 5.5.2](#), we show hardness of entangled homomorphism to the diamond graph.

5.5.1 Hardness of entangled colouring

Theorem 5.5.1. There exists a cominon homomorphism $\text{Pol}^+(K_n) \rightarrow \text{Pol}^+(3\text{SAT})$.

Proof. Due to [\[49\]](#), we know that 3SAT admits an oracular algebraic commutativity gadget. Therefore, [Theorem 4.3.4](#) implies that $\text{Pol}^+(3\text{SAT}) = \text{Pol}(3\text{SAT})$. Further, from the classical theory, $\text{pol}(3\text{SAT})$ consists only of the projection maps $\pi_i^k : \{0, 1\}^k \rightarrow \{0, 1\}$. Then, $\text{Mor}^{o+}(3\text{SAT}^k, 3\text{SAT})$ is generated by the functions $f_i^k : \{\pi_1^k, \dots, \pi_k^k\} \rightarrow \mathbb{C}$ defined as $f_i^k(\pi_j^k) = \delta_{i,j}$. Define $\Xi_k : \text{Mor}^{o+}(3\text{SAT}^k, 3\text{SAT}) \rightarrow \text{Mor}^{o+}(K_n^k, K_n)$ with respect to this generating set as $\Xi_k(f_i^k) = \Pi_i^k$, where Π_i^k is as defined in [Theorem 4.4.1](#), with the superscript k added for clarity. Since $\{\Pi_i^k\}_{i=1}^k$ is a central PVM, Ξ_k is clearly a $*$ -homomorphism. To show Ξ is a cominon homomorphism, it remains to show that it intertwines the minor operations. In fact, noting that in $\text{Mor}^{o+}(3\text{SAT}^k, 3\text{SAT})$, $p_{a,a}$ corresponds to the function $p_{a,a}(\pi_i^k) = \delta_{\pi_i^k(a),a} = \delta_{a_i,a}$, we have that $p_{a,a} = \sum_{i \in [k]: a_i=a} f_i^k$. In $\text{Mor}^+(K_n^k, K_n)$, $\Pi_i^k = \sum_{a \in K_n} p_{x,a}p_{y,a}$, for any $\mathbf{x}, \mathbf{y} \in K_n^k$ that agree only on the i -th entry. Let $\pi : [m] \rightarrow [k]$. We aim to show that $\Xi_m \circ \Phi_\pi = \Phi_\pi \circ \Xi_k$, which we do by showing it on the generators $p_{a,a}$. On the left hand side

$$\Xi_m(\Phi_\pi(p_{a,a})) = \Xi_m(p_{a \circ \pi, a}) = \sum_{i \in [m]: a_{\pi(i)}=a} \Xi_m(f_i^m) = \sum_{i \in [m]: a_{\pi(i)}=a} \Pi_i^m = \Pi_{(a \circ \pi)^{-1}(a)}^m,$$

and on the right hand side

$$\Phi_\pi(\Xi_k(p_{a,a})) = \sum_{i \in [k]: a_i = a} \Phi_\pi(\Pi_i^k) = \sum_{i \in [k]: a_i = a} \sum_{x \in K_n} p_{\mathbf{x}_i \circ \pi, x} p_{\mathbf{y}_i \circ \pi, y},$$

where $\mathbf{x}_i, \mathbf{y}_i \in K_n^m$ agree only at the i -th position. We know that $\sum_{x \in K_n} p_{\mathbf{x}_i \circ \pi, x} p_{\mathbf{y}_i \circ \pi, y} = \Pi_{\pi^{-1}(i)}^m$, and therefore that $\Phi_\pi(\Xi_k(p_{a,a})) = \Pi_{(a \circ \pi)^{-1}(a)}^m$ as well, as wanted. ■

Corollary 5.5.2. Let $n \geq 3$. There exists $s < 1$ such that $\text{MIP}_{no}^q(K_n)_{1,s}$ and ${}^e\text{MIP}_{no}^q(K_n)_{1,s}$ are RE-complete via polynomial-time reductions from the halting problem; and $\text{MIP}_{no}^{co}(K_n)_{1,s}$ is coRE-complete via a polynomial-time reduction from the cohalting problem.

Proof. The proof follows immediately by applying [Lemma 5.4.1](#) to get q - and co -stable minion homomorphisms, and then [Theorem 5.3.12](#) to get a reduction from oracular entangled 3SAT. Hardness of 3SAT follows by applying [Theorem 4.1.6](#) using the commutativity gadget of [\[49\]](#). ■

5.5.2 Hardness of homomorphism to the diamond graph

Lemma 5.5.3. Let A and B be σ -structures that are homomorphism equivalent, *i.e.* there exist homomorphisms $f : A \rightarrow B$ and $g : B \rightarrow A$. Then, there exist cominon homomorphisms $\Xi : \text{Pol}^{o+}(A) \rightarrow \text{Pol}^{o+}(B)$ and $\Xi : \text{Pol}^+(A) \rightarrow \text{Pol}^+(B)$.

Proof. We show the oracular case; the non-oracular case is identical. Let $\Xi_n : \text{Mor}^{o+}(B^n, B) \rightarrow \text{Mor}^{o+}(A^n, A)$ be defined on the generators as $\Xi_n(p_{\mathbf{b},b}) = \sum_{a \in f^{-1}(b)} p_{g(\mathbf{b}),a}$, where $g(\mathbf{b}) = (g(b_1), \dots, g(b_n))$. Ξ_n is a $*$ -homomorphism as $\Xi_n = \hat{f} \circ \text{id}_{\text{Mor}^{o+}(B^n, B)} \circ \hat{g}^{\otimes n}$ (where we recall that the notation $\hat{\cdot}$ from [Definition 4.4.6](#) maps classical homomorphisms to characters). It remains to show that it intertwines the minor homomorphisms. Let $\pi : [n] \rightarrow [m]$. Then $\Xi_n(\Phi_\pi(p_{\mathbf{b},b})) = \sum_{a \in f^{-1}(b)} p_{g(\mathbf{b} \circ \pi),a}$ and $\Phi_\pi(\Xi_m(p_{\mathbf{b},b})) = \sum_{a \in f^{-1}(b)} p_{g(\mathbf{b}) \circ \pi, a}$. But for each $i \in [n]$, $g(\mathbf{b} \circ \pi)_i = g(b_{\pi(i)}) = (g(\mathbf{b}) \circ \pi)_i$, giving the wanted equality. ■

Lemma 5.5.4. Let D be the diamond graph (see [Fig. 4.3](#)). There exist cominon homomorphisms $\text{Pol}^{o+}(D) \rightarrow \text{Pol}^{o+}(K_3)$ and $\text{Pol}^+(D) \rightarrow \text{Pol}^{o+}(K_3)$.

Proof. D is homomorphism equivalent to K_3 , so there exist cominon homomorphisms $\text{Pol}^{o+}(D) \rightarrow \text{Pol}^{o+}(K_3)$ and $\text{Pol}^+(D) \rightarrow \text{Pol}^+(K_3)$. Note also that $\text{Pol}^+(K_3) = \text{Pol}^{o+}(K_3)$ as $\text{pol}^+(K_3) = \text{pol}(K_3)$, so the second homomorphism is in fact $\text{Pol}^+(D) \rightarrow \text{Pol}^{o+}(K_3)$. ■

This implies the same hardness results as in the previous section.

5.6 Equivalence of the constraint-variable and constraint-constraint frameworks

In this section, we show the equivalence between the oracular constraint-constraint and constraint-variable frameworks for entangled CSPs. In [Sections 5.6.1](#) and [5.6.2](#) we show the two directions of the implication.

5.6.1 Reduction from constraint-variable to constraint-constraint

Definition 5.6.1. The **oracular constraint-constraint weighted polymorphism cominion** of A is $\text{Pol}^{occ}(A) = \{\text{Mor}_{\mathbf{u}}^{occ}(A^n, A)\}_n$ with minor homomorphisms $\Phi_\pi : \text{Mor}^{occ}(A^m, A) \rightarrow \text{Mor}^{occ}(A^n, A)$ for $\pi : [n] \rightarrow [m]$ given on the generators as $\Phi_\pi(\Phi_{R,\mathbf{a},\mathbf{a}}) = \Phi_{R,(\mathbf{a}_1 \circ \pi, \dots, \mathbf{a}_{\text{ar}(R)} \circ \pi), \mathbf{a}}$.

Lemma 5.6.2. $\text{Pol}^{occ}(A)$ is well-defined as a weighted cominion and approximates $\text{Pol}^{o+}(A)$.

This is proved identically to [Lemmata 5.3.3](#) and [5.3.5](#).

Theorem 5.6.3. Let Q be a finite measured tensor monoid. Let A be a σ -structure and B, Y be τ -structures. Let π be a probability distribution on the constraints of Y . Let X be the σ -structure constructed in [Theorem 2.10.2](#). Suppose that there exists a Q -stable minion homomorphism $\xi : \text{Pol}^{occ}(A) \rightarrow \text{Pol}^{ocv}(B)$. Then, there exists a (connected) probability distribution π' on pairs of constraints of X such that, for any $\varepsilon > 0$, there exists $\delta > 0$ (depending only on A, B, ξ , and ε) such that, for any tracial state τ on $\text{Mor}_{\pi'}^{occ}(X, A)$ over Q , if $\text{def}(\tau) \leq \delta$, then there exists a tracial state τ' on $\text{Mor}_{\pi}^{ocv}(Y, B)$ over Q such that $\text{def}(\tau') \leq \varepsilon$.

Proof. Define the probability distribution π' as follows. Let $C = \sum_{S \in \sigma} |S^{A^{\mathfrak{D}(B)}}|$ be the number of constraints in $A^{\mathfrak{D}(B)}$, and for each $R \in \tau$, let $C_R = \sum_{S \in \sigma} |S^{A^{R^B}}|$ be the number of constraints in A^{R^B} . Write $\pi(y) = \sum_{\substack{R \in \tau, \mathbf{y} \in R^Y \\ i \in [\text{ar}(R)] : y_i = y}} \frac{\pi(R, \mathbf{y})}{\text{ar}(R)}$. For $S, S' \in \sigma$, $\mathbf{x} \in S^X$, and $\mathbf{x}' \in S'^X$, let

$$\pi'((S, \mathbf{x}), (S', \mathbf{x}')) = \begin{cases} \frac{\pi(R, \mathbf{y})}{2C_R^2} & x_i = [(R, \mathbf{y}, \mathbf{a}_i)] \wedge x'_j = [(R, \mathbf{y}, \mathbf{a}'_j)] \\ \frac{\pi(y)}{2C^2} & x_i = [(y, \mathbf{a}_i)] \wedge x'_j = [(y, \mathbf{a}'_j)]. \end{cases}$$

As it is uniform on the pairs of constraints coming from the same copy of $A^{\mathfrak{D}(B)}$ or A^{R^B} , this probability distribution is connected. Let τ be a tracial state on $\text{Mor}_{\pi'}^{occ}(X, A)$ over Q

such that $\text{def}(\tau) \leq \delta$. Analogously to the gapless case, let the embedding $*$ -homomorphisms $\Phi_y : \text{Mor}_{\mathbf{u}}^{\text{occ}}(A^{\mathfrak{D}(B)}, A) \rightarrow \text{Mor}_{\pi'}^{\text{occ}}(X, A)$ and $\Phi_{R,y} : \text{Mor}_{\mathbf{u}}^{\text{occ}}(A^{R^B}, A) \rightarrow \text{Mor}_{\pi'}^{\text{occ}}(X, A)$, and we have that $\Phi_{R,y} \circ \Phi_{\pi_i} = \Phi_{y_i}$. It is also direct to see that

$$\text{def}(\tau) = \sum_{R \in \tau, y \in R^Y} \frac{\pi(R, y)}{2} \text{def}(\tau \circ \Phi_{R,y}) + \sum_{y \in Y} \frac{\pi(y)}{2} \text{def}(\tau \circ \Phi_y).$$

As in [Theorem 5.3.11](#), the proof does not make use of the structure of the defect on A , so the result follows in the same way as in [Theorem 5.3.10](#). $\blacksquare$

Theorem 5.6.4. Let A be a σ -structure and B be a τ -structure, let Q be a finite measured tensor monoid, and let $s < 1$. If there exists a Q -stable minion homomorphism $\text{Pol}^{\text{occ}}(A) \rightarrow \text{Pol}^{\text{ocv}}(B)$, then there exists $s' < 1$ such that there are polynomial-time reductions from $\text{CSP}_{\text{ocv}}^Q(B)_{1,s}$ to $\text{CSP}_{\text{occ}}^Q(A)_{1,s'}$ and from ${}^e\text{CSP}_{\text{ocv}}^Q(B)_{1,s}$ to ${}^e\text{CSP}_{\text{occ}}^Q(A)_{1,s'}$; and if B is well-connected, from $\text{MIP}_{\text{ocv}}^Q(B)_{1,s}$ to $\text{MIP}_{\text{occ}}^Q(A)_{1,s'}$ and from ${}^e\text{MIP}_{\text{ocv}}^Q(B)_{1,s}$ to ${}^e\text{MIP}_{\text{occ}}^Q(A)_{1,s'}$.

The proof follows in the same way as [Theorem 5.3.12](#).

Corollary 5.6.5. Let A be a σ -structure, let $0 \leq s \leq 1$, and let Q be a tensor monoid. Then, there exists $0 \leq s' \leq 1$ such that $s' = 1 \iff s = 1$ for which there are polynomial time reductions from $\text{CSP}_{\text{ocv}}^Q(A)_{1,s}$ to $\text{CSP}_{\text{occ}}^Q(A)_{1,s'}$ and ${}^e\text{CSP}_{\text{ocv}}^Q(A)_{1,s}$ to ${}^e\text{CSP}_{\text{occ}}^Q(A)_{1,s'}$; and if A is well-connected, from $\text{MIP}_{\text{ocv}}^Q(A)_{1,s}$ to $\text{MIP}_{\text{occ}}^Q(A)_{1,s'}$ and ${}^e\text{MIP}_{\text{ocv}}^Q(A)_{1,s}$ to ${}^e\text{MIP}_{\text{occ}}^Q(A)_{1,s'}$.

This corollary implies that the oracular constraint-variable form of a CSP has complexity at most that of the oracular constraint-constraint form.

Proof. Since both $\text{Pol}^{\text{ocv}}(A)$ and $\text{Pol}^{\text{occ}}(A)$ approximate $\text{Pol}^{o+}(A)$, we have by [Lemma 5.4.1](#) that there is a Q -stable minion homomorphism $\text{Pol}^{\text{occ}}(A) \rightarrow \text{Pol}^{\text{ocv}}(A)$. Hence, [Theorem 5.6.4](#) gives us the wanted reductions. $\blacksquare$

5.6.2 Reduction from constraint-constraint to constraint-variable

Lemma 5.6.6 (Lemma III.7 in [\[49\]](#)). Let X and A be σ -structures and let π be a symmetric probability distribution on pairs of constraints of X . Let $\pi'(R, \mathbf{x}) = \sum_{S \in \sigma, y \in S^x} \pi((R, \mathbf{x}), (S, \mathbf{y}))$ and let $L = \max_{R \in \sigma} \text{ar}(R)$. Then, there exists a $4L$ -homomorphism $\alpha : \text{Mor}_{\pi}^{\text{occ}}(X, A) \rightarrow \text{Mor}_{\pi'}^{\text{ocv}}(X, A)$.

Theorem 5.6.7. Let A be a σ -structure, let $0 \leq s \leq 1$, and let Q be a tensor monoid. Then, there exists $0 \leq s' \leq 1$ such that $s' = 1 \iff s = 1$ for which there are polynomial time reductions from $\text{CSP}_{occ}^Q(A)_{1,s}$ to $\text{CSP}_{ocv}^Q(A)_{1,s'}$, ${}^e\text{CSP}_{occ}^Q(A)_{1,s}$ to ${}^e\text{CSP}_{ocv}^Q(A)_{1,s'}$, $\text{MIP}_{occ}^Q(A)_{1,s}$ to $\text{MIP}_{ocv}^Q(A)_{1,s'}$, and ${}^e\text{MIP}_{occ}^Q(A)_{1,s}$ to ${}^e\text{MIP}_{ocv}^Q(A)_{1,s'}$.

This theorem implies that the oracular constraint-constraint form of a CSP has complexity at most that of the oracular constraint-variable form.

Proof. We show that reduction $\text{MIP}_{occ}^Q(A)_{1,s}$ to $\text{MIP}_{ocv}^Q(A)_{1,s'}$. The others are essentially identical. Let $s' = 1 - \frac{1-s}{4L}$. Suppose M is an instance of $\text{MIP}_{occ}^Q(A)_{1,s}$, which corresponds to a σ -structure $X(M)$ and a probability distribution π_M on pairs of constraints of $X(M)$. Let M' be the probabilistic Turing machine that samples from a probability distribution $\pi_{M'}$ on the constraints of $X(M)$ by sampling a pair of constraints from π_M and sampling one uniformly at random. We claim $M \mapsto M'$ is the wanted reduction. Clearly, the mapping is polynomial-time. Next, if M is a yes instance, then $\inf_{\tau} \text{def}(\tau) = 0$, where the infimum is over tracial states on $\text{Mor}_{\pi_M}^{occ}(X(M), A)$ over Q . Since $C_u^*(\text{Mor}_{\pi_M}^{occ}(X(M), A) / \langle\langle \mu_{occ, \pi_M} \rangle\rangle) \cong C_u^*(\text{Mor}_{\pi_{M'}}^{ocv}(X(M), A) / \langle\langle \mu_{ocv, \pi_{M'}} \rangle\rangle)$ by connectedness of the distribution, this implies that the infimum of the defects of the tracial states on $\text{Mor}_{\pi_{M'}}^{ocv}(X(M), A)$ over Q is also 0. Hence M' is a yes instance of $\text{MIP}_{ocv}^Q(A)_{1,s'}$. Now, suppose that M is a no instance. Then $\inf_{\tau} \text{def}(\tau) > 1 - s$, where the infimum is over tracial states on $\text{Mor}_{\pi_M}^{occ}(X(M), A)$ over Q . Then, using [Lemma 5.6.6](#), for every tracial state τ on $\text{Mor}_{\pi_{M'}}^{ocv}(X(M), A)$, $\tau \circ \alpha$ is a tracial state on $\text{Mor}_{\pi_M}^{occ}(X(M), A)$ satisfying $\text{def}(\tau \circ \alpha) \leq 4L \text{def}(\tau)$. Taking infima, we find that $\inf_{\tau} \text{def}(\tau) > \frac{1-s}{4L}$, where the infimum is over tracial states on $\text{Mor}_{\pi_{M'}}^{ocv}(X(M), A)$ over Q . Hence M' is a no instance of $\text{MIP}_{ocv}^Q(A)_{1,s'}$. ■

Chapter 6

Hardness of binary games

In this chapter, we show that nonlocal games where each of the players answer a single bit (binary nonlocal games) capture the full complexity of MIP^* , even for a family of games qualitatively very close to XOR games. In [Section 6.1](#), we review the NP-hardness proof of classical XOR games, which we will adapt in the quantum setting. In [Section 6.2](#), we introduce XOR games, the class of binary games for which we show hardness. In [Sections 6.3](#) and [6.4](#), we show hardness for the gapless and gapped tilted XOR game problems, respectively. In [Section 6.5](#), we discuss the implications for hardness of approximation.

6.1 Classical hardness of XOR games

The NP-completeness of approximating the classical values of XOR games is ascertained in [\[72\]](#) via a gadget reduction from 3LIN games.

Recall that solving linear equations over the boolean field $\{0, 1\}$ with n variables per equation is the CSP with associated boolean relational structure $n\text{LIN}$ with relations $L_{b,n}^{n\text{LIN}} = \{\mathbf{a} \in \mathbb{Z}_2^n \mid a_1 + \dots + a_n = b\}$ for $b \in \{0, 1\}$. It was shown by Håstad that, for all $\varepsilon > 0$, $\text{CSP}_{ocv}^c(3\text{LIN})_{1-\varepsilon, \frac{5}{6}+\varepsilon}$ is NP-complete [\[72\]](#). This gives optimal hardness of approximation as there is always a strategy that wins with probability $\frac{5}{6}$. First, note that, for any 3-linear system, at least half of the equations can be simultaneously satisfied. Then, we take the strategy where Alice outputs the value of the assignment for all satisfied equations, and the assignment with one bit flipped for all unsatisfied equations; and Bob outputs the value of the assignment.

XOR games can also be seen as CSPs. In fact, the problem of deciding if an XOR game has classical value $\geq c$ or $< s$ is exactly ${}_g\text{CSP}_{no}^d(2\text{LIN})_{c,s}$. The reduction from 3LIN to XOR games essentially relies on the following game.

Definition 6.1.1. Let $X_0, X_1 \subseteq \{0, 1\}^3$ be the sets of bit strings of even and odd parity, respectively. Then, the **cube game** is the XOR game $G^\square$ with question sets $X = X_0$ and $Y = X_1$, and game matrix $H_{x,y} = \frac{1}{16}$ for $y = \neg x$ and $H_{x,y} = -\frac{1}{16}$ otherwise.

In [72], this game was originally phrased as a system of equations, but it is natural to represent it as an XOR game. In this work, Håstad attributes the construction to G. Sorkin.

Lemma 6.1.2 ([72]). The value $\mathfrak{w}_c(G^\square) = \frac{3}{4}$. If $p(a, b|x, y) = \delta_{a,g(x)}\delta_{b,h(y)}$ is an optimal deterministic strategy, then $g(011) + g(101) + g(110) = g(000)$. If a deterministic strategy is not optimal, then its value is $\leq \frac{5}{8}$.

To finish this section, we sketch the reduction from 3LIN games to XOR games due to [72]. Let X be a $\sigma_{3\text{LIN}}$ -structure and let π be a probability distribution on the constraints of X . We construct an XOR game as follows. We replace each constraint of the form $(L_{0,3}, \mathbf{x})$ by a copy of the bipartite cube game, identifying the variables x_1, x_2 , and x_3 in the equation with 011, 101, and 110 in the cube, respectively. For constraints of the form $(L_{1,3}, \mathbf{x})$, we do the same except we flip the sign of the relations involving 000 — this is equivalent to flipping the parity of the variable 000 in those cubes. We also identify all the variables corresponding to 000 between the cubes. Since the value of an assignment to an $\sigma_{2\text{LIN}}$ -structure is invariant under negation, we may assume that, for any classical strategy, the variable 000 is assigned value 0.

Now, if (X, π) is a yes instance of $\text{CSP}_{ocv}^c(3\text{LIN})_{1-\varepsilon, \frac{5}{6}+\varepsilon}$, there exists a classical assignment where $1 - O(\varepsilon)$ of the equations are satisfied, which implies that $1 - O(\varepsilon)$ of the cubes may be won with probability $\frac{3}{4}$. Hence the value of the new game is $\geq \frac{3}{4} - O(\varepsilon)$. Conversely, suppose (X, π) is such that the XOR game constructed as above has value $\geq \frac{11}{16} + \varepsilon$ — to ascertain soundness of the reduction, it suffices to see that the classical value of $G_{ocv}(X, 3\text{LIN}, \pi)$ is $\geq \frac{5}{6} + O(\varepsilon)$. Since the value of the cube when the corresponding 3-linear equation is not satisfied is at most $\frac{10}{16}$, the lower bound in the winning probability implies that at least $\frac{1}{2} + O(\varepsilon)$ of the equations are satisfied. For a satisfied equation, the value of the corresponding term in the LCS game is 1, and for an unsatisfied equation, we can choose a strategy where, by making Alice answer consistently on 2 of the 3 variables, the value is $\frac{2}{3}$. Putting these together, we find a strategy for the original 3LIN instance with value $\frac{5}{6} + O(\varepsilon)$.

6.2 Tilted XOR games

In this section, we define tilted XOR games and study some of their properties. In [Section 6.2.1](#), we introduce tilted XOR games and their associated decision problems. In [Section 6.2.2](#), we define the tilted bipartite cube game, which we use in proof of hardness of tilted XOR games.

6.2.1 Definitions and basic properties

Definition 6.2.1. A **tilted XOR game** is a nonlocal game $\mathbf{G} = (X, Y, A, B, \mu, V)$ where $A = B = \{0, 1\}$ and X contains a distinguished element $\perp$, and there exists a function $f : X \times Y \rightarrow \{0, 1\}$ such that

$$V(a, b|x, y) = \begin{cases} \delta_{b, f(\perp, y)} & x = \perp \\ \delta_{a+b, f(x, y)} & \text{else.} \end{cases}$$

This corresponds to an XOR game with some additional questions where only Bob is asked for a bit; however, it can also be thought of an XOR game with the additional condition that, whenever $x = \perp$, the extra condition that $a = 0$ must be met.

By swapping the roles of Alice and Bob, we can see that the variant of tilted XOR games where Bob receives the distinguished question is completely equivalent. One may further generalise tilted XOR games to a **two-sided** variant where both Alice and Bob may receive the distinguished question. However, we do not make use of this variant in this work.

We can express the value of a quantum strategy $p(a, b|x, y) = \langle \psi | P_a^x \otimes Q_b^y | \psi \rangle$ for G using the associated quantum correlation matrix $C_{x,y} = \langle \psi | A_x \otimes B_y | \psi \rangle$ with $A_x = P_0^x - P_1^x$ and $B_y = Q_0^y - Q_1^y$. Without loss of generality, we may suppose that $A_\perp = I$, since Alice's

answer on the distinguished question does not affect the value. Then, the value of p is

$$\begin{aligned}
\mathfrak{w}(\mathbf{G}, p) &= \sum_{y \in Y} \pi(\perp, y) \langle \psi | I \otimes Q_{f(\perp, y)}^y | \psi \rangle + \sum_{\substack{x \in X \setminus \{\perp\}, y \in Y \\ a, b \in \{0, 1\}: a+b=f(x, y)}} \pi(x, y) \langle \psi | P_a^x \otimes Q_b^y | \psi \rangle \\
&= \frac{1}{2} \sum_{y \in Y} \pi(\perp, y) \langle \psi | A_\perp \otimes (I + (-1)^{f(\perp, y)} B_y) | \psi \rangle \\
&\quad + \frac{1}{4} \sum_{x \in X \setminus \{\perp\}, y \in Y} \pi(x, y) \sum_{a \in \{0, 1\}} \langle \psi | (I + (-1)^a A_x) \otimes (I + (-1)^{a+f(x, y)} B_y) | \psi \rangle \\
&= \frac{1}{2} + \frac{1}{2} \sum_{x \in X, y \in Y} (-1)^{f(x, y)} \pi(x, y) \langle \psi | A_x \otimes B_y | \psi \rangle.
\end{aligned}$$

As for an XOR game, let the game matrix be $H_{x, y} = (-1)^{f(x, y)} \pi(x, y)$ and the bias of a correlation matrix $\mathfrak{b}(G, C) = \sum_{x \in X, y \in Y} H_{x, y} C_{x, y}$.

