

A Quantum Unique Games Conjecture

Hamoon Mousavi ✉

Simons Institute for Theoretical Computer Science, University of California, Berkeley, CA, USA

Taro Spirig ✉ 

QMATH, Department of Mathematical Sciences, University of Copenhagen, Denmark

Abstract

After the NP-hardness of computational problems such as 3SAT and MaxCut was established, a natural next step was to explore whether these problems remain hard to approximate. While the quantum nonlocal games extensions of some of these problems are known to be hard – indeed undecidable – their inapproximability remains largely unresolved. In this work, we introduce definitions for the quantum extensions of Label-Cover and Unique-Label-Cover. We show that these problems play a similarly crucial role in studying the inapproximability of quantum constraint satisfaction problems as they do in the classical setting.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory; Theory of computation → Quantum complexity theory

Keywords and phrases hardness of approximation, quantum computing, noncommutative constraint satisfaction problems, Fourier analysis

Digital Object Identifier 10.4230/LIPIcs.ITCS.2025.76

Related Version *Full Version:* <https://arxiv.org/abs/2409.20028> [38]

Funding *Hamoon Mousavi:* Supported by a Simons-CIQC postdoctoral fellowship through NSF QLCI Grant No. 2016245.

Taro Spirig: Supported by the European Union under the Grant Agreement No 101078107, QInteract and VILLUM FONDEN via Villum Young Investigator grant (No 37532) and the QMATH Centre of Excellence (Grant No 10059).

Acknowledgements We thank Henry Yuen for valuable discussions on the topics explored in this paper. We thank Eric Culf and Michael Chapman for helpful discussions.

1 Introduction

1.1 Motivations

The Label-Cover problem [1, 8, 16], a well-known constraint satisfaction problem (CSP), has been central in the study of hardness of approximation. According to the PCP theorem [2, 3, 4, 5, 8, 15, 19, 45], it is NP-hard to approximate Label-Cover within any additive constant. This result, combined with gap-preserving reductions from Label-Cover, has led to optimal inapproximability results for several other key problems, including 3XOR and 3SAT, due largely to the work of Håstad [24, 25].

An instance of Label-Cover consists of a bipartite graph with left vertex set U and right vertex set V , a large label set (also called an alphabet) \mathcal{L} , and maps $\pi_{u,v} : \mathcal{L} \rightarrow \mathcal{L}$ associated with each edge (u, v) . These maps are usually called projections. A *labelling* $a : U \cup V \rightarrow \mathcal{L}$ is said to satisfy an edge (u, v) if it respects the associated projection map, meaning if $\pi_{u,v}(a(u)) = a(v)$. The *value of the instance* is the maximum fraction of satisfied edges over all possible labelings.



© Hamoon Mousavi and Taro Spirig;

licensed under Creative Commons License CC-BY 4.0

16th Innovations in Theoretical Computer Science Conference (ITCS 2025).

Editor: Raghu Meka; Article No. 76; pp. 76:1–76:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In quantum information theory, the quantum extension of Label-Cover, known as *entangled projection games* [14], is a special case of *nonlocal games* [10, 17].¹ The primary distinction lies in the definition of labeling. A *quantum labeling* generalizes probabilistic labeling. In classical probabilistic labeling, each vertex is assigned a probability distribution over labels. In quantum labeling, each edge is assigned a distribution over pairs of labels – one for each vertex incident on the edge.

More formally, a quantum labeling assigns quantum measurements to each vertex, with the measurement outcomes corresponding to labels. When both vertices of an edge are measured, the resulting probability distribution determines the likelihood that the edge is satisfied. The *quantum value* of the instance is the maximum expected fraction of satisfied edges over all possible quantum labelings. A formal definition is given in Section 3.1 of the full version [38].

In the quantum setting, the role of the PCP theorem is played by the $\text{MIP}^* = \text{RE}$ theorem of Ji et al. [29]: it is RE-hard to approximate the quantum value of Label-Cover within any additive constant.² Now, as with the classical case, one might expect to derive optimal inapproximability results for other quantum constraint satisfaction problems. O’Donnell and Yuen [41], for instance, extended Håstad’s results to the optimal inapproximability of quantum 3XOR. So far, the classical and quantum theories of inapproximability have followed parallel trajectories.

In the classical development, after the initial success of Håstad, researchers discovered some of the limitations of Label-Cover. To address these, Khot [32] introduced a modification to the Label-Cover problem, imposing a uniqueness condition on the projection maps. This modified problem is known as the Unique-Label-Cover and it remains an open question whether it is hard to approximate. However, assuming that Unique-Label-Cover is hard to approximate, researchers showed optimal inapproximability results for a variety of other problems, including MaxCut [34]. This assumption is known as the Unique Games Conjecture (UGC) [32].

A striking divergence arises between the classical and quantum theories of inapproximability with the introduction of Unique-Label-Cover. Kempe et al. [31] showed that the quantum value of Unique-Label-Cover is efficiently approximable. This implies that a *direct* quantum analogue of the UGC does not hold. As a result, none of the classical consequences of the UGC can be straightforwardly extended to the quantum domain. This presents a significant challenge to the development of a comprehensive quantum theory for hardness of approximation.

For instance, understanding the complexity of the anti-ferromagnetic Heisenberg model – a quantum generalization of MaxCut and a problem of great significance in physics and quantum information – remains an open question. While this problem is known to be QMA-hard, the complexity of approximating it is still unresolved.³ Hwang et al. [26] made significant progress on this question, but their results rely on the classical UGC, which limits them to showing NP-hardness rather than QMA-hardness.

Finding the right candidate for a quantum extension of UGC is key for advancing the theory of inapproximability in the quantum setting. This work proposes such a candidate.

