My research is in probability theory, mathematical physics, and functional analysis. More specifically: I work on problems in stochastic analysis, diffusion processes on Lie groups, random matrices, and free probability; quantum field theory, Yang-Mills theory, and quantum information theory; logarithmic Sobolev inequalities, heat kernel analysis, and holomorphic and subharmonic function spaces.

I have 19 published or accepted papers, and 3 preprints [1]—[22]; all 22 papers are available on the arXiv or my webpage. 15 of these papers are from the last 4 years, and 12 are from the last 2 years. The published papers have appeared (or been accepted) in the Annals of Probability, Probability Theory and Related Fields, Communications in Mathematical Physics, the Journal of Functional Analysis, Letters in Mathematical Physics, the Transactions of the AMS, IMRN, AIHP, the Canadian Journal of Mathematics, and others.

I currently hold an NSF CAREER Award, DMS-1254807, for the period 2013-2018. Previously, I was awarded two NSF grants: DMS-0701162 for 2007–2010, and DMS-1001894 for 2010–2013. I also held the Hellman Foundation Fellowship for 2011-2012.

I have advised 5 PhD students, two of whom graduated previously. Natasha Bltivić is now a Lecturer (tenure track Assistant Professor) at Lancaster University in the UK. David Zimmermann is a Dickson Instructor at the University of Chicago. Of my three current PhD students, the most senior is Ching Wei Ho, who will graduate in 2018, and already has one accepted paper the Journal of Functional Analysis.

Following are descriptions of the general directions of my research programs in probability theory, mathematical physics, and functional analysis, with brief summaries of some of my papers.

1 Probability Theory

My work in probability theory is centered on random matrices, where I apply tools from stochastic analysis and functional analysis to understand the large-N limits of various statistics associated to diffusion processes in $\mathbb{M}_N$ (the state space of $N \times N$ complex matrices).

The most classical ensembles of random matrices are the GOE_N and GUE_N (Gaussian Orthogonal / Unitary Ensembles), of symmetric real / complex Hermitian matrices with i.i.d. Gaussian entries above the main diagonal. A complete study of many spectral statistics (the empirical eigenvalue distribution, the largest eigenvalue, the spacing between eigenvalues and their k-point correlation functions, etc.) was carried out from the 1950s through 1990s. The empirical spectral distribution converges almost surely to Wigner's semicircle law $\frac{1}{2\pi}\sqrt{(4-x^2)_+}\,dx$, with Gaussian fluctuations of order $\frac{1}{N}$. The largest eigenvalue converges almost surely to the edge of the law, 2, with $O(N^{-2/3})$ fluctuations described by the Tracy–Widom law, discovered first in this context, but now known to be of great importance in many different structures in probability theory and statistical physics. In the parallel non-Hermitian setting where *all* the entires of the matrix Z_N are i.i.d. Gaussians with variance $\frac{1}{N}$ (called the Ginibre ensemble), the empirical spectral distribution converges a.s. to the uniform probability measure on the unit disk, with Gaussian fluctuations of order $\frac{1}{N}$.

A large fraction of the work done in random matrix theory in the last two decades is the "universality program", which aims to prove that the above features (and others) don't fundamentally depend on the Gaussian distributions. There are two kinds of generalizations: *Wigner ensembles*, where the entries are still independent but not Gaussian, and *invariant ensembles* where the joint law of entries has a density whose logarithm is the negative of a convex potential function, but not necessarily quadratic. In wide generality, it is now understood precisely when the above strong laws of large numbers and central limit theorems apply.

In my view, there is a third natural direction to generalize the Gaussian ensembles. In a GUE_N, one may take the entries to be Brownian motions, to bring to bear the tools of stochastic analysis. (The eigenvalues

then form an interacting particle system known as the Dyson Brownian motion.) This is a diffusion process in the space of Hermitian matrices; multiplying by i, it becomes the standard Brownian motion in the *Lie algebra* $\mathfrak{u}_N$ of the unitary group $\mathbb{U}_N$. Similarly, the Ginibre ensemble Z_N may be thought of as an endpoint of Brownian motion on the Lie algebra $\mathbb{M}_N = \mathfrak{gl}_N$ of the Lie group $\mathbb{GL}_N$ of invertible matrices in $\mathbb{M}_N$. It is therefore natural to investigate the large-N behavior of the usual random matrix statistics for the Brownian motions on the *Lie groups* $\mathbb{U}_N$ and $\mathbb{GL}_N$: the classical Gaussian cases then being the infinitesimal version. My papers [10, 13, 14, 15, 19] answer several questions about these Brownian motions, and my papers [8, 22] address other random matrix models generalizing the classical Gaussian ensembles. The following paragraphs highlight some of the results in these papers.