Definition 6.2.2. The **XOR relaxation** of a tilted XOR game $\mathbf{G}$ is the XOR game $\mathbf{G}_{XOR}$ with the same distribution and predicate $V_{XOR}(a, b|x, y) = \delta_{a+b, f(x, y)}$, *i.e.* the distinguished question $\perp$ is treated as a usual question in an XOR game.

Lemma 6.2.3. Let $\mathbf{G}$ be a tilted XOR game. Then, $\mathfrak{w}_c(\mathbf{G}_{XOR}) = \mathfrak{w}_c(\mathbf{G})$ and $\mathfrak{w}_q(\mathbf{G}_{XOR}) \geq \mathfrak{w}_q(\mathbf{G})$.

Note that for the two-sided variant, we can only guarantee $\mathfrak{w}_c(\mathbf{G}_{XOR}) \geq \mathfrak{w}_c(\mathbf{G})$ in general, as for the quantum value.

Proof. For the upper bound, note that via the representation of the tilted XOR game bias in terms of correlation matrices, every strategy for $\mathbf{G}$ corresponds to a strategy for $\mathbf{G}_{XOR}$ where Alice always outputs 0 on question $\perp$. Thus the classical strategies for $\mathbf{G}$ form a subset of the classical strategies for $\mathbf{G}_{XOR}$, and the quantum strategies for $\mathbf{G}$ form a subset of the quantum strategies for $\mathbf{G}_{XOR}$, giving the upper bounds.

For the reverse inequality in the classical case, let $p(a, b|x, y) = \delta_{g(x), a} \delta_{h(y), b}$ be an optimal classical strategy for $\mathbf{G}_{XOR}$. Then, define the classical strategy p' for $\mathbf{G}$ via the functions $g'(x) = g(x) + g(\perp)$ and $h'(y) = h(y) + g(\perp)$. Since $g'(x) + h'(y) = g(x) + h(y)$ and $g'(\perp) = 0$, this is a strategy for $\mathbf{G}$ with the same value as the value of p for $\mathbf{G}_{XOR}$. So, the optimal classical value of $\mathbf{G}$ upper bounds the optimal classical value of $\mathbf{G}_{XOR}$. $\blacksquare$

As for XOR games, tilted XOR games fit within the framework of constraint system games. In the same way as XOR games correspond to linear systems with two variables per equation, tilted XOR games correspond to linear systems with one or two variables per equation.

Lemma 6.2.4. Let TXOR be the boolean relational structure with linear relations $L_{b,2}^{\text{TXOR}}$ and $L_{b,\perp}^{\text{TXOR}} = \{(0, b), (1, b)\}$ for $b \in \{0, 1\}$. Then, for every tilted XOR game $\mathbf{G}$, there exists a relational structure Z and a probability distribution π on the constraints of Z such that $\mathbf{w}_Q(\mathbf{G}) = \mathbf{w}_Q(\mathbf{G}_{no}(Z, \text{TXOR}, \pi))$. There is an analogous value-preserving mapping in the converse direction.

This correspondence allows the computational problem of approximating the value of a tilted XOR game to be phrased as an entangled CSP. We will see from the proof that the correspondence maps between succinctly-presented tilted XOR games and succinctly-presented relational structures efficiently as well. We are interested in the problems ${}_g\text{CSP}_{no}^Q(\text{TXOR})_{c,s}$ and ${}_g\text{MIP}_{no}^Q(\text{TXOR})_{c,s}$.

Proof. Let $\mathbf{G} = (X, Y, S, \{0, 1\}, \{0, 1\}, \mu, V)$ be a tilted XOR game (where we are also allowing for side information). Then, we construct a σ_{TXOR} -structure Z by taking the domain to be $\mathfrak{D}(Z) = X \sqcup Y$ and the relations to be $L_{b,2}^Z = \{(x, y) \in X \times Y \mid x \neq \perp \wedge \exists s : f(x, y, s) = b\}$ and $L_{b,\perp}^Z = \{(\perp, y) \mid \exists s : f(\perp, y, s) = b\}$; and let the probability distribution on the constraints to be $\pi(L_{b,2}, (x, y)) = \sum_{s: f(x, y, s) = b} \mu(x, y)$ for $x \neq \perp$ and otherwise $\pi(L_{b,\perp}, (\perp, y)) = \sum_{s: f(\perp, y, s) = b} \mu(\perp, y)$. It is clear that this mapping preserves the value.

For the converse, if we have a σ_{TXOR} -structure Z with probability distribution π on the constraints, we construct a tilted XOR game with side information as follows. Let $\mathbf{G} = (X, Y, S, \{0, 1\}, \{0, 1\}, \mu, V)$ be the tilted XOR game where $X = \mathfrak{D}(Z) \sqcup \{\perp\}$; $Y = \mathfrak{D}(Z)$; $S = \sigma_{\text{TXOR}}$; $\mu(x, y, L_{b,2}) = \pi(L_{b,2}, (x, y))$ if $x \neq \perp$ and $\mu(\perp, y, L_{b,\perp}) = \sum_x \pi(L_{b,\perp}, (x, y))$; and $f(x, y, L_{b,2}) = b$ and $f(\perp, y, L_{b,\perp}) = b$. Since, for constraints $L_{b,\perp}$, the winning condition does not depend on Alice's output, it follows that this mapping also preserves the value. ■

It follows from the hardness of approximating the classical value of XOR games [72] and Lemma 6.2.3 that for all $\varepsilon > 0$, ${}_g\text{CSP}_{no}^c(\text{TXOR})_{\frac{3}{4}-\varepsilon, \frac{11}{16}+\varepsilon}$ is NP-complete. Also, for $c > s$, ${}_g\text{CSP}_{no}^q(\text{TXOR})_{c,s} \in \text{RE}$ as the problem consists of approximating the value of a nonlocal game to constant precision.

6.2.2 The tilted bipartite cube game

In this section, we study the structure of a tilted XOR arising from the bipartite cube game, which we will make use of as a gadget to reduce from $\text{CSP}_{ocv}^q(3\text{LIN})$. This replaces the cube game of Definition 6.1.1 in the classical reduction from 3LIN to XOR games.

Definition 6.2.5. The **tilted bipartite cube game** is the tilted XOR game $\mathbf{G}^\diamond$ with the same data as $\mathbf{G}^\square$ from Definition 6.1.1 and distinguished variable 000.

It is clear from the definition that $\mathbf{G}_{XOR}^\diamond = \mathbf{G}^\square$. Also, it follows by [Lemma 6.2.3](#) that $\mathfrak{w}_c(\mathbf{G}^\diamond) = \mathfrak{w}_c(\mathbf{G}^\square) = \frac{3}{4}$. Now, we bound the quantum value of $\mathbf{G}^\diamond$.

Lemma 6.2.6. The quantum value $\mathfrak{w}_q(\mathbf{G}^\square) = \frac{3}{4}$.

Since $\mathfrak{w}_q(\mathbf{G}^\diamond)$ is sandwiched in between $\mathfrak{w}_q(\mathbf{G}^\square)$ and $\mathfrak{w}_c(\mathbf{G}^\diamond) = \frac{3}{4}$, the lemma implies $\mathfrak{w}_q(\mathbf{G}^\diamond) = \frac{3}{4}$ as well.

Proof. It is clear that $\mathfrak{w}_q(\mathbf{G}^\square) \geq \mathfrak{w}_c(\mathbf{G}^\square) = \frac{3}{4}$. We use the dual SDP [\(2.5\)](#) discussed in [Section 2.12.2](#) to upper bound the quantum value of $\mathbf{G}^\square$. Ordering the elements of X_0 as (000, 011, 101, 110) and the elements of X_1 as (111, 100, 010, 001) the game matrix is

$$H = \frac{1}{16} \begin{bmatrix} 1 & -1 & -1 & -1 \\ -1 & 1 & -1 & -1 \\ -1 & -1 & 1 & -1 \\ -1 & -1 & -1 & 1 \end{bmatrix}.$$

Taking the square, we find that

$$H^2 = \frac{1}{16^2} \begin{bmatrix} 4 & 0 & 0 & 0 \\ 0 & 4 & 0 & 0 \\ 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & 4 \end{bmatrix} = \frac{1}{64} I,$$

which implies that the eigenvalues of H are $\pm \frac{1}{8}$. Then, the eigenvalues of $B = \frac{1}{2} \begin{bmatrix} 0 & H \\ H^T & 0 \end{bmatrix}$ are $\pm \frac{1}{16}$. Therefore, $\frac{1}{16} I \geq B$, giving that $a_x = b_y = \frac{1}{8}$ is a feasible point for the dual SDP [\(2.5\)](#) for $\mathbf{G}^\square$, giving an upper bound on the bias $\mathfrak{b}_q(\mathbf{G}^\square) \leq \frac{1}{2} \sum_{x \in X_0} a_x + \frac{1}{2} \sum_{y \in X_1} b_y = \frac{1}{2}$. This gives the upper bound on the value $\mathfrak{w}_q(\mathbf{G}^\square) \leq \frac{1}{2} + \frac{1}{2} \cdot \frac{1}{2} = \frac{3}{4}$. ■

6.3 Gapless hardness of tilted XOR games

In this section, we show hardness for the gapless entangled CSP associated to tilted XOR games. In [Section 6.3.1](#), we fully characterise the optimal strategies of the tilted bipartite cube game, and use this in [Section 6.3.2](#) to get the hardness reduction.




Figure 6.1: Schematic representation of the tilted bipartite cube game $\mathbf{G}^\diamond$. The gray vertices are questions to Alice and the white vertices are questions to Bob; the red double edges are inequality relations and the blue single edges are equality relations.

6.3.1 Optimal strategies of the tilted bipartite cube game

Lemma 6.3.1. Let $C_{x,y} = \langle \psi | A_x \otimes B_y | \psi \rangle$ be an optimal quantum correlation matrix for $\mathbf{G}^\diamond$. Suppose that the observables A_x and B_y are unitaries, and write H_ψ for the support of $|\psi\rangle$ on Alice's system. Then, A_{011} , A_{101} , and A_{110} commute, preserve H_ψ , and $A_{011}A_{101}A_{110} = 1$ on H_ψ .

Proof. Since $\mathfrak{w}_q(\mathbf{G}^\square) = \mathfrak{w}_q(\mathbf{G}^\diamond)$, C is an optimal quantum correlation matrix for $\mathbf{G}^\square$. Let $|u_x\rangle = (A_x \otimes I)|\psi\rangle$ and $|v_y\rangle = (I \otimes B_y)|\psi\rangle$. As the inner products $\langle u_x | v_y \rangle$ are all real, we can restrict to the real inner product space generated by $|u_x\rangle$ and $|v_y\rangle$ without changing the value of the correlation. As such, it follows by [Theorem 2.12.1](#) that $\sum_{x \in X} H_{x,y} |u_x\rangle = \frac{1}{8} |v_y\rangle$ and $\sum_{y \in Y} H_{x,y} |v_y\rangle = \frac{1}{8} |u_x\rangle$, as the optimal row and column biases are $a_x = b_y = \frac{1}{8}$ by [Lemma 6.2.6](#). Therefore,

$$\begin{aligned} \sum_{x \in X} H_{x,y} (A_x \otimes 1) |\psi\rangle &= \frac{1}{8} (1 \otimes B_y) |\psi\rangle, \\ \sum_{y \in Y} H_{x,y} (1 \otimes B_y) |\psi\rangle &= \frac{1}{8} (A_x \otimes 1) |\psi\rangle. \end{aligned}$$

Taking the square,

$$\frac{1}{64} |\psi\rangle = \frac{1}{64} (1 \otimes B_y)^2 |\psi\rangle = \frac{1}{8} \sum_x H_{x,y} (A_x \otimes B_y) |\psi\rangle = \left(\sum_x H_{x,y} (A_x \otimes 1) \right)^2 |\psi\rangle.$$

Since for every $|v\rangle \in H_\psi$, there exists $|w\rangle \in H_B$ such that $(1 \otimes \langle w |) |\psi\rangle = |v\rangle$, we have that $A_x |v\rangle = 8 \sum_{y \in Y} H_{x,y} (1 \otimes \langle w | B_y) |\psi\rangle \in H_\psi$. Hence A_x preserves H_ψ , and as A_x is Hermitian,

A_x also preserves $H_\psi^\perp$. As such, we restrict to H_ψ in the following. Then, $(\sum_x H_{x,y} A_x)^2 = \frac{1}{64}$. Writing e_i for the canonical basis vectors in $\{0, 1\}^3$, we have for any $y \in \{0, 1\}^3$ of odd parity that

$$(A_{\neg y} - A_{y+e_1} - A_{y+e_2} - A_{y+e_3})^2 = 4,$$

as $|H_{x,y}| = \frac{1}{16}$. Since C is a quantum correlation matrix for the tilted XOR game $\mathbf{G}^\diamond$, we may assume that $A_{000} = 1$. Writing $A_1 = A_{011}$, $A_2 = A_{101}$, and $A_3 = A_{110}$, we get the four relations

$$\begin{aligned} (I - A_1 - A_2 - A_3)^2 &= 4, \\ (-I + A_1 - A_2 - A_3)^2 &= 4, \\ (-I - A_1 + A_2 - A_3)^2 &= 4, \\ (-I - A_1 - A_2 + A_3)^2 &= 4. \end{aligned}$$

Adding the first two relations gives

$$\begin{aligned} 8I &= ((1 - A_1) - (A_2 + A_3))^2 + (-(1 - A_1) - (A_2 + A_3))^2 \\ &= 2((1 - A_1)^2 + (A_2 + A_3)^2) \\ &= 2(1 - 2A_1 + 1 + 1 + A_2A_3 + A_3A_2 + 1) \\ &= 2(4 - 2A_1 + \{A_2, A_3\}), \end{aligned}$$

which can be rearranged to give $\{A_2, A_3\} = 2A_1$. By symmetry of the relations, we get the other anticommutation relations $\{A_1, A_2\} = 2A_3$ and $\{A_1, A_3\} = 2A_2$. Consider the following:

$$\begin{aligned} 4A_1 &= \{A_2, 2A_3\} \\ &= \{A_2, \{A_1, A_2\}\} \\ &= A_2(A_1A_2 + A_2A_1) + (A_1A_2 + A_2A_1)A_2 \\ &= 2(A_2A_1A_2 + A_1). \end{aligned}$$

Rearranging, we find $A_2A_1A_2 - A_1 = 0$, and then multiplying by A_2 gives the commutation relation $[A_1, A_2] = 0$. In the same way, we find the other commutation relations $[A_1, A_3] = [A_2, A_3] = 0$. Hence, Alice's observables commute. To finish the proof, we see using the commutation and anticommutation relations that

$$A_1A_2A_3 = \frac{1}{2}A_1\{A_2, A_3\} = A_1^2 = 1. \quad \blacksquare$$

6.3.2 Reduction from 3LIN

Theorem 6.3.2. The problems ${}_g\text{CSP}_{no}^q(\text{TXOR})_{\frac{3}{4}, \frac{3}{4}}$ and ${}_g\text{CSP}_{no}^q(\text{TXOR})_{\frac{3}{4}, \frac{3}{4}}$ are undecidable.

It follows that the succinct versions of these problems are also undecidable. This gives the first complexity-theoretic separation between XOR games and tilted XOR games: the analogues of these problems for XOR games, ${}_g\text{CSP}_{no}^q(2\text{LIN})_{c,c}$ and ${}_g\text{CSP}_{no}^q(2\text{LIN})_{c,c}$ with $1 > c > \frac{1}{2}$, are contained in **PSPACE**. This is due to the fact that these problems correspond to the exact feasibility of a semidefinite program, which is a special case of the existential theory of the reals, known to be decidable in **PSPACE** [35].

The proof of the theorem follows by a reduction from gapless entangled 3LIN, whose undecidability was shown in [109]. The following theorem is identical to [Theorem 4.6.2\(i\)](#), but phrased in the CSP notation rather than solution group notation.

Theorem 6.3.3 ([109]). The problems ${}_g\text{CSP}_{ocv}^q(3\text{LIN})_{1,1}$ and ${}_g\text{CSP}_{ocv}^q(3\text{LIN})_{1,1}$ are undecidable.

Proof of [Theorem 6.3.2](#). We show the case of ${}_g\text{CSP}_{no}^q(\text{TXOR})_{\frac{3}{4}, \frac{3}{4}}$ here; the other part of the theorem is identical. Let (X, π) be an instance of ${}_g\text{CSP}_{ocv}^q(3\text{LIN})_{1,1}$. We construct a tilted XOR game analogously to how the XOR game is constructed in the reduction of [72], sketched in [Section 6.1](#). We replace each constraint of the form $(L_{0,3}, \mathbf{x})$ by a copy of the tilted bipartite cube game, identifying the variables x_1, x_2 , and x_3 in the equation with 011, 101, and 110 in the cube, respectively. For constraints of the form $(L_{1,3}, \mathbf{x})$, we do the same except we flip the sign of the relations involving 000 — this is equivalent to flipping the parity of the variable 000 in those cubes. We also identify all the variables corresponding to 000 between the cubes, and take this to be the distinguished variable. Call this instance Y .

First, suppose that (X, π) is a yes instance. Then, by [43, 42], there exists a d -dimensional operator solution to the linear system for some $d \in \mathbb{N}$ — corresponding to a representation of $\text{Mor}^{o+}(X, 3\text{LIN})$. Explicitly, for each $x \in A$, there exists a binary observable A_x , such that for all $\mathbf{x} \in L_{b,3}^X$, A_{x_1}, A_{x_2} , and A_{x_3} commute and satisfy $A_{x_1}A_{x_2}A_{x_3} = (-1)^b$. Now we construct a strategy for the tilted XOR game corresponding to Y with winning probability $\frac{3}{4}$. For the copy of the tilted bipartite cube game coming from the constraint $(L_{b,3}, \mathbf{x})$, assign the observables $A_{x_1}, A_{x_2}, A_{x_3}$ to the vertices 011, 101, and 110 of the tilted cube game, and

take as usual $A_\perp = 1$. Also, we take Bob's observables on this cube to be

$$\begin{aligned} B_{000} &= \frac{1}{2}(-1 + \bar{A}_{x_1} + \bar{A}_{x_2} + \bar{A}_{x_3}) \\ B_{100} &= \frac{1}{2}(+1 - \bar{A}_{x_1} + \bar{A}_{x_2} + \bar{A}_{x_3}) \\ B_{010} &= \frac{1}{2}(+1 + \bar{A}_{x_1} - \bar{A}_{x_2} + \bar{A}_{x_3}) \\ B_{001} &= \frac{1}{2}(+1 + \bar{A}_{x_1} + \bar{A}_{x_2} - \bar{A}_{x_3}), \end{aligned}$$

where the complex conjugate is with respect to a choice of canonical basis. Finally, we take the shared state to be the maximally-entangled state $|\psi\rangle = \frac{1}{\sqrt{d}} \sum_{k=1}^d |k\rangle \otimes |k\rangle$. Then, the quantum correlation matrix of this strategy is $C_{x,y} = \langle \psi | A_x \otimes B_y | \psi \rangle = \text{tr}(A_x \bar{B}_y)$. It is direct to verify that this strategy wins with probability $\frac{3}{4}$ on each cube, and thus wins with probability $\frac{3}{4}$ at the tilted XOR game of Y .

For soundness, suppose that a strategy for the tilted XOR game of Y wins with probability $\geq 3/4$. Note that by Naimark dilation, we may assume that all of the observables are unitary. Since the tilted XOR game is a probabilistic mixture of edge-disjoint tilted bipartite cube games, it follows that the strategy achieves optimal winning probability $3/4$ on each of these. Therefore, by [Lemma 6.3.1](#), for each constraint $(L_{b,3}, \mathbf{x})$, the corresponding observables A_{x_1} , A_{x_2} , and A_{x_3} commute and satisfy $A_{x_1} A_{x_2} A_{x_3} = (-1)^b$. It follows that the observables give a perfect strategy for $\mathbb{G}_{ocv}(X, 3\text{LIN}, \pi)$. $\blacksquare$

6.4 Gapped hardness of tilted XOR games

In this section, we show hardness of the gapped entangled CSP associated to tilted XOR games. In [Section 6.4.1](#), we tighten the XOR game structure theorem of [\[108\]](#). In [Section 6.4.2](#), we use this theorem to characterise near-optimal strategies for the tilted bipartite cube game. In [Section 6.4.3](#), we show the hardness reduction.

6.4.1 Improved XOR game structure theorem

In this section, we prove a tighter average-case version of [Theorem 2.12.1](#) [\[108, Theorem 3.1\]](#) when there is a symmetry condition on the XOR game.

Lemma 6.4.1. Let $\mathbf{G}$ be an XOR game such that the question sets $|X| = |Y|$, and suppose that all the optimal row and column biases are equal $a_x = b_y =: \alpha$. Then, for any ε -optimal vector correlation matrix $C_{x,y} = \langle u_x | v_y \rangle$,

$$\sum_{x \in X} \left\| \sum_{y \in Y} H_{x,y} |v_y\rangle - \alpha |u_x\rangle \right\|^2 \leq 8(\alpha + 2\mathbf{b}_q(G))\varepsilon.$$

By symmetry, we also have that

$$\sum_{y \in Y} \left\| \sum_{x \in X} H_{x,y} |u_x\rangle - \alpha |v_y\rangle \right\|^2 \leq 8(\alpha + 2\mathbf{b}_q(G))\varepsilon.$$

Note that the bipartite cube game satisfies the conditions of this lemma.

Proof. First, write $n = |X|$ and $\beta = \mathbf{b}_q(G)$; we know $\beta = n\alpha$. By definition, there exists some $\varepsilon' \leq 2\varepsilon$ such that $\sum_{x,y} H_{x,y} \langle u_x | v_y \rangle = \beta - \varepsilon'$. Then, we expand the left-hand side from the lemma statement to get that

$$\begin{aligned} \sum_x \left\| \sum_y H_{x,y} |v_y\rangle - \alpha |u_x\rangle \right\|^2 &= \sum_x \left(\alpha^2 - 2\alpha \sum_y H_{x,y} \langle u_x | v_y \rangle + \left\| \sum_y H_{x,y} |v_y\rangle \right\|^2 \right) \\ &= -\alpha\beta + 2\alpha\varepsilon' + \sum_x \left\| \sum_y H_{x,y} |v_y\rangle \right\|^2 \end{aligned}$$

Now, we focus on bounding the last term. Following the argument of [108, Theorem 3.1], let $|u'_x\rangle$ be the normalisation of $\sum_y H_{x,y} |v_y\rangle$ (if $\sum_y H_{x,y} |v_y\rangle = 0$, we take $|u'_x\rangle$ to be an arbitrary unit vector). Since

$$\sum_{x,y} H_{x,y} \langle u'_x | v_y \rangle = \sum_x \left\| \sum_y H_{x,y} |v_y\rangle \right\| \geq \sum_{x,y} H_{x,y} \langle u_x | v_y \rangle,$$

the vector correlation $C'_{x,y} = \langle u'_x | v_y \rangle$ is an ε -optimal strategy as well, with bias $\beta - \varepsilon'_B$ for some $\varepsilon'_B \leq \varepsilon' \leq 2\varepsilon$. Then,

$$\begin{aligned} \sum_x \left\| \sum_y H_{x,y} |v_y\rangle \right\|^2 &= \sum_x \left\| \sum_y H_{x,y} |v_y\rangle - \alpha |u'_x\rangle - \alpha |u'_x\rangle \right\|^2 \\ &= \sum_x \left(\alpha^2 - 2\alpha \sum_y H_{x,y} \langle u'_x | v_y \rangle + 2\alpha^2 + \left\| \sum_y H_{x,y} |v_y\rangle - \alpha |u'_x\rangle \right\|^2 \right) \\ &= \alpha\beta + 2\alpha\varepsilon'_B + \sum_x \left(\sum_y H_{x,y} \langle u'_x | v_y \rangle - \alpha \right)^2. \end{aligned}$$

As in the SDP formulation, write M for the $2n \times 2n$ Gram matrix of the vectors $|u'_x\rangle$ and $|v_y\rangle$ over x and y , and let

$$S = \frac{1}{2} \begin{bmatrix} \alpha I & -H \\ -H^T & \alpha I \end{bmatrix}.$$

We have that $\varepsilon'_B = \text{Tr}(SM)$, $\alpha I \geq S \geq 0$, and $\|M\| \leq \text{Tr}(M) \leq 2n$. Then,

$$\begin{aligned} \sum_x \left(\sum_y H_{x,y} \langle u'_x | v_y \rangle - \alpha \right)^2 &= 4 \sum_{x \in X} (SM)_{x,x}^2 \\ &\leq 4 \|SM\|_F^2 = 4 \text{Tr}(SM^2 S) \\ &\leq 4 \|S\| \|M\| \text{Tr}(SM) \\ &\leq 4(\alpha)(2n)\varepsilon'_B = 8\beta\varepsilon'_B. \end{aligned}$$

Putting it all together,

$$\begin{aligned} \sum_x \left\| \sum_y H_{x,y} |v_y\rangle - \alpha |u_x\rangle \right\|^2 &\leq -\alpha\beta + 2\alpha\varepsilon' + (\alpha\beta + 2\alpha\varepsilon'_B + 8\beta\varepsilon'_B) \\ &\leq 8(\alpha + 2\beta)\varepsilon. \end{aligned} \quad \blacksquare$$

6.4.2 Application to the tilted cube game

Now, we prove a robust version of the [Lemma 6.3.1](#).

Proposition 6.4.2. Let $C_{x,y} = \langle \psi | A_x \otimes B_y | \psi \rangle$ be an ε -optimal quantum correlation matrix for $\mathbf{G}^\diamond$ such that the observables are unitaries. Then,

$$\left\| \left((A_{011}A_{101}A_{110} - I) \otimes I \right) | \psi \rangle \right\| \leq 68\sqrt{13}\varepsilon.$$

By symmetry, the same upper bound holds independent of the order of the observables.