¹ We primarily adopt the CSP terminology, while referencing certain naming conventions from the nonlocal games literature for the convenience of readers familiar with that field. See Section 1.2 for a discussion of our terminology.

² Recall that RE is the class of recursively enumerable languages. Stating that a problem is RE-hard is another way of saying that it is undecidable.

³ Recall that QMA is the quantum analogue of NP.

We begin by noting that Kempe et al.’s algorithm for approximating the quantum value of Unique-Label-Cover applies only to some of many possible quantum values that are defined in the literature on nonlocal games. Examples of quantum values include, but are not limited to, the tensor-product, commuting-operator, synchronous, and oracularizable-synchronous values. Label-Cover remains hard across all these models [29, 47, 48], whereas Unique-Label-Cover is efficiently approximable in only some of them [31].

We choose a model where Unique-Label-Cover might still be hard to approximate. This choice forms the basis of our *quantum Unique Games Conjecture* (qUGC). We give an informal overview in Section 1.3. Formal definitions can be found in Section 3.2 of the full version [38].

We then show how classical reductions from Unique-Label-Cover can be “quantized” to prove inapproximability results for quantum CSPs. We illustrate this approach with 2-Lin [32] and MaxCut [34], and we believe the same tools can be applied to Raghavendra’s seminal work on CSPs [44]. An informal overview is given in Section 1.4. The proof ideas are sketched in Section 1.5. Detailed statements and proofs are provided in Sections 4 and 5 of the full version [38].

We conclude this section with a few remarks.

Relationship between UGC and qUGC. The cases of 3XOR, 2-Lin, and MaxCut could be instances of a broader principle: any reduction from Label-Cover (or its variants such as Unique-Label-Cover, Smooth-Label-Cover, and 2-to-2 Label-Cover) might be adaptable to the quantum setting. If true, this would imply that the classical UGC implies qUGC. We explore this possibility and related open questions in Section 1.6.

Variants of qUGC. Section 1.3 and Section 3.2 in the full version [38] present several variants of the quantum Unique Games Conjecture. These variants may help in proving hardness for other models, such as the anti-ferromagnetic Heisenberg model or CSPs investigated in [12]. These topics are further detailed as Problems 4, 7, and 8 in Section 1.6.

Alternative Approaches to UGC. Some of the variants of quantum UGC proposed in this paper may be easier to prove hardness for than the classical UGC itself. Proving hardness for any of these variants could strengthen confidence in the UGC. Conversely, some of our UGC variants may be easier to disprove, and doing so could offer insights into disproving UGC. Figure 5, illustrates the reductions between variants of Unique-Label-Cover.

Universality. Label-Cover is exhibiting a remarkable *universality*, remaining hard across all the different quantum and classical models, see Figure 4. In contrast, the brittleness of Unique-Label-Cover aptly mirrors the difficulty in resolving the UGC. From another perspective, however, the constraints encountered in the search for a quantum analogue of the UGC could serve as a guiding principle in identifying the most appropriate model for quantum value from a computational standpoint.

This paper builds on our previous work on noncommutative CSPs [12], exploring how noncommutativity reshapes the landscape of CSPs. One of our motivations is the direct applications of this research to quantum information. We are also equally interested in whether a broader perspective on CSPs might uncover unifying principles that could provide new insights into classical CSPs. Our remarks on alternative approaches to UGC and universality are examples of the insights we are pursuing.

1.2 Quantum Value

We illustrate the concept of *quantum value* of a CSP instance using the example of MaxCut. For formal definitions refer to Section 3 in the full version [38]. Before proceeding, we briefly outline the connection to *nonlocal games*; however, no prior knowledge in quantum information is required to follow this work.

Nonlocal Games. The types of values we introduce for CSPs correspond directly to the values for *nonlocal games* previously studied in quantum information. Every CSP instance can be viewed as a nonlocal game and vice versa and the correspondence between their values is straightforward. This relationship is detailed in Section 10.2 of [12]. The physical, mathematical, and computational considerations for the definitions of the various types of values are extensively explored in the literature on nonlocal games [6, 9, 10, 20, 23, 29, 30, 35, 37, 42, 43, 46, 47, 48, 50, 51]. For readers familiar with this area, Table 1 summarizes the correspondences relevant to this paper. Please also refer to the discussion on previous work at the end of this section for our choice of terminology.

■ **Table 1** Correspondence between types of values for CSPs and nonlocal games. Formal definitions for the values in the left column are given in Section 3 of the full version [38].


CSP terminology	Nonlocal game terminology
classical	classical synchronous
noncommutative [12]	quantum synchronous [43, 23, 35, 37]
quantum	oracularizable quantum synchronous [23, 29, 37]

MaxCut. The objective of MaxCut is to partition the vertices of a given graph $G = (V, E)$ into two subsets, such that the number of edges crossing the partition is maximized (as illustrated in Figure 1).

Classical Value. A (classical) partition can be represented as an assignment $x : V \rightarrow \{-1, +1\}$. The value of this partition, meaning the number of edges that cross the partition, is given by $\sum_{(i,j) \in E} \frac{1 - x_i x_j}{2}$. The (classical) value of the instance, denoted $\omega_c(G)$, is the maximum partition value over all possible (classical) partitions:

$$\text{maximize: } \sum_{(i,j) \in E} \frac{1 - x_i x_j}{2} \tag{1}$$

subject to: $x_i \in \{-1, +1\}$.



■ **Figure 1** An instance of Max-Cut and an example of a partition (cut).

