

AN INTRODUCTION TO THE THEORY OF GRAPH LIMITS

MARC DE FONTNOUVELLE

ABSTRACT. The theory of graph limits enables us to analyze discrete combinatorial objects, such as graphs, through the study of continuous objects such as *graphons*. In this paper, we define a notion of *quasirandom* graph sequences and present a surprising and elegant property of such sequences. These ideas naturally lead us toward the development of the theory of graph limits. We then use this rich theory to prove the aforementioned property of quasirandom graph sequences.

CONTENTS

1. Introduction: On Randomness	1
1.1. Graph Limits	3
2. Graphs and Sequences	3
3. Kernels and Graphons	4
4. Kernel Operators	8
5. An Infinitary Proof	9
Acknowledgments	12
References	13

1. INTRODUCTION: ON RANDOMNESS

Much of modern mathematics deals with the notion of randomness. Probabilists study random walks; computer scientists develop random algorithms; logicians analyse random graphs. But what makes a mathematical object – say a number, graph, or function – random?

For a finite mathematical object – a real number, a finite graph, a function between finite sets – this question can not admit any reasonable answer. But for infinite *sequences* of finite objects, a notion of randomness that matches our intuition can often be defined. For example, consider the collection of infinite binary sequences. We would like to develop a notion of randomness under which sequences such as 0101010101... are not considered random, whereas sequences such as 0111010011..., whose digits have no discernable pattern, are considered random.

It turns out that this is a surprisingly difficult task: many definitions of random sequences that appear intuitive either admit sequences whose digits have a discernible pattern, or admit no sequences at all. It was not until the 1960s, when the logician Per Martin-Löf had the idea to use the theory of computation to formalize the notion of a “randomness test,” that a suitable definition of random sequences

was found. We refer the reader to Chapter 9 of [2] for further discussion on this topic, including a definition of Martin-Löf randomness.

Curiously, when we consider sequences of finite *graphs*, rather than sequences of bits, the situation becomes somewhat simpler. In 1989, Fan Chung, Ronald Graham, and Richard Wilson developed an intuitive and versatile theory of “quasirandom” sequences of graphs [3]. By a sequence of graphs, we mean a sequence (G_n) such that for each $n \in \mathbb{N}$, G_n is a finite graph. Let us define $v(G_n)$ to be the number of vertices in G_n , and let us only consider those graph sequences (G_n) for which $v(G_n) \rightarrow \infty$. For simplicity of notation in the forthcoming definitions, let us further assume that $v(G_n) = n$ for all n .

The theory of quasirandom graph sequences relies upon the notion of a graph homomorphism. A graph homomorphism is a map between the vertex sets of two graphs which preserves edges (we will define this notion formally later, at the beginning of Section 2). We can now define a quasirandom graph sequence as follows:

Definition 1.1 (Quasirandom graph sequence). A graph sequence $G(n)$ is *quasirandom* with density $p \in (0, 1)$ if for every finite graph H , the number of homomorphisms of H into G_n is asymptotically equal to $p^{e(H)} n^{v(H)}$.

Intuitively, this definition stipulates that the edge density of the graphs in (G_n) is asymptotically p – and furthermore, that the number of appearances of any subgraph H in G_n asymptotically depends only on the number of vertices and edges in H , rather than on the specific structure of H . A sequence of graphs which is constructed in such a way as to either generate or avoid the presence of any particular subgraph will *not* be quasirandom according to this definition.

This definition is rather intuitive and elegant, especially in comparison to the definition of random sequences. One of the most remarkable aspects of this definition is the variety of equivalent ways in which it can be formulated. In fact, each of the following is an equivalent definition of a quasirandom sequence of graphs: ¹

Definition 1.2. A graph sequence $G(n)$ is *quasirandom* with density $p \in (0, 1)$ if the number of edges contained within any set of $n/2$ nodes is asymptotically equal to $pn^2/8$.

Definition 1.3. A graph sequence $G(n)$ is *quasirandom* with density $p \in (0, 1)$ if the number of edges in G_n is asymptotically equal to $pn^2/2$, and the number of 4-cycles is asymptotically equal to $p^4 n^4/8$.