Proof. Note that Alice's observable $A_{000} = 1$. Then, in the notation of [Lemma 6.4.1](#), $\alpha = \frac{1}{8}$ and $\beta = \frac{3}{4}$, which implies that

$$\sum_y \left\| \sum_x H_{x,y} (A_x \otimes I) | \psi \rangle - \frac{1}{8} (I \otimes B_y) | \psi \rangle \right\|^2 \leq 8 \left(\frac{1}{8} + 2\frac{3}{4} \right) \varepsilon = 13\varepsilon. \quad (6.1)$$

Now, let $|\psi\rangle = \sum_i \sqrt{p_i} |a_i\rangle \otimes |b_i\rangle$ be the Schmidt decomposition of $|\psi\rangle$, and take $\lambda : H_B \rightarrow H_A$ to be the operator $\lambda = \sum_i \sqrt{p_i} |a_i\rangle \langle b_i|$. Then, $\lambda\lambda^* = \psi_A$ and $\lambda^*\lambda = \psi_B$, the marginals of $|\psi\rangle$ on H_A and H_B , respectively. We can rewrite (6.1) in terms of the Frobenius norm as

$$\sum_y \left\| \sum_x H_{x,y} A_x \lambda - \frac{1}{8} \lambda \overline{B}_y \right\|_F^2 \leq 13\varepsilon,$$

where the complex conjugate is with respect to the basis $\{|b_i\rangle\}_i$. Writing $h_{x,y} = 16H_{x,y}$, we have $|h_{x,y}| = 1$ and

$$\sum_y \left\| \sum_x h_{x,y} A_x \lambda - 2\lambda \overline{B}_y \right\|_F^2 \leq 2^8 \cdot 13\varepsilon.$$

Next, we remove the dependence on $\overline{B}_y$:

$$\begin{aligned} \left\| \left(\sum_x h_{x,y} A_x \right)^2 \lambda - 4\lambda \right\|_F &\leq \left\| \left(\sum_x h_{x,y} A_x \right)^2 \lambda - 2 \sum_x h_{x,y} A_x \lambda \overline{B}_y \right\|_F + \left\| 2 \sum_x h_{x,y} A_x \lambda \overline{B}_y - 4\lambda \right\|_F \\ &\leq \left\| \sum_x h_{x,y} A_x \right\| \left\| \sum_x h_{x,y} A_x \lambda - 2\lambda \overline{B}_y \right\|_F + 2 \left\| \sum_x h_{x,y} A_x \lambda - 2\lambda \overline{B}_y \right\|_F \left\| \overline{B}_y \right\| \\ &\leq 6 \left\| \sum_x h_{x,y} A_x \lambda - 2\lambda \overline{B}_y \right\|_F, \end{aligned}$$

and therefore

$$\sum_y \left\| \left(\sum_x h_{x,y} A_x \right)^2 \lambda - 4\lambda \right\|_F^2 \leq 36 \sum_y \left\| \sum_x h_{x,y} A_x \lambda - 2\lambda \overline{B}_y \right\|_F^2 \leq 2^{10} \cdot 117\varepsilon.$$

Now, as in the proof of [Lemma 6.3.1](#), let $A_1 = A_{011}$, $A_2 = A_{101}$, and $A_3 = A_{110}$. Then we expand to see that

$$\begin{aligned} &\left\| (1 - A_1 - A_2 - A_3)^2 \lambda - 4\lambda \right\|_F^2 + \left\| (-1 + A_1 - A_2 - A_3)^2 \lambda - 4\lambda \right\|_F^2 \\ &+ \left\| (-1 - A_1 + A_2 - A_3)^2 \lambda - 4\lambda \right\|_F^2 + \left\| (-1 - A_1 - A_2 + A_3)^2 \lambda - 4\lambda \right\|_F^2 \leq 2^{10} \cdot 117\varepsilon. \end{aligned}$$

Using the facts that

$$\begin{aligned} (1 - A_1 - A_2 - A_3)^2 + (-1 + A_1 - A_2 - A_3)^2 &= 2(4 - 2A_1 + \{A_2, A_3\}) \\ (-1 - A_1 + A_2 - A_3)^2 + (-1 - A_1 - A_2 + A_3)^2 &= 2(4 + 2A_1 - \{A_2, A_3\}), \end{aligned}$$

we find that

$$\begin{aligned}
\|\{A_2, A_3\}\lambda - 2A_1\lambda\|_F^2 &= \frac{1}{8}\|(1 - A_1 - A_2 - A_3)^2\lambda + (-1 + A_1 - A_2 - A_3)^2\lambda - 8\lambda\|_F^2 \\
&\quad + \frac{1}{8}\|(-1 - A_1 + A_2 - A_3)^2\lambda + (-1 - A_1 - A_2 + A_3)^2\lambda - 8\lambda\|_F^2 \\
&\leq \frac{1}{4}\|(1 - A_1 - A_2 - A_3)^2\lambda - 4\lambda\|_F^2 + \frac{1}{4}\|(-1 + A_1 - A_2 - A_3)^2\lambda - 4\lambda\|_F^2 \\
&\quad + \frac{1}{4}\|(-1 - A_1 + A_2 - A_3)^2\lambda - 4\lambda\|_F^2 + \frac{1}{4}\|(-1 - A_1 - A_2 + A_3)^2\lambda - 4\lambda\|_F^2 \\
&\leq 2^8 \cdot 117\varepsilon.
\end{aligned}$$

By symmetry, the other anticommutation relations hold approximately as well, with the same error. Next, let $\tilde{B}_1 = \frac{1}{2}(\bar{B}_{100} - \bar{B}_{010} - \bar{B}_{001} - \bar{B}_{111})$, $\tilde{B}_2 = \frac{1}{2}(\bar{B}_{010} - \bar{B}_{100} - \bar{B}_{111} - \bar{B}_{001})$, and $\tilde{B}_3 = \frac{1}{2}(\bar{B}_{001} - \bar{B}_{111} - \bar{B}_{100} - \bar{B}_{010})$. By construction, $\|\tilde{B}_i\| \leq 2$ and by [Lemma 6.4.1](#), $\sum_i \|A_i\lambda - \lambda\tilde{B}_i\|_F^2 \leq 2^6 \cdot 13\varepsilon$. Now, note that

$$\begin{aligned}
A_1A_2A_3 - I &= \frac{1}{2}(A_1(\{A_2, A_3\} - 2A_1) + A_1[A_2, A_3]) \\
&= \frac{1}{2}A_1(\{A_2, A_3\} - 2A_1) - \frac{1}{4}A_1A_2(\{A_2, \{A_2, A_3\}\} - 4A_3) \\
&= \frac{1}{4}A_1(\{A_2, A_3\} - 2A_1) - \frac{1}{2}A_1A_2(\{A_2, A_1\} - 2A_3) - \frac{1}{4}A_1A_2(\{A_2, A_3\} - 2A_1)A_2
\end{aligned}$$

Then, we can use this to bound

$$\begin{aligned}
\|(A_1A_2A_3 - I)\lambda\|_F &\leq \frac{1}{4}\|A_1(\{A_2, A_3\} - 2A_1)\lambda\|_F + \frac{1}{2}\|A_1A_2(\{A_2, A_1\} - 2A_3)\lambda\|_F \\
&\quad + \frac{1}{4}\|A_1A_2(\{A_2, A_3\} - 2A_1)A_2\lambda\|_F \\
&\leq \frac{1}{4}\|(\{A_2, A_3\} - 2A_1)\lambda\|_F + \frac{1}{2}\|(\{A_2, A_1\} - 2A_3)\lambda\|_F \\
&\quad + \frac{1}{4}\|(\{A_2, A_3\} - 2A_1)\lambda\tilde{B}_2\|_F + \frac{1}{4}\|(\{A_2, A_3\} - 2A_1)(A_2\lambda - \lambda\tilde{B}_2)\|_F \\
&\leq \frac{3}{4}\|(\{A_2, A_3\} - 2A_1)\lambda\|_F + \frac{1}{2}\|(\{A_2, A_1\} - 2A_3)\lambda\|_F + \|A_2\lambda - \lambda\tilde{B}_2\|_F \\
&\leq \left(\frac{3}{4} + \frac{1}{2}\right)2^4\sqrt{117\varepsilon} + 2^3\sqrt{13\varepsilon} \\
&= 68\sqrt{13\varepsilon}.
\end{aligned}$$

■

Now, we show that the relation shown for near-optimal strategies for the tilted bipartite cube game in [Proposition 6.4.2](#) gives rise to operators that are near-perfect for a single 3LIN relation. This is necessary in the proof of soundness.

Lemma 6.4.3. Let $p(a, b|x, y) = \langle \psi | P_a^x \otimes Q_b^y | \psi \rangle$ be an ε -optimal correlation for $\mathbb{G}^\diamond$ such that the players' measurements are PVMs. Then, there exists a state $|\phi\rangle \in H_A \otimes H_A$ that depends only on $|\psi\rangle$, and a POVM $\{\Pi_a\}_{a \in \{0,1\}^3} \subseteq \mathcal{B}(H_A)$ such that

$$\frac{1}{3} \sum_{\substack{a \in \{0,1\}^3 \\ a_1 + a_2 + a_3 = 0}} \langle \phi | \Pi_a \otimes (P_{a_1}^{011} + P_{a_2}^{101} + P_{a_3}^{110}) | \phi \rangle \geq 1 - C\varepsilon,$$

where $C = 301814$.

Proof. We use the same notation as in [Proposition 6.4.2](#), and note that Alice and Bob's observables are $A_x = P_0^x - P_1^x$ and $B_y = Q_0^y - Q_1^y$, respectively. Fix an orthonormal basis $\{|k\rangle\}_k$ of H_A , and let $|\phi\rangle = \sum_k |k\rangle \otimes \psi_A^{1/2} |k\rangle$. Let $\Pi_a = \overline{P_{a_3}^3 P_{a_2}^2 P_{a_1}^1 P_{a_2}^2 P_{a_3}^3}$. Write

$$w = \frac{1}{3} \sum_{\substack{a \in \{0,1\}^3 \\ a_1 + a_2 + a_3 = 0}} \langle \phi | \Pi_a \otimes (P_{a_1}^1 + P_{a_2}^2 + P_{a_3}^3) | \phi \rangle.$$

We will lower bound this by upper bounding

$$1 - w = \frac{1}{3} \sum_{\substack{a \in \{0,1\}^3 \\ a_1 + a_2 + a_3 = 0}} \langle \phi | \Pi_a \otimes (P_{\neg a_1}^1 + P_{\neg a_2}^2 + P_{\neg a_3}^3) | \phi \rangle + \sum_{\substack{a \in \{0,1\}^3 \\ a_1 + a_2 + a_3 = 1}} \langle \phi | \Pi_a \otimes 1 | \phi \rangle.$$

First, we want to show that A_i approximately commutes with $\psi_A^{1/2}$. Due to [Lemma 6.4.1](#), $\|A_i \lambda - \lambda \tilde{B}_i\|_F \leq 8\sqrt{13\varepsilon}$. with the goal of removing $\tilde{B}_i$, we first round it to an order-2 unitary. We have

$$\|\lambda(1 - \tilde{B}_i^2)\|_F \leq \|A_i^2 \lambda - A_i \lambda \tilde{B}_i\|_F + \|A_i \lambda \tilde{B}_i - \lambda \tilde{B}_i^2\|_F \leq 3\|A_i \lambda - \lambda \tilde{B}_i\|_F \leq 24\sqrt{13\varepsilon}.$$

Now, let $C_i = \text{sgn}(\tilde{B}_i)$, which is an order-2 unitary by construction. As in [Lemma 3.1.14](#), we have that

$$(1 - \tilde{B}_i^2)^2 = (C_i - \tilde{B}_i)^2 (C_i + \tilde{B}_i)^2 \geq (C_i - \tilde{B}_i)^2,$$

so $\|\lambda(C_i - \tilde{B}_i)\|_F \leq \|\lambda(1 - \tilde{B}_i^2)\|_F \leq 24\sqrt{13\varepsilon}$. Therefore, we can replace $\tilde{B}_i$ with C_i and find

$$\|A_i\lambda - \lambda C_i\|_F \leq \|A_i\lambda - \lambda\tilde{B}_i\|_F + \|\lambda(C_i - \tilde{B}_i)\|_F \leq 32\sqrt{13\varepsilon}.$$

To finish this step, we employ the Araki-Yamagami inequality [3], which gives a tight modulus of continuity for the operator absolute value with respect to the Frobenius norm:

$$\||S| - |T|\|_F \leq \sqrt{2}\|S - T\|_F.$$

Since $|\lambda^* A_i| = \sqrt{A_i \lambda \lambda^* A_i} = A_i \psi_A^{1/2} A_i$ and $|C_i \lambda^*| = \sqrt{\lambda C_i^2 \lambda^*} = \psi_A^{1/2}$,

$$\|[A_i, \psi_A^{1/2}]\|_F = \|A_i \psi_A^{1/2} A_i - \psi_A^{1/2}\|_F = \||\lambda^* A_i| - |C_i \lambda^*|\|_F \leq \sqrt{2}\|\lambda^* A_i - C_i \lambda^*\|_F \leq 32\sqrt{26\varepsilon}.$$

Next, we show that the A_i approximately commute. In fact, using the result of [Proposition 6.4.2](#),

$$\begin{aligned} \|[A_1, A_2]\psi_A^{1/2}\|_F &= \|(A_1 A_2 - A_2 A_1)\lambda\|_F \\ &\leq \|(A_1 A_2 - A_3)\lambda\|_F + \|(A_2 A_1 - A_3)\lambda\|_F \\ &= \|(A_3 A_1 A_2 - 1)\lambda\|_F + \|(A_3 A_2 A_1 - 1)\lambda\|_F \\ &\leq 136\sqrt{13\varepsilon}. \end{aligned}$$

By symmetry, this bound holds for the other two commutators as well. These commutation relations also apply to the PVM elements: $\|[P_a^i, \psi_A^{1/2}]\|_F \leq 16\sqrt{26\varepsilon}$ and $\|[P_a^i, P_b^j]\psi_A^{1/2}\|_F \leq 34\sqrt{13\varepsilon}$. Now, we bound the terms of $1 - w$ one by one. First, note that, for any i ,

$$\langle \phi | \Pi_a \otimes P_{\neg a_i}^i | \phi \rangle = \text{Tr}[\bar{\Pi}_a \psi_A^{1/2} P_{\neg a_i}^i \psi_A^{1/2}] = \|P_{a_1}^1 P_{a_2}^2 P_{a_3}^3 \psi_A^{1/2} P_{\neg a_i}^i\|_F^2.$$

If $i = 3$,

$$\begin{aligned} \|P_{a_1}^1 P_{a_2}^2 P_{a_3}^3 \psi_A^{1/2} P_{\neg a_3}^3\|_F &\leq \|P_{a_1}^1 P_{a_2}^2 P_{a_3}^3 \psi_A^{1/2}\|_F + \|[P_{a_3}^3, \psi_A^{1/2}]\|_F \\ &\leq 16\sqrt{26\varepsilon}; \end{aligned}$$

if $i = 2$,

$$\begin{aligned} \|P_{a_1}^1 P_{a_2}^2 P_{a_3}^3 \psi_A^{1/2} P_{\neg a_2}^2\|_F &\leq \|P_{a_1}^1 P_{a_3}^3 P_{a_2}^2 \psi_A^{1/2} P_{\neg a_2}^2\|_F + \|[P_{a_2}^2, P_{a_3}^3]\psi_A^{1/2}\|_F \\ &\leq \|P_{a_1}^1 P_{a_3}^3 P_{a_2}^2 \psi_A^{1/2}\|_F + \|[P_{a_2}^2, \psi_A^{1/2}]\|_F + \|[P_{a_2}^2, P_{a_3}^3]\psi_A^{1/2}\|_F \\ &\leq 16\sqrt{26\varepsilon} + 34\sqrt{13\varepsilon} = 2(8\sqrt{2} + 17)\sqrt{13\varepsilon}; \end{aligned}$$

and if $i = 1$,

$$\begin{aligned}
\left\| P_{a_1}^1 P_{a_2}^2 P_{a_3}^3 \psi_A^{1/2} P_{\neg a_1}^1 \right\|_F &\leq \left\| P_{a_1}^1 P_{a_2}^2 \psi_A^{1/2} P_{a_3}^3 P_{\neg a_1}^1 \right\|_F + 16\sqrt{26\varepsilon} \\
&\leq \left\| P_{a_1}^1 P_{a_2}^2 \psi_A^{1/2} P_{\neg a_1}^1 P_{a_3}^3 \right\|_F + 16\sqrt{26\varepsilon} + 34\sqrt{13\varepsilon} \\
&\leq \left\| P_{a_2}^2 P_{a_1}^1 \psi_A^{1/2} P_{\neg a_1}^1 P_{a_3}^3 \right\|_F + 16\sqrt{26\varepsilon} + 68\sqrt{13\varepsilon} \\
&\leq \left\| P_{a_2}^2 P_{a_1}^1 P_{\neg a_1}^1 \psi_A^{1/2} P_{a_3}^3 \right\|_F + 32\sqrt{26\varepsilon} + 68\sqrt{13\varepsilon} \\
&= 4(8\sqrt{2} + 17)\sqrt{13\varepsilon}.
\end{aligned}$$

For the final term,

$$\begin{aligned}
\sum_{\substack{\mathbf{a} \in \{0,1\} \\ a_1+a_2+a_3=1}} \langle \phi | \Pi_{\mathbf{a}} \otimes 1 | \phi \rangle &= \sum_{\substack{\mathbf{a} \in \{0,1\} \\ a_1+a_2+a_3=1}} \text{Tr} \left(P_{a_3}^3 P_{a_2}^2 P_{a_1}^1 P_{a_2}^2 P_{a_3}^3 \psi_A \right) \\
&= \frac{1}{32} \sum_{\substack{\mathbf{a} \in \{0,1\} \\ a_1+a_2+a_3=1}} \text{Tr} \left[(1 + (-1)^{a_3} A_3) (1 + (-1)^{a_2} A_2) (1 + (-1)^{a_1} A_1) (1 + (-1)^{a_2} A_2) (1 + (-1)^{a_3} A_3) \psi_A \right] \\
&= \frac{1}{32} \text{Tr} \left[(4 - 4A_1 A_2 A_3 + 4 - 4A_2 A_1 A_3 + 4 - 4A_3 A_1 A_2 + 4 - A_3 A_2 A_1) \psi_A \right] \\
&= \frac{1}{8} \left(\|(A_1 A_2 A_3 - 1)\lambda\|_F^2 + \|(A_2 A_1 A_3 - 1)\lambda\|_F^2 \right) \\
&\leq \frac{1}{4} (68\sqrt{13\varepsilon})^2 = 4 \cdot 13 \cdot 17^2 \varepsilon = 15028\varepsilon.
\end{aligned}$$

Putting everything together, we find the bound

$$\begin{aligned}
1 - w &\leq \frac{4}{3} \left((16\sqrt{26\varepsilon})^2 + (2(8\sqrt{2} + 17)\sqrt{13\varepsilon})^2 + (4(8\sqrt{2} + 17)\sqrt{13\varepsilon})^2 \right) + 15028\varepsilon \\
&= \frac{4}{3} \left(6656\varepsilon + 260(417 + 272\sqrt{2})\varepsilon \right) + 15028\varepsilon \\
&\leq 301814\varepsilon.
\end{aligned}$$

■

6.4.3 Proof of the reduction

In this section, we show the following theorem

Theorem 6.4.4. There exist $1 > c > s > \frac{1}{2}$ such that ${}_g\text{MIP}_{no}^q(\text{TXOR})_{c,s}$ and ${}_g\text{MIP}_{no}^q(\text{TXOR})_{c,s}$ are RE-complete, via polynomial-time reductions from the halting problem.

This implies immediately that the non-succinct versions of the problems are also RE-complete. From the proof, we may choose explicit values as $c = \frac{3}{4} - \varepsilon$ for any $\varepsilon > 0$, and $s = \frac{3}{4} - 10^{-8}$. We prove the result via a reduction from 3LIN, in an analogous way to [72]. We rely on the following result.

Theorem 6.4.5 ([112]). For all $\varepsilon > 0$, ${}_g\text{MIP}_{no}^q(3\text{LIN})_{1-\varepsilon, \frac{119}{120}+\varepsilon}$ and ${}_g\text{MIP}_{no}^q(3\text{LIN})_{1-\varepsilon, \frac{119}{120}+\varepsilon}$ are RE-complete, via polynomial-time reductions from the halting problem.

We also need a converse of Lemma 6.4.3: operators satisfying a single 3LIN relation near-perfectly induce a near-optimal strategy for the tilted bipartite cube game. This is necessary for the proof of completeness.

Lemma 6.4.6. Let $|\psi\rangle \in H_A \otimes H_B$ be a quantum state, and let $\{P_a^u\}_{a \in \{0,1\}} \subseteq \mathcal{B}(H_B)$ for $u \in \{011, 101, 110\}$ and $\{\Pi_a\}_{a \in \{0,1\}^3} \subseteq \mathcal{B}(H_A)$ be PVMs satisfying

$$\frac{1}{3} \sum_{\substack{a \in \{0,1\}^3 \\ a_1+a_2+a_3=0}} \langle \psi | \Pi_a \otimes (P_{a_1}^{011} + P_{a_2}^{101} + P_{a_3}^{110}) | \psi \rangle \geq 1 - \varepsilon.$$

Then, there exist PVMs $\{Q_a^u\}_{a \in \{0,1\}} \subseteq \mathcal{B}(H_A)$ for $u \in \{001, 010, 100, 111\}$ such that $p(a, b|x, y) = \langle \psi | Q_b^y \otimes P_a^x | \psi \rangle$ is a $\frac{3}{8}\varepsilon$ -optimal correlation for the tilted bipartite cube game.

Proof. Without loss of generality we may suppose that $\Pi_a = 0$ if $a_1 + a_2 + a_3 = 1$ — we can change Alice's PVM to one satisfying this condition without decreasing the value. Now write $\Pi_a^{011} = \sum_{a_2, a_3} \Pi_{aa_2a_3}$, $\Pi_a^{101} = \sum_{a_1, a_3} \Pi_{a_1aa_3}$, and $\Pi_a^{110} = \sum_{a_1, a_2} \Pi_{a_1a_2a}$. Let $A_u = P_0^u - P_1^u$ and $C_u = \Pi_0^u - \Pi_1^u$ for all $u \in \{011, 101, 110\}$; let $A_{000} = I$ and $C_{000} = C_{011}C_{101}C_{110}$; and define $B_v = \frac{1}{2}(C_{\neg v} - C_{v+e_1} - C_{v+e_2} - C_{v+e_3})$ for all $v \in X_1$. Since the C_u are commuting order-2 unitaries satisfying $C_{000}C_{011}C_{101}C_{110} = 1$, the B_u are also order-2 unitaries. Next, writing C for the correlation matrix corresponding to p , note that

$$\begin{aligned} 1 - \varepsilon &= \frac{1}{3} \sum_{\substack{u \in \{011, 101, 110\} \\ a \in \{0,1\}}} \langle \psi | \Pi_a^u \otimes P_a^u | \psi \rangle \\ &= \frac{1}{6} \sum_{u \in X_0} (1 + \langle \psi | C_u \otimes A_u | \psi \rangle) - \frac{1}{6} (1 + \langle \psi | C_{000} \otimes 1 | \psi \rangle) \\ &= \frac{1}{2} + \frac{1}{6} \sum_{u \in X_0} \langle \psi | \frac{1}{2} (B_{\neg u} - B_{u+e_1} - B_{u+e_2} - B_{u+e_3}) \otimes A_u | \psi \rangle \\ &\quad - \frac{1}{6} \sum_{a \in \{0,1\}^3} (-1)^{a_1+a_2+a_3} \langle \psi | \Pi_a \otimes 1 | \psi \rangle \\ &= \frac{1}{2} + \frac{4}{3} \mathbf{b}(\mathbf{G}^\diamond, C) - \frac{1}{6} \langle \psi | I \otimes I | \psi \rangle = \frac{1}{3} + \frac{4}{3} \mathbf{b}(\mathbf{G}^\diamond, C). \end{aligned}$$

Hence, we can rearrange and find that $\mathbf{b}(\mathbf{G}^\diamond, C) \geq \frac{1}{2} - \frac{3}{4}\varepsilon$, as wanted. $\blacksquare$

Now we can pass to the proof of the main theorem.

Proof of Theorem 6.4.4. We show the case of ${}^e\text{MIP}_{no}^q(\text{TXOR})_{c,s}$ here; the other part of the theorem is identical. Let M be an instance of ${}^e\text{MIP}_{ocv}^q(3\text{LIN})_{1-\varepsilon,s}$. We construct an instance of ${}^e\text{MIP}_{no}^q(\text{TXOR})_{c',s'}$, where c' and s' will be specified later, similarly to Theorem 6.3.2, by constructing a Turing machine M' sampling constraints Y constructed from $X(M)$. M' samples a constraint $(L_{b,3}, \mathbf{x})$ from M , and then samples a uniformly random constraint from the tilted bipartite cube game, identifying 011, 101, and 110 with x_1, x_2 , and x_3 , respectively; and flipping the signs of the edges adjacent to 000 if $b = 1$.

First, we show completeness. Suppose there exists a quantum correlation p such that $\mathbf{w}(\mathbf{G}_{ocv}(X(M), 3\text{LIN}, \pi_M), p) \geq 1 - \varepsilon$. Using Naimark dilation, we may assume it is projective. Let $1 - \varepsilon_{b,x}$ be the value on the constraint $(L_{b,3}, \mathbf{x})$. Then, by constructing a strategy p' for G using Lemma 6.4.6, the value of the corresponding cube is $\geq \frac{3}{4} - \frac{3}{8}\varepsilon_{b,x}$, and therefore $\omega(G, p') \geq \frac{3}{4} - \frac{3}{8}\varepsilon =: c'$.

Next, we show soundness. Suppose there exists a quantum correlation p for the game $\mathbf{G}_{no}(X(M'), \text{TXOR}, \pi_{M'})$ such that $\mathbf{w}(\mathbf{G}_{no}(X(M'), \text{TXOR}, \pi_{M'}), p) \geq \frac{3}{4} - \delta$. As above, we can suppose the strategy is projective due to Naimark dilation. Then, via Lemma 6.4.3, there exists a strategy p' for $\mathbf{G}_{ocv}(X(M), 3\text{LIN}, \pi_M)$ such that $\mathbf{w}(\mathbf{G}_{ocv}(X(M), 3\text{LIN}, \pi_M), p') \geq 1 - 4 \cdot 10^5 \delta$. Then, as we want $s = 1 - 4 \cdot 10^5 \delta$, we take $s' = \frac{3}{4} - \delta = \frac{3}{4} - \frac{1-s}{4 \cdot 10^5}$.

Putting these together, the construction maps ${}^e\text{MIP}_{ocv}^q(3\text{LIN})_{1-\varepsilon,s}$, which is RE-complete by Theorem 6.4.5, to ${}^e\text{MIP}_{no}^q(\text{TXOR})_{\frac{3}{4} - \frac{3}{8}\varepsilon, \frac{3}{4} - \frac{1-s}{4 \cdot 10^5}}$. $\blacksquare$

6.5 Hardness of approximation

In this section, we show that the best known techniques for approximating the classical value of an XOR game based on the quantum value, due to [41], extend to approximation algorithms for the quantum value of tilted XOR games. We also study the range of values where Theorem 6.4.4 implies approximating the quantum value of a tilted XOR game is hard.

Lemma 6.5.1. Let $\mathbf{G}$ be a tilted XOR game. Then, there exists a classical correlation p for $\mathbf{G}$ such that $\mathbf{w}(\mathbf{G}, p) \geq \frac{1}{2}$.

This also implies that ${}_g\text{CSP}_{no}^c(\text{TXOR})_{c,s}$ and ${}_g\text{CSP}_{no}^q(\text{TXOR})_{c,s}$ with $c \leq \frac{1}{2}$ are trivial, as every instance is a yes instance.