Noncommutative Value. A *noncommutative partition* assigns an observable to each vertex, i.e. $X : V \rightarrow \text{Obs}(\mathcal{H})$, where \mathcal{H} is some finite-dimensional Hilbert space. Observables are unitary operators with eigenvalues in the set $\{-1, +1\}$ (refer to Section 2.2 in the full version [38] for background on quantum measurements and observables). If we substitute observables for ± 1 in (1), we obtain the noncommutative value of the instance, denoted $\omega_{\text{nc}}(G)$:

$$\begin{aligned} \text{maximize: } & \sum_{(i,j) \in E} \frac{1 - \text{tr}(X_i X_j)}{2} \\ \text{subject to: } & X_i \in \text{Obs}(\mathcal{H}), \end{aligned} \quad (2)$$

where tr is the dimension-normalized trace on \mathcal{H} . Crucially, the optimization is over all finite-dimensional \mathcal{H} . When \mathcal{H} is one-dimensional, we recover the classical value, therefore $\omega_{\text{nc}}(G) \geq \omega_c(G)$.

The noncommutative value has been the focus of our previous work on CSPs [12]. The main focus of this paper, however, is the *quantum value* of CSPs, which we define next.

Quantum Value. A *quantum partition* lies somewhere in between a classical and noncommutative partition: A noncommutative partition X is *quantum* if X_i and X_j commute whenever (i, j) is an edge. In quantum mechanics, commuting observables can be simultaneously measured, emphasizing the physical motivation for this additional commutation relation.

The *quantum value*, denoted $\omega_q(G)$, is therefore obtained from (2) by imposing the extra commutation relations:

$$\begin{aligned} \text{maximize: } & \sum_{(i,j) \in E} \frac{1 - \text{tr}(X_i X_j)}{2} \\ \text{subject to: } & X_i \in \text{Obs}(\mathcal{H}), \\ & X_i X_j = X_j X_i \text{ for all } (i, j) \in E. \end{aligned} \quad (3)$$

Since every classical assignment is a quantum assignment and every quantum assignment is a noncommutative assignment, we obtain the chain of inequalities $\omega_{\text{nc}}(G) \geq \omega_q(G) \geq \omega_c(G)$.

Using the same approach outlined here, we can define the classical and quantum values for all CSPs (refer to Section 3 in the full version [38] for the relevant definitions). However, the noncommutative value is defined only for CSPs where each constraint involves only two variables.⁴

We assume $\omega_c, \omega_q, \omega_{\text{nc}}$ are normalized appropriately so that they are between 0 and 1. So for example in (1), (2), and (3), the normalization amounts to dividing the objective function by the number of edges.

Remark. It is important to note that the commutativity requirement in a quantum assignment does not imply that all observables in the assignment commute with each other, nor does it mean that all observables can be diagonalized in the same basis. Commutativity is not a transitive property: every operator commutes with the identity operator, but not every

⁴ This restriction is related to a similar issue for synchronous value of nonlocal games: with current definitions in the literature, the quantum synchronous value of a three-player nonlocal game reduces to the classical synchronous value. See also the remark after Definition 13 in the full version [38].

pair of operators commute with each other. For example, in MaxCut, if two vertices are at a distance of two or more, their observables in a quantum assignment need not commute.

Previous Work. The concept of noncommutative constraint satisfaction problems studied in this work traces its roots to the binary constraint system games introduced by Cleve and Mittal [11]. To the best of our knowledge, the first work to introduce a concept closely related to our notion of quantum assignment – referred to as oracularizable strategies for entangled nonlocal games – is by Ito et al. [27]. Oracularization is also a widely used technique in the study of classical nonlocal games (see references in [27]). Helton et al. [23] later examined oracularizable strategies from a mathematical perspective, employing the term locally commuting algebras (see Sections 7 and 8 of their paper).

Our definition of quantum assignment is the direct translation (for CSPs) of the *oracularizable quantum synchronous strategies* in [37]. The definition in [37] itself dates back to the breakthrough results [40, 29]. In this work, we refer to these assignments as *quantum* because the requirement of *commutativity* stems from *simultaneous measurability in quantum mechanics*. In contrast, we refer to assignments as *noncommutative* when there is no commutation requirement. This terminology originates from our earlier work on noncommutative CSPs [12].⁵

1.3 Quantum Unique Games Conjecture

Before moving on to our conjecture, we first briefly review decision problems and reductions between them. Readers familiar with the terminology can skip the next section.

1.3.1 Reductions

In the context of hardness of approximation, it is often easier to work with a version of a computational problem known as the *decision problem with a promise*. Let \mathcal{P} be any computational problem for which $\omega_c, \omega_q, \omega_{nc}$ are defined. For example \mathcal{P} could be MaxCut, Label-Cover, or Unique-Label-Cover. For every $s, t \in \{c, q, nc\}$ and $\zeta, \delta > 0$, the *decision problem* $\mathcal{P}_{s,t}(1 - \zeta, \delta)$ is defined as follows: given an instance $\phi \in \mathcal{P}$ that is promised to be either

- $\omega_s(\phi) \geq 1 - \zeta$, called a “yes” instance, or
- $\omega_t(\phi) < \delta$, called a “no” instance,


decide whether the instance is a “yes” or “no” instance. When $s = t$, we simply write $\mathcal{P}_s(1 - \zeta, \delta)$ to denote $\mathcal{P}_{s,s}(1 - \zeta, \delta)$.

A *reduction* from $\mathcal{P}_{s,t}(\gamma, \delta)$ to $\mathcal{P}'_{s',t'}(\gamma', \delta')$ is any efficient map $\mathcal{P} \rightarrow \mathcal{P}'$ that is

- complete: it maps “yes” instances to “yes” instances, and
- sound: it maps “no” instances to “no” instances.

As a simple example, for every choice of $\zeta, \delta > 0$, the decision problem $\mathcal{P}_{q,nc}(1 - \zeta, \delta)$ reduces to $\mathcal{P}_{nc,nc}(1 - \zeta, \delta)$ via the identity map. Soundness is trivial, and completeness holds because $\omega_{nc}(\phi) \geq \omega_q(\phi)$ for all instances ϕ . This is denoted by $\mathcal{P}_{q,nc} \rightarrow \mathcal{P}_{nc,nc}$ (we often drop the completeness and soundness parameters in our informal discussions in this introduction). There are several other trivial reductions between decision versions of \mathcal{P} , and they are illustrated in Figure 2.