Definition 1.4. A graph sequence $G(n)$ is *quasirandom* with density $p \in (0, 1)$ if for any two disjoint sets A and B of vertices, the number of edges between A and B is $p|A||B| + o(n^2)$.

Definition 1.5. A graph sequence $G(n)$ is *quasirandom* with density $p \in (0, 1)$ if degrees of all but finitely many vertices are asymptotically pn , and the codegrees (number of common neighbors of two nodes) of all but finitely many pairs of vertices are asymptotically $p^2 n$.

¹These definitions are drawn from Chapter 1 of [1].

Perhaps the most surprising of these equivalences is that between Definitions 1.1 and 1.3. If a sequence of graphs contains asymptotically “random” numbers of edges and 4-cycles, it must also contain asymptotically “random” numbers of every other finite graph!

1.1. Graph Limits.

In defining a quasirandom graph sequence (G_n) (Definition 1.1), we stipulated that the number of homomorphisms of every finite graph H into G_n must be asymptotically equal to a certain scalar multiple of $n^{v(H)}$, which is in turn the number of *total* homomorphisms from H into G_n . Furthermore, we required that the value of this scalar multiple depend only on $e(H)$. It is natural to ask what would happen if we removed this second requirement. That is, what if we considered sequences (G_n) such that the number of homomorphisms of H into G_n is asymptotically equal to some scalar multiple of $n^{v(H)}$ for each H – but imposed no restriction on what this scalar multiple can be?

It is precisely this idea that underlies the theory of *graph limits*. In the following section, we will formalize the notion of a “convergent” sequence of graphs. Furthermore, we will specify a class of objects called “graphons” *to which* sequences of graphs converge. Whereas graphs are defined over a discrete set of vertices, graphons are defined over a continuous domain.

Precisely because of their continuous nature, graphons turn out to be a very powerful tool for studying and analyzing finite graphs. Theorems and properties about graphs are often easier to state and prove in the space of graphons – and proofs in the space of graphons can often be transferred back into the space of graphs. As an example of this proof technique, we will use the theory of graphons to prove the aforementioned property of quasirandom graphs – that is, the equivalence between Definitions 1.1 and 1.3.

2. GRAPHS AND SEQUENCES

Let us begin by formally defining graphs and homomorphisms therebetween. A graph G is given by a set $V(G)$ of vertices together with a set $E(G) \subset V(G)^2$ of edges between vertices. We will write $v(G)$ to denote $|V(G)|$, and $e(G)$ to denote $|E(G)|$. In a *simple* graph G , edges are undirected, and loops are forbidden. That is, for all $x, y \in V(G)$, $(x, y) \in E(G) \Leftrightarrow (y, x) \in E(G)$, and for all $x \in V(G)$, $(x, x) \notin E(G)$.

In a *weighted* graph G , every vertex $x \in V(G)$ is assigned a positive real weight $\alpha_x(G)$, and every edge $(x, y) \in E(G)$ is assigned a real weight $\beta_{(x, y)}(G)$. Furthermore, loops are permitted, i.e., we may have $(x, x) \in E(G)$ for some $x \in V(G)$. In this paper, graphs will be assumed to be simple unless otherwise specified. Given a subset $S \subset V(G)$, we will write $G[S]$ to denote the subgraph of G induced by S . That is, the edges of $G[S]$ are all those edges $(x, y) \in E(G)$ such that $x, y \in S$.

A *homomorphism* is a map between graphs that preserves edges. That is, a function f is a homomorphism between graphs G and H if for all $x, y \in V(G)$, $(x, y) \in E(G) \Rightarrow (f(x), f(y)) \in E(H)$. We write $G \rightarrow H$ to denote the existence of a homomorphism $f : G \rightarrow H$. For instance, letting K_n denote the complete graph on n vertices, $K_n \rightarrow G$ means that G contains a clique of size n .