Proof. We define a classical correlation by sampling a random deterministic correlation $p(a, b|x, y) = \delta_{a,g(x)}\delta_{b,h(y)}$ for g and h uniformly random functions. Then, the bias is ≥ 0 , giving that the value of the strategy is $\geq \frac{1}{2}$. ■

Lemma 6.5.2. Let $\gamma : [1/2, 1] \rightarrow [1/2, 1]$ be a monotone increasing function such that for all XOR games $\mathbf{G}$, $\mathfrak{w}_q(\mathbf{G}) \leq \gamma(\mathfrak{w}_c(\mathbf{G}))$. Then, if $c > \gamma(s)$, there exist polynomial-time algorithms to decide ${}_g\text{CSP}_{no}^q(\text{TXOR})_{c,s}$ and ${}_g\text{CSP}_{no}^c(\text{TXOR})_{c,s}$.

We need only consider the γ on the domain $[1/2, 1]$ as $\mathfrak{w}_c(\mathbf{G}) \geq \frac{1}{2}$ by the previous lemma.

Proof. Consider the following algorithm:

1. On input $\mathbf{G}$ compute $\mathfrak{w}_q(\mathbf{G}_{XOR})$ to precision $\frac{1}{2}(c - \gamma(s))$.
2. If $\mathfrak{w}_q(\mathbf{G}_{XOR}) \geq c$ (within the bounds of the precision), output YES and else output NO.

Since $\mathfrak{w}_q(\mathbf{G}_{XOR})$ can be approximated by an SDP, this algorithm is polynomial-time in the instance size and $\log(1/(c - \gamma(s)))$. Now, if $\mathfrak{w}_q(\mathbf{G}) \geq c$, then $\mathfrak{w}_q(\mathbf{G}_{XOR}) \geq \mathfrak{w}_q(\mathbf{G}) \geq c$, so the algorithm outputs YES correctly. On the other hand, if $\mathfrak{w}_q(\mathbf{G}) < s$, then, $\mathfrak{w}_c(\mathbf{G}_{XOR}) = \mathfrak{w}_c(\mathbf{G}) < s$, so using the assumption on γ , $\mathfrak{w}_q(\mathbf{G}_{XOR}) < \gamma(s) < c$. As such, the algorithm outputs NO correctly.


The proof for the classical case proceeds identically. ■

The next theorem provides an explicit form of the function γ needed above.

Theorem 6.5.3 ([41]). There exists a monotone increasing function $\gamma : [1/2, 1] \rightarrow [1/2, 1]$ such that $\mathfrak{w}_q(\mathbf{G}) \leq \gamma(\mathfrak{w}_c(\mathbf{G}))$ for all XOR games $\mathbf{G}$, where

$$\gamma(x) = \begin{cases} \frac{1}{2} + \kappa\left(x - \frac{1}{2}\right) & \frac{1}{2} \leq x \leq \gamma_0 \\ \gamma_1 x & \gamma_0 < x \leq \gamma_2 \\ \sin^2\left(\frac{\pi}{2}x\right) & \gamma_2 < x \leq 1, \end{cases}$$

where $\kappa \approx 1.7822$ is an upper bound on Grothendieck's constant, $\gamma_0 = \frac{\kappa-1}{2(\kappa-\gamma_1)} \approx 0.60730$ is the point where the two line segments intersect, and $\gamma_1 \approx 1.1382$ and $\gamma_2 \approx 0.74202$ are such that $\gamma_1 x$ is tangent to $\sin^2\left(\frac{\pi}{2}x\right)$ at $0 < \gamma_2 < 1$.



(a) Hardness of ${}_g\text{CSP}_{no}^c(\text{TXOR})_{c,s}$, following from Lemmata 6.5.1, 6.5.2 and 6.5.4 and Theorem 6.5.3.

(b) Hardness of ${}_g\text{CSP}_{no}^q(\text{TXOR})_{c,s}$, following from Lemmata 6.5.1, 6.5.2 and 6.5.4 and Theorems 6.3.2, 6.4.4 and 6.5.3.

Figure 6.2: Hardness of tilted XOR games. The green region indicates problems solvable in polynomial time, the red and purple regions indicate hard instances, the white region indicates problems where the complexity is unknown, and the gray region indicates $s > c$ where the problem makes no sense. The RE-complete region is not to scale (though it is the interior of a triangle with positive area).

On the other hand, the hardness results Theorems 6.3.2 and 6.4.4 give rise to a set of completeness and soundness parameters where the games are hard to approximate.

Lemma 6.5.4. For any $c \geq s$, the decision problem ${}_g\text{CSP}_{no}^q(\text{TXOR})_{c,s}$ reduces to the problem ${}_g\text{CSP}_{no}^q(\text{TXOR})_{c',s'}$ for any (c', s') that is a convex combination of (c, s) , $(1, 1)$ and $(1/2, 1/2)$, as long as the weight of (c, s) is nonzero. The same holds for the classical problems.

Proof. The first reduction follows simply by taking a convex combination with the XOR games with winning probability 1 and $\frac{1}{2}$. ■

The results of this section are illustrated in Fig. 6.2.

Chapter 7

Approximation algorithms for entangled CSPs

In this chapter, we develop approximation algorithms for entangled CSPs. In [Section 7.1](#), we introduce the optimisation problems we are interested in approximating. In [Section 7.2](#), we describe our approximation algorithm. In [Section 7.3](#), we analyse the algorithm and find its approximation ratio. In [Section 7.4](#), we study a variant of the algorithm which guarantees finite-dimensional strategies.

7.1 Graph colouring and Max-Cut problems

As noted in [Section 2.11.2](#), the CSP of a σ -structure A can be expressed as an optimisation problem as follows. Given a σ -structure X , maximise the number of constraints that may be satisfied by a function $f : \mathfrak{D}(X) \rightarrow \mathfrak{D}(A)$.

Definition 7.1.1. The optimisation problem corresponding to the relational structure K_k (the complete graph on k vertices) is called Max- k -Cut.

An instance of of Max- k -Cut is given by a graph G , and the optimal value of G is the value of the nonlocal game $\mathsf{G}_{no}(G, K_k, \mathfrak{u})$ under symmetric classical strategies, times the number of edges.

The name Max- k -Cut arises from the fact that the optimisation corresponds to maximising the number of edges cut when splitting the graph into k parts.

Definition 7.1.2. The **entangled Max- k -Cut problem** is the problem of, given a graph G , maximising the value of $\mathbf{G}_{no}(G, K_k, \mathfrak{u})$ (times the number of edges) under symmetric tracial quantum strategies. This problem is denoted $\text{Max-}k\text{-Cut}^q$.

$\text{Max-}k\text{-Cut}^q$ is the optimisation problem corresponding to the problem $\text{CSP}_{no}^q(K_k)_{c,s}$. The objective function for $\text{Max-}k\text{-Cut}$ has a useful expression in terms of observables.

Lemma 7.1.3. Given a graph G , $\text{Max-}k\text{-Cut}$ for this graph is the problem

$$\begin{aligned} & \text{maximise} \quad \sum_{\{i,j\} \in E(G)} \left(1 - \frac{1}{k} \sum_{s=0}^{k-1} x_i^{-s} x_j^s \right) \\ & \text{subject to } x_i^k = 1, \ x_i \in \mathbb{C} \ \forall i \in V(G). \end{aligned}$$

$\text{Max-}k\text{-Cut}^q$ for G is the problem

$$\begin{aligned} & \text{maximise} \quad \sum_{\{i,j\} \in E(G)} \left(1 - \frac{1}{k} \sum_{s=0}^{k-1} \langle X_i^s, X_j^s \rangle_{\text{tr}} \right) \\ & \text{subject to } X_i^* X_i = X_i^k = 1, \ X_i \in \mathcal{B}(\mathbb{C}^d) \ \forall i \in V(G), \ d \in \mathbb{N}, \end{aligned}$$

where $\langle \cdot, \cdot \rangle_{\text{tr}}$ is the inner product induced by the normalised trace.

Note that for simplicity our presentation here does not consider different weights on the edges, but it can account for that by having multiple edges between pairs of vertices.

Proof. Given a PVM $\{P_s^i\}_{s=0}^{k-1}$, we take its observable to be $X_i = P_0^i + \omega P_1^i + \dots + \omega^{k-1} P_{k-1}^i$, where ω is a primitive k -th root of unity. We have $P_s^i = \frac{1}{k} \sum_{t=0}^{k-1} \omega^{-st} X_i^t$. Then, the term corresponding to colouring edge $\{i, j\}$ correctly can be expanded in terms of the observables as

$$\sum_{s \neq s'} \text{tr}(P_s^i P_{s'}^j) = 1 - \sum_{s=0}^{k-1} \text{tr}(P_s^i P_s^j) = 1 - \frac{1}{k^2} \sum_{s,t,t'=0}^{k-1} \omega^{-s(t+t')} \text{tr}(X_i^t X_j^{t'}) = 1 - \frac{1}{k} \sum_{t=0}^{k-1} \text{tr}(X_i^{-t} X_j^t),$$

which gives immediately the wanted expression. ■

To finish this section, we note a surprising distinction between the classical and quantum versions of $\text{Max-}k\text{-Cut}$. For any instance G , the corresponding Max-2-Cut^q optimisation problem is an XOR game, whose value is easy to approximate (though as discussed in [Chapter 6](#), the classical hardness can be somewhat reinstated in the entangled setting).

That gives an approximation algorithm with approximation ratio 1, in stark contrast with classical setting, where the optimal approximation ratio is the Goemans-Williamson constant ≈ 0.878 , and this is tight assuming the unique games conjecture [80].

In this chapter, we study an approximation algorithm for Max- k -Cut^q. In Table 7.1, we compare the approximation ratios of our algorithm to the best known approximation ratios for Max- k -Cut. Note that these values may be compared, as every classical strategy for a nonlocal game is also a quantum strategy, and hence every approximation algorithm for Max- k -Cut is also an approximation algorithm for Max- k -Cut^q.

Max- k -Cut	Classical	Entangled
$k = 2$	.878 [67]	1 [114, 41]
$k = 3$	.836 [66]	.8649
$k = 4$	.857 [54]	.8642
$k = 5$	.876 [54]	.8746
$k = 10$	.926 [54]	.9195

Table 7.1: Approximation ratios for classical and entangled Max- k -Cut for various k . The unattributed values are from our work.

Max- k -Cut is a special case of optimisation problems consisting of linear equality and inequality relations in two variables.

Definition 7.1.4. Max-Lin(k) is the optimisation problem arising from the CSP of the relational structure with domain $\mathbb{Z}_k$, and relations $R_{=,c} = \{(x, y) \in \mathbb{Z}_k^2 \mid x - y = c\}$ and $R_{\neq,c} = \{(x, y) \in \mathbb{Z}_k^2 \mid x - y \neq c\}$. The entangled analogue is Max-Lin(k)^q.

Max- k -Cut corresponds to the special case of instances where all but the relation $R_{\neq,0}$ is empty. We will also apply our approximation algorithm to **homogeneous** Max-Lin(k), or HMax-Lin(k), which is the special case where only the $R_{=,0}$ and $R_{\neq,0}$ are nonempty, *i.e.* where the linear relations are all homogeneous.

Lemma 7.1.5. HMax-Lin(k) for an instance H with edge sets $E_ =$ and $E_{\neq}$ corresponding to relations $R_{=,0}^H$ and $R_{\neq,0}^H$, respectively, is the problem

$$\begin{aligned}
& \text{maximise} \quad \sum_{\{i,j\} \in E_ =} \frac{1}{k} \sum_{s=0}^{k-1} x_i^{-s} x_j^s + \sum_{\{i,j\} \in E_{\neq}} \left(1 - \frac{1}{k} \sum_{s=0}^{k-1} x_i^{-s} x_j^s \right) \\
& \text{subject to} \quad x_i^k = 1 \text{ and } x_i \in \mathbb{C} \ \forall i \in V(G).
\end{aligned}$$

HMax-Lin(k)^q for this instance is the problem

$$\begin{aligned} & \text{maximise} \quad \sum_{\{i,j\} \in E=} \frac{1}{k} \sum_{s=0}^{k-1} \langle X_i^s, X_j^s \rangle_{\text{tr}} + \sum_{\{i,j\} \in E_{\neq}} \left(1 - \frac{1}{k} \sum_{s=0}^{k-1} \langle X_i^s, X_j^s \rangle_{\text{tr}} \right) \\ & \text{subject to} \quad X_i^* X_i = X_i^k = 1, \quad X_i \in \mathcal{B}(\mathbb{C}^d) \quad \forall i \in V(G), \quad d \in \mathbb{N}. \end{aligned}$$

The proof is identical to the proof of [Lemma 7.1.3](#). Applying our algorithm to HMax-Lin(k) gives the approximation ratios in [Table 7.2](#).

HMax-Lin(k)	Classical	Entangled
$k = 2$	.878 [67]	1 [114]
$k = 3$	.793 [66]	.864
$k = 4$	.708	.813
$k = 5$	.629	.738
$k = 10$	.383	.475

Table 7.2: Approximation ratios for classical and entangled HMax-Lin(k) for various k . The unattributed classical values are by applying the algorithm for Max- k -Cut to HMax-Lin(k), and the unattributed entangled values are from our work.

7.2 An approximation algorithm

In this section, we present our approximation algorithm for Max- k -Cut^q. It is modelled on successful approximation algorithms for the classical analogues of the problem: an SDP relaxation followed by a rounding step. In our case, the rounding takes two steps. The first is a vector-to-unitary construction to express the SDP solution in terms of unitaries, and the second is a unitary rounding to turn the unitaries into observables. These three parts are presented in [Sections 7.2.1 to 7.2.3](#), respectively, and the algorithm is presented in [Section 7.2.4](#)

7.2.1 SDP relaxation

In this section, we construct an SDP relaxation of Max- k -Cut^q. Since every classical strategy for a nonlocal game is also a quantum strategy, Max- k -Cut^q is a relaxation of Max- k -Cut. As such any SDP relaxation of Max- k -Cut^q is an SDP relaxation of Max- k -Cut.

The SDP we consider the same as the relaxation studied in the classical setting in [66, 94]. This gives best known approximation algorithm for Max-3-Cut, but not for $k > 3$. The improvement for $k \geq 4$ is given by an SDP relaxation that seems harder to apply to the quantum setting [54].

Lemma 7.2.1. The following SDP is given a relaxation of both Max- k -Cut and Max- k -Cut^q for an instance G .

$$\begin{aligned} & \text{maximise } \frac{k-1}{k} \sum_{\{i,j\} \in E(G)} (1 - X_{i,j}) \\ & \text{subject to } X_{i,i} = 1 \ \forall i \in V(G) \\ & \quad X_{i,j} \geq -\frac{1}{k-1} \ \forall i, j \in V(G) \\ & \quad X \in \mathcal{B}(\mathbb{R}^{V(G)}), X \geq 0. \end{aligned}$$

Proof. To show that is SDP is in fact a relaxation of Max- k -Cut^q, it suffices to see that every feasible point of Max- k -Cut^q gives rise to a feasible point of the SDP with the same value. Let $\{X_i\}_{i \in V(G)}$ be a feasible point, *i.e.* a collection of order- k unitaries. Then, we take $X_{i,j} = \frac{1}{k-1} \sum_{s=1}^{k-1} \langle X_i^s, X_j^s \rangle_{\text{tr}}$. First, we know that X is positive semidefinite as it is the Gram matrix of the vectors $(X_i, X_i^2, \dots, X_i^{k-1})$ for $i \in V(G)$. Next, $X_{i,i} = \frac{1}{k-1} \sum_{s=1}^{k-1} \text{tr}(X_i^{-s} X_i^s) = 1$. We know from the proof of Lemma 7.1.3 that $\sum_{s=0}^{k-1} \langle X_i^s, X_j^s \rangle_{\text{tr}} \geq 0$, and hence $X_{i,j}$ is real and $X_{i,j} \geq -\frac{1}{k-1}$ for all i, j . As such, X is a feasible point of the SDP. To finish, we see that X has the same value as $\{X_i\}_i$:

$$\begin{aligned} \frac{k-1}{k} \sum_{\{i,j\} \in E(G)} (1 - X_{i,j}) &= \frac{k-1}{k} \sum_{\{i,j\} \in E(G)} \left(1 - \frac{1}{k-1} \sum_{s=1}^{k-1} \langle X_i^s, X_j^s \rangle_{\text{tr}} \right) \\ &= \sum_{\{i,j\} \in E(G)} \left(1 - \frac{1}{k} \sum_{s=0}^{k-1} \langle X_i^s, X_j^s \rangle_{\text{tr}} \right). \quad \blacksquare \end{aligned}$$

In an identical way, we get an SDP relaxation of HMax-Lin(k).

Lemma 7.2.2. The following SDP is given a relaxation of both HMax-Lin(k) and

HMax-Lin(k) ^{q} for an instance H .

$$\begin{aligned}
& \text{maximise } \frac{1}{k} \sum_{\{i,j\} \in E=} (1 + (k-1)X_{i,j}) + \frac{k-1}{k} \sum_{\{i,j\} \in E\neq} (1 - X_{i,j}) \\
& \text{subject to } X_{i,i} = 1 \ \forall i \in V(G) \\
& \quad X_{i,j} \geq -\frac{1}{k-1} \ \forall i, j \in V(G) \\
& \quad X \in \mathcal{B}(\mathbb{R}^{V(G)}), \ X \geq 0.
\end{aligned}$$

7.2.2 Vector-to-unitary construction

A vector-to-unitary construction is a mapping from vectors to matrices that sends unit vectors to unitaries. We make use of this to convert the solution X of the SDP relaxation to a feasible point of Max- k -Cut ^{q} , by seeing X as the Gram matrix of some unit vectors and mapping those to unitaries. For now, we will not worry about the order of the unitaries — this will be achieved by the rounding step in the next section.

We make use of the vector-to-unitary construction used by Tsirelson [114] to show that every vect strategy for an XOR nonlocal game, which can be seen as a feasible point of an SDP relaxation, gives rise to a quantum strategy

Definition 7.2.3. A set of d **Weyl-Brauer operators** is a set of hermitian unitaries $\sigma_1, \dots, \sigma_d$ satisfying the anticommutation relations $\sigma_i \sigma_j = -\sigma_j \sigma_i$ for all $i \neq j$.

A set of d Weyl-Brauer operators admits a representation in dimension $2^{O(d)}$ by taking

$$\begin{aligned}
\sigma_1 &= X \otimes I \otimes I \otimes I \otimes \dots \otimes I \\
\sigma_2 &= Y \otimes I \otimes I \otimes I \otimes \dots \otimes I \\
\sigma_3 &= Z \otimes X \otimes I \otimes I \otimes \dots \otimes I \\
\sigma_4 &= Z \otimes Y \otimes I \otimes I \otimes \dots \otimes I \\
\sigma_5 &= Z \otimes Z \otimes X \otimes I \otimes \dots \otimes I \\
&\vdots
\end{aligned}$$

where X, Y, Z are the Pauli matrices. In fact, any set of d Weyl-Brauer operators is unitarily equivalent to this set.

Definition 7.2.4. The **Weyl-Brauer vector-to-unitary construction** is the linear map $U_{WB} : \mathbb{R}^d \rightarrow \mathbb{M}_{2O(d)}$ defined as $\mathbf{v} \mapsto v_1\sigma_1 + v_2\sigma_2 + \dots + v_d\sigma_d$, where $\sigma_1, \dots, \sigma_d$ is a set of d Weyl-Brauer operators.

Lemma 7.2.5 ([114]). U_{WB} is a vector-to-unitary construction into the hermitian matrices satisfying

$$\langle U_{WB}(\mathbf{v}), U_{WB}(\mathbf{w}) \rangle_{\text{tr}} = \langle \mathbf{v}, \mathbf{w} \rangle.$$

Proof. Since $U_{WB}(\mathbf{v})$ is trivially hermitian for real $\mathbf{v}$, and for unit vector $\mathbf{v}$, it is unitary as

$$U_{WB}(\mathbf{v})^2 = \sum_{i,j=1}^d v_i v_j \sigma_i \sigma_j = \sum_{i=1}^d v_i^2 \sigma_i^2 + \sum_{i < j} (v_i v_j - v_j v_i) \sigma_i \sigma_j = \sum_{i=1}^d v_i^2 = 1.$$

Finally, noting that $\text{tr}(\sigma_i \sigma_j) = -\text{tr}(\sigma_j \sigma_i) = 0$ for $i \neq j$, we have that

$$\langle U_{WB}(\mathbf{v}), U_{WB}(\mathbf{w}) \rangle_{\text{tr}} = \sum_{i,j=1}^d v_i w_j \text{tr}(\sigma_i \sigma_j) = \sum_{i=1}^d v_i w_i = \langle \mathbf{v}, \mathbf{w} \rangle. \quad \blacksquare$$

As a consequence of the hardness of approximation of entangled k -colouring for $k > 2$, there cannot be an analogous vector-to-unitary construction U that maps real unit vectors to order- k unitaries satisfying $\langle \mathbf{v}, \mathbf{w} \rangle = \frac{1}{k-1} \sum_{s=1}^{k-1} \langle U(\mathbf{v})^s, U(\mathbf{w})^s \rangle_{\text{tr}}$. Since perfect approximation in this way is impossible, we need to make use of a rounding step.

7.2.3 Rounding unitaries

In the final step of our approximation algorithm, we round the unitaries constructed by the Weyl-Brauer vector-to-unitary construction to order- k unitaries. For simplicity, we round to the nearest order- k unitary. This unitary has a simple form, demonstrated by the following lemma.

Lemma 7.2.6. Let $\Omega_k : S^1 \rightarrow S^1$ be the function that maps the input to the nearest k -th root of unity (with respect to complex modulus). Given a unitary $U \in \mathcal{U}_d$, the nearest order- k unitary with respect to the little Frobenius norm $\|\cdot\|_{\text{tr}}$ is $\Omega_k(U)$.

This is special case of the Hoffman-Wielandt theorem [74].

Proof. Write W for a closest order- k unitary to U . Let $U = \sum_{\lambda} \lambda P_{\lambda}$ and $W = \sum_{i=1}^k \omega_k^i \Pi_i$ be the spectral decompositions. By unitarity, $\|U - W\|_{\text{tr}}^2 = 2 - 2 \text{Re}\langle U, W \rangle_{\text{tr}}$, and then expanding

$$\begin{aligned} \text{Re}\langle U, W \rangle_{\text{tr}} &= \sum_{\lambda, i} \text{Re}(\lambda^* \omega_k^i) \text{tr}(P_{\lambda} \Pi_i) \\ &\leq \sum_{\lambda} \max_i \text{Re}(\lambda^* \omega_k^i) \text{tr}(P_{\lambda}), \end{aligned}$$

since $\sum_i \text{tr}(P_{\lambda} \Pi_i) = \text{tr}(P_{\lambda})$. On the other hand, we have that $\text{Re}(\lambda^* \Omega_k(\lambda)) = \max_i \text{Re}(\lambda^* \omega_k^i)$, so since $\Omega_k(U) = \sum_{\lambda} \Omega_k(\lambda) P_{\lambda}$,

$$\text{Re}\langle U, \Omega_k(U) \rangle = \sum_{\lambda} \text{Re}(\lambda^* \Omega_k(\lambda)) \text{tr}(P_{\lambda}) = \sum_{\lambda} \max_i \text{Re}(\lambda^* \omega_k^i) \text{tr}(P_{\lambda}).$$

Hence, $\text{Re}\langle U, W \rangle_{\text{tr}} \leq \text{Re}\langle U, \Omega_k(U) \rangle$, so we may take $W = \Omega_k(U)$. ■

7.2.4 The algorithm

We present the approximation algorithm for Max- k -Cut ^{q} .

- 1 Solve the SDP relaxation of [Lemma 7.2.1](#) to get X .
- 2 Extract unit vectors $\mathbf{x}_i \in \mathbb{R}^{V(G)}$ such that $X_{i,j} = \langle \mathbf{x}_i, \mathbf{x}_j \rangle$.
- 3 Run the Weyl-Brauer vector-to-unitary construction to obtain unitary operators $X_i = U_{WB}(\mathbf{x}_i) \in \mathcal{U}_M$ for some M .
- 4 Sample a unitary U from the Haar measure on $\mathcal{U}_{MN}$.
- 5 Let $\tilde{X}_i = \Omega_k(U(X_i \otimes I_N))$ be the closest order- k unitary to $U(X_i \otimes I_N)$.
- 6 Output the set of operators $\tilde{X}_i$ as the solution.

Algorithm 1: Approximation algorithm for Max- k -Cut ^{q}

We note two important aspects of the rounding that were not mentioned in [Section 7.2.3](#). First, we multiply the matrices constructed from the vector-to-unitary construction by a Haar-random unitary. This serves to distribute the eigenvalues of the constructed unitary evenly while preserving the inner products. Also, we increase the dimension of the unitary before applying the Haar-random unitary and rounding. This step is important to allow analysis of the algorithm based on free probability, which holds in the limit of infinite dimension.

We can also apply the same algorithm to $\text{HMax-Lin}(k)^q$.

- 1 Solve the SDP relaxation of [Lemma 7.2.2](#) to get X .
- 2 Extract unit vectors $\mathbf{x}_i \in \mathbb{R}^{V(G)}$ such that $X_{i,j} = \langle \mathbf{x}_i, \mathbf{x}_j \rangle$.
- 3 Run the Weyl-Brauer vector-to-unitary construction to obtain unitary operators $X_i = U_{WB}(\mathbf{x}_i) \in \mathcal{U}_M$ for some M .
- 4 Sample a unitary U from the Haar measure on $\mathcal{U}_{MN}$.
- 5 Let $\tilde{X}_i = \Omega_k(U(X_i \otimes I_N))$ be the closest order- k unitary to $U(X_i \otimes I_N)$.
- 6 Output the set of operators $\tilde{X}_i$ as the solution.

Algorithm 2: Approximation algorithm for $\text{HMax-Lin}(k)^q$

7.3 Analysis of the algorithm

In this section, we analyse the algorithm described in the previous section. In [Section 7.3.1](#), we define the notion of a relative distribution, which underlies our technique. In [Section 7.3.2](#), we relate the relative distribution to the approximation ratio via an integral formula. In [Section 7.3.3](#), we apply the results of the previous two sections to our approximation algorithm. In [Section 7.3.4](#), we prove a technical lemma needed in the previous section.

7.3.1 Relative distribution

First, we formally define the relative distribution.

Definition 7.3.1. Let $A, B \in \mathcal{U}_d$. Let $A = \sum_s e^{i\phi_s} \Pi_{A,s}$ and $B = \sum_t e^{i\psi_t} \Pi_{B,t}$ be the spectral decompositions of A and B . Define the **weight measure** $w_{A,B} : \mathcal{B}([0, 2\pi)) \rightarrow \mathbb{R}_{\geq 0}$ as

$$w_{A,B}(E) = \sum_{s,t: \psi_t - \phi_s \in E} \langle \Pi_{A,s}, \Pi_{B,t} \rangle_{\text{tr}}. \quad (7.1)$$

Since the phases ψ_t and ϕ_s are considered as angles, the expression $\psi_t - \phi_s \in E$ is modulo 2π . It follows immediately from the definition that $w_{A,B}$ is a well-defined measure, since it can be simply expressed as positive linear combination of the Dirac delta measures on $[0, 2\pi)$ (Eq. [2.2](#))

$$w_{A,B} = \sum_{s,t} \langle \Pi_{A,s}, \Pi_{B,t} \rangle_{\text{tr}} \delta_{\psi_t - \phi_s}.$$

Also, we can see the weight measure is a probability distribution since $w_{A,B}([0, 2\pi)) = \sum_{s,t} \langle \Pi_{A,s}, \Pi_{B,t} \rangle = 1$.