⁵ Our decision to use *quantum synchronous strategies*, both here and in [12], is driven by the syntactical simplicity they provide compared to tensor-product strategies. While we believe that all of our results can be extended to the tensor-product model, we leave a more detailed exploration of this for future work.



■ **Figure 2** Trivial Reductions: if a problem in this diagram is hard (for some complexity class), then every problem reachable from it is at least as hard.

Lastly, we use the abbreviations LC for Label-Cover and ULC for Unique-Label-Cover.

1.3.2 Label-Cover

The PCP theorem states that the classical value of Label-Cover is NP-hard to approximate, or more precisely:

► **Theorem 1** (PCP Theorem). *For every $\delta > 0$, there is a large enough alphabet over which $\text{LC}_c(1, \delta)$ is NP-hard.*

Since we have the trivial reductions $\text{LC}_{c,c} \rightarrow \text{LC}_{q,c} \rightarrow \text{LC}_{nc,c}$, we conclude that both $\text{LC}_{q,c}$ and $\text{LC}_{nc,c}$ are at least NP-hard. This is illustrated in Figure 3.

What can we say about the other variants of Label-Cover in Figure 3? The noncommutative analogue of the PCP theorem is the $\text{MIP}^* = \text{RE}$ theorem. This theorem states that the noncommutative value of Label-Cover is RE-hard to approximate. In fact [29] proves a stronger theorem:

► **Theorem 2** ($\text{MIP}^* = \text{RE}$, informal). *For every $\delta > 0$, the decision problem $\text{LC}_{q,nc}(1, \delta)$ is RE-hard.*

This immediately implies that the problems in the last two rows in Figure 3 are RE-hard. Our full knowledge of the complexity landscape of Label-Cover is summarized in Figure 4.


1.3.3 Unique-Label-Cover

A Unique-Label-Cover instance is a special case of Label-Cover instance where every projection map is a bijection (i.e. a permutation), see Section 3.2 in the full version [38] for further details. The well-known Unique Games Conjecture asserts that it is hard to approximate the classical value of Unique-Label-Cover:


► **Conjecture 3** (UGC [32], informal). *For all $\zeta, \delta > 0$, the decision problem $\text{ULC}_c(1 - \zeta, \delta)$ is NP-hard.*

Kempe et al. [31] provided an approximation algorithm for the noncommutative value of Unique-Label-Cover (see Theorem 19 in the full version [38] for details), but this result does not address the quantum value. Thus, we propose the following quantum analogue of the Unique Games Conjecture (see Conjecture 20 in the full version [38] for the precise statement):

► **Conjecture 4** (qUGC, informal). *For all $\zeta, \delta > 0$, the decision problem $\text{ULC}_q(1 - \zeta, \delta)$ is RE-hard.*



■ **Figure 3** Implication of the PCP theorem: problems in the third column are NP-hard.



■ **Figure 4** The complexity landscape of Label-Cover.

Assuming both UGC and qUGC, along with the algorithm from Kempe et al., the complexity landscape of Unique-Label-Cover can be summarized as shown in Figure 5.

Assuming that any one of $ULC_{q,q}$, $ULC_{q,c}$, $ULC_{nc,q}$, $ULC_{nc,c}$ is RE-hard constitutes a variant of the qUGC. The strongest among these variants is the conjecture concerning $ULC_{q,q}$. As discussed in Section 3.2 in the full version [38], there are many (indeed, infinitely many) variants of qUGC. When we refer to qUGC without additional qualifiers, we are specifically referring to the conjecture stated above.

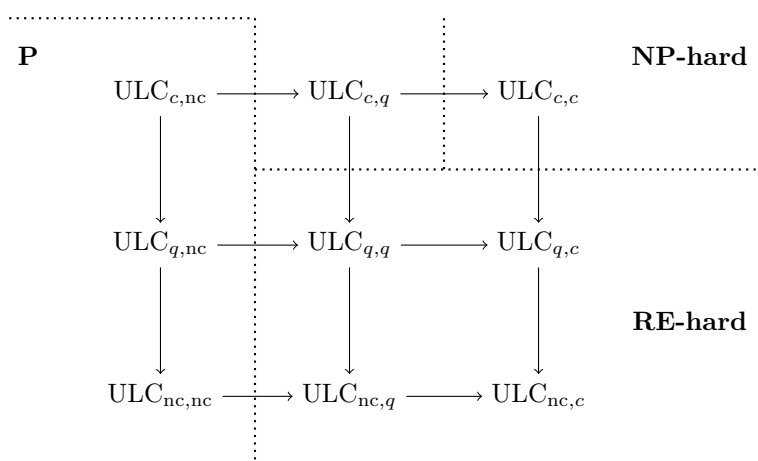
Given the reductions shown in Figure 5, it is likely easier to prove hardness for $ULC_{q,c}$ or $ULC_{nc,c}$ than to prove either UGC or qUGC. Conversely, devising a polynomial-time approximation algorithm for $ULC_{c,q}$ seems more achievable than disproving either UGC or qUGC. Note that $ULC_{c,q}$ is in NP (since $ULC_{c,c}$ is in NP).

1.3.4 Applications to Hardness of Approximation

The result by O’Donnell and Yuen [41] on the quantum value of 3XOR is stated in the next theorem.