We have seen that the definition of quasirandom graphs involves the *number* of homomorphisms from a given graph H into G_n ; this will also be a central notion in the theory of graph limits. Let $\text{Hom}(G, H)$ be the set of homomorphisms between G and H , and let $\text{hom}(G, H)$ be the number of such homomorphisms, i.e., $\text{hom}(G, H) = |\text{Hom}(G, H)|$. It will often be useful to normalize $\text{hom}(G, H)$ in order to define a *homomorphism density* between two graphs. In particular, we define:

$$t(G, H) = \frac{\text{hom}(G, H)}{v(H)^{v(G)}}$$

Note that there are $v(H)^{v(G)}$ total homomorphisms from G to H ; as such, $t(G, H)$ is the probability that a map $f : G \rightarrow H$ chosen uniformly at random is a homomorphism.

We are now ready to formalize the notion of a convergent sequence of finite graphs:

Definition 2.1. A sequence of graphs (G_n) with $v(G_n) \rightarrow \infty$ is *convergent* if for every finite graph F , the sequence of homomorphism densities $t(F, G_n)$ converges.

A basic example of a convergent sequence of graphs is the sequence of complete graphs (K_n) . For any n and any finite graph F , every map $f : F \rightarrow K_n$ is a homomorphism. Therefore $t(F, G_n) = 1$ for all n and F , and so the sequence $(t(F, G_n))$ is convergent. It is also clear that any quasirandom graph sequence is convergent: if a sequence of graphs $G(n)$ satisfies Definition 1.1, then $(t(F, G_n))$ converges to $p^{e(F)}$ for all finite graphs F .

Another equivalent way to define convergent graph sequences is through the notion of *sampling*. Given finite graphs F and G , let $p(F, G)$ be the probability that a randomly selected subset $S \subset V(G)$ of size $|S| = v(F)$ induces the subgraph F . We may then define a sequence of graphs (G_n) to be convergent if for every finite graph F , the sequence $(p(F, G_n))$ converges. This definition is perhaps more natural than Definition 2.1, and in some instances, more useful. However, the two definitions are in fact equivalent; we refer the reader to Chapter 11 of [1] for the justification of this equivalence.

3. KERNELS AND GRAPHONS

Once we have defined the notion of a convergent sequence of graphs, a natural question arises: what does a sequence of graphs converge to? For instance, given the notion of Cauchy convergence on the rational numbers \mathbb{Q} , the real numbers \mathbb{R} emerge as the numbers *to which* sequences of rational numbers converge. In other words, real numbers are the limit objects of sequences of rational numbers. Can we extend this same concept to sequences of graphs? That is, can we define some limit object of a sequence of graphs that somehow captures all the important information

contained within the sequence?

At first glance, it may seem as though the limit object of a sequence of finite graphs should be a countably infinite graph. Given a convergent sequence of graphs (G_n) , one candidate for a limit object might be a countably infinite graph G that captures the information about homomorphism densities contained in the sequence. In particular, for every finite graph F , we might like to have:

$$t(F, G) = \lim_{n \rightarrow \infty} t(F, G_n)$$

At this point, a problem emerges: when G is an infinite graph, $t(F, G)$ is no longer well defined. Perhaps we can avoid this issue by using the alternative characterization of $t(F, G)$ that we previously described – that is, as the probability that a randomly chosen map from F to G is a homomorphism. But here we encounter an even more fundamental problem. A probability distribution must be countably additive – so there is no uniform probability distribution on a countably infinite set! Therefore, the notion of a “randomly chosen map” from F to G can not be well defined.

Even if we could somehow circumvent this issue, countably infinite graphs would not be sufficiently complex to capture all the information we would like to about convergent sequences of graphs. We refer the reader to Chapter 11 of [1], and to Corollary 11.15 in particular, for further discussion of this topic.

It turns out that the ideal limit object of a convergent sequence of graphs is a *weighted* graph on the interval $[0, 1]$. This construction avoids the aforementioned problem, as there does exist a uniform probability distribution on $[0, 1]$. It also provides us with enough complexity to capture the important features of a convergent graph sequence, in a sense that we will soon make more precise. These limit objects are formally defined as functions, and are called *graphons*, a contraction of the words “graph” and “function.”