Definition 7.3.2. Let $A, B \in \mathcal{U}_d$ and $U \sim \eta_{\mathcal{U}_d}$ the Haar measure. The **relative distribution** of (A, B) is the probability measure $\Delta_{A,B} : \mathcal{B}([0, 2\pi)) \rightarrow \mathbb{R}_{\geq 0}$ such that

$$\Delta_{A,B}(E) = \int_{\mathcal{U}_d} w_{UA,UB}(E) dU. \quad (7.2)$$

Lemma 7.3.3. The relative distribution is a well-defined probability distribution.

Proof. First, to see that $\Delta_{A,B}(E)$ is well-defined for all E , we need that $U \mapsto w_{UA,UB}(E)$ is a measurable function. In fact, it is the composition of the multiplication $U \mapsto (UA, UB)$, which is continuous; the diagonalisation $(X, Y) \mapsto (\{\phi_s\}, \{\psi_t\}, \{\Pi_{X,s}\}, \{\Pi_{Y,t}\})$ where $X = \sum_s e^{i\phi_s} \Pi_{X,s}$ and $Y = \sum_t e^{i\psi_t} \Pi_{Y,t}$, which is continuous almost everywhere; and the map $(\{\phi_s\}, \{\psi_t\}, \{\Pi_{X,s}\}, \{\Pi_{Y,t}\}) \mapsto \sum_{s,t: \phi_s - \psi_t \in E} \langle \Pi_{X,s}, \Pi_{Y,t} \rangle_{\text{tr}}$, which is measurable. Therefore, the Haar integral $\int_{\mathcal{U}_d} w_{UA,UB}(E) dU$ is well-defined.

Next, since $w_{UA,UB}(E) \geq 0$ for all U , $\Delta_{A,B}(E) \geq 0$. Also, $\Delta_{A,B}$ is countably additive by monotone convergence theorem. As such, $\Delta_{A,B}$ is also a measure. Finally, $\Delta_{A,B}([0, 2\pi)) = \int_{\mathcal{U}_d} w_{UA,UB}([0, 2\pi)) dU = 1$, so it is normalised, giving a distribution. $\blacksquare$

Although $w_{A,B}$ is not absolutely continuous with respect to the Lebesgue measure Λ on $[0, 2\pi)$, the relative distribution $\Delta_{A,B}$ in fact is. For any measurable subset $E \subseteq [0, 2\pi)$, if E has measure 0, the subset of $\mathcal{U}_d$ whose eigenvalues are in E must also have measure 0. As such $w_{UA,UB}(E) \neq 0$ with probability 0, so $\Delta_{A,B}(E) = 0$. Therefore the Radon-Nikodym derivative $p_{\Delta_{A,B}} = \frac{d\Delta_{A,B}}{d\Lambda}$ exists — this is the PDF of the relative distribution.

Lemma 7.3.4. The characteristic function of the relative distribution is

$$\chi_{\Delta_{A,B}}(n) = \int_{\mathcal{U}_d} \text{tr}[U^{-n}(UA^*B)^n] dU.$$

Proof. First, we find the characteristic function of the weight measure:

$$\begin{aligned} \chi_{w_{A,B}}(n) &= \sum_{s,t} \langle \Pi_{A,s}, \Pi_{B,t} \rangle_{\text{tr}} \chi_{\delta_{\psi_t - \phi_s}}(n) = \sum_{s,t} \langle \Pi_{A,s}, \Pi_{B,t} \rangle_{\text{tr}} e^{in(\psi_t - \phi_s)} \\ &= \left\langle \sum_s e^{in\phi_s} \Pi_{A,s}, \sum_t e^{in\psi_t} \Pi_{B,t} \right\rangle_{\text{tr}} = \langle A^n, B^n \rangle_{\text{tr}} = \text{tr}[A^{-n} B^n]. \end{aligned}$$

This immediately gives, by integrating over the Haar measure, that the characteristic function of the relative distribution is

$$\chi_{\Delta_{A,B}}(n) = \int \text{tr}[(UA)^{-n}(UB)^n]dU.$$

To finish, we can simplify the expression using Haar invariance. We replace U by UA^* to get $\chi_{\Delta_{A,B}}(n) = \int \text{tr}[U^{-n}(UA^*B)^n]dU$, as wanted. $\blacksquare$

In order to motivate the theorem below, we can compute some of the values of the characteristic function. First, note that $\chi_{\Delta_{A,B}}(-n) = \chi_{\Delta_{A,B}}(n)^*$ since $\Delta_{A,B}$ is real-valued. As such, we need only compute the values of the characteristic function for $n \geq 0$. Next, it's easy to calculate the characteristic function for small n . Letting $C = A^*B$ and $\lambda = \text{tr}(C) = \text{tr}(A^*B)$,

$$\begin{aligned}\chi_{\Delta_{A,B}}(0) &= \int \text{tr}[U^{-0}(UC)^0]dU = \int dU = 1 \\ \chi_{\Delta_{A,B}}(1) &= \int \text{tr}[U^*UC]dU = \int \text{tr}(C)dU = \lambda \\ \chi_{\Delta_{A,B}}(2) &= \int \text{tr}[(U^*)^2(UC)^2]dU = \text{tr}\left[C \int UCU^*dU\right] = \text{tr}[C \text{tr}(C)] = \lambda^2.\end{aligned}$$

It is also feasible, though far less direct, to calculate the characteristic function at $n = 3$. We get

$$\chi_{\Delta_{A,B}}(3) = \lambda^3 + \frac{\text{tr}(C^3) + \lambda^3 - 2 \text{tr}(C^2)/d}{d^2 - 1}.$$

We see here that, unlike for smaller n , the value is not simply λ^n . However, it does converge there as $d \rightarrow \infty$. On the other hand, the distribution with characteristic function λ^n for $n \geq 0$ and $(\lambda^*)^{-n}$ for $n < 0$ is exactly the wrapped Cauchy distribution $\Delta_\lambda = \mathcal{W}(\arg \lambda, -\ln |\lambda|)$ (see [Section 2.5.3](#)).

Theorem 7.3.5 (Cauchy Law). Let $A, B \in \mathcal{U}_d$ and let $\lambda = \text{tr}(A^*B)$. In the limit of $m \rightarrow \infty$, the relative distribution $\Delta_{A \otimes I_m, B \otimes I_m}$ converges in distribution to Δ_λ .

The proof makes use of techniques from free probability ([Section 2.8](#)).

Proof. Let $A_m = A^*B \otimes I_m$, and let U_m be a dm -dimensional Haar random unitary. Then, by [Theorem 2.8.1](#), we get $U_m, A_m \rightarrow u, a$ ($*$ -dist) where u and a are free. As such, for every $n \geq 0$

$$\chi_{\Delta_{A \otimes I_m, B \otimes I_m}}(n) \rightarrow \tau(u^{-n}(ua)^n),$$




Figure 7.1: PDF of the wrapped Cauchy distribution Δ_λ for a selection of values of λ .

as $m \rightarrow \infty$. This is a mixed $*$ -moment of u, a , so we can use freeness to compute it in terms of the $*$ -moments of u and a . To rewrite this in terms of free variables of expectation 0, note that $\tau(a) = \lim_{m \rightarrow \infty} \text{tr}(A_m) = \lambda$. Then

$$\tau((ua)^n u^{-n}) = \tau((u(a - \lambda + \lambda))^n u^{-n}) = \sum_{e_1, \dots, e_n \in \{0,1\}} \lambda^{n - \sum e_i} \tau(u(a - \lambda)^{e_1} \cdots u(a - \lambda)^{e_n} u^{-n})$$

If the e_i are not all 0, then the argument of τ is a product of free variables each with expectation 0, so that term does not contribute. That is, we get

$$\begin{aligned} \tau((ua)^n u^{-n}) &= \lambda^n \tau(1) + \sum_{k=1}^n \lambda^{n-k} \sum_{n_1, \dots, n_k > 0, \sum n_k \leq n} \tau(u^{n_1}(a - \lambda) \cdots u^{n_k}(a - \lambda) u^{-\sum n_k}) \\ &= \lambda^n + \sum_{k=1}^n \lambda^{n-k} \sum_{n_1, \dots, n_k > 0, \sum n_k \leq n} 0 = \lambda^n, \end{aligned}$$

since $\tau(a - \lambda) = 0$ and $\tau(u^n) = 0$ for all $n \neq 0$. This gives exactly the characteristic function of the wrapped Cauchy distribution. Using Theorem 2.5.1, $\Delta_{A \otimes I_m, B \otimes I_m}$ converges in distribution to Δ_λ . $\blacksquare$

7.3.2 Integral formula

In this section, we make use of the relative distribution introduced in the previous section to calculate the expected value of polynomials of rounded operators. We show that these expected values can be expressed as an integral of a function, which we call the **fidelity** function, depending only on the polynomial, with respect to the relative distribution (Definition 7.3.2).

In the discussion below, for every k we work with noncommutative $*$ -polynomials of the form $P(x, y) = \sum_{s,t=0}^{k-1} c_{s,t}(x^*)^s y^t$. In our applications, the variables x, y will always be substituted with order- k unitaries. Because of this, we also assume that the indices s, t in $c_{s,t}$ are considered modulo k , *i.e.* we let $c_{s+ak, t+bk} = c_{s,t}$ for all integers a, b .

Definition 7.3.6. Let $k \in \mathbb{N}$. For a noncommutative $*$ -polynomial of the form $P(x, y) = \sum_{s,t=0}^{k-1} c_{s,t}(x^*)^s y^t$, the **fidelity function** is the map $\text{fid}_{k,P} : [0, 2\pi) \rightarrow \mathbb{C}$ defined as

$$\text{fid}_{k,P}(\theta) = \frac{1}{2\pi} \int_0^{2\pi} \sum_{s,t=0}^{k-1} c_{s,t} \Omega_k(e^{i\varphi})^{-s} \Omega_k(e^{i(\varphi+\theta)})^t d\varphi, \quad (7.3)$$

where Ω_k is the rounding to the closest k -th root of unity (Lemma 7.2.6), and the integral is with respect to the Lebesgue measure over $[0, 2\pi)$.

When k is clear from the context we let $\text{fid}_P := \text{fid}_{k,P}$.

Theorem 7.3.7. Let $A, B \in \mathcal{U}_d$ be two unitaries. For any noncommutative $*$ -polynomial $P(x, y) = \sum_{s,t=0}^{k-1} c_{s,t}(x^*)^s y^t$,

$$\int P(\Omega_k(UA), \Omega_k(UB)) dU = \int \text{fid}_P(\theta) d\Delta_{A,B}(\theta). \quad (7.4)$$

Proof. Let $UA = \sum_i \alpha_{U,i} P_{U,i}$ and $UB = \sum_j \beta_{U,j} Q_{U,j}$ be the spectral decompositions. This gives that

$$\begin{aligned} \int P(\Omega_k(UA), \Omega_k(UB)) dU &= \int \sum_{s,t} c_{s,t} \langle \Omega_k(UA)^s, \Omega_k(UB)^t \rangle_{\text{tr}} dU \\ &= \sum_{i,j} \sum_{s,t} c_{s,t} \int \Omega_k(\alpha_{U,i})^{-s} \Omega_k(\beta_{U,j})^t \langle P_{U,i}, Q_{U,j} \rangle_{\text{tr}} dU \end{aligned}$$

Using Haar invariance, we can multiply U by any $\zeta \in S^1$, which sends (α_i, β_j) to $(\zeta\alpha_i, \zeta\beta_j)$. In particular, we can multiply by a randomly distributed phase, giving that

$$\begin{aligned} \int P(\Omega_k(UA), \Omega_k(UB))dU &= \frac{1}{2\pi} \int_0^{2\pi} \sum_{i,j,s,t} c_{s,t} \int \Omega_k(e^{i\varphi} \alpha_{U,i})^{-s} \Omega_k(e^{i\varphi} \beta_{U,j})^t \langle P_{U,i}, Q_{U,j} \rangle_{\text{tr}} dU d\varphi \\ &= \int \sum_{i,j,s,t} c_{s,t} \frac{1}{2\pi} \int_0^{2\pi} \Omega_k(e^{i\varphi})^{-s} \Omega_k(e^{i\varphi} \alpha_{U,i}^* \beta_{U,j})^t d\varphi \langle P_{U,i}, Q_{U,j} \rangle_{\text{tr}} dU \\ &= \int \sum_{i,j} \text{fid}_P(\arg(\alpha_{U,i}^* \beta_{U,j})) \langle P_{U,i}, Q_{U,j} \rangle_{\text{tr}} dU. \end{aligned}$$

Now, using the weight measure $w_{UA,UB}$ (Definition 7.3.1), we can view the sum over the eigenvalues as an integral with respect to this measure:

$$\sum_{i,j} \text{fid}_P(\arg(\alpha_{U,i}^* \beta_{U,j})) \langle P_{U,i}, Q_{U,j} \rangle_{\text{tr}} = \int \text{fid}_P(\theta) dw_{UA,UB}(\theta).$$

Finally, using the definition of the relative distribution (Definition 7.3.2), we get that

$$\int \int \text{fid}_P(\theta) dw_{UA,UB}(\theta) dU = \int \text{fid}_P(\theta) d\Delta_{X,Y}(\theta). \quad \blacksquare$$

Due to the results Section 7.3.1, we know that the relative distribution becomes a wrapped Cauchy distribution in the limit of large dimension. In this case, we can exactly compute the integral formula. First, we need a simpler expression of the fidelity function.

Lemma 7.3.8. The fidelity function defined in Definition 7.3.6 simplifies to

$$\text{fid}_P(\theta) = \sum_{t=0}^{k-1} c_{t,t} \omega^{t \lfloor \frac{k}{2\pi} \theta \rfloor} \left(1 + (\omega^t - 1) \left(\frac{k}{2\pi} \theta - \left\lfloor \frac{k}{2\pi} \theta \right\rfloor \right) \right). \quad (7.5)$$

Its Fourier coefficients are $\widehat{\text{fid}}_P(0) = 2\pi c_{0,0}$ and for $n \neq 0$,

$$\widehat{\text{fid}}_P(n) = \frac{2k^2}{\pi n^2} \sin^2\left(\frac{\pi n}{k}\right) c_{n,n}. \quad (7.6)$$

Proof. Consider the integral $f_{s,t}(\theta) = \int_0^{2\pi} \Omega_k(e^{i\varphi})^{-s} \Omega_k(e^{i(\varphi+\theta)})^t d\varphi$. This simplifies as

$$\begin{aligned} f_{s,t}(\theta) &= \sum_{r=0}^{k-1} \int_{-\frac{\pi}{k} + \frac{2\pi}{k}r}^{\frac{\pi}{k} + \frac{2\pi}{k}r} \Omega_k(e^{i\varphi})^{-s} \Omega_k(e^{i(\varphi+\theta)})^t d\varphi \\ &= \sum_{r=0}^{k-1} \omega^{(t-s)r} \int_{-\frac{\pi}{k}}^{\frac{\pi}{k}} \Omega_k(e^{i\varphi})^{-s} \Omega_k(e^{i(\varphi+\theta)})^t d\varphi \\ &= k \delta_{s,t} \int_{-\frac{\pi}{k}}^{\frac{\pi}{k}} \Omega_k(e^{i(\varphi+\theta)})^t d\varphi. \end{aligned}$$

Now, $\Omega_k(e^{i(\varphi+\theta)})^t$ is a piecewise constant function (*i.e.* a step function), with the value changing at $x_\theta = \frac{2\pi}{k} \lfloor \frac{k}{2\pi} \theta \rfloor - \theta + \frac{\pi}{k}$. More precisely, on the interval $\varphi \in (-\frac{\pi}{k}, x_\theta]$, $\Omega_k(e^{i(\varphi+\theta)})^t = \Omega_k(e^{i(\theta-\frac{\pi}{k})})^t = \omega^{t \lfloor \frac{k}{2\pi} \theta \rfloor}$ and on its complement $\varphi \in (x_\theta, \frac{\pi}{k}]$, $\Omega_k(e^{i(\varphi+\theta)})^t = \Omega_k(e^{i(\theta+\frac{\pi}{k})})^t = \omega^{t(\lfloor \frac{k}{2\pi} \theta \rfloor + 1)}$. Therefore, $f_{s,t}(\theta) = k\delta_{s,t}\omega^{t \lfloor \frac{k}{2\pi} \theta \rfloor} \left(x_\theta + \frac{\pi}{k} + \omega^t \left(\frac{\pi}{k} - x_\theta \right) \right)$. Putting the terms together,

$$\text{fid}_P(\theta) = \frac{1}{2\pi} \sum_{s,t=0}^{k-1} c_{s,t} f_{s,t}(\theta) = \frac{k}{2\pi} \sum_{t=0}^{k-1} c_{t,t} \omega^{t \lfloor \frac{k}{2\pi} \theta \rfloor} \left(\frac{2\pi}{k} + (\omega^t - 1) \left(\frac{\pi}{k} - x_\theta \right) \right).$$

To find the Fourier coefficients, consider the Fourier coefficients of $f_{t,t}$. They are

$$\hat{f}_{t,t}(0) = k \int_0^{2\pi} \int_{-\frac{\pi}{k}}^{\frac{\pi}{k}} \Omega_k(e^{i(\theta+\varphi)})^t d\varphi d\theta = 2k\pi\delta_{t,0} \int_{-\frac{\pi}{k}}^{\frac{\pi}{k}} d\varphi = 4\pi^2\delta_{t,0},$$

and for $n \neq 0$

$$\begin{aligned} \hat{f}_{t,t}(n) &= k \int_0^{2\pi} \int_{-\frac{\pi}{k}}^{\frac{\pi}{k}} \Omega_k(e^{i(\theta+\varphi)})^t e^{-in\theta} d\varphi d\theta \\ &= k \int_0^{2\pi} \int_{-\frac{\pi}{k}}^{\frac{\pi}{k}} \Omega_k(e^{i\theta})^t e^{-in(\theta-\varphi)} d\varphi d\theta \\ &= 2k \frac{\sin\left(\frac{\pi n}{k}\right)}{n} \int_0^{2\pi} \Omega_k(e^{i\theta})^t e^{-in\theta} d\theta \\ &= 2k \frac{\sin\left(\frac{\pi n}{k}\right)}{n} \int_{-\frac{\pi}{k}}^{\frac{\pi}{k}} e^{-in\theta} d\theta \sum_{r=0}^{k-1} \omega^{r(t-n)} \\ &= \begin{cases} 4k^2 \frac{\sin^2\left(\frac{\pi n}{k}\right)}{n^2}, & t = n \pmod{k}, \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

Plugging these into $\widehat{\text{fid}}_P(n) = \frac{1}{2\pi} \sum_{t=0}^{k-1} c_{t,t} \hat{f}_{t,t}(n)$ gives the wanted result. ■

Proposition 7.3.9. For $\lambda \in \mathbb{C}$ with $|\lambda| \leq 1$ and $P(x, y) = \sum_{s,t} c_{s,t} (x^*)^s y^t$ we have

$$\int \text{fid}_P(\theta) d\Delta_\lambda(\theta) = c_{0,0} + \frac{k^2}{\pi^2} \sum_{n=1}^{\infty} \frac{\sin^2\left(\frac{\pi n}{k}\right)}{n^2} (c_{-n,-n} \lambda^n + c_{n,n} (\lambda^*)^n). \quad (7.7)$$

Proof. By Parseval's theorem (Eq. 2.1),

$$\int \text{fid}_P(\theta) d\Delta_\lambda(\theta) = \frac{1}{2\pi} \sum_{n \in \mathbb{Z}} \widehat{\text{fid}}_P(-n) \chi_{\Delta_\lambda}(n).$$

The characteristic function of the wrapped Cauchy distribution Δ_λ is $\chi_{\Delta_\lambda}(n) = \lambda^n$ for $n \geq 0$ and $\chi_{\Delta_\lambda}(n) = (\lambda^*)^{-n}$ for $n < 0$. Using this and the previous lemma,

$$\int \text{fid}_P(\theta) d\Delta_\lambda(\theta) = c_{0,0} + \frac{1}{2\pi} \sum_{n=1}^{\infty} \frac{2k^2}{\pi n^2} \sin^2\left(\frac{\pi n}{k}\right) (c_{-n,-n} \lambda^n + c_{n,n} (\lambda^*)^n),$$

giving the result. ■

7.3.3 Application to Max- k -Cut

Since the objective function of Max- k -Cut is $\sum_{\{i,j\} \in E(G)} \left(1 - \frac{1}{k} \sum_{s=0}^{k-1} x_i^{-s} x_j^s\right)$, we can find the relevant fidelity function (Definition 7.3.6) by considering the noncommutative $*$ -polynomial $P_k(x, y) = \frac{1}{k} \sum_{s=0}^{k-1} (x^*)^s y^s$.

Corollary 7.3.10 (of Lemma 7.3.8). For $P_k(x, y) = \frac{1}{k} \sum_{s=0}^{k-1} (x^*)^s y^s$, the fidelity function $\text{fid}_k := \text{fid}_{k, P_k}$ is

$$\text{fid}_k(\theta) = \begin{cases} 1 - \frac{k}{2\pi} \theta & 0 \leq \theta < \frac{2\pi}{k} \\ 0 & \frac{2\pi}{k} \leq \theta < 2\pi - \frac{2\pi}{k} \\ 1 - \frac{k}{2\pi} (2\pi - \theta) & 2\pi - \frac{2\pi}{k} \leq \theta < 2\pi \end{cases} \quad (7.8)$$

Proof. First, $\text{fid}_k(\theta) = \frac{1}{k} \sum_{t=0}^{k-1} \omega^{t \lfloor \frac{k}{2\pi} \theta \rfloor} \left(1 + (\omega^t - 1) \left(\frac{k}{2\pi} \theta - \lfloor \frac{k}{2\pi} \theta \rfloor\right)\right)$ by Lemma 7.3.8. If $\lfloor \frac{k}{2\pi} \theta \rfloor \neq 0, k-1$, the terms all cancel and $\text{fid}_k(\theta) = 0$. Else, if $\lfloor \frac{k}{2\pi} \theta \rfloor = 0$, we have

$$\text{fid}_k(\theta) = \frac{1}{k} \sum_t \left(1 + (\omega^t - 1) \frac{k}{2\pi} \theta\right) = 1 - \frac{k}{2\pi} \theta.$$

Finally, if $\lfloor \frac{k}{2\pi} \theta \rfloor = k-1$, we have

$$\text{fid}_k(\theta) = \frac{1}{k} \sum_t \left(\omega^{-t} + (1 - \omega^{-t}) \left(\frac{k}{2\pi} \theta - (k-1)\right)\right) = \frac{k}{2\pi} \theta - k + 1$$

Putting these three cases together gives the result. ■

To make use of the fidelity, we use the integral formula of Theorem 7.3.7 to get that

$$\mathbb{E}(P_k(\tilde{X}_i, \tilde{X}_j)) = \int \text{fid}_k(\theta) d\Delta_{X_i \otimes I_N, X_j \otimes I_N}(\theta).$$




Figure 7.2: The fidelity function fid_k for selected values of k .

Due to [Theorem 7.3.5](#), the relative distribution $\Delta_{X_i \otimes I_N, X_j \otimes I_N}$ tends to a wrapped Cauchy distribution (in distribution). Further, since we increase the dimension of the X_i while preserving their $*$ -moments (*i.e.* inner products of the form $\langle X_i^s, X_j^s \rangle_{\text{tr}}$) by tensoring on the identity $X_i \mapsto X_i \otimes I_N$ for larger and larger N , we can assume that we are working in the limit of large dimension, and hence replace the relative distribution by the wrapped Cauchy distribution.

Corollary 7.3.11 (Of [Proposition 7.3.9](#)). For λ real we have

$$\int \text{fid}_k(\theta) d\Delta_\lambda(\theta) = \frac{1}{k} + \frac{2k}{\pi^2} \sum_{n=1}^{\infty} \frac{\sin^2(\pi n/k)}{n^2} \lambda^n = \frac{1}{k} + \frac{k}{\pi^2} \text{Re}[\text{Li}_2(\lambda) - \text{Li}_2(\lambda e^{i\frac{2\pi}{k}})], \quad (7.9)$$

where $\text{Li}_2(x) = \sum_{n=1}^{\infty} \frac{x^n}{n^2}$ is the dilogarithm.

Proof. With $c_{n,n} = \frac{1}{k}$, for all n , in [Proposition 7.3.9](#), the integral is expressed as the power series

$$\int \text{fid}_k(\theta) d\Delta_\lambda(\theta) = \frac{1}{k} + \frac{k^2}{\pi^2} \sum_{n=1}^{\infty} \frac{\sin^2\left(\frac{\pi n}{k}\right)}{n^2} \left(\frac{1}{k} \lambda^n + \frac{1}{k} (\lambda^*)^n \right) = \frac{1}{k} + \frac{2k}{\pi^2} \sum_{n=1}^{\infty} \frac{\sin^2\left(\frac{\pi n}{k}\right)}{n^2} \lambda^n.$$

The closed form in terms of dilogarithms can be found by expanding $\sin^2\left(\frac{\pi n}{k}\right) = \frac{1}{2} \text{Re}\left[1 - e^{i\frac{2\pi n}{k}}\right]$ in the series. ■

We need the following technical lemma.

Lemma 7.3.12. The function

$$\lambda \in [-1, 1] \mapsto \frac{1 - \frac{1}{k} - \frac{k}{\pi^2} \operatorname{Re} \left[\operatorname{Li}_2(\lambda) - \operatorname{Li}_2(\lambda e^{i\frac{2\pi}{k}}) \right]}{1 - \lambda} \quad (7.10)$$

is increasing for all $k \geq 3$.

The proof of the above is elementary but quite involved, so we defer it to [Section 7.3.4](#). Now, we can find the approximation ratio of our algorithm for Max- k -Cut^q.

Theorem 7.3.13. The approximation ratio of [Algorithm 1](#) in the case of Max- k -Cut^q for $k \geq 3$ tends to

$$1 - \frac{1}{k} - \frac{k}{\pi^2} \operatorname{Re} \left[\operatorname{Li}_2\left(-\frac{1}{k-1}\right) - \operatorname{Li}_2\left(-\frac{e^{i\frac{2\pi}{k}}}{k-1}\right) \right] \quad (7.11)$$

in the limit of large dimension N .

Computing this function gives the values presented in [Table 7.1](#). We can alternatively write this as the series

$$1 - \frac{1}{k} - \frac{2k}{\pi^2} \sum_{n=1}^{\infty} \frac{\sin^2(\pi n/k)}{n^2(1-k)^n}.$$

Note the rapid convergence of this series due to the exponential term $(1-k)^n$ in the denominator.