1.1 [10, 19] Unitary Brownian Motion: Hard Edge, The Jacobi Process, and Liberation

The standard Brownian motion $U_t = U_t^N$ on $\mathbb{U}_N$ is the diffusion process determined by the left-invariant Riemannian metric given by the Hilbert–Schmidt inner product $\langle X, Y\rangle = N\mathrm{Re}\,\mathrm{Tr}(XY^*)$ on $\mathbb{M}_N$. The empirical eigenvalue distribution has a large-N limit (proven by Philippe Biane and Eric Rains independently in the late 1990s) ν_t, which is a probability measure on the unit circle having many qualitative properties in common with Wigner's semicircle law: it is supported in a strict sub-arc for $t < 4$. The main result of our paper [19] is that the spectrum has a "hard edge": the largest angle eigenvalue converges almost surely to the edge of the support of ν_t. We also prove a multi-time version of this theorem called "strong convergence", which must be measured in terms of operator norm of polynomials in the endpoints of the process, due to noncommutativity. The proof relies on a precise estimate on the growth of moments of the empirical eigenvalue distribution, which we prove with a coupling argument between U_t^{2N} and a block diagonal process composed of two independent copies of U_t^N, and many involved sharp estimates using the representation theory of $\mathbb{U}_N$ and the Schur–Weyl duality (to transport computations to the more manageable symmetric group).

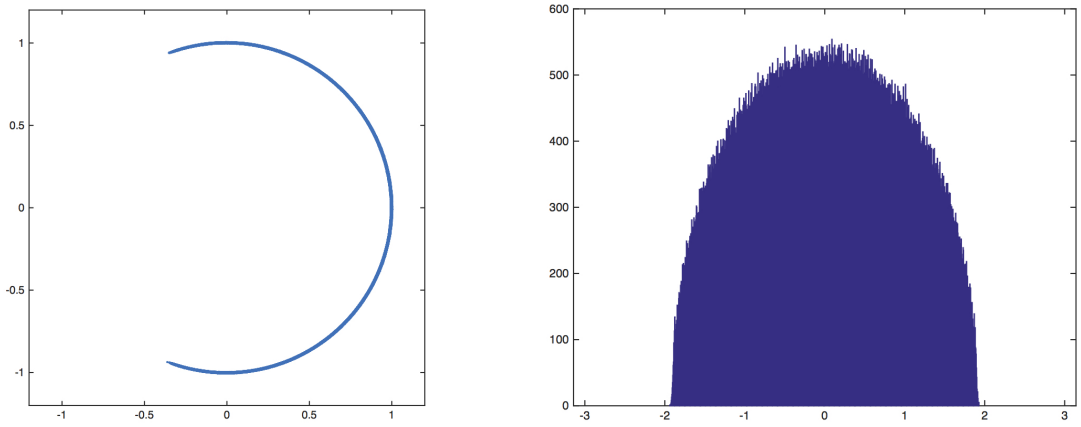

Figure 1: The spectrum of the unitary Brownian motion U_t^N with $N = 400$ and $t = 1$ (1000 trials). On the left is a plot of the eigenvalues, while on the right is a 1000-bin histogram of their complex arguments. The argument range of the data is $[-1.9392, 1.9291]$, as compared to the predicted large-N limit range (to four digits) $[-1.9132, 1.9132]$.