Definition 3.1. A *graphon* W is a symmetric, Lebesgue measurable function

$$W : [0, 1]^2 \rightarrow [0, 1].$$

If we consider $W(x, y) = W(y, x)$ to be the “weight” of the edge connecting x and y , we can indeed consider a graphon as a weighted graph on $[0, 1]$. It will also be useful to generalize graphons by defining *kernels*, whose images are bounded, but not necessarily within the interval $[0, 1]$:

Definition 3.2. A *kernel* W is a symmetric, Lebesgue measurable, bounded function

$$W : [0, 1]^2 \rightarrow \mathbb{R}.$$

We will write \mathbb{W} to denote the space of kernels, and \mathbb{W}_0 to denote the space of graphons. We will also write J to denote the kernel which is the constant function 1: that is, $J(x, y) = 1$ for all $x, y \in [0, 1]$. We can then write pJ to denote the kernel which is the constant function p .

We now wish to assign graphons as the limits of graph sequences in the way previously discussed. That is, given a sequence of graphs $G(n)$, its limit should be a graphon W such that

$$t(F, W) = \lim_{n \rightarrow \infty} t(F, G_n)$$

for every finite graph F . We will first need to extend the notion of homomorphism densities to graphons – and then verify that a graphon satisfying the above equation can indeed be found for every convergent graph sequence (G_n) .

Definition 3.3. Let W be a kernel and F a finite graph, and for ease of notation, let us write $V = V(F)$ and $E = E(F)$. We define the homomorphism density of F in W by:

$$t(F, W) = \int_{[0,1]^V} \prod_{(i,j) \in E} W(x_i, x_j) \prod_{i \in V} dx_i$$

This definition is an infinite analogue of homomorphism densities into finite weighted graphs. As a basic example, we have

$$t(K_2, W) = \int_{[0,1]^2} W(x, y) \, dx \, dy$$

which we recognize as the average value of W on the domain $[0, 1]^2$. For any finite graph G , $t(K_2, G)$ gives the edge density of G . If we consider G as a function $g : V(G)^2 \rightarrow \{0, 1\}$, where $g(x, y) = 1$ if and only if $(x, y) \in E(G)$, then the edge density of G corresponds to the average value of g on $V(G)^2$. We thus see that in this case, Definition 3.3 provides a natural infinite analogue of homomorphism densities into finite graphs.

The use of graphons as limit objects for sequences of graphs is justified by the following critical property:

Theorem 3.4. *For every convergent sequence of graphs (G_n) , there exists a graphon W such that $t(F, G_n)$ converges to $t(F, W)$ for every finite graph F .*

The proof of this theorem relies on the notion of *cut distance*, which is a metric on the space of graphs and kernels – that is, a way of defining the distance between any two graphs or kernels. We will write $d_{\square}(G, G')$ to denote the cut distance between two graphs, and $d_{\square}(W, W')$ for the cut distance between two graphons. The cut distance is rather difficult to define formally, and will not play a central role in this paper; we therefore omit its formal definition, and refer the reader to Chapter 8 of [1]. For our purposes, the most important property of the cut distance is the following:

Theorem 3.5. *The space of graphons is compact with respect to the cut distance. That is, (W, d_{\square}) is compact.*

In particular, because the cut distance is bounded between 0 and 1, every sequence of graphons contains a subsequence that converges to a limit graphon with respect to the cut distance. The other useful feature of the cut distance is that if the cut distance between two graphs (or graphons) is small, then so is the difference

between their homomorphism densities for any graph F . This result is known as the counting lemma:

Lemma 3.6 (Counting Lemma). *If F, G , and G' are simple graphs, then:*

$$|t(F, G) - t(F, G')| \leq e(F)d_{\square}(G, G').$$

Similarly, if F is a simple graph and W and W' are graphons, then:

$$|t(F, W) - t(F, W')| \leq e(F)d_{\square}(W, W').$$

Using these two properties of the cut distance, we can outline the proof of Theorem 3.4.

Proof of Theorem 3.4 (sketch). Let (G_n) be a convergent sequence of graphs. For each graph G_n in the sequence, we can define a corresponding graphon W_{G_n} by partitioning the interval $[0, 1]$ into $v(G_n)$ intervals of equal length. Let us call these intervals $J_1, \dots, J_{v(G_n)}$, and associate each interval with one of the vertices of $G(n)$. We then "connect" two intervals in the graphon if and only if their corresponding vertices in G_n are connected. That is, for all (x, y) with $x \in J_i$ and $y \in J_j$, we set $W(x, y) = 1$ if the vertices of G_n associated to J_i and J_j are connected, and $W(x, y) = 0$ otherwise. This definition ensures that for any simple graph F , $t(F, W_{G_n}) = t(F, G_n)$.