Proof. Let $\{\tilde{X}_i\}$ be the feasible solution of a Max- k -Cut^q instance G created in [Algorithm 1](#). To calculate the expected value of our solution, by linearity of the expectation, we just need to calculate $\mathbb{E}\left(\frac{1}{k} \sum_{s=0}^{k-1} \langle \tilde{X}_i^s, \tilde{X}_j^s \rangle_{\text{tr}}\right)$ for every i, j . By [Theorem 7.3.7](#) it holds that

$$\frac{1}{k} \mathbb{E}\left(\sum_{s=0}^{k-1} \langle \tilde{X}_i^s, \tilde{X}_j^s \rangle_{\text{tr}}\right) = \int \operatorname{fid}_k(\theta) d\Delta_{X_i, X_j}(\theta). \quad (7.12)$$

Therefore the contribution to the solution by edge $\{i, j\}$ is $1 - \int \operatorname{fid}_k(\theta) d\Delta_{X_i, X_j}(\theta)$. So the ratio of the value of the expected solution to the value of the SDP is at least the minimum of

$$\frac{k}{k-1} \frac{1 - \int \operatorname{fid}_k(\theta) d\Delta_{X_i, X_j}(\theta)}{1 - \langle X_i, X_j \rangle_{\text{tr}}}$$

over all edges $\{i, j\}$. Note that we can always exclude the edges for which $\langle X_i, X_j \rangle_{\text{tr}} = 1$ from this analysis because the contribution of such edges to the SDP value is zero and the numerator is always nonnegative.

Finally note that in all feasible solutions $X_{i,j} = \langle \mathbf{x}_i, \mathbf{x}_j \rangle$ to the SDP relaxation we have $\langle \mathbf{x}_i, \mathbf{x}_j \rangle \geq -\frac{1}{k-1}$. Thus in the worst case over all instances, the quantity above is never less than

$$\min_{\substack{\lambda \in [-\frac{1}{k-1}, 1) \\ X, Y: \langle X, Y \rangle_{\text{tr}} = \lambda}} \frac{k}{k-1} \frac{1 - \int \text{fid}_k(\theta) d\Delta_{X,Y}(\theta)}{1 - \lambda}$$

In the limit of large dimension, by [Theorem 7.3.5](#), the relative distribution tends to the wrapped Cauchy distribution. As such, using [Corollary 7.3.11](#), the minimization becomes

$$\min_{\lambda \in [-\frac{1}{k-1}, 1)} \frac{k}{(k-1)(1-\lambda)} \left(1 - \frac{1}{k} - \frac{k}{\pi^2} \text{Re} \left[\text{Li}_2(\lambda) - \text{Li}_2(\lambda e^{i\frac{2\pi}{k}}) \right] \right)$$

Finally, using [Lemma 7.3.12](#), the function we are minimizing over is increasing, and as such the minimum is attained at $\lambda = -\frac{1}{k-1}$. $\blacksquare$

The proof of the theorem generalizes to all homogeneous CSPs.

Theorem 7.3.14. The approximation ratio of [Algorithm 2](#) for homogeneous NC-Max-Lin(k) for $k \geq 3$, in the limit of large dimension, is the minimum of

$$1 - \frac{1}{k} - \frac{k}{\pi^2} \text{Re} \left[\text{Li}_2\left(-\frac{1}{k-1}\right) - \text{Li}_2\left(-\frac{e^{i\frac{2\pi}{k}}}{k-1}\right) \right]$$

and

$$\min_{\lambda \in [-\frac{1}{k-1}, 1)} \frac{1 + \frac{k^2}{\pi^2} \text{Re} \left[\text{Li}_2(\lambda) - \text{Li}_2(\lambda e^{i\frac{2\pi}{k}}) \right]}{1 + (k-1)\lambda},$$

over all $\lambda \in [-\frac{1}{k-1}, 1)$

The proof is the same as for [Theorem 7.3.13](#) except that, in the case of equation constraints, the ratio for a given $\lambda = \langle X_i, X_j \rangle_{\text{tr}}$ is

$$\frac{k \int \text{fid}_k(\theta) \Delta_{X_i, X_j}(\theta) d\theta}{1 + (k-1)\lambda} \rightarrow \frac{1 + \frac{k^2}{\pi^2} \text{Re} \left[\text{Li}_2(\lambda) - \text{Li}_2(\lambda e^{i\frac{2\pi}{k}}) \right]}{1 + (k-1)\lambda},$$

which is not a monotone function, so the minimisation cannot be simplified in the same way as NC-Max- k -Cut.

This theorem gives the results presented in [Table 7.2](#).

7.3.4 Minimising the ratio

Lemma 7.3.15. Let $f : [-1, 1] \rightarrow \mathbb{R}$ be a smooth function. If $f(1) \leq 1$ and f is convex, then $g(x) := \frac{1-f(x)}{1-x}$ is increasing.

Proof. Since f is smooth and convex, $f''(x) \geq 0$ for all x . Therefore, $h(x) := 1 - f(x) - (1-x)f'(x)$ is decreasing, as its derivative $h'(x) = -f'(x) + f'(x) - (1-x)f''(x) = -(1-x)f''(x) \leq 0$. Also, we have that $h(1) = 1 - f(1) \geq 0$. Putting these together, $h(x) \geq 0$ for all $x \in [-1, 1]$. The derivative of g is $g'(x) = \frac{-f'(x)(1-x) + (1-f(x))}{(1-x)^2} = \frac{h(x)}{(1-x)^2} \geq 0$, and as such g is increasing. ■

Lemma 7.3.16. Let $f : [-1, 1] \rightarrow \mathbb{R}$ be $f(x) = \frac{1}{k} + \frac{k}{\pi^2} \operatorname{Re}[\operatorname{Li}_2(x) - \operatorname{Li}_2(xe^{i\frac{2\pi}{k}})]$ for $k \geq 3$. Then $f(1) = 1$.

Proof. Consider the Fourier series of the function $p(x) = x(1-x)$ on the interval $[0, 1]$. The Fourier coefficients are

$$\hat{p}(n) = \int_0^1 x(1-x)e^{-2\pi i n x} dx = -\frac{1}{2\pi n^2},$$

for $n \neq 0$, and $\hat{p}(0) = \frac{1}{6}$. Thus, $x(1-x) = \frac{1}{6} - \frac{1}{\pi^2} \sum_{n=1}^{\infty} \frac{\cos(2\pi n x)}{n^2}$. In particular, as $\frac{1}{k} \in [0, 1]$, $-\sum_{n=1}^{\infty} \frac{\cos(2\pi n/k)}{n^2} = \pi^2 \left(\frac{1}{k} \left(1 - \frac{1}{k} \right) - \frac{1}{6} \right)$. Using this and the fact that $\sum_{n=1}^{\infty} \frac{1}{n^2} = \frac{\pi^2}{6}$,

$$f(1) = \frac{1}{k} + \frac{k}{\pi^2} \sum_{n=1}^{\infty} \frac{1 - \cos(2\pi n/k)}{n^2} = \frac{1}{k} + \frac{k}{\pi^2} \left[\frac{\pi^2}{6} + \frac{\pi^2}{k} \left(1 - \frac{1}{k} \right) - \frac{\pi^2}{6} \right] = 1. \quad \blacksquare$$

Lemma 7.3.17. Let $f : [-1, 1] \rightarrow \mathbb{R}$ be as in the previous lemma. Then, f is convex.

Proof. It suffices to show that the second derivative of f is positive. Since the dilogarithm has integral expression $\operatorname{Li}_2(x) = -\int_0^x \frac{\ln(1-t)}{t} dt$, the first derivative of f is

$$f'(x) = \frac{k}{\pi^2} \operatorname{Re} \left[\frac{-\ln(1-x) + \ln(1-xe^{i\frac{2\pi}{k}})}{x} \right] = \frac{k}{2\pi^2} \frac{1}{x} \ln \left(1 + C \frac{x}{(1-x)^2} \right),$$

where $C := 4 \sin^2\left(\frac{\pi}{k}\right) \in (0, 3]$. Differentiating again,

$$f''(x) = \frac{k}{2\pi^2} \left[-\frac{1}{x^2} \ln \left(1 + C \frac{x}{(1-x)^2} \right) + C \frac{1+x}{x(1-x)((1-x)^2 + Cx)} \right].$$

This is positive if and only if the multiple by $2\pi^2 x^2/k$, $g(x) := -\ln\left(1 + C\frac{x}{(1-x)^2}\right) + C\frac{x(1+x)}{(1-x)((1-x)^2+Cx)}$ is. To show g is positive, we need only show that it is positive at the boundary and at any inflection point. At $x = -1$,

$$g(-1) = -\ln\left(1 - \frac{C}{4}\right) \geq 0,$$

as $C \leq 3$ so $1 - \frac{C}{4} \leq 1$. In the limit $x \rightarrow 1$, $-\ln\left(1 + C\frac{x}{(1-x)^2}\right) \sim 2\ln(1-x)$ and $C\frac{x(1+x)}{(1-x)((1-x)^2+Cx)} \sim \frac{2}{1-x}$, so $\lim_{x \rightarrow 1} g(x) = +\infty$. To find the values of g at the inflection points, consider the derivative

$$g'(x) = \frac{Cx(2x^3 + Cx^2 + (2C - 6)x - C + 4)}{(1-x)^2((1-x)^2 + Cx)^2};$$

the inflection points are at the values of x where $g'(x) = 0$. First, we have a zero at $x = 0$, in which case $g(0) = 0 \geq 0$. I claim that g' has no other zero in the interval $[-1, 1]$, and hence no other inflection points we must consider. In fact, the only zeros of g' are the zeros of $q(x) := 2x^3 + Cx^2 + (2C - 6)x - C + 4$. We can show that q only has one real root at some $x \leq -1$ by showing that it is positive at its inflection points, and that its concave inflection point is at $x = -1$. To do so, $q'(x) = 6x^2 + 2Cx + 2C - 6$, which has roots -1 and $1 - \frac{C}{3}$. Thus, the values of q at its inflection points are $q(-1) = -2 + C - (2C - 6) - C + 4 = 8 - 2C \geq 0$, and $q(1 - C/3) = \frac{C}{27}(C^2 - 18C + 54) \geq 0$ as its smallest nonzero root is $9 - 3\sqrt{3} > 3$ (seeing C as a variable). As such, the only inflection point of g in $[-1, 1]$ is at $x = 0$, where it is positive, so $g(x) \geq 0$ for all $x \in [-1, 1]$, and so f is convex. ■

Lemma 7.3.18 (Restatement of Lemma 7.3.12). The function

$$\lambda \in [-1, 1] \mapsto \frac{1 - \frac{1}{k} - \frac{k}{\pi^2} \operatorname{Re}\left[\operatorname{Li}_2(\lambda) - \operatorname{Li}_2(\lambda e^{i\frac{2\pi}{k}})\right]}{1 - \lambda} \quad (7.13)$$

is increasing for all $k \geq 3$.

Proof. Let $f(x) = \frac{1}{k} + \frac{k}{\pi^2} \operatorname{Re}\left[\operatorname{Li}_2(x) - \operatorname{Li}_2(xe^{i\frac{2\pi}{k}})\right]$. Then, using Lemma 7.3.15, if $f(1) \leq 1$ and f is convex, we get that $\frac{1-f(x)}{1-x}$ is increasing. To get these two properties, we can use Lemmas 7.3.16 and 7.3.17. ■

7.4 Dimension-efficient algorithm

The implementation of the algorithm for entangled Max- k -Cut and HMax-Lin(k) introduced in Section 7.2.4 is hindered by the fact that the known approximation ratio is only attained

in the limit of large dimension, with no information on the rate of convergence. In this section, we construct a variant of the algorithm that achieves the same approximation ratios in dimension exponential in the instance size. In [Section 7.4.1](#), we develop group-theoretic methods which emulate the behaviour of the infinite-dimensional limit of the relative distribution studied in the previous section, in finite dimension. In [Section 7.4.2](#), we apply this to a modification of the approximation algorithm.

7.4.1 Generalised Weyl-Brauer group and strong isometry

Lemma 7.4.1. Let H be a Hilbert space and let $U : \mathbb{R}^d \rightarrow \mathcal{U}_H$, $\mathbf{v} \mapsto \sum_{i=1}^d v_i \sigma_i$ be a linear map. U maps unit vectors to unitaries if and only if σ_i is a unitary for all i and $\sigma_i^* \sigma_j = -\sigma_j^* \sigma_i$ for all $i \neq j$.

We call the property satisfied by the σ_i ***-anticommutation**.

Proof. First, suppose that the σ_i are *-anticommuting unitaries. We have for $i \neq j$ that $\sigma_i \sigma_j^* = \sigma_j \sigma_i^* \sigma_i \sigma_j^* = -\sigma_j \sigma_i^* \sigma_j \sigma_i^* = -\sigma_j \sigma_i^*$. Then,

$$\begin{aligned} U(\mathbf{v})^* U(\mathbf{v}) &= \sum_{i,j} v_i v_j \sigma_i^* \sigma_j = \sum_i v_i^2 \sigma_i^* \sigma_i + \sum_{i < j} v_i v_j (\sigma_i^* \sigma_j + \sigma_j^* \sigma_i) = \|\mathbf{v}\|^2, \\ U(\mathbf{v}) U(\mathbf{v})^* &= \sum_{i,j} v_i v_j \sigma_i \sigma_j^* = \sum_i v_i^2 \sigma_i \sigma_i^* + \sum_{i < j} v_i v_j (\sigma_i \sigma_j^* + \sigma_j \sigma_i^*) = \|\mathbf{v}\|^2, \end{aligned}$$

so U is a vector-to-unitary construction.

Conversely, if U is a vector-to-unitary construction, write $\mathbf{e}_i = (0, \dots, 0, 1, 0, \dots, 0)$ for the i -th canonical basis vector. We have first that $U(\mathbf{e}_i) = \sigma_i$, so σ_i must be unitary for all i . Also, for any $i \neq j$,

$$2 = U(\mathbf{e}_i + \mathbf{e}_j)^* U(\mathbf{e}_i + \mathbf{e}_j) = \sigma_i^* \sigma_i + \sigma_j^* \sigma_j + \sigma_i^* \sigma_j + \sigma_j^* \sigma_i,$$

giving that we must have $\sigma_i^* \sigma_j + \sigma_j^* \sigma_i = 0$ as well. ■

For more general vector-to-unitary constructions than the Weyl-Brauer construction of [Definition 7.2.4](#), we make use of a generalisation of the Weyl-Brauer operators.

Definition 7.4.2. The **generalized Weyl-Brauer (GWB) group with n generators** is

$$\text{GWB}_n = \langle \sigma_1, \dots, \sigma_n, J : J^2, [\sigma_i, J], J(\sigma_i^{-1} \sigma_j)^2 \forall i \neq j \rangle. \quad (7.14)$$

Any unitary representation of GWB_n sending $J \mapsto -1$ gives a set of n operators satisfying $*$ -anticommutation relation. Conversely, any set of unitary operators satisfying $*$ -anticommutation provides a representation of the group.

Definition 7.4.3. Let $\pi : \text{GWB}_n \rightarrow \mathcal{B}(H)$ be a unitary representation such that $\pi(J) = -1$. The π -**vector-to-unitary construction** is the linear map $U_\pi : \mathbb{R}^n \rightarrow \mathcal{B}(H)$, $\mathbf{x} \mapsto \sum_i x_i \pi(\sigma_i)$.

We can study properties of general vector-to-unitary constructions by studying the representation theory of GWB_n . The goal will be to construct representations π of GWB_n such that the associated vector-to-unitary construction has the **strong isometry property** $\langle U_\pi(\mathbf{v})^s, U_\pi(\mathbf{w})^s \rangle_{\text{tr}} = \langle \mathbf{v}, \mathbf{w} \rangle^s$.

First, we work at the level of the group algebra, which allows to see that the strong isometry property is in fact possible. However, we must first note that GWB_n has infinite cardinality, which makes the analysis more involved: there is a surjective homomorphism $\text{GWB}_n \rightarrow \mathbb{Z}$ such that $J \mapsto 0$ and $\sigma_i \mapsto 1$. In order to get a handle on the structure of GWB_n , we consider an alternate presentation. In this presentation, the $*$ -anticommutation relation $\sigma_i^{-1} \sigma_j = J \sigma_j^{-1} \sigma_i$ can be replaced by the standard anticommutation relation, and hence we can reason about this group using intuition from Pauli matrices.

Lemma 7.4.4. The group GWB_n admits the presentation

$$\text{GWB}_n = \langle p_1, \dots, p_{n-1}, c, J : J^2, Jp_i^2, [p_i, J], [c, J], J(p_i p_j)^2 \forall i \neq j \rangle. \quad (7.15)$$

The new generators are related to the original ones by $c = \sigma_1$ and $p_i = \sigma_1^{-1} \sigma_{i+1}$.

We call this the **Pauli presentation** of GWB_n because the generators p_i could be represented by Pauli matrices; consider the representation that sends $c \mapsto 1$, $J \mapsto -1$ and generators p_i to pairwise anticommuting and anti-Hermitian Pauli matrices. In fact it is possible to show that the subgroup $P_n = \langle p_1, \dots, p_{n-1} \rangle$ is isomorphic to some subgroup of index at most 2 of the Pauli group on $\lfloor \frac{n-1}{2} \rfloor$ qubits.

Proof. Let $c = \sigma_1$ and $p_i = \sigma_1^{-1} \sigma_{i+1}$. This is evidently a set of generators for GWB_n . It suffices to rewrite the relations of GWB_n in terms of these generators. First, note that the order relation $J^2 = 1$ remains the same, and the commutation relations $[\sigma_i, J] = 1$ become commutation with the new generators $[c, J] = 1$ and $[p_i, J] = 1$. The final relations $J(\sigma_i^{-1} \sigma_j)^2 = 1$, if $i = 1, j > 1$, become $1 = Jp_{j-1}^2$; we get the same relation for $j = 1, i > 1$; and for $i, j \neq 1$, $1 = J(cp_{i-1}p_{j-1}^{-1}c^{-1})^2 = cJ(p_{i-1}p_{j-1})^2c^{-1}$, which is the last relation $J(p_{i-1}p_{j-1})^2 = 1$. ■

The quotient by the subgroup $\langle J \rangle$ is the free product $\text{GWB}_n / \langle J \rangle \cong \mathbb{Z} * \mathbb{Z}_2^{n-1}$. So, there are no relations between c and p_i in $\text{GWB}_n / \langle J \rangle$. This also implies that $c, p_i, p_i^{-1}p_j \notin \{1, J\}$ in GWB_n unless $i = j$. Finally, since the quotient has trivial centre, $\langle J \rangle$ is the centre of GWB_n .

The subgroup $C = \langle c, J \rangle$ is isomorphic to $\mathbb{Z} \times \mathbb{Z}_2$, and $\langle J \rangle$ is a subgroup of both C and P_n . So, by the definition of the amalgamated product (see [Section 2.3.1](#)), we arrive at the characterisation

$$\text{GWB}_n \cong (\mathbb{Z} \times \mathbb{Z}_2) *_{\mathbb{Z}_2} P_n$$

where the amalgamated product is over the subgroups $\mathbb{Z}_2 \cong \langle J \rangle$.

Now, let λ be the left regular representation of GWB_n on $H = \ell^2 \text{GWB}_n$. The group von Neumann algebra $L(\text{GWB}_n)$ of GWB_n is the SOT closure of $\lambda(\text{GWB}_n)$ in $\mathcal{B}(H)$.

To construct a representation where $J \mapsto -1$, we can take the quotient algebra

$$\mathcal{M} = L(\text{GWB}_n) / \langle\langle 1 + J \rangle\rangle.$$

The element $p = (1 + J)/2$ is a projection in $L(\text{GWB}_n)$. Since the ideals $\langle\langle 1 + J \rangle\rangle = \langle\langle p \rangle\rangle$, we have that $\mathcal{M} = L(\text{GWB}_n) / \langle\langle p \rangle\rangle$, and since p is central, we have $\mathcal{M} \cong (1 - p)L(\text{GWB}_n)$.

Definition 7.4.5. Let $\bar{\lambda} : \text{GWB}_n \rightarrow \mathcal{B}((1 - p)H)$ be the representation induced by the left regular representation. The **infinite-dimensional vector-to-unitary construction** is the $\bar{\lambda}$ -vector-to-unitary construction $U_{\bar{\lambda}}$.

Write $\bar{\sigma}_i = \bar{\lambda}(\sigma_i)$ and $u_x = U_{\bar{\lambda}}(\mathbf{x})$. First, we see that the $\bar{\sigma}_i$ $*$ -anticommute, and hence $U_{\bar{\lambda}}$ is a vector-to-unitary construction by [Lemma 7.4.1](#).

Since $L(\text{GWB}_n)$ is the group von Neumann algebra of a discrete group, it has a tracial state τ that extends the canonical trace on the group algebra $\mathbb{C}[\text{GWB}_n]$. The trace of any group element other than the identity is 0. Now the algebra $(1 - p)L(\text{GWB}_n)$ has a trace

$$\frac{\tau|_{(1-p)L(\text{GWB}_n)}}{\tau(1-p)},$$

where $\tau|_{(1-p)L(\text{GWB}_n)}$ is the restriction of τ to the compression $(1 - p)L(\text{GWB}_n)$. Thus, on $\mathcal{M}$, this becomes the tracial state

$$\tau_{\mathcal{M}}(\bar{a}) = 2\tau((1 - p)a) = \tau((1 - J)a),$$

since $\tau(p) = \tau(\frac{1+J}{2}) = \frac{1}{2}$, where $\bar{a}$ is the element of the quotient algebra with representative a .

The trace $\tau_{\mathcal{M}}$ gives rise to an inner product on $\mathcal{M}$, which allows $\mathcal{M}$ to be completed to the Hilbert space $\ell^2 \mathcal{M} = (1 - p)H$. We show that the infinite-dimensional vector-to-unitary construction is an isometry with respect to this inner product.

Lemma 7.4.6. For all vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$,

$$\tau_{\mathcal{M}}(u_{\mathbf{x}}^* u_{\mathbf{y}}) = \langle \mathbf{x}, \mathbf{y} \rangle \quad (7.16)$$

Proof. We can simply calculate the inner product using the Pauli presentation of GWB_n : writing $p_0 = 1$,

$$\tau_{\mathcal{M}}(u_{\mathbf{x}} u_{\mathbf{y}}) = \sum_{i,j=1}^n x_i y_j \tau((1-J)\sigma_i^{-1}\sigma_j) = \sum_{i,j=1}^n x_i y_j \tau((1-J)p_{i-1}^{-1}p_{j-1}) = \sum_{i=1}^n x_i y_i = \langle \mathbf{x}, \mathbf{y} \rangle,$$

since $p_{i-1}^{-1}p_{j-1} \in \{1, J\}$ in the group GWB_n if and only if $i = j$. ■

In fact, we have the stronger property characterising the inner products of powers of the unitaries.

Proposition 7.4.7 (Strong isometry). For all vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ and $m \in \mathbb{N}$,

$$\tau_{\mathcal{M}}((u_{\mathbf{x}}^m)^* u_{\mathbf{y}}^m) = \langle \mathbf{x}, \mathbf{y} \rangle^m. \quad (7.17)$$

Note that by traciality of the state and the fact that the vectors are real, we get that $\tau_{\mathcal{M}}((u_{\mathbf{x}}^m)^* u_{\mathbf{y}}^m) = \langle \mathbf{x}, \mathbf{y} \rangle^{|m|}$ for all $m \in \mathbb{Z}$.

Proof. We calculate the inner product as in the previous lemma. Writing again $p_0 = 1$, we have

$$\begin{aligned} \tau_{\mathcal{M}}((u_{\mathbf{x}}^m)^* u_{\mathbf{y}}^m) &= \sum_{\substack{i_1, \dots, i_m \\ j_1, \dots, j_m}} x_{i_1} y_{j_1} \cdots x_{i_m} y_{j_m} \tau((1-J)\sigma_{i_m}^{-1} \cdots \sigma_{i_1}^{-1} \sigma_{j_1} \cdots \sigma_{j_m}) \\ &= \sum_{\substack{i_1, \dots, i_m \\ j_1, \dots, j_m}} x_{i_1} y_{j_1} \cdots x_{i_m} y_{j_m} \tau((1-J)p_{i_m-1}^{-1} c^{-1} \cdots p_{i_2-1} c^{-1} p_{i_1-1}^{-1} p_{j_1-1} c p_{i_2-1} \cdots c p_{j_m-1}) \\ &= \sum_{i_1, \dots, i_m} x_{i_1} y_{i_1} \cdots x_{i_m} y_{i_m} = \langle \mathbf{x}, \mathbf{y} \rangle^m, \end{aligned}$$

as, for $(i_1, \dots, i_m) \neq (j_1, \dots, j_m)$, $p_{i_m-1}^{-1} c^{-1} \cdots p_{i_2-1} c^{-1} p_{i_1-1}^{-1} p_{j_1-1} c p_{i_2-1} \cdots c p_{j_m-1} \notin \{1, J\}$ owing to the fact that it is not identity in the quotient $\text{GWB}_n / \langle J \rangle$ since there are no relations between the p_i and c there. ■

Now, we extend these ideas to construct finite-dimensional representations, by adding relations to GWB_n to obtain a finite group GWB_n^k with order- k generators σ_i . We then study the representation theory of GWB_n^k and give a vector-to-unitary construction that preserves the strong isometry property up to order k .

We construct the group GWB_n^k in two steps. We first introduce an intermediate variant of GWB_n .

Definition 7.4.8. Let $f_{i,j,t} \in \mathbb{Z}_2$ for $i, j \in \{1, \dots, n-1\}$ and $t \in \mathbb{Z}_k$ be such that $f_{i,i,0} = 0$ and $f_{i,j,t} = f_{j,i,-t}$. Define $G_{n,f}^k$ as the group

$$\langle p_1, \dots, p_{n-1}, c, J : J^2, c^k, Jp_i^2, [p_i, J], [c, J], J^{f_{i,j,t}}[p_i, c^{-t}p_jc^t] \rangle. \quad (7.18)$$

The difference between this and GWB_n is the addition of the order- k relation $c^k = e$ and the commutation relations $J^{f_{i,j,t}}[p_i, c^{-t}p_jc^t] = 1$. We show that this group is finite by introducing a normal form for its elements using the commutation relations in the presentation. Let $\leq$ denote the lexicographic order on $\mathbb{Z} \times \mathbb{Z}_k$. For $\alpha \in \mathbb{M}_{[n-1] \times \mathbb{Z}_k}(\mathbb{Z}_2)$, write

$$p^\alpha = \prod_{(i,t) \in [n-1] \times \mathbb{Z}_k} c^{-t} p_i^{\alpha_{i,t}} c^t,$$

where the terms are ordered with $\leq$. We also use the shorthand notation $p_{i,t} = c^{-t}p_i c^t$.

Lemma 7.4.9. Every element of $G_{n,f}^k$ is of the form $J^b c^s p^\alpha$ for some $(b, s, \alpha) \in \mathbb{Z}_2 \times \mathbb{Z}_k \times \mathbb{M}_{[n-1] \times \mathbb{Z}_k}(\mathbb{Z}_2)$.

We call $J^b c^s p^\alpha$ the **normal form** of an element in $G_{n,f}^k$.