One nice application of these ideas is to the Jacobi ensemble: a third rung in the classical random matrix ensembles. Making it into a matrix-valued process as above, it can be realized in the following form: $J_t^N = J_t = QU_tPU_t^*Q$ for fixed projection matrixes $P, Q \in \mathbb{M}_N$. This encodes geometry as follows: the

eigenvalues of J_t are (trigonometric polynomials in) the principle angles between the range spaces of Q and $P_t = U_t P U_t^*$; i.e. J_t measures the "matrix-valued angle" between two subspaces when one is rotated by a unitary Brownian motion. In [10], we studied the large-N limit of the empirical eigenvalue distribution of J_t^N, when the projections P, Q also have an appropriate large-N limit (and, for technical reasons, have range spaces of dimension $N/2$). We proved that J_t has an almost sure limit empirical spectral distribution μ_t. Ut is supported in $[0, 1]$, has a real analytic density on the interior of its support; and the measure evolves in a smooth way (in particular the boundary of the support evolves smoothly). The techniques used involve stochastic analysis, operator theory, nonlinear complex PDE, and complex analysis to develop the appropriate subordination theory for the Cauchy transform of μ_t. The strong convergence results from our paper [19] then show that the empirical eigenvalue distribution of the Jacobi process J_t^N has the same qualitative properties for sufficiently large N.

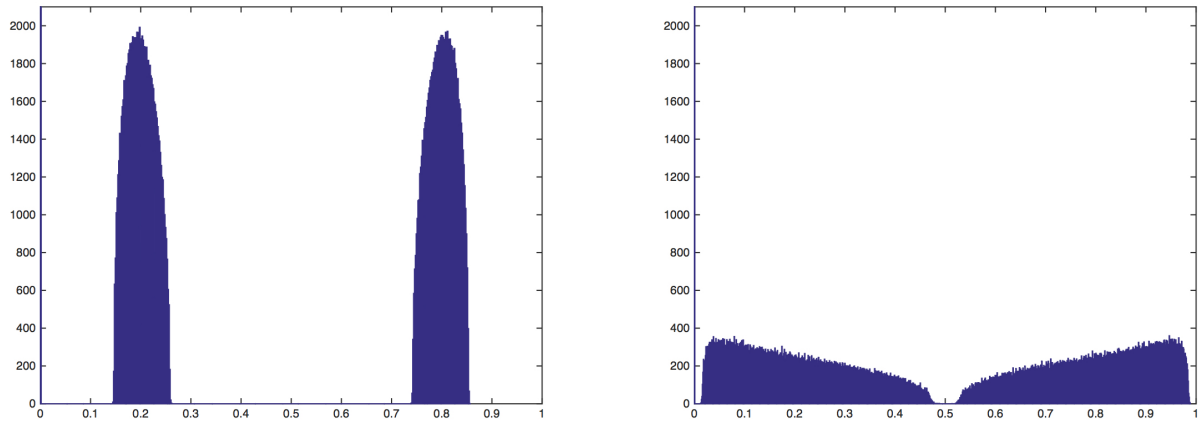

Figure 2: An example of the eigenvalues of the Jacobi process J_t^N with $N = 400$, initial data concentrated at 0.2 and 0.8, and $t = 0.01$ (on the left) and $t = 0.25$ (on the right). (1000 trials.)

1.2 [13, 14, 15] Brownian Motions on $\mathbb{GL}_N$

The standard Brownian motion on $\mathbb{GL}_N$ is determined by the same Hilbert–Schmidt metric as above. In [14], I studied a two-parameter family of diffusions $(B_{r,s}^N(t))_{t\geq 0}$ (where $r, s > 0$) on $\mathbb{GL}_N$ with distribution invariant under conjugation by unitaries (a condition closely tied to the existence of a large-N limit). The standard Brownian motion is the special case $r = s = 1/2$. These processes are difficult to analyze from a random matrix perspective, as they are a.s. not normal matrices for all positive time (as I prove in the paper). I proved that $B_{r,s}^N(t)$ has a large-N limit (in terms of finite-dimensional distributions) which can be realized as a stochastic process taking values in the operators on an infinite-dimensional Hilbert space. This, in particular, solved an outstanding conjecture of Philippe Biane's from 1995. The techniques used include stochastic calculus and concentration of measure bounds that are of independent interest and were also used in my other work (including [9] and [13]).