We can then consider the sequence of graphons (W_{G_n}) . By Theorem 3.5, this sequence has a subsequence $(W_{G_{n_k}})$ which converges to a limit graphon W . Then by applying Lemma 3.6, for any simple graph F we have:

$$|t(F, G_{n_k}) - t(F, W)| = |t(F, W_{G_{n_k}}) - t(F, W)| \leq e(F)d_{\square}(W_{G_{n_k}}, W) \rightarrow 0$$

Therefore, $t(F, G_{n_k}) \rightarrow t(F, W)$. By Definition 2.1, the sequence $t(F, G_n)$ is convergent. Therefore, because $t(F, G_n)$ contains a subsequence which converges to $t(F, W)$, the sequence itself must also converge to $t(F, W)$. As this result holds for every simple graph F , our proof of Theorem 3.4 is complete. \square

Let us examine several simple examples of graphons which act as the limit object as a sequence of graphs. The previously discussed sequence of complete graphs (K_n) converges to the constant graphon J . Indeed, we see from Definition 3.3 that for all simple graphs F , $t(F, J) = 1$. More generally, we see that for any simple graph F and any real p , $t(F, pJ) = p^{e(F)}$. Therefore, any graph sequence which is quasirandom with density p (according to Definition 1.1) converges to the graphon pJ . It is not immediately clear that a graph sequence satisfying the other given definitions of quasirandomness – for example, a graph sequence satisfying Definition 1.3 – should also converge to pJ . We will prove this result later in this paper – and in fact, this result will be the key to demonstrating the equivalence between Definitions 1.1 and 1.3.

One graph sequence which converges to a somewhat more interesting graphon is the sequence of complete bipartite graphs (K_n^n) . This sequence converges to a sort of "bipartite graphon" W , given by:

$$W(x, y) = \begin{cases} 1 & x \leq 1/2 \leq y \\ 1 & y \leq 1/2 \leq x \\ 0 & \text{otherwise} \end{cases}$$

To illustrate the limit graphons of more complex graph sequences, it is often helpful to use diagrams known as “pixel pictures.” We refer the reader to Section 11.4 of [1] for information on these diagrams, and for a discussion of the limits of more complex graph sequences.

4. KERNEL OPERATORS

In the first section of this paper, we stated that theorems can often be proven more easily in the space of graphons than in the space of graphs. One of the primary reasons for this is that because graphons, unlike graphs, are functions on a continuous domain, we can often analyse them using the tools of functional analysis. One such tool is the theory of operators between Lebesgue measurable functions:

Definition 4.1. Given a kernel $W \in \mathbb{W}$, we define the corresponding *kernel operator* $T_W : L_1[0, 1] \rightarrow L_\infty[0, 1]$ by:

$$(T_W f)(x) = \int_0^1 W(x, y) f(y) dy.$$

An *eigenfunction* of T_W is a function $f \in L_1[0, 1]$ such that for some corresponding *eigenvalue* $\lambda \in \mathbf{R}$, $T_W f = \lambda f$. That is, $(T_W f)(x) = \lambda f(x)$ for all $x \in [0, 1]$.

If we consider T_W as a map from $L_2[0, 1]$ to $L_2[0, 1]$, it has a bounded Hilbert-Schmidt norm, and is therefore a Hilbert-Schmidt operator. We refer the interested reader to [4] for further information on the rich theory of Hilbert-Schmidt operators. For our purposes, it will suffice to know that T_W has a *discrete spectrum*. That is, it has a countable set of nonzero eigenvalues $\{\lambda_1, \lambda_2, \dots\}$ such that $\lambda_n \rightarrow 0$.