► **Theorem 5.** *For every $\delta > 0$, the decision problem $3XOR_q(1 - \delta, \frac{1}{2} + \delta)$ is RE-hard.*

This result was originally proven via a reduction from LC_{nc} . However, a reduction from LC_q suffices with some straightforward modifications. The hardness of LC_q follows directly from the $MIP^* = RE$ theorem (as shown in Figure 4). Since a direct proof of LC_q is likely easier than the full proof of $MIP^* = RE$ [29], we propose Problem 5 in Section 1.6.



■ **Figure 5** The complexity landscape of Unique-Label-Cover assuming both UGC and qUGC.



(a) Classical value [25].

(b) Quantum value [41].

■ **Figure 6** Approximability of classical and quantum values of 3XOR: the vertical axis denotes approximation ratios between 0 and 1, with points of sharp transition in hardness (0.5 in both cases) indicated.

This theorem mirrors Håstad’s celebrated result on the classical value of 3XOR [25]. Our current understanding of approximability of 3XOR is summarized in Figure 6.

In Sections 4 and 5 in the full version [38], we prove similar results for 2-Lin and MaxCut. Our reductions are from ULC_q rather than LC_q . An informal overview of our MaxCut result is provided in the next section.

1.4 Result on MaxCut


How well can the classical, quantum, and noncommutative values of MaxCut be approximated? The current knowledge, including our result, is summarized in Figure 7.

The well-known Goemans-Williamson algorithm [21] for the classical value of MaxCut achieves an approximation ratio of $\alpha_{\text{gw}} \approx 0.878$. More specifically, they showed that $\omega_c(G) \geq \alpha_{\text{gw}} \omega_{\text{sdp}}(G)$, where $\omega_{\text{sdp}}(G)$ is the optimal value of a certain semidefinite programming (SDP) relaxation for MaxCut. Meanwhile, Tsirelson [50, 51] showed that $\omega_{\text{nc}}(G) = \omega_{\text{sdp}}(G)$. Combining these results yields the following chain of inequalities:


$$\alpha_{\text{gw}} \omega_{\text{sdp}}(G) \leq \omega_c(G) \leq \omega_q(G) \leq \omega_{\text{nc}}(G) = \omega_{\text{sdp}}(G). \tag{4}$$

From this chain of inequalities (see also Figure 8), we infer that the Goemans-Williamson algorithm is also a 0.878-approximation algorithm for the quantum value of MaxCut. But can we do better for the quantum value?

76:10 A Quantum Unique Games Conjecture



■ **Figure 7** Approximability of MaxCut across different types of values.



■ **Figure 8** The MaxCut interval.

Under the assumption of the quantum Unique Games Conjecture introduced in this paper, we show:

The classical Goemans-Williamson algorithm is not only the best efficient approximation algorithm for the quantum value; it is the best approximation algorithm for the quantum value—regardless of efficiency.

► **Theorem 6** (Theorem 31 in the full version [38], informal). *Assuming the quantum Unique Games Conjecture, it is RE-hard to approximate the quantum value of MaxCut to any ratio better than α_{gw} .*

Compare this with the following optimal inapproximability result from Khot et al. [34]:

► **Theorem 7** (KKMO, informal). *Assuming the Unique Games Conjecture, it is NP-hard to approximate the classical value of MaxCut to any ratio better than α_{gw} .*


Integrality Gap. Similar to the Goemans-Williamson SDP relaxation, the quantum value—though undecidable—serves as a relaxation to the classical value. A natural question arises: how close is this quantum relaxation to the classical value? Can the quantum value outperform the SDP value in approximating the classical value?

Under the assumption of the quantum Unique Games Conjecture, we show that the SDP and quantum relaxations are incomparable. This is illustrated in Figure 9 and formally established in Theorem 35 in the full version [38]. Constructing examples that exhibit each of the extremes shown in Figure 9 remains an open problem.

1.5 Proof Ideas

In the conventional NP proof system for Label-Cover, the certificate is simply the list of labels for the vertices. This is what we called *classical assignment* in Section 1.2. The verifier’s task in this proof system is to check that the assigned labels satisfy the edge constraints.

Håstad’s gap-preserving reduction from Label-Cover (used, for instance, in proving the hardness of 3SAT and 3XOR) is best viewed as the construction of a more efficient version of this NP proof system, known as a probabilistically checkable proof (PCP). The PCP certificate for Label-Cover is an encoding of the classical assignment in binary using an error-correcting code, commonly referred to as the Long-Code, pioneered in [7]. In addition



■ **Figure 9** There exist instances realizing both extremes depicted in this picture. That is, among the integrality gap instances of MaxCut (see Feige and Schechtman [18]), i.e. those instances with $\omega_c(G) = \alpha_{\text{gw}}\omega_{\text{sdp}}(G)$, there are instances for which $\omega_q(G) = \omega_c(G)$ and instances for which $\omega_q(G) = \omega_{\text{sdp}}(G)$.

to verifying that the labels in the encoding satisfy the edge constraints, the PCP verifier must also check that the proof consists of valid codewords. The soundness of this PCP relies on Fourier analysis. This was one of the key insights of Håstad’s celebrated work.

In our case, we are reducing Unique-Label-Cover to 2-Lin (Section 4 in the full version [38]) and MaxCut (Section 5 in the full version [38]), but the core principles remain the same. Our reductions are based on the classical reductions by Khot [32] and KKMO [34], respectively. There are, however, several key differences, which we discuss in the rest of this section. For concreteness, we focus on the reduction from Unique-Label-Cover to MaxCut.

Completeness. This step requires a generalization of the Long-Code to encode quantum assignments. This is a straightforward generalization of the Long-Code to the operator setting and it is due to O’Donnell and Yuen [41]. The main challenge in proving completeness is ensuring that the encoded quantum assignment satisfies the commutation requirements of a quantum MaxCut assignment.

Soundness. The proof of soundness, as is customary, is done contrapositively. We begin with a quantum assignment for the MaxCut instance that achieves a large value, and from this, we construct a quantum assignment that has a sufficiently large value in the Unique-Label-Cover instance.