Taking these ideas quantitatively further, in [13] I studied the rates and regularity of weak a.s. convergence of the empirical spectral measure of the unitary Brownian motion U_t^N, and the corresponding empirical measure of *singular values* of the $\mathbb{GL}_N$ Brownian motions I introduced in [14]. To be precise: if ν_t^N is the empirical measure of eigenvalues and its limit is ν_t, the question is: for which test functions f does it follow

that

$$\int_{\mathbb{R}} f\, d\nu_t^N \to \int_{\mathbb{R}} f\, d\nu_t, \quad a.s. \tag{1}$$

and at what rate? Due to the non-normality in the $\mathbb{GL}_N$ case, the largest class of test functions which can be used for matrix computations are analytic functions on a neighborhood of the spectrum, with appropriate decay in the coefficients to guarantee convergence. For this class, I proved that the variance is always $O(1/N^2)$, giving a convergence rate of $O(1/N)$. In the unitary case, where more spectral tools are available, I used Fourier methods to improve the known results to the best yet-known: if f is in the Sobolev space H_p on the circle for $p > 1$, then the rate of convergence is at least $O(1/N^{p-1/2})$ for $p < 3/2$, and $O(1/N)$ for $p \geq 3/2$. (The $p = 3/2$ case coincides with Lipschitz test functions, reproducing the best previously know results.)

In [15], we turned to the study of fluctuations of these diffusions. We give a unified analysis for the $\mathbb{U}_N$ and $\mathbb{GL}_N$ cases, with a class of test functions that generalizes polynomials to "trace polynomials": for example $f(U) = \mathrm{Tr}(U^2)\mathrm{Tr}(U^3) - 2\mathrm{Tr}(U)^2$. Restricting to the context of (1), this is a generalization of the linear statistics of the ensembles. We prove that, for any times $t_1, \ldots, t_k > 0$, the rescaled fluctuations

$$N[f(B_{t_1}^N, \ldots, B_{t_k}^N) - \mathbb{E}f(B_{t_1}^N, \ldots, B_{t_k}^N)]$$

converges in distribution to a multivariate centered Gaussian as $N \to \infty$, and we give a precise quantitative description of the covariance. (Here the B_t^N can be taken among the two-parameter $\mathbb{GL}_N$ Brownian motions from [14], or the unitary Brownian motion.) This generalizes the 2010 work of Thierry Lévy and Mylene Maïda, where they computed the fluctuations for the unitary case, only for a single time t. In this case, it is of note that the (complicated) covariance we compute can be shown to converge, as $t \to \infty$, to the square Sobolev $H_{1/2}$ norm of f, which aligns with the 2001 work of Diaconis and Evans on the fluctuations of eigenvalues of Haar distributed unitary matrices.

1.3 [8, 22] Further Generalizations of Wigner Matrices

I wish to highlight two more papers that deal with asymptotic random matrix theory, in this case back in the Hermitian category generalizing Wigner matrices. In both cases, the techniques and results involve novel combinations of random matrices with tools from functional analysis: namely stochastic analysis and logarithmic Sobolev inequalities.

In [8], we generalized the "fourth moment phenomenon" to the GUE_N context (in the large-N limit). The *fourth moment theorem*, proved originally by Nualart and Peccati, deals with the distributions of random variables in a fixed order of Wiener chaos. Namely, if X_k is a sequence of multiple Wiener integerals of a fixed order, with standardized mean and variance, then (X_k) converges in distribution to the standard normal distribution if and only if $\mathbb{E}(X_k^4) \to 3$; i.e. iff the fourth moment converges to the fourth moment of the normal distribution. This remarkable theorem has spawned a whole industry of fourth moment phenomena in different contexts, with connections to Stein's method, and significant applications to Malliavin calculus. Our work in [8] was devoted to proving the fourth moment theorem in the large in limit of the GUE_N, in the context of "Wigner chaos". Here, the theorem takes the same form: a sequence X_k if "Wigner integrals" (defined in analogy to Wiener integrals but in the operator theoretic large-N limit of free probability) of a fixed order converges to a *semicircular* distribution if and only if the fourth moment converges to 2, the fourth moment of the semicircle law. We used a combination of methods both combinatorial (the diagrammatic approach to the method of moments) and analytic (Malliavin calculus and Fourier analysis to estimate concentration of Wasserstein-like metrics) to prove the theorem, and gave several applications to Malliavin calculus and a transfer principle between the Wigner and Wiener chaos expansions. Notably, we proved that no random variable in higher orders of Wigner chaos can be semicircular; this is the first (and to date only) concrete distributional result about the highly noncommutative Wigner chaos.