Moreover, T_W has a *spectral decomposition*. That is, the kernel W is the limit of the following series of eigenvalues and eigenfunctions of T_W :

$$(4.2) \quad W(x, y) \sim \sum_k \lambda_k f_k(x) f_k(y)$$

where f_k is the eigenfunction corresponding to the eigenvalue λ_k such that $\|f_k\|_2 = 1$. The \sim symbol indicates that equality is conditional on the convergence of the series. It will also be important to note that if we define the product of kernels U and W by

$$(U \circ W)(x, y) = \int_0^1 U(x, z) W(z, y) dz$$

then we have that for any two kernels U and W ,

$$(4.3) \quad T_{U \circ W} = T_U T_W.$$

5. AN INFINITARY PROOF

As an illustration of the power of this spectral decomposition – and of the theory of graph limits more generally – we will now prove the previously discussed property of quasirandom graph sequences. That is, we will prove the equivalence between Definition 1.1 and Definition 1.3. Having introduced the notion of homomorphism densities, we can now state these definitions more elegantly. In particular, we can reformulate Definition 1.1 as:

Definition 1.1. A graph sequence (G_n) with $v(G_n) \rightarrow \infty$ is quasirandom with density p if for every finite graph F , $t(F, G_n) \rightarrow p^{e(F)}$.

We will also use the notation C_4 to denote the 4-cycle graph. The equivalence which we seek to prove can thus be stated as:

Theorem 5.1. *If (G_n) is a sequence of graphs such that $v(G_n) \rightarrow \infty$, $t(K_2, G_n) \rightarrow p$, and $t(C_4, G_n) \rightarrow p^4$, then (G_n) is quasirandom with density p .*

Our proof of this theorem will follow a pattern that is common among proofs in the theory of graph limits. We will first translate this theorem into the space of kernels and prove it in this space, assisted by the tools of functional analysis. We will then use the theory of graph limits to translate our proof back into the space of graph sequences.

In order to translate this theorem into the space of kernels, we begin a proof by contradiction. Suppose there exists a sequence of graphs (G_n) such that $v(G_n) \rightarrow \infty$, $t(K_2, G_n) \rightarrow p$, and $t(C_4, G_n) \rightarrow p^4$, but (G_n) is *not* quasirandom with density p . Then there exists some graph F such that $t(F, G_n) \not\rightarrow p^{e(F)}$. We can then select a subsequence (G_{n_k}) such that $t(F, G_{n_k}) \rightarrow c \neq p^{e(F)}$. Furthermore, by Theorem 3.5, we can select a convergent subsequence of (G_{n_k}) .

Let W be the limit graphon of this subsequence. Then by definition, $t(K_2, W) = p$, $t(C_4, W) = p^4$, and $t(F, W) = c$. Observe that W can not be the constant function pJ , because $t(F, pJ) = p^{e(F)} \neq c$. In order to prove Theorem 5.1, it will therefore suffice to prove the following lemma about kernels:

Lemma 5.2. *If W is a kernel such that $t(K_2, W) = p$ and $t(C_4, W) = p^4$ for some real number p , then $W = pJ$.*

To prove this lemma, we will need to introduce the concept of *labeled* graphs, and the related notions of labeled graph homomorphisms and homomorphism densities. A k -labeled graph is a graph G in which k of the vertices of $V(G)$ are labeled by $1, \dots, k$. In this proof, we will make use of labeled path graphs. Let us therefore write P_n^\bullet to denote the path on n vertices with one endnode labeled, and $P_n^{\bullet\bullet}$ to denote the path on n vertices with both endnodes labeled.

For a k -labeled graph F , a simple graph G , and any vertices $x_1, \dots, x_k \in V(G)$, let $\text{Hom}_{x_1 \dots x_k}(F, G)$ be the set of homomorphisms from F to G which send the k labeled nodes of F to the vertices x_1, \dots, x_k in G . As before, let $\text{hom}_{x_1 \dots x_k}(F, G) = |\text{Hom}_{x_1 \dots x_k}(F, G)|$, and let us define the homomorphism density by normalizing this quantity:

$$t_{x_1 \dots x_k}(F, G) = \frac{\text{hom}_{x_1 \dots x_k}(F, G)}{v(H)^{v(G)}}$$