Proof. Elements of $G_{n,f}^k$ are arbitrary products of terms of the form $c^{-t}p_i c^t$ and powers of J and c . Since J commutes with everything we can always bring terms that are powers of J to the beginning. Similarly, powers of c can always be moved to the beginning because $(c^{-t}p_i c^t)c^s = c^s(c^{-(t+s)}p_i c^{t+s})$. Finally, using the commutation relations of $G_{n,f}^k$, we have $[c^{-s}p_i c^s, c^{-t}p_j c^t] = J^{f_{i,j,t-s}}$, and therefore elements of the form $c^{-t}p_i c^t$ for all i and t , can always be moved past each other, which changes the exponent of J . ■

Since the index set $\mathbb{Z}_2 \times \mathbb{Z}_k \times \mathbb{M}_{[n-1] \times \mathbb{Z}_k}(\mathbb{Z}_2)$ is of size $2k \cdot 2^{k(n-1)}$, we conclude that $|G_{n,f}^k| \leq 2k \cdot 2^{k(n-1)}$. In fact, we can prove equality.

Lemma 7.4.10. $|G_{n,f}^k| = 2k \cdot 2^{k(n-1)}$. In particular every element in $G_{n,f}^k$ has a unique normal form.

Proof. Let X be the set of all $2k \cdot 2^{k(n-1)}$ normal form words $J^b c^s p^\alpha$. We let π be the action of the group $G_{n,f}^k$ on the set X defined on the generators so that

$$\begin{aligned}\pi_J(J^b c^s p^\alpha) &= J^{b+1} c^s p^\alpha, \\ \pi_c(J^b c^s p^\alpha) &= J^b c^{s+1} p^\alpha, \\ \pi_{p_i}(J^b c^s p^\alpha) &= J^{b+\alpha_{i,s} + \sum_{(j,t) \leq (i,s)} \alpha_{j,t} f_{i,j,t-s}} c^s p^{\alpha + E_{i,s}},\end{aligned}$$

where $E_{i,s}$ is the standard basis matrix that is 1 at the entry (i, s) and 0 everywhere else. To show that π indeed extends to an action of the group $G_{n,f}^k$ on X we need to show that every relator in the presentation acts as the identity map on X . We can directly verify that

$$\begin{aligned}\pi_{J^2}(J^b c^s p^\alpha) &= J^b c^s p^\alpha \\ \pi_{c^k}(J^b c^s p^\alpha) &= J^b c^s p^\alpha \\ \pi_{p_i J}(J^b c^s p^\alpha) &= \pi_{J p_i}(J^b c^s p^\alpha) \\ \pi_{cJ}(J^b c^s p^\alpha) &= \pi_{Jc}(J^b c^s p^\alpha)\end{aligned}$$

Similarly we have

$$\pi_{J p_i^2}(J^b c^s p^\alpha) = \pi_J(J^{b+1} c^s p^\alpha) = J^b c^s p^\alpha.$$

Finally, a tedious but straightforward calculation verifies that

$$\pi_{[p_i, c^{-t} p_j c^t]}(J^b c^s p^\alpha) = J^{b+\delta_{(i,s) \leq (j,s+t)} f_{i,j,t} + \delta_{(j,s+t) \leq (i,s)} f_{j,i,-t}} c^s p^\alpha,$$

which is $J^{b+f_{i,j,t}} c^s p^\alpha$, as desired, owing to our assumptions on the coefficients $f_{i,j,t}$.

Let $\epsilon = J^0 c^0 p^0 \in X$ denote the empty word and note that $J^b c^s p^\alpha = \pi_{J^b c^s p^\alpha}(\epsilon)$ are all distinct words in X . Therefore distinct normal form words in X are also distinct as elements of $G_{n,f}^k$. We conclude that $|G_{n,f}^k| = |X|$. $\blacksquare$

For the rest of this discussion fix $f_{i,j,t}$ as follows:

$$f_{i,j,t} = \begin{cases} 1 & i \neq j, t = 0, \\ 1 & i \leq j, t = 1, \\ 1 & i \geq j, t = -1, \\ 0 & \text{otherwise.} \end{cases} \quad (7.19)$$

With this choice of $f_{i,j,t}$, we have that $G_{n,f}^k$ is a quotient of GWB_n .

For every $i \in [n-1]$ let $r_i := J^k p_i (c^{-1} p_i c) (c^{-2} p_i c^2) \cdots (c^{-(k-1)} p_i c^{k-1})$. Let H be the subgroup generated by $r_1, \dots, r_{n-1}$.

Lemma 7.4.11. H is a central subgroup, and thus is trivially normal in $G_{n,f}^k$.

Proof. Using the commutation relations, we can write

$$\begin{aligned} Jr_i J^{-1} &= r_i \\ c^{-1} r_i c &= J^k (c^{-1} p_i c) \cdots (c^{-(k-1)} p_i c^{k-1}) p_i = J^3 p_i (c^{-1} p_i c) \cdots (c^{-(k-1)} p_i c^{k-1}) = r_i \\ p_j r_i p_j^{-1} &= J^{k+\delta_{i \neq j} + \delta_{i \leq j} + \delta_{j \leq i}} p_i (c^{-1} p_i c) \cdots (c^{-(k-1)} p_i c^{k-1}) = r_i. \end{aligned}$$

Definition 7.4.12. The **order- k generalized Weyl-Brauer group with n generators** is the quotient group $\text{GWB}_n^k = G_{n,f}^k / H$.

Since GWB_n^k is also a quotient of GWB_n , all the representation of GWB_n^k are also representations of GWB_n .

In the same way as for GWB_n , GWB_n^k can be presented in terms of generators $\sigma_1 = c$ and $\sigma_{i+1} = c p_i$ for $i \in [n-1]$. In particular, the relations $\sigma_i^k = 1$ hold in GWB_n^k . This follows by expressing the relation $r_{i-1} = 1$ in terms of these generators.

Theorem 7.4.13. The group GWB_n^k is finite of order $2k \cdot 2^{(k-1)(n-1)}$. Further, the words p^α for $\alpha \in \mathbb{M}_{[n-1] \times \mathbb{Z}_k}(\mathbb{Z}_2)$ with $\alpha_{i,k-1} = 0$, for all $i \in [n-1]$, are equal to 1 or J if and only if $\alpha = 0$.

Note that in this theorem, the non-trivial words are chosen to not involve any of the elements $p_{i,k-1}$ due to the relation $p_{i,k-1} = J^{k-1} p_{i,0} \cdots p_{i,k-2}$ for all $i \in [n-1]$.

Proof. Since $\text{GWB}_n^k = G_{n,f}^k / H$ we have $|\text{GWB}_n^k| = |G_{n,f}^k| / |H|$. We calculate $|H|$. First note that from the commutation relations of $G_{n,f}^k$, we have

$$[c^{-s} p_i c^s, c^{-t} p_i c^t] = c^{-s} [p_i, c^{-(t-s)} p_i c^{t-s}] c^s = \begin{cases} J & \text{if } t - s = 1, k-1 \pmod{k}, \\ 1 & \text{otherwise.} \end{cases}$$

Now a repeated application of these commutation relations gives that $r_i^2 = 1$ for all $i \in [n-1]$. Therefore $|H| \leq 2^{n-1}$ as it is generated by $n-1$ commuting generators of order 2. However by [Lemma 7.4.10](#) we know that all the normal form words are distinct in $G_{n,f}^k$. Thus $\prod_{i=1}^{n-1} r_i^{e_i}$ for $e_i \in \{0, 1\}$ are all distinct, and thus $|H| = 2^{n-1}$. This gives the result $|\text{GWB}_n^k| = |G_{n,f}^k| / |H| = 2k \cdot 2^{k(n-1)} / 2^{n-1} = 2k \cdot 2^{(k-1)(n-1)}$.

For the second statement, none of the p^α considered are contained in H or JH , so they are not equal to 1 or J in the quotient group GWB_n^k . ■

Corollary 7.4.14. There exists a finite-dimensional representation π of GWB_n^k of dimension $k \cdot 2^{(k-1)(n-1)}$, where $\pi(J) = -1$ and $\text{tr}(\pi(p^\alpha)) = 0$ for all nonzero $\alpha \in \mathbb{M}_{[n-1] \times \mathbb{Z}_k}(\mathbb{Z}_2)$ such that $\alpha_{i,k-1} = 0$ for all $i \in [n-1]$.

Proof. The proof is similar to the infinite-dimensional case, except now the Hilbert space $H = \ell^2 \text{GWB}_n^k = \mathbb{C}[\text{GWB}_n^k]$ is finite-dimensional of dimension $2k \cdot 2^{(n-1)(k-1)}$, and the von Neumann group algebra is simply $\mathbb{C}[\text{GWB}_n^k]$. The group algebra $\mathbb{C}[\text{GWB}_n^k]$ acts by left multiplication on H , therefore it embeds into $\mathcal{B}(H)$. Let tr_H denote the canonical normalised trace on $\mathcal{B}(H)$.

Notice that $q = (1 - J)/2$ is a projection with trace $\text{tr}_H(q) = \frac{1}{2}$. Therefore $H_0 = qH$ is a subspace of dimension $k \cdot 2^{(n-1)(k-1)}$. Let $\pi : \mathbb{C}[\text{GWB}_n^k] \rightarrow \mathcal{B}(H_0)$, be defined as $\pi(x) = qx$. Since q is a central projection, it is easy to verify that π is a representation and that $\pi(J) = -1$.

Finally notice that $\text{tr}_H(\pi(x)) = \frac{1}{2} \text{tr}_H((1 - J)x)$, so $\text{tr}_{H_0} = 2 \text{tr}_H$ is a tracial state on $\mathcal{B}(H_0)$. Now since $p^\alpha \notin \{1, J\}$ by Theorem 7.4.13, $\text{tr}_{H_0}(\pi(p^\alpha)) = 0$. $\blacksquare$

For simplicity from here on we identify σ_i with its image under the representation $\pi(\sigma_i)$.

Corollary 7.4.15 (Finite-dimensional strong isometry). Let π be as in the theorem. The vector-to-unitary construction U_π is strongly isometric up to order k , *i.e.* it satisfies

$$\text{tr}[U_\pi(\mathbf{x})^{-s} U_\pi(\mathbf{y})^s] = \langle \mathbf{x}, \mathbf{y} \rangle^s,$$

for all $s = 0, \dots, k-1$ and $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$.

Proof. Writing $p_0 = 1$, we can expand

$$\begin{aligned} & \text{tr}[U_\pi(\mathbf{x})^{-s} U_\pi(\mathbf{y})^s] \\ &= \sum_{\substack{i_1, \dots, i_s \\ j_1, \dots, j_s}} x_{i_1} y_{j_1} x_{i_2} y_{j_2} \cdots x_{i_s} y_{j_s} \text{tr} \left((1 - J) \sigma_{i_s}^{-1} \cdots \sigma_{i_2}^{-1} \sigma_{i_1}^{-1} \sigma_{j_1} \sigma_{j_2} \cdots \sigma_{j_s} \right) \\ &= \sum_{\substack{i_1, \dots, i_s \\ j_1, \dots, j_s}} x_{i_1} y_{j_1} x_{i_2} y_{j_2} \cdots x_{i_s} y_{j_s} \text{tr} \left((1 - J) p_{i_s-1}^{-1} c^{-1} \cdots p_{i_2-1}^{-1} c^{-1} p_{i_1-1}^{-1} c^{-1} c p_{j_1-1} c p_{j_2-1} \cdots c p_{j_s-1} \right) \\ &= \sum_{\substack{i_1, \dots, i_s \\ j_1, \dots, j_s}} x_{i_1} y_{j_1} x_{i_2} y_{j_2} \cdots x_{i_s} y_{j_s} \text{tr} \left((1 - J) p_{i_s-1,0}^{-1} p_{i_s-1-1,1}^{-1} \cdots p_{i_1-1,s-1}^{-1} p_{j_1-1,s-1} \cdots p_{j_s-1,0} \right), \end{aligned}$$

where in the last line we used the shorthand notation $p_{i,t} = c^{-t}p_i c^t$. By Theorem 7.4.13

$$p_{i_s-1,0}^{-1} \cdots p_{i_1-1,s-1}^{-1} p_{j_1-1,s-1} \cdots p_{j_s-1,0} \in \{1, J\}$$

if and only if $(i_1, \dots, i_l) = (j_1, \dots, j_l)$. Thus

$$\text{tr}[U_\pi(\mathbf{x})^{-s} U_\pi(\mathbf{y})^s] = \sum_{i_1, \dots, i_s} x_{i_1} y_{i_1} \cdots x_{i_s} y_{i_s} = \langle \mathbf{x}, \mathbf{y} \rangle^s. \quad \blacksquare$$

7.4.2 The algorithm and its analysis

Basing ourselves on the algebraic results of the previous section, we consider the following algorithm for Max- k -Cut^q.

- 1 Solve the SDP relaxation of Lemma 7.2.1 to get X .
- 2 Extract unit vectors $\mathbf{x}_i \in \mathbb{R}^{V(G)}$ such that $X_{i,j} = \langle \mathbf{x}_i, \mathbf{x}_j \rangle$.
- 3 Run the $m \cdot 2^{(m-1)(n-1)}$ -dimensional π -vector-to-unitary construction given by the representation of GWB_n^m from Corollary 7.4.14 to obtain unitary operators $X_i = U_\pi(\mathbf{x}_i)$.
- 4 Sample a random phase $\zeta \in S^1$.
- 5 Let $\tilde{X}_i = \Omega_k(\zeta X_i)$ be the closest order- k unitary to ζX_i .
- 6 Output the set of operators $\tilde{X}_i$ as the solution.

Algorithm 3: Dimension-efficient approximation algorithm for Max- k -Cut^q.

We can again analyse this algorithm by means of the relative distribution method introduced in Section 7.3. In this case, the relative distribution is defined as follows.

Definition 7.4.16. Let $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ be unit vectors and π be a representation of GWB_n . Define the **relative distribution** induced by $\mathbf{x}, \mathbf{y}$, and π as the measure $\Delta_{\mathbf{x}, \mathbf{y}}^\pi : \mathcal{B}([0, 2\pi)) \rightarrow \mathbb{R}_{\geq 0}$ such that

$$\Delta_{\mathbf{x}, \mathbf{y}}^\pi(E) = \frac{1}{2\pi} \int_0^{2\pi} w_{e^{i\theta} U_\pi(\mathbf{x}), e^{i\theta} U_\pi(\mathbf{y})}(E) d\theta. \quad (7.20)$$

where the weight measure w is defined in Definition 7.3.1.

First, note that since the distribution of the $\{\zeta X_i\}$ is phase-invariant, the result of Theorem 7.3.7 still holds for this algorithm. We restate it in this context.

Lemma 7.4.17. Let $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ be unit vectors and π be a representation of GWB_n . For any noncommutative $*$ -polynomial $P(x, y) = \sum_{s,t=0}^{k-1} c_{s,t}(x^*)^s y^t$,

$$\frac{1}{2\pi} \int P(\Omega_k(e^{i\theta} U_\pi(\mathbf{x})), \Omega_k(e^{i\theta} U_\pi(\mathbf{y}))) d\theta = \int \text{fid}_P(\theta) d\Delta_{\mathbf{x}, \mathbf{y}}^\pi(\theta), \quad (7.21)$$

where fid_P is as defined in [Definition 7.3.6](#).

In [Algorithm 1](#), the relative distribution approached the wrapped Cauchy in distribution in the limit of large dimension. Here, the sense in which the relative distribution approaches the wrapped Cauchy distribution is different. Due the strong isometry property of GWB_n^m , the relative distribution matches the wrapped Cauchy distribution exactly for small moments.

Lemma 7.4.18. Let $\mathbf{x}, \mathbf{y} \in R^n$, and let π be the representation of GWB_n^m from [Corollary 7.4.14](#). Write $\lambda = \langle \mathbf{x}, \mathbf{y} \rangle$. Then, the characteristic function of $\Delta_{\mathbf{x}, \mathbf{y}}^\pi$ satisfies $\chi_{\Delta_{\mathbf{x}, \mathbf{y}}^\pi}(k) = \chi_{\Delta_\lambda}(k)$ for $k = -m, \dots, m$.

Proof. First, as in [Lemma 7.3.4](#), the characteristic function of the weight distribution is $\chi_{w_{A,B}}(k) = \text{tr}[A^{-k} B^k]$. Therefore,

$$\chi_{\Delta_{\mathbf{x}, \mathbf{y}}^\pi}(k) = \frac{1}{2\pi} \int \text{tr}[(e^{i\theta} U_\pi(\mathbf{x}))^{-k} (e^{i\theta} U_\pi(\mathbf{y}))^k] d\theta = \text{tr}[U_\pi(\mathbf{x})^{-k} U_\pi(\mathbf{y})^k].$$

To finish, note that $\chi_{\Delta_{\mathbf{x}, \mathbf{y}}^\pi}(-n) = \chi_{\Delta_{\mathbf{x}, \mathbf{y}}^\pi}(n)$, and due to [Corollary 7.4.15](#), for $k = 0, \dots, m$,

$$\chi_{\Delta_{\mathbf{x}, \mathbf{y}}^\pi}(k) = \langle \mathbf{x}, \mathbf{y} \rangle^k = \chi_{\Delta_\lambda}(k). \quad \blacksquare$$

Next, we need a lemma that will allow us to exclude the cases where the approximation given by the low-order moments of the relative distribution is insufficient.

Lemma 7.4.19. There exists $\varepsilon > 0$ independent of k and m such that for all $\lambda = \langle \mathbf{x}, \mathbf{y} \rangle \geq 1 - \varepsilon$, the approximation ratio

$$\frac{k}{k-1} \frac{1 - \int \text{fid}_k(\theta) \Delta_{\mathbf{x}, \mathbf{y}}^\pi(\theta) d\theta}{1 - \lambda} \geq 1,$$

for large enough m .

As noted in the proof below, we may take any $\varepsilon \leq \frac{1}{4}$ and $m \geq 4$.

Proof. We need to split the proof into a few cases, dependent on k . However, all the cases have the same general idea. Since we know $\int e^{il\theta} d\Delta_{x,y}^\pi(\theta) = \lambda^{|l|}$ for $|l| < m$, if we can lower bound $1 - \text{fid}_k(\theta) \geq \sum_{l=0}^{m-1} a_l \cos(l\theta)$ for all θ , then, we get the lower bound on the ratio

$$\frac{k}{k-1} \frac{1 - \int \text{fid}_k(\theta) d\Delta_{x,y}^\pi(\theta)}{1 - \lambda} \geq \frac{k}{k-1} \frac{\sum_{l=0}^{m-1} a_l \lambda^l}{1 - \lambda}.$$

As such, by maximising this near $\lambda = 1$, we can find a lower bound satisfying the lemma.

First, consider the case $k \geq 7$. Consider the function $f(\theta) = \frac{1}{2}(1 - \cos(3\theta))$. For $k \geq 7$, it is easy to see that $1 - \text{fid}_k(\theta) \geq f(\theta)$ for all θ . Then, by the strong isometry property, we know that since $m \geq 3$,

$$\frac{1 - \int \text{fid}_k(\theta) d\Delta_{x,y}^\pi(\theta)}{1 - \lambda} \geq \frac{\frac{1}{4} \int 2 - e^{3i\theta} - e^{-3i\theta} d\Delta_{x,y}^\pi(\theta)}{1 - \lambda} = \frac{1 - \lambda^3}{2(1 - \lambda)} = \frac{1}{2}(1 + \lambda + \lambda^2).$$

This is greater than or equal to 1 for all $\lambda \geq \frac{3}{4}$, so we can take $\varepsilon = \frac{1}{4}$.

Next, consider $k = 5, 6$. We can take a similar function $f(\theta) = \frac{1}{2}(1 - \cos(2\theta))$, which lower bounds $1 - \text{fid}_k(\theta)$ in these cases. This gives

$$\frac{k}{k-1} \frac{1 - \int \text{fid}_k(\theta) d\Delta_{x,y}^\pi(\theta)}{1 - \lambda} \geq \frac{k}{2(k-1)} \frac{1 - \lambda^2}{1 - \lambda} \geq \frac{6}{10}(1 + \lambda),$$

which is ≥ 1 for $\lambda \geq \frac{2}{3}$, so $\varepsilon = \frac{1}{3}$.

For $k = 4$, we can take the lower bound function $f(\theta) = (\sqrt{2} - \frac{3}{4}) - (\sqrt{2} - 1) \cos(\theta) - \frac{1}{4} \cos(2\theta)$. Then, we have $\frac{k}{k-1} \frac{1 - \int \text{fid}_k(\theta) d\Delta_{x,y}^\pi(\theta)}{1 - \lambda} \geq \frac{4}{3} \left(\sqrt{2} - \frac{3}{4} + \frac{1}{4} \lambda \right)$, which is ≥ 1 for $\lambda \geq 6 - 4\sqrt{2}$, so we can take $\varepsilon = \frac{1}{2}$.

Finally, for $k = 3$, we can take the lower bound function $f(\theta) = \frac{3}{5} - \frac{1}{2} \cos(\theta) - \frac{1}{10} \cos(2\theta)$, giving $\frac{k}{k-1} \frac{1 - \int \text{fid}_k(\theta) d\Delta_{x,y}^\pi(\theta)}{1 - \lambda} \geq \frac{3}{2} \left(\frac{3}{5} + \frac{1}{10} \lambda \right)$. Then, the approximation ratio is ≥ 1 for $\lambda \geq \frac{2}{5}$. Putting all the cases together, we can take $\varepsilon = \frac{1}{4}$ for any $m \geq 4$. $\blacksquare$

Theorem 7.4.20. The approximation ratio of [Algorithm 3](#) in the case of NC-Max- k -Cut for fixed $k \geq 3$ is

$$1 - \frac{1}{k} - \frac{k}{\pi^2} \text{Re} \left[\text{Li}_2 \left(-\frac{1}{k-1} \right) - \text{Li}_2 \left(-\frac{e^{i\frac{2\pi}{k}}}{k-1} \right) \right] + O\left(\frac{1}{m}\right) \quad (7.22)$$

Proof. As in Theorem 7.3.13, the approximation ratio is the minimum, over unit vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ such that $\lambda = \langle \mathbf{x}, \mathbf{y} \rangle \in [-\frac{1}{k-1}, 1)$, of

$$\frac{k}{k-1} \frac{1 - \int \text{fid}_k(\theta) \Delta_{\mathbf{x}, \mathbf{y}}^\pi(\theta) d\theta}{1 - \lambda} = \frac{k}{k-1} \frac{1 - \frac{1}{2\pi} \sum_{n \in \mathbb{Z}} \widehat{\text{fid}}_k(n) \chi_{\Delta_{\mathbf{x}, \mathbf{y}}^\pi}(n)}{1 - \lambda}.$$

Note first that, using the previous lemma, for large enough m , there exists a neighbourhood of 1, of radius ε independent of m , where the approximation ratio is greater than 1. Hence, we need only consider the minimisation of λ on the interval $[-\frac{1}{k-1}, 1 - \varepsilon)$. Next, due to the order- m strong isometry property (Corollary 7.4.14), we have that for $|n| < m$, $\chi_{\Delta_{\mathbf{x}, \mathbf{y}}^\pi}(n) = \lambda^{|n|}$; for the other values of n , we have by definition that $|\chi_{\Delta_{\mathbf{x}, \mathbf{y}}^\pi}(n)| \leq 1$ and that $\chi_{\Delta_{\mathbf{x}, \mathbf{y}}^\pi}(n) \in \mathbb{R}$. As such, the difference

$$\begin{aligned} \left| \sum_{n \in \mathbb{Z}} \widehat{\text{fid}}_k(n) \chi_{\Delta_{\mathbf{x}, \mathbf{y}}^\pi}(n) - \sum_{n \in \mathbb{Z}} \widehat{\text{fid}}_k(n) \lambda^{|n|} \right| &= \left| 2 \sum_{n=m}^{\infty} \text{Re}(\widehat{\text{fid}}_k(n)) (\chi_{\Delta_{\mathbf{x}, \mathbf{y}}^\pi}(n) - \lambda^{|n|}) \right| \\ &\leq 4 \sum_{n=m}^{\infty} \frac{k}{\pi n^2} \left(1 - \cos\left(\frac{2\pi n}{k}\right) \right) \\ &\leq \frac{8k}{\pi} \sum_{n=m}^{\infty} \frac{1}{n^2} \leq \frac{8k}{\pi} \int_{m-1}^{\infty} \frac{1}{x^2} dx = \frac{8k}{\pi(m-1)}. \end{aligned}$$

Hence, due to Corollary 7.3.11, the ratio for any λ is lower-bounded by

$$\begin{aligned} \frac{k}{k-1} \frac{1 - \frac{1}{2\pi} \sum_{n \in \mathbb{Z}} \widehat{\text{fid}}_k(n) \chi_{\Delta_{\mathbf{x}, \mathbf{y}}^\pi}(n)}{1 - \lambda} &\geq \frac{k}{k-1} \frac{1 - \frac{1}{k} - \frac{k}{\pi^2} \text{Re}[\text{Li}_2(\lambda) - \text{Li}_2(\lambda e^{i\frac{2\pi}{k}})]}{1 - \lambda} - \frac{4k}{\pi^2(m-1)} \\ &\geq \frac{k}{k-1} \frac{1 - \frac{1}{k} - \frac{k}{\pi^2} \text{Re}[\text{Li}_2(\lambda) - \text{Li}_2(\lambda e^{i\frac{2\pi}{k}})]}{1 - \lambda} - \frac{4k^2}{\pi^2(k-1)(m-1)\varepsilon}. \end{aligned}$$

Since this is simply the ratio in Theorem 7.3.13, up to additive constant, we can minimise in the same way using Lemma 7.3.12 to get the wanted result. $\blacksquare$

Although we do not analyse it here, it is also possible to extend this algorithm to all $\text{HMax-Lin}(k)^q$.

Chapter 8

Discussion

In this thesis, we studied the behaviour of constraint satisfaction problems presented as nonlocal games. First, we defined the various ways CSPs can embed into nonlocal games. CSP games can be oracular or non-oracular, in the sense that the players may receive constraints and variables, or just variables, respectively, as questions. Then, these games naturally give rise to decision problems, where the winning probability of the game is taken as representing the number of constraints satisfied. We generally restricted to tracial correlations, which give a nice algebraic description for the strategies. We showed that, assuming perfect completeness, it suffices to work with tracial correlations for all CSP games.

First, we studied how to construct reductions between entangled CSPs. As a major technical tool, we made use of commutativity gadgets, which allow all (local) classical reductions to a given CSP to be quantised. Since there are a variety of different types of entangled CSP, we required a variety of different types of commutativity gadget to facilitate reductions between them. We fully characterised when commutativity gadgets for a CSP in any framework exist, using the structure of the polymorphism algebra of the associated relational structure. Using this characterisation, we were able to show a variety of equivalences, in particular reducing each of the robust commutativity gadgets to one of the exact cases. Conversely, we were able to show separations between some of the classes of gadget. Finally, we showed that it is in general undecidable to decide if a relational structure admits a non-oracular commutativity gadget.