The classical proof constructs a probabilistic labeling based on the Fourier coefficients of the MaxCut assignment. Given a MaxCut assignment f , the labeling samples a subset of labels S with probability proportional to the squared Fourier coefficients, $\hat{f}(S)^2$. It then probabilistically selects a label $a \in S$ for each vertex. In the quantum setting, these squared Fourier weights are positive semidefinite operators. Defining $P^a := \sum_{S:a \in S} \frac{1}{|S|} \hat{f}(S)^2$ for each label a results in a *POVM measurement* for each vertex in the Unique-Label-Cover instance (see Section 2.2 in the full version [38] for the definition of POVMs). This is also due to [41].

These measurements satisfy the commutation relations required for a quantum assignment. However, they are not projective measurements, as required by definition of quantum assignment. Naimark’s dilation theorem (see, for example, Watrous [52]) is typically used in such cases to convert POVMs into projective measurements. However, due to the commutation requirements (and the use of tracial states), Naimark’s theorem is not applicable. Therefore, we need to projectivize these POVMs while preserving the commutation relations, which we accomplish using the projectivization lemma proved in Section 6 in the full version [38].

Classically, the proof of soundness relies on a deep result from Boolean Fourier analysis, known as the Majority-Is-Stablest [36], which relates Boolean functions to their Fourier weights. If the operators in our assignments were simultaneously diagonalizable, we could apply this theorem directly, as in the classical proof. However, the inherent noncommutativity of quantum assignments prevents us from doing so, forcing us to take a new approach.

To explain the problem and our resolution at a high level, recall that our goal is to construct a quantum labeling for an instance of Unique-Label-Cover. Let u be a vertex from the left-vertex set, and let N_u denote its neighbors in the right-vertex set. In the operator setting, while each neighbor's operator individually commutes with u 's, the operators assigned to the neighbors may not commute with each other. This noncommutativity is a fundamental obstacle.

We first encountered this issue earlier, when attempting to view quantum assignments as extensions of probabilistic assignments. In a classical probabilistic assignment, each vertex u is assigned a single probability measure P_u over the set of labels. In contrast, a quantum assignment associates two probability measures, $P_{u,e}$ and $P_{v,e}$, with each edge $e = (u, v)$. The crux of the issue is that different edges incident on u – say, e and e' – can impose vastly different probability measures on u , which may conflict with each other.

This same phenomenon arises in the proof of soundness, except it manifests not in the probability measures, but in the Fourier coefficients of the quantum assignment. Classical soundness proofs organize around sets of neighbors, but this approach is too coarse to be applicable to the operator setting. The key to resolving this issue is to structure the proof around individual edges; this allows us to manage the noncommutative nature of the operator assignments without the sort of conflicts we mentioned earlier.

1.6 Future Directions

There are many open problems concerning the quantum value of CSPs and the conjectures we have proposed regarding Unique-Label-Cover:

1. **Quantizing Classical Reductions:** The works of O'Donnell and Yuen [41], as well as this paper, establish straightforward translations of some of the most well-known classical inapproximability results to the quantum setting. Could this indicate a more general phenomenon? More precisely, can any classical gap-preserving reduction from Label-Cover (or its variants like Unique-Label-Cover or Smooth-Label-Cover) be extended to a reduction for the quantum value?
2. **UGC versus qUGC:** What is the relationship between the classical UGC and the quantum UGC? If UGC is proven via a reduction from Label-Cover, can this reduction be “quantized” to prove qUGC? A positive answer to the previous question would also resolve this one affirmatively.
3. **Variants of qUGC:** In the classical setting, several formulations of UGC have been shown to be equivalent (see Section 3 in [33] and references therein). Could similar equivalences hold for the quantum variants discussed in this paper? See Section 3.2 in the full version [38] for further discussion of these qUGC variants.
4. **Noncommutative CSPs:** This work focuses primarily on the quantum value of CSPs. Could variants of qUGC help resolve hardness of approximation questions for the noncommutative value? The noncommutative value, which also plays an important role in quantum information, exhibits some intriguing behavior. As shown in Figure 7, unlike the classical or quantum value, the noncommutative value of MaxCut can be computed in polynomial time. However, it is known that the noncommutative value of Max-3-Cut is undecidable [28, 29]. [12] provides a 0.864-approximation algorithm for the noncommutative value of Max-3-Cut. We suspect that a variant of qUGC could prove useful in establishing the optimality of this algorithm.
5. **Weaker Version of $\text{MIP}^* = \text{RE}$:** Is there a direct way to prove the hardness of approximating the quantum value of Label-Cover without relying on $\text{MIP}^* = \text{RE}$? One potential approach is to quantize Dinur's proof of the classical PCP theorem [3, 2, 45, 13].

The commutation relations in LC_q could make this quantization feasible. Could this weaker result also have some of the implications of $MIP^* = RE$ such as the resolution of the Connes Embedding Problem?

6. **Integrality Gap Instances for MaxCut:** What is the smallest ratio between the classical and quantum values across all MaxCut instances? Assuming quantum UGC, we show this ratio is the Goemans-Williamson constant $\alpha_{gw} \approx 0.878$ (Theorem 35 in the full version [38]). Can we construct an instance attaining this ratio?
7. **Label-Cover for QMA:** Approximating the classical value of Label-Cover is NP-complete (PCP theorem). Similarly, approximating the noncommutative and quantum values of Label-Cover is RE-complete ($MIP^* = RE$ theorem). What version of Label-Cover or Unique-Label-Cover naturally captures QMA? Could further constraints on the variations of Label-Cover and Unique-Label-Cover explored in this paper lead to suitable candidates? A positive resolution to this problem would answer the Quantum Games PCP Conjecture [39].
8. **Hardness of Approximation in Hamiltonian Complexity:** In quantum information theory, quantum MaxCut often refers to the anti-ferromagnetic Heisenberg model, a special case of the local Hamiltonian problem. Although [26] attempts to prove UGC hardness for the anti-ferromagnetic Heisenberg model, reductions from UGC only establish NP-hardness. Could variations of qUGC yield stronger results?
While the inapproximability of the anti-ferromagnetic Heisenberg model is one of our motivations for proposing the quantum UGC, our current methods do not yet apply to the inapproximability of local Hamiltonian problems.
9. **Gadget Reductions:** In classical complexity theory, gadgets [25, 49] are widely used to demonstrate hardness of approximation. Can this method be “quantized”? [22, 28] are examples of gadget reductions in the quantum setting.