Finally, in the very recent preprint [22], we consider real symmetric random matrices whose entries are *not* necessarily i.i.d. There is a vast literature on *band matrices*, where the identical distribution requirement is softened (for example just to uniform integrability), but in almost all cases, the studied models still have independent upper triangular entries. Some more recent work (using tools of operator-valued free probability) has considered block matrices, where the blocks are all independent. The existing techniques can only handle the case where the block sizes in the ensemble X_N do not grow with N (or, dually, where there are a fixed number of $O(N)$-sized blocks). As such examples demonstrate, the limit empirical spectral distribution is typically not semicircular; it depends in a complicated way on the correlation structure of X_N.

In [22], we used concentration techniques to prove that the empirical spectral distributions of such ensembles still converge to their means (in probability) for block sizes up to $o(\log N)$; and, more generally, the blocks need not be square and can be distributed at random. To be precise: we proved that if the centered entries of X_N are uniformly integrable, and can be partitioned into groups each with size no more than some bound $d_N = o(\log N)$ with all groups independent of each other, then the empirical spectral distribution of $\frac{1}{\sqrt{N}}X_N$ converges in probability to its mean. The principle tool used to prove these results is a new concentration method we call *mollified log Soboelv inequalities*. We proved earlier (published in David Zimmermann's thesis paper in J. Funct. Anal.) the following: if X is any bounded random variable, and G is a standard normal random variable independent from X, then the law of $X + tG$ satisfies a log Sobolev inequality for all $t > 0$. (See (5) below for the definition of a log Sobolev inequality.) This can be used to give a totally new proof of Wigner's semicircle law. For the present application, we needed to prove a multivariate version of this theorem (random vectors X and G) with bounds on the log Sobolev constant as a function of t and the dimension. Then the ensemble X_N can be mollified by adding t times an independent GOE_N, and letting $t \downarrow 0$ as $N \uparrow \infty$ appropriately to take advantage of the concentration of measure the log Sobolev inequality provides. This work has been submitted to a prestigious probability journal.

2 Mathematical Physics

I have an undergraduate degree in physics, and my PhD thesis was in mathematical physics (related to Fermion quantum fields and coherent states). While my research program evolved more towards functional analysis and probability theory afterward, I have recently returned to working on many problems in and related to mathematical physics. My papers [1, 9, 16, 17, 18, 20, 21, **?**] are all mathematical physics papers; three of them are published or accepted in CMP, and three others are currently in submission to mathematical physics journals. Below, I briefly describe the contents of six of these papers, evenly divided into three themes.

2.1 [16, 17] 2-Dimensional Yang–Mills Theory

The Yang–Mills field theory describes a random connection on a principal bundle with a compact structure group K. In two dimensions, the theory is tractable and is extensively studied. In particular, on the plane, it is possible to use gauge fixing to make the measure Gaussian, opening the door to rigorous calculations via stochastic calculus. The typical observables studied in the theory are *Wilson loop* random variables: if L is a loop in the base manifold, the associated Wilson loop is $W_L = \mathrm{Tr}(\mathrm{hol}(L))$, where $\mathrm{hol}(L)$ is the holonomy of the random connection around L, and the trace is taken in a fixed representation of K.

One tool presented in the physics literature for computing Wilson loop expectations is the *Makeenko–Migdal equation*. When the underlying manifold is the plane and the structure group is $K = \mathbb{U}_N$, it can be stated in the following form. Let L be a loop with a simple crossing at a point ν. There are four regions of the complement of L in the plane adjacent to ν; call them F_1 through F_4 (in counter-clockwise order). Denote the area of the face F_j as $t_j > 0$ (for the sake of simplicity we presently assume $t_1, \ldots, t_4$ are finite). We may

also decompose L into two concatenated loops: L_1 starting at ν and going around until the first return to ν, and then continuing with L_2 until reaching ν again. See Figure 3 below.