Next, we studied reductions between CSPs that do not make use of commutativity gadgets. We developed a theory of stable minion homomorphisms, which allowed the strongest classical reduction between CSPs to be lifted to the setting of gapped entangled

CSPs. This allowed us to show hardness of CSPs that do not admit commutativity gadgets.

We used these techniques to show that a selection of entangled CSPs are undecidable. Notably, we showed that entangled k -colouring is undecidable for all $k \geq 3$, in both the oracular and non-oracular settings. The hardness of 3-colouring and oracular k -colouring can be shown to be undecidable using commutativity gadgets. On the other hand, non-oracular k -colouring does not admit a commutativity gadget, but we were able to show hardness using a minion homomorphism-based reduction, by working out a full description of the polymorphisms of the complete graphs.

For CSPs without perfect completeness, we showed that a simple variant of XOR games, the tilted XOR games, suffice to capture the undecidability of the quantum value of nonlocal games. This shows that binary nonlocal games are in general undecidable, despite being decidable with perfect completeness. This also indicates how special the class of XOR games is, as even the slight perturbation of asking some questions to only one of the players immediately sends the complexity from polynomial-time to undecidable.

Finally, we studied approximation algorithms for entangled CSPs. Our main focus was on approximation algorithms for entangled Max- k -Cut problems, which are the optimisation analogues of k -colouring. We devised and analysed an algorithm inspired by the Goemans-Williamson algorithm for Max-Cut. We showed that our algorithm improves on the classical algorithm for entangled Max-3-Cut and Max-4-Cut. Our approach to the analysis makes use of the relative distribution, which allows us to connect the value to a distribution over random matrices, and hence make use of ideas from free probability to analyse the performance of the algorithm.

8.1 Related work

A line of work has also investigated what happens to the classically polynomial-time solvable CSPs with entanglement. Slofstra [109, 110] showed that gapless LIN is undecidable with entanglement, which is a significant break with the classical setting, where it is easy. If gapped entangled LIN were also undecidable, that would imply the existence of a non-hyperlinear group; however, this seems infeasible to achieve via reductions from other CSPs, as the entangled homomorphism algebra for a linear system game always admits a tracial state if it is nontrivial [99]. Further, Atserias, Kolaitis, and Severini [12] show that this is the only boolean CSP where a classically easy problem can become quantumly-hard, in a strong way: for any classically-easy boolean CSP that is not LIN, an instance admits a quantum assignment if and only if it admits a classical assignment. Hence the entangled CSP is

solvable in polynomial time. Recently, Bulatov and Živný [33] showed that this extends to all classically-easy CSPs. If a CSP is bounded-width, meaning it cannot simulate linear systems over some finite field in logspace, its entangled CSP is also solvable in polynomial time. Further, they show that non-bounded width CSPs can simulate linear systems over some finite field in the entangled setting as well. This indicates that easy CSPs becoming hard in the entangled setting relies entirely on the hardness of linear systems.

The study of CSPs is intertwined with the study of unique games. A nonlocal game is unique if for each one of Alice’s answers, exactly one of Bob’s answers is correct, and vice versa. The unique games conjecture asks whether the classical value of unique games is NP-hard to approximate for all gaps < 1 , and has many significant consequences in theoretical computer science [79, 80, 104]. Surprisingly, due to a result of Kempe, Regev, and Toner [78], the unique games conjecture no longer holds in the entangled setting — the quantum value of a unique game becomes easy to approximate beyond a certain constant gap. Recent work of Mousavi and Spirig [90] has attempted to slightly modify the definition of entangled unique games by restricting to locally-commuting strategies, in order to reinstate a plausible unique games conjecture.

CSP nonlocal games are games with a question size that grows in the instance size, but an answer size that stays constant. Conversely, Natarajan and Zhang [93] have shown that nonlocal games with constant question size and growing answer size are also undecidable. In recent work, Mančinska, Spaas, Spirig, and Vernooij [85] reduce this to the setting of independent set games, where the players have to convince the referee that a graph has an independent set of a given size. This is qualitatively similar to CSP games, as it corresponds to a graph homomorphism from a fixed graph to a graph that varies with the instance. However, it has the interesting property that the classical problem is solvable in polynomial time but the entangled problem is undecidable.

8.2 Open questions

This work leaves many interesting questions still open. We present a few of them here.

- **Do entangled CSPs satisfy a dichotomy theorem similar to classical CSPs?** The evidence in this work suggests that is the case, as we are able to lift the classical reductions to quantum-sound reductions in many cases. Also, we know that there is a dichotomy theorem in some restricted cases: for example, for oracular gapless commuting-operator boolean CSPs, every problem is either **coRE**-complete or solvable

in polynomial time [49, 33] — however the line between the easy and hard instances shifts, as entangled LIN is hard. On the other hand, the fact that commutativity gadgets may not exist hints that there may be cases where the reductions fail quantumly.

- **Are there classically-hard CSPs that become easy with entanglement, even with perfect completeness?** The closest behaviour known to this is for XOR games, but we know that both the classical and quantum problems with perfect completeness are easy. We show that a class of games qualitatively close to XOR games captures the full hardness of MIP^* , indicating that XOR games may be a very special case. Nevertheless, we show that despite being significantly harder in general, there are CSPs that admit an approximation algorithm with better approximation ratio in the entangled setting.
- **What is the minimal NP-complete CSP that does not admit a commutativity gadget?** Due to [49], every NP-complete boolean CSP admits an oracular commutativity gadget. We show that there exists a CSP with domain size 4 that admits neither an oracular nor a non-oracular commutativity gadget. The only NP-complete graph CSP with domain size 3 is 3-colouring, and that admits a commutativity gadget. It remains open whether every NP-complete CSP on with alphabet size 3 admits a commutativity gadget.
- **Can we separate or identify the remaining classes of commutativity gadget?** We have shown that robust quantum, quantum approximate, and commuting operator commutativity gadgets are each equivalent to an exact gadget; we have separated oracular from non-oracular gadgets; and we know that quantum gadgets are distinct, and identically for quantum approximate gadgets under the assumption of a non-hyperlinear group. It remains to separate algebraic gadgets from commuting operator gadgets; and all the subclasses of oracular gadget from each other. It remains possible that some of these classes of commutativity gadget could also be identified.
- **Is it also undecidable to decide the existence of an oracular commutativity gadget?** We showed that it is undecidable to decide the existence of any non-oracular commutativity gadget, but our techniques rely on simplifying the structure of the polymorphisms using the complete graph, which cannot be applied in the oracular setting without making all the polymorphisms classical.
- **Is there a form of reduction based on the symmetries of relational structures that captures hardness for all entangled CSPs?** The strongest form of symmetry-based reduction known to exist is the reduction based on minion homomorphisms.

For classical CSPs, this suffices to capture all the hard CSPs. However, we know that this type of reduction is insufficient in the entangled setting, as 3LIN and 3SAT are both hard but there is no minion homomorphism from the entangled polymorphisms of 3LIN to those of 3SAT, because they are all classical.

- **What is the optimal approximation algorithm for entangled k -colouring?**

We give an algorithm that improves on the classical approximation ratio for entangled 3- and 4-colouring. However, this algorithm is no longer an improvement over the classical one for $k \geq 5$, so it stands to reason that there may be an improvement to be made at least there. On the other hand, we know entangled k -colouring is undecidable for all $k \geq 3$ with some gap, so this limits the possible approximation ratios.

References

- [1] Samson Abramsky, Rui Soares Barbosa, Nadish de Silva, and Octavio Zapata. The Quantum Monad on Relational Structures. In Kim G. Larsen, Hans L. Bodlaender, and Jean-Francois Raskin, editors, *42nd International Symposium on Mathematical Foundations of Computer Science (MFCS 2017)*, volume 83 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 35:1–35:19, Dagstuhl, Germany, 2017. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [2] Antonio Acín, Serge Massar, and Stefano Pironio. Randomness versus nonlocality and entanglement. *Physical review letters*, 108(10):100402, 2012.
- [3] Huzihiro Araki and Shigeru Yamagami. An inequality for Hilbert-Schmidt norm. *Communications in Mathematical Physics*, 81(1):89–96, 1981.
- [4] P. K. Aravind. Quantum mysteries revisited again. *American Journal of Physics*, 72(10):1303–1307, 2004.
- [5] Alex Arkhipov. Extending and characterizing quantum magic games. *arXiv preprint arXiv:1209.3819*, 2012.
- [6] Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.
- [7] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM (JACM)*, 45(3):501–555, 1998.
- [8] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *Journal of the ACM (JACM)*, 45(1):70–122, 1998.
- [9] Vahid R. Asadi, Richard Cleve, Eric Culf, and Alex May. Linear gate bounds against natural functions for position-verification. *Quantum*, 9:1604, 2025.

- [10] Vahid R Asadi, Eric Culf, and Alex May. Rank lower bounds on non-local quantum computation. In *16th Innovations in Theoretical Computer Science Conference (ITCS 2025)*, pages 11–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2025.
- [11] Alain Aspect, Jean Dalibard, and Gérard Roger. Experimental test of Bell’s inequalities using time-varying analyzers. *Physical review letters*, 49(25):1804, 1982.
- [12] Albert Atserias, Phokion G Kolaitis, and Simone Severini. Generalized satisfiability problems via operator assignments. In *International Symposium on Fundamentals of Computation Theory*, pages 56–68. Springer, 2017.
- [13] Sheldon Axler. *Measure, integration & real analysis*. Springer, 2020.
- [14] László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. *Computational complexity*, 1(1):3–40, 1991.
- [15] Libor Barto, Jakub Bulín, Andrei Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. *Journal of the ACM (JACM)*, 68(4):1–66, 2021.
- [16] John S Bell. On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika*, 1(3):195, 1964.
- [17] Mihir Bellare, Oded Goldreich, and Madhu Sudan. Free bits, PCPs, and nonapproximability—towards tight results. *SIAM Journal on Computing*, 27(3):804–915, 1998.
- [18] Archishna Bhattacharyya, Anne Broadbent, and Eric Culf. The uncloneable bit exists. *arXiv preprint arXiv:2603.08916*, 2026.
- [19] Archishna Bhattacharyya and Eric Culf. Uncloneable encryption from decoupling. *Nature Physics*, 22:315–318, 2026.
- [20] Julien Bichon. Hopf-Galois objects and cogroupoids. *Rev. Un. Mat. Argentina*, 55(2):11–69, 2014.
- [21] Bruce Blackadar. *Operator algebras theory of C^* -algebras and von Neumann algebras*. Encyclopaedia of mathematical sciences, v. 122. Operator algebras and non-commutative geometry ; 3. Springer, Berlin, 2006.
- [22] V.G. Bodnarchuk, L.A. Kaluzhnin, V.N. Kotov, and B.A. Romov. Galois theory for Post algebras. I. *Cybernetics*, 5(3):243–252, 1969.

- [23] V.G. Bodnarchuk, L.A. Kaluzhnin, V.N. Kotov, and B.A. Romov. Galois theory for Post algebras. II. *Cybernetics*, 5(5):531–539, 1969.
- [24] Max Born. Quantenmechanik der Stoßvorgänge. *Zeitschrift für physik*, 38(11):803–827, 1926.
- [25] Pierre Botteron, Anne Broadbent, Eric Culf, Ion Nechita, Clément Pellegrini, and Denis Rochette. Towards unconditional uncloneable encryption. *Quantum*, 10:2157, 2026.
- [26] Brian Harold Brandsen and Charles Jean Joachain. *Quantum Mechanics*. Pearson, Harlow, 2nd edition, 2000.
- [27] Michael Brannan, Eric Culf, and Matthijs Vernooij. Quantum expanders and property (T) discrete quantum groups. *Journal of Mathematical Physics*, 67(3):032201, 2026.
- [28] Ola Bratteli and Derek W. Robinson. *Operator Algebras and Quantum Statistical Mechanics II*. Springer, Berlin, Heidelberg, 1st edition, 1981.
- [29] Anne Broadbent, Eric Culf, and Denis Rochette. Optimal untelegraphable encryption and implications for uncloneable encryption. *IACR Communications in Cryptology*, 3(2), 2026.
- [30] Nathaniel P. Brown and Narutaka Ozawa. *C*-algebras and finite-dimensional approximations*. Graduate studies in mathematics, v. 88. American Mathematical Society, Providence, R.I, 2008.
- [31] Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Physical review letters*, 87(16):167902, 2001.
- [32] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 319–330. IEEE, 2017.
- [33] Andrei A. Bulatov and Stanislav Živný. Satisfiability of commutative vs. non-commutative CSPs. In *52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025)*, pages 37–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2025.
- [34] Peter J. Cameron, Ashley Montanaro, Michael W. Newman, Simone Severini, and Andreas Winter. On the quantum chromatic number of a graph. *Electronic Journal of Combinatorics*, 14(1):#R81, 2007.

- [35] John Canny. Some algebraic and geometric computations in PSPACE. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, STOC '88, page 460–467, New York, NY, USA, 1988. Association for Computing Machinery.
- [36] Valerio Capraro, Martino Lupini, and Vladimir Pestov. *Introduction to Sofic and Hyperlinear Groups and Connes' Embedding Conjecture*, volume 2136 of *Lecture Notes in Mathematics*. Springer Nature, Cham, 1st edition, 2015.
- [37] Hubie Chen and Benoit Larose. Asking the metaquestions in constraint tractability. *ACM Transactions on Computation Theory (TOCT)*, 9(3):1–27, 2017.
- [38] Lorenzo Ciardo, Gideo Joubert, and Antoine Mottet. Quantum polymorphisms and commutativity gadgets. *arXiv preprint arXiv:2511.23445*, 2025.
- [39] John F Clauser, Michael A Horne, Abner Shimony, and Richard A Holt. Proposed experiment to test local hidden-variable theories. *Physical review letters*, 23(15):880, 1969.
- [40] Richard Cleve, Eric Culf, and Aviv Taller. XOR games at full tilt: The hardness of binary nonlocal games. *arXiv preprint arXiv:2607.06876*, 2026.
- [41] Richard Cleve, Peter Høyer, Benjamin Toner, and John Watrous. Consequences and limits of nonlocal strategies. In *Proceedings. 19th IEEE Annual Conference on Computational Complexity, 2004.*, pages 236–249, 2004.
- [42] Richard Cleve, Li Liu, and William Slofstra. Perfect commuting-operator strategies for linear system games. *Journal of Mathematical Physics*, 58(1), 2017.
- [43] Richard Cleve and Rajat Mittal. Characterization of binary constraint system games. In *Automata, Languages, and Programming (ICALP 2014)*, pages 320–331. Springer, 2014.
- [44] Richard Cleve, William Slofstra, Falk Unger, and Sarvagya Upadhyay. Perfect parallel repetition theorem for quantum XOR proof systems. *Computational Complexity*, 17(2):282–299, 2008.
- [45] Eric Culf. Almost perfect strategies for projection games are approximately tracial. *arXiv preprint arXiv:2603.14746*, 2026.
- [46] Eric Culf. The complexity of entangled graph colouring via polymorphisms, 2026. In preparation.

- [47] Eric Culf, Josse van Dobben de Bruyn, Matthijs Vernooij, and Peter Zeman. Existence and nonexistence of commutativity gadgets for entangled CSPs. *arXiv preprint arXiv:2509.07835*, 2025.
- [48] Eric Culf, Josse van Dobben van Dobben de Bruyn, and Peter Zeman. Quantum polymorphism characterisation of commutativity gadgets in all quantum models. *arXiv preprint arXiv:2604.01408*, 2026.
- [49] Eric Culf and Kieran Mastel. RE-completeness of entangled constraint satisfaction problems. In *2025 IEEE 66th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 2194–2230. IEEE, 2025.
- [50] Eric Culf, Kieran Mastel, Connor Paddock, and Taro Spirig. The quantum smooth label cover problem is undecidable. In *53rd International Colloquium on Automata, Languages, and Programming (ICALP 2026)*, pages 71–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2026.
- [51] Eric Culf and Arthur Mehta. New approaches to complexity via quantum graphs. *Quantum Information & Computation*, 25(2025):453–487, 2025.
- [52] Eric Culf, Hamoon Mousavi, and Taro Spirig. Approximation algorithms for noncommutative CSPs. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 920–929. IEEE, 2024.
- [53] Louis de Broglie. Waves and quanta. *Nature*, 112(2815):540–540, 1923.
- [54] Etienne de Klerk, Dmitrii V. Pasechnik, and Joost P. Warners. On approximate graph colouring and MAX-k-CUT algorithms based on the θ -function. *Journal of Combinatorial Optimization*, 8:267–294, 2004.
- [55] Mikael de la Salle. Orthogonalization of positive operator valued measures. *Comptes Rendus. Mathématique*, 360(G5):549–560, 2022.
- [56] Irit Dinur. The PCP theorem by gap amplification. *Journal of the ACM (JACM)*, 54(3):12-es, 2007.
- [57] Irit Dinur, David Steurer, and Thomas Vidick. A parallel repetition theorem for entangled projection games. *computational complexity*, 24(2):201–254, 2015.
- [58] Josse van Dobben de Bruyn, David E Roberson, and Simon Schmidt. Asymmetric graphs with quantum symmetry. *Proceedings of the London Mathematical Society*, 131(5):e70098, 2025.

- [59] Josse van Dobben de Bruyn, David E Roberson, and William Slofstra. Future work (in preparation), 2026. Personal communication.
- [60] Yangjing Dong, Honghao Fu, Anand Natarajan, Minglong Qin, Haochen Xu, and Penghui Yao. The computational advantage of MIP* vanishes in the presence of noise. volume 72, pages 1–78. ACM New York, NY, 2025.
- [61] A. Einstein, B. Podolsky, and N. Rosen. Can quantum-mechanical description of physical reality be considered complete? *Phys. Rev.*, 47:777–780, May 1935.
- [62] Honghao Fu, Kieran Mastel, and Xingjian Zhang. Succinct perfect zero-knowledge for MIP*. *arXiv preprint arXiv:2503.04517*, 2025.
- [63] Viktor Galliard and Stefan Wolf. Pseudo-telepathy, entanglement, and graph colorings. In *Proceedings IEEE International Symposium on Information Theory*,, page 101. IEEE, 2002.
- [64] David Geiger. Closed systems of functions and predicates. *Pacific journal of mathematics*, 27(1):95–100, 1968.
- [65] Andrew Gleason. Measures on the closed subspaces of a Hilbert space. *Indiana University Mathematics Journal*, 6(4):885–893, 1957.
- [66] Michel X Goemans and David Williamson. Approximation algorithms for MAX-3-CUT and other problems via complex semidefinite programming. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 443–452, 2001.
- [67] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, nov 1995.
- [68] Isaac Goldbring. *Model theory of operator algebras*. De Gruyter Series in Logic and Its Applications Series ; v.11. Walter de Gruyter GmbH, Berlin/Boston, 1st edition, 2023.
- [69] Rudolf Haag. *Local Quantum Physics: Fields, Particles, Algebras*. Theoretical and Mathematical Physics. Springer, Berlin, Heidelberg, 2nd edition, 2012.
- [70] Uffe Haagerup. The standard form of von Neumann algebras. *Mathematica Scandinavica*, 37(2):271–283, 1975.

- [71] Uffe Haagerup. L^p -spaces associated with an arbitrary von Neumann algebra. In *Algebres d'opérateurs et leurs applications en physique mathématique (Proc. Colloq., Marseille, 1977)*, volume 274, pages 175–184, 1979.
- [72] Johan Håstad. Some optimal inapproximability results. *Journal of the ACM (JACM)*, 48(4):798–859, 2001.
- [73] Teiko Heinosaari, Daniel Reitzner, and Peter Stano. Notes on joint measurability of quantum observables. *Foundations of Physics*, 38:1133–1147, 2008.
- [74] A. J. Hoffman and H. W. Wielandt. The variation of the spectrum of a normal matrix. *Duke Mathematical Journal*, 20(1):37 – 39, 1953.
- [75] Zhengfeng Ji. Binary constraint system games and locally commutative reductions. *arXiv preprint arXiv:1310.3794*, 2013.
- [76] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. $MIP^* = RE$. *Communications of the ACM*, 64(11):131–138, 2021.
- [77] Abraham Karrass and Donald Solitar. The subgroups of a free product of two groups with an amalgamated subgroup. *Transactions of the American Mathematical Society*, 150(1):227–255, 1970.
- [78] Julia Kempe, Oded Regev, and Ben Toner. Unique games with entangled provers are easy. *SIAM Journal on Computing*, 39(7):3207–3229, 2010.
- [79] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings 17th IEEE Annual Conference on Computational Complexity*, pages 25–, 2002.
- [80] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM Journal on Computing*, 37(1):319–357, 2007.
- [81] Se-Jin Kim, Vern Paulsen, and Christopher Schafhauser. A synchronous game for binary constraint systems. *Journal of Mathematical Physics*, 59(3), 2018.
- [82] Stephen Cole Kleene. Introduction to metamathematics. 1952.
- [83] Junqiao Lin. $MIP^{\text{co}} = \text{co}RE$. *arXiv preprint arXiv:2510.07162*, 2025.
- [84] Laura Mančinska and David E. Roberson. Quantum homomorphisms. *Journal of Combinatorial Theory. Series B*, 118:228–267, 2016.

- [85] Laura Mančinska, Pieter Spaas, Taro Spirig, and Matthijs Vernooij. Gap-preserving reductions and RE-completeness of independent set games. *arXiv preprint arXiv:2505.05253*, 2025.
- [86] Amine Marrakchi and Mikael de la Salle. Almost synchronous correlations and Tomita-Takesaki theory. *arXiv preprint arXiv:2307.08129*, 2023.
- [87] Kieran Mastel and William Slofstra. Two prover perfect zero knowledge for MIP*. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC 2024)*, pages 991–1002, 2024.
- [88] N. David Mermin. Simple unified form for the major no-hidden-variables theorems. *Physical Review Letters*, 65(27):3373–3376, 1990.
- [89] Hamoon Mousavi, Seyed Sajjad Nezhadi, and Henry Yuen. Nonlocal games, compression theorems, and the arithmetical hierarchy. In *Proceedings of the 54th annual ACM SIGACT symposium on theory of computing*, pages 1–11, 2022.
- [90] Hamoon Mousavi and Taro Spirig. A quantum unique games conjecture. In *16th Innovations in Theoretical Computer Science Conference (ITCS 2025)*, pages 76–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2025.
- [91] Francis J Murray and John von Neumann. On rings of operators. IV. *Annals of Mathematics*, 44(4):716–808, 1943.
- [92] Mark A Naimark. On a representation of additive operator set functions. In *Dokl. Akad. Nauk SSSR*, volume 41, pages 373–375, 1943.
- [93] Anand Natarajan and Tina Zhang. Quantum free games. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1603–1616, 2023.
- [94] Alantha Newman. Complex semidefinite programming and Max-k-Cut. In *1st Symposium on Simplicity in Algorithms (SOSA 2018)*, pages 13–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2018.
- [95] Alexandru M. Nica and Roland Speicher. *Lectures on the combinatorics of free probability*. London Mathematical Society lecture note series ; 335. Cambridge University Press, Cambridge, 2006.
- [96] Michael A Nielsen and Isaac L Chuang. *Quantum computation and quantum information*, volume 2. Cambridge university press Cambridge, 2001.

- [97] Narutaka Ozawa. About the Connes embedding conjecture: Algebraic approaches. *Japanese Journal of Mathematics. 3rd Series*, 8(1):147–183, 2013.
- [98] Connor Paddock. Rounding near-optimal quantum strategies for nonlocal games to strategies using a maximally entangled state. *Journal of Mathematical Physics*, 66(12):122204, 2025.
- [99] Connor Paddock and William Slofstra. Satisfiability problems and algebras of boolean constraint system games. *Illinois Journal of Mathematics*, 69(1):81–107, 2025.
- [100] Vern I Paulsen, Simone Severini, Daniel Stahlke, Ivan G Todorov, and Andreas Winter. Estimating quantum chromatic numbers. *Journal of Functional Analysis*, 270(6):2188–2222, 2016.
- [101] Asher Peres. Incompatible results of quantum measurements. *Physics Letters A*, 151(3):107–108, 1990.
- [102] Piotr Podleś. Symmetries of quantum spaces. Subgroups and quotient spaces of quantum $SU(2)$ and $SO(3)$ groups. *Communications in mathematical physics*, 170(1):1–20, 1995.
- [103] Sandu Popescu and Daniel Rohrlich. Quantum nonlocality as an axiom. *Foundations of Physics*, 24(3):379–385, 1994.
- [104] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, STOC '08*, page 245–254, New York, NY, USA, 2008. Association for Computing Machinery.
- [105] Thomas J Schaefer. The complexity of satisfiability problems. In *Proceedings of the 10th Annual ACM Symposium on Theory of Computing (STOC'78)*, pages 216–226, 1978.
- [106] Simon Schmidt. *Quantum automorphism groups of finite graphs*. PhD thesis, Universität des Saarlandes, 2020.
- [107] Erwin Schrödinger. An undulatory theory of the mechanics of atoms and molecules. *Physical review*, 28(6):1049, 1926.
- [108] William Slofstra. Lower bounds on the entanglement needed to play XOR non-local games. *Journal of mathematical physics*, 52(10):102202, 2011.

- [109] William Slofstra. The set of quantum correlations is not closed. *Forum of Mathematics, Pi*, 7:e1:1–e1:41, 2019.
- [110] William Slofstra. Tsirelson’s problem and an embedding theorem for groups arising from non-local games. *Journal of the American Mathematical Society*, 33(1):1–56, 2020.
- [111] Masamichi Takesaki. *Theory of operator algebras II*, volume 125. Springer, 2003.
- [112] Aviv Tallor and Thomas Vidick. Approximating the quantum value of an LCS game is RE-hard. *arXiv preprint arXiv:2507.22444*, 2025.
- [113] Marianne Terp. Lp spaces associated with von Neumann algebras. *Notes, Math. Institute, Copenhagen Univ*, 3(4):5, 1981.
- [114] Boris Tsirelson. Quantum analogues of the Bell inequalities. the case of two spatially separated domains. *Journal of Soviet Mathematics*, 36(4):557–570, 1987.
- [115] Alan M. Turing. On computable numbers, with an application to the Entscheidungsproblem. *J. of Math*, 58(345-363):5, 1936.
- [116] Lieven Vandenbergh and Stephen Boyd. Semidefinite programming. *SIAM Review*, 38(1):49–95, 1996.
- [117] Thomas Vidick. Almost synchronous quantum correlations. *Journal of mathematical physics*, 63(2), 2022.
- [118] Thomas Vidick and John Watrous. Quantum proofs. *Foundations and Trends in Theoretical Computer Science*, 11(1-2):1–215, 2016.
- [119] Shuzhou Wang. Quantum symmetry groups of finite spaces. *Communications in Mathematical Physics*, 195(1):195–211, 1998.
- [120] John Watrous. Advanced topics in quantum information theory, 2020. Lecture notes.
- [121] Stanisław L Woronowicz. Compact matrix pseudogroups. *Communications in Mathematical Physics*, 111(4):613–665, 1987.
- [122] Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. *Journal of the ACM*, 67(5):30:1–30:78, 2020.