In the same way that we previously extended the notion of homomorphism densities to kernels, we can do the same for homomorphism densities of labeled graphs. Let $W \in \mathbb{W}$ be a kernel, and F be a k -labeled graph. For ease of notation, let us write $V = V(F)$, $E = E(F)$, and let $V_0 = V \setminus [k]$ be the set of unlabeled vertices of F . For any $x_1, \dots, x_k \in [0, 1]$, we define

$$t_{x_1 \dots x_k}(F, W) = \int_{x \in [0, 1]^{V_0}} \prod_{(i, j) \in E} W(x_i, x_j) \prod_{i \in V_0} dx_i.$$

For example, let K_2^\bullet denote the graph with two connected vertices, one of which is labeled. Then for any graphon W and any $x \in [0, 1]$, $t_x(K_2^\bullet, W)$ gives the average weight of edges between x and all other points in the interval $[0, 1]$.² This is a sort of infinite analogue of the degree of a vertex x in a simple graph G . Indeed, for any $x \in V(G)$, $t_x(K_2^\bullet, G)$ is equal to the degree of x . We thus see that in this case, the above definition gives a natural extension of labeled homomorphism densities into the space of graphons.

It will often be useful to consider labeled homomorphism densities as functions. In the case of graphs, we can consider $t_{x_1 \dots x_k}(F, G)$ as a function from $V(G)^k$ to $[0, 1]$, while in the case of kernels, we can consider $t_{x_1 \dots x_k}(F, W)$ as a function from $[0, 1]^k$ to $[0, 1]$. This will allow us to use the following lemma, upon which the proof of Lemma 5.2 hinges:

Lemma 5.3. *The set of Lebesgue measurable functions from $[0, 1]^2$ to $[0, 1]$ forms an inner product space, with inner product given by*

$$\langle f, g \rangle = \int_{[0, 1]^2} f(x, y)g(x, y) \, dx dy.$$

We omit the proof of this lemma, which involves checking that the above definition satisfies symmetry, linearity, and positive definiteness.

Let F and G be any 2-labeled graphs, and suppose we take the inner product $\langle t_{xy}(F, W), t_{xy}(G, W) \rangle$ according to the above definition. This inner product will equal $t(FG, W)$, where FG denotes the graph resulting from “gluing” F and G together at the two labeled vertices. More formally, FG is defined by taking the disjoint union of F and G , and identifying vertices with the same label. We refer the reader to Section 4.2 of [1] for further information (and visualization!) of gluing products.

For instance, if we glue two copies of $P_3^{\bullet\bullet}$ together at the labeled vertices, the resulting graph is C_4 . And indeed, the following series of substitutions verifies that $\langle t_{xy}(P_3^{\bullet\bullet}, W), t_{xy}(P_3^{\bullet\bullet}, W) \rangle = t(C_4, W)$:

²Here again, we are considering W as a weighted graph on the interval $[0, 1]$.

$$\begin{aligned}
\langle t_{xy}(P_3^{\bullet\bullet}, W), t_{xy}(P_3^{\bullet\bullet}, W) \rangle &= \int_{[0,1]^2} t_{xy}(P_3^{\bullet\bullet}, W) t_{xy}(P_3^{\bullet\bullet}, W) dx dy \\
&= \int_{[0,1]^2} \int_{[0,1]} W(x, z) W(z, y) dz \int_{[0,1]} W(x, w) W(w, y) dw dx dy \\
&= \int_{[0,1]^4} W(x, z) W(z, y) W(y, w) W(w, x) dx dz dy dw \\
&= t(C_4, W)
\end{aligned}$$

In a similar way, gluing together two copies of P_2^\bullet yields the graph P_3 . Therefore, we will have:

$$(5.4) \quad \langle t_x(P_2^\bullet), W \rangle, t_x(P_2^\bullet), W \rangle = t(P_3, W)$$

which may be verified through a similar series of substitutions. We are now ready to prove Lemma 5.2, thereby proving the property of quasirandom graph sequences discussed in the beginning of this paper.