Organization of the Full Version of the Paper [38]

In Section 2, we present notations, definitions, and some Fourier analytic results that underpin this work. Section 2.1 specifically addresses the notations used throughout the paper. In Section 2.2, we introduce the concept of quantum measurement and discuss several definitions related to projective measurements. Section 2.3 reviews concepts such as reductions and decision problems, and Section 2.4 summarizes results from Fourier analysis.

Section 3.1 offers the most general definition of CSPs necessary for this work. It also covers the classical, quantum, and noncommutative values of a CSP instance, providing definitions which are crucial for the remainder of the paper. In Section 3.2, we discuss variants of quantum UGC. Section 4 contains the proof of our theorem on the quantum value of 2-Lin, and Section 5 presents the proof of our theorem on the quantum value of MaxCut. Finally, two technical lemmas are given in Sections 6 and 7.

References

- 1 Sanjeev Arora, László Babai, Jacques Stern, and Z Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. *Journal of Computer and System Sciences*, 54(2):317–331, 1997. doi:10.1006/jcss.1997.1472.
- 2 Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, May 1998. doi:10.1145/278298.278306.

- 3 Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, January 1998. doi:10.1145/273865.273901.
- 4 László Babai, Lance Fortnow, Leonid A. Levin, and Mario Szegedy. Checking computations in polylogarithmic time. In *Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing*, STOC '91, pages 21–32, New York, NY, USA, 1991. Association for Computing Machinery. doi:10.1145/103418.103428.
- 5 László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1(1):3–40, March 1991. doi:10.1007/BF01200056.
- 6 John Stewart Bell. On the Einstein Podolsky Rosen paradox. *Physics*, 1(3):195–200, 1964. doi:10.1103/PhysicsPhysiqueFizika.1.195.
- 7 Mihir Bellare, Oded Goldreich, and Madhu Sudan. Free bits, PCPs and non-approximability-towards tight results. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*, pages 422–431, 1995. doi:10.1109/SFCS.1995.492573.
- 8 Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson. Multi-prover interactive proofs: How to remove intractability assumptions. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, STOC '88, pages 113–131, New York, NY, USA, 1988. ACM. doi:10.1145/62212.62223.
- 9 John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt. Proposed experiment to test local hidden-variable theories. *Phys. Rev. Lett.*, 23:880–884, October 1969. doi:10.1103/PhysRevLett.23.880.
- 10 Richard Cleve, Peter Hoyer, Benjamin Toner, and John Watrous. Consequences and limits of nonlocal strategies. In *Proceedings of the 19th IEEE Annual Conference on Computational Complexity*, CCC '04, pages 236–249, Washington, DC, USA, 2004. IEEE Computer Society. doi:10.1109/CCC.2004.9.
- 11 Richard Cleve and Rajat Mittal. Characterization of binary constraint system games. In *International Colloquium on Automata, Languages, and Programming (ICALP) 2012*, pages 320–331, 2012. doi:10.1007/978-3-662-43948-7_27.
- 12 Eric Culf, Hamoon Mousavi, and Taro Spirig. Approximation algorithms for noncommutative constraint satisfaction problems, 2023. doi:10.48550/arXiv.2312.16765.
- 13 Irit Dinur. The pcp theorem by gap amplification. *J. ACM*, 54(3):12–es, June 2007. doi:10.1145/1236457.1236459.
- 14 Irit Dinur, David Steurer, and Thomas Vidick. A parallel repetition theorem for entangled projection games. In *2014 IEEE 29th Conference on Computational Complexity (CCC)*, pages 197–208, 2014. doi:10.1109/CCC.2014.28.
- 15 Uriel Feige, Shafi Goldwasser, Laszlo Lovász, Shmuel Safra, and Mario Szegedy. Interactive proofs and the hardness of approximating cliques. *J. ACM*, 43(2):268–292, March 1996. doi:10.1145/226643.226652.
- 16 Uriel Feige and Joe Kilian. Two-prover protocols—low error at affordable rates. *SIAM Journal on Computing*, 30(1):324–346, 2000. doi:10.1137/S0097539797325375.
- 17 Uriel Feige and László Lovász. Two-prover one-round proof systems: their power and their problems (extended abstract). In *Proceedings of the Twenty-Fourth Annual ACM Symposium on Theory of Computing*, STOC '92, pages 733–744, New York, NY, USA, 1992. Association for Computing Machinery. doi:10.1145/129712.129783.
- 18 Uriel Feige and Gideon Schechtman. On the optimality of the random hyperplane rounding technique for max cut. *Random Structures & Algorithms*, 20(3):403–440, 2002. doi:10.1002/rsa.10036.
- 19 Lance Fortnow, John Rempel, and Michael Sipser. On the power of multi-prover interactive protocols. *Theoretical Computer Science*, 134(2):545–557, 1994. doi:10.1016/0304-3975(94)90251-8.
- 20 Tobias Fritz. Tsirelson's problem and Kirchberg's conjecture. *Reviews in Mathematical Physics*, 24(05):1250012, 2012. doi:10.1142/S0129055X12500122.