Proof of Lemma 5.2. We are given that $t(C_4, W) = p^4$. Using the above property, and using Lemma 5.3 to apply the Cauchy-Schwarz inequality, we have:

$$p^4 = t(C_4, W) = \langle t_{xy}(P_3^{\bullet\bullet}), t_{xy}(P_3^{\bullet\bullet}) \rangle \geq \langle t_{xy}(P_3^{\bullet\bullet}), J \rangle^2.$$

Considering the term on the right side of the inequality, we have:

$$\langle t_{xy}(P_3^{\bullet\bullet}), J \rangle = \int_{[0,1]^2} t_{xy}(P_3^{\bullet\bullet}) dx dy = \int_{[0,1]^3} W(x, z) W(z, y) dx dy dz = t(P_3, W).$$

So we conclude that

$$p^4 \geq t(P_3, W)^2$$

with equality holding if and only if $t_{xy}(P_3^{\bullet\bullet})$ is a constant function. But by Equation 5.4 and another application of the Cauchy-Schwarz inequality, we have:

$$t(P_3, W)^2 = \langle t_x((P_2^\bullet), W), t_x((P_2^\bullet), W) \rangle \geq \langle t_x((P_2^\bullet), W), J \rangle^4.$$

Again considering the term on the right side of the inequality, we have:

$$\langle t_x((P_2^\bullet), W), J \rangle = \int_0^1 t_x(P_2^\bullet) dx = \int_{[0,1]^2} W(x, y) dx dy = t(K_2, W).$$

So we may conclude that

$$p^4 \geq t(K_2, W)^4$$

with equality holding if and only if $t_{xy}(P_3^{\bullet\bullet})$ is a constant function. But we are given that $t(K_2, W) = p$ – so equality *does* hold, and $t_{xy}(P_3^{\bullet\bullet})$ must be a constant function!

By integrating $t_{xy}(P_3^{\bullet\bullet})$ over the domain $[0, 1]^2$, we see that it must be the constant function with value p^2 . But observe that:

$$t_{xy}(P_3^{\bullet\bullet}) = \int_0^1 W(x, z)W(z, y)dz = (W \circ W)(x, y).$$

So we have in fact shown that the kernel $W \circ W$ is the constant function p^2 . Our goal, however, was to show that W is the constant function p . To reach this conclusion, we will make use of the functional analysis tools introduced in the previous section.

Consider the operator $T_{W \circ W}$, which is given by:

$$(T_{W \circ W}f)(x) = \int_0^1 (W \circ W)(x, y)f(y) dy = p^2 \int_0^1 f(y) dy.$$

We see that $T_{W \circ W}$ maps any function f to a constant function. Therefore, only a constant function can be an eigenfunction of $T_{W \circ W}$. More specifically, $T_{W \circ W}$ has a single nonzero eigenvalue p^2 , with corresponding normalized eigenfunction $f = 1$.

But by Equation 4.3, $T_{W \circ W} = T_W^2$. Therefore, W has a single nonzero eigenvalue $\pm p$ with the same eigenfunction. Then by the spectral decomposition of T_W (cf. Equation 4.2), W is either the constant function p or the constant function $-p$. But we have:

$$t(K_2, W) = p \neq -p = t(K_2, -pJ)$$

which rules out the latter option. We conclude that W is the constant function p , i.e., $W = pJ$, thus concluding the proof. \square

ACKNOWLEDGMENTS

I would like to thank my mentor, Olga Medrano Martin del Campo, for being an invaluable source of guidance, knowledge, and support over the past two summers. Thank you as well to Professor Alexander Razborov for stimulating my interest in continuous combinatorics and graph limits, and providing guidance in my study of the central proof presented in this paper. Finally, I would like to thank Peter May, together with all the professors who volunteered their time to give lectures over the summer, for putting together an excellent research program. This is my third and final summer as a student at the REU, and each consecutive summer has exposed me to fascinating new fields of mathematics. The teaching and mentorship I've received through the REU has deepened my maturity as a mathematician, and has also inspired my own endeavors to teach mathematics to younger students.

REFERENCES

- [1] Lovász, László. *Large Networks and Graph Limits*. Providence: American Mathematical Society, 2012.
- [2] Hirschfeldt, Denis. *Algorithmic Randomness and Complexity*. New York: Springer, 2010.
- [3] Chung, Fan; Graham, Ronald; Wilson, Richard. “Quasi-random graphs”. *Combinatorica* 9, p345–362, 1989.
- [4] Conway, John. *A Course in Functional Analysis*. New York: Springer-Verlag, 1990.