- 21 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, November 1995. doi:10.1145/227683.227684.
- 22 Samuel J. Harris. Universality of Graph Homomorphism Games and the Quantum Coloring Problem. *Annales Henri Poincare*, 25(10):4321–4356, 2024. doi:10.1007/s00023-024-01422-5.
- 23 John William Helton, Kyle P Meyer, Vern I Paulsen, and Matthew Satriano. Algebras, synchronous games and chromatic numbers of graphs. *arXiv preprint*, 2017. doi:10.48550/arXiv.1703.00960.
- 24 Johan Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. In *Proceedings of 37th Conference on Foundations of Computer Science*, pages 627–636, 1996. doi:10.1109/SFCS.1996.548522.
- 25 Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, July 2001. doi:10.1145/502090.502098.
- 26 Yeongwoo Hwang, Joe Neeman, Ojas Parekh, Kevin Thompson, and John Wright. *Unique Games hardness of Quantum Max-Cut, and a conjectured vector-valued Borell’s inequality*, pages 1319–1384. Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2023. doi:10.1137/1.9781611977554.ch48.
- 27 Tsuyoshi Ito, Hirotada Kobayashi, and Keiji Matsumoto. Oracularization and two-prover one-round interactive proofs against nonlocal strategies. *2009 24th Annual IEEE Conference on Computational Complexity*, pages 217–228, 2008. doi:10.1109/CCC.2009.22.
- 28 Zhengfeng Ji. Binary constraint system games and locally commutative reductions, 2013. doi:10.48550/arXiv.1310.3794.
- 29 Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. $MIP^* = RE$. *Commun. ACM*, 64(11):131–138, October 2021. doi:10.1145/3485628.
- 30 Marius Junge, Miguel Navascues, Carlos Palazuelos, David Perez-Garcia, Volker B. Scholz, and Reinhard F. Werner. Connes’ embedding problem and tsirelson’s problem. *Journal of Mathematical Physics*, 52(1):012102, 2011. doi:10.1063/1.3514538.
- 31 Julia Kempe, Oded Regev, and Ben Toner. Unique games with entangled provers are easy. *SIAM Journal on Computing*, 39(7):3207–3229, 2010. doi:10.1137/090772885.
- 32 Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings 17th IEEE Annual Conference on Computational Complexity*, pages 25–, 2002. doi:10.1109/CCC.2002.1004334.
- 33 Subhash Khot. On the unique games conjecture. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS’05)*, pages 3–, 2005. doi:10.1109/SFCS.2005.61.
- 34 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. doi:10.1137/S0097539705447372.
- 35 Se-Jin Kim, Vern Paulsen, and Christopher Schafhauser. A synchronous game for binary constraint systems. *Journal of Mathematical Physics*, 59(3):032201, 2018. doi:10.1063/1.4996867.
- 36 Elchanan Mossel, Ryan O’Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: Invariance and optimality. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS’05)*, pages 21–30, 2005. doi:10.1109/SFCS.2005.53.
- 37 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*, STOC 2022, pages 1–11, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3519935.3519949.
- 38 Hamoon Mousavi and Taro Spirig. A quantum unique games conjecture, 2024. doi:10.48550/arXiv.2409.20028.
- 39 Anand Natarajan and Chinmay Nirkhe. The status of the quantum pcg conjecture (games version), 2024. doi:10.48550/arXiv.2403.13084.

76:16 A Quantum Unique Games Conjecture

- 40 Anand Natarajan and John Wright. $\text{NEEXP} \subseteq \text{MIP}^*$. In *IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 510–518, 2019. doi:10.48550/arXiv.1904.05870.
- 41 Ryan O’Donnell and Henry Yuen. Private communication, 2024.
- 42 Narutaka Ozawa. About the Connes embedding conjecture: Algebraic approaches. *Jpn. J. Math.*, 8:147–183, 2013. doi:10.1007/s11537-013-1280-5.
- 43 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. doi:10.1016/j.jfa.2016.01.010.
- 44 Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, STOC ’08*, pages 245–254, New York, NY, USA, 2008. Association for Computing Machinery. doi:10.1145/1374376.1374414.
- 45 Ran Raz. A parallel repetition theorem. In *Proceedings of the Twenty-Seventh Annual ACM Symposium on Theory of Computing, STOC ’95*, pages 447–456, New York, NY, USA, 1995. Association for Computing Machinery. doi:10.1145/225058.225181.
- 46 Volkher B Scholz and Reinhard F Werner. Tsirelson’s problem, 2008. doi:10.48550/arXiv.0812.4305.
- 47 William Slofstra. The set of quantum correlations is not closed. In *Forum of Mathematics, Pi*, volume 7. Cambridge University Press, 2019. doi:10.1017/fmp.2018.3.
- 48 William Slofstra. Tsirelson’s problem and an embedding theorem for groups arising from non-local games. *Journal of the American Mathematical Society*, 2019. doi:10.1090/jams/929.
- 49 Luca Trevisan, Gregory B. Sorkin, Madhu Sudan, and David P. Williamson. Gadgets, approximation, and linear programming. *SIAM Journal on Computing*, 29(6):2074–2097, 2000. doi:10.1137/S0097539797328847.
- 50 Boris Tsirelson. Quantum analogues of the Bell inequalities. the case of two spatially separated domains. *Journal of Soviet Mathematics*, 36(4):557–570, 1987. doi:doi:10.1007/BF01663472.
- 51 Boris Tsirelson. Some results and problems on quantum Bell-type inequalities. *Hadronis Journal Supplement*, 8:320–331, 1993. URL: <https://www.tau.ac.il/~tsirel/download/hadron.pdf>.
- 52 John Watrous. *The Theory of Quantum Information*. Cambridge University Press, 2018. doi:10.1017/9781316848142.