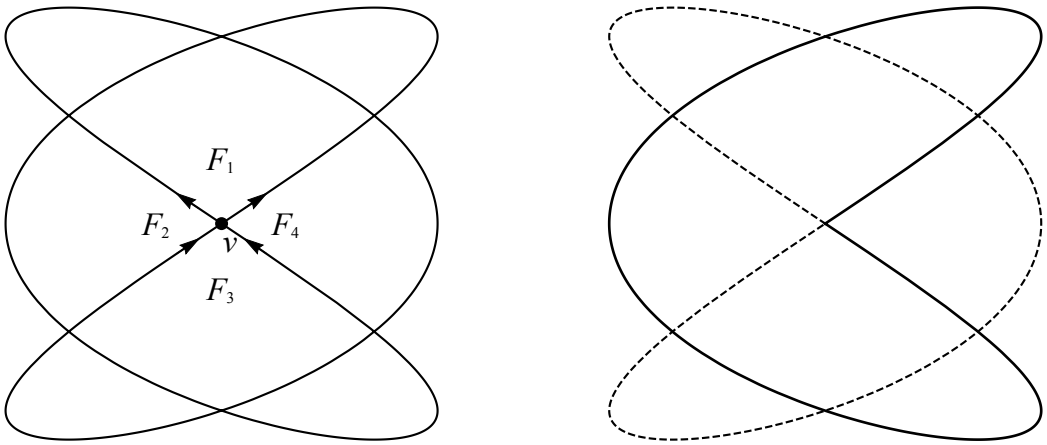

Figure 3: A typical loop L for the Makeenko–Migdal equation on the left, with its decomposition into L_1 (solid line) and L_2 (dashed line) on the right.

The Makeenko–Migdal equation then asserts that, as the loop L is perturbed (heuristically) through the collection of smooth loops with simple crossing at ν, the associated Wilson loops satisfy

$$\left(\frac{\partial}{\partial t_1} - \frac{\partial}{\partial t_2} + \frac{\partial}{\partial t_3} - \frac{\partial}{\partial t_4}\right) \mathbb{E}(W_L) = \mathbb{E}(W_{L_1} W_{L_2}). \tag{2}$$

If the right-hand side of (2) were instead the product of expectations $\mathbb{E}(W_{L_1})\mathbb{E}(W_{L_2})$, one could use the equation to iteratively reduce the computation of any Wilson loop expectation into solving a linear system of ODEs. This is what happens in the large-N limit (where, on the plane, Yang–Mills theory is supposed to converge to an object called the *Master Field*, introduced by Singer), where concentration of measure results similar to those I developed in [13, 14] show that $\mathbb{E}(W_{L_1} W_{L_2}) - \mathbb{E}(W_{L_1})\mathbb{E}(W_{L_2}) \to 0$ as $N \to \infty$.

The Makeenko–Migdal equation (2) was introduced (in a different form) by Makeenko and Migdal in 1979, and proved using heuristic methods that are challenging to make rigorous. Recently, Thierry Lévy and Antoine Dahlqvist gave independent rigorous proofs of the equation. Their methods rely on the technology of Markovian holonomy fields developed by Lévy, and use complicated arguments converting the problem to an algebraic form in the Brauer group. They also rely on the existence of an unbounded face for the complement of L, and so their arguments do not generalize beyond the plane.

In [16], we gave three new proofs of a generalized form the Makeenko–Migdal equation for any compact structure group K, and for a wider class of observables than Wilson loops. All three proofs are of a similar nature: they view the equation as the infinitesimal form of an invariance principle for the Yang–Mills connection. In particular, they only require that the random variables involved satisfy a symmetry we call *extended gauge invariance* with respect to any simple crossing in the loop. The proofs themselves are then an exercise in heat kernel analysis on the group K, and are quite elementary; in fact, the shortest of the three is slightly more than one page long (as compared to dozens of pages in the first rigorous proof by Lévy). What's more, two of the proofs given are "local" in nature: they do not rely on any structure (like an unbounded face) that is not in a neighborhood of the chosen crossing ν. As a result, without a lot of additional effort, these proofs generalize to prove the exact same Makeenko–Migdal equation (2) holds for the Yang–Mills theory over *any* compact surface, even with boundary (under a variety of boundary conditions in the Yang–Mills equation). We proved this in [17]. Both papers have recently been accepted for publication in CMP.

2.2 [9, 20] The Segal–Bargmann (Coherent State) Transform

The classical Segal–Bargmann transform is a unitary isomorphism that interchanges the holomorphic and Schrödinger representations. Introduced independently by its namesakes in that context, a finite-dimensional version of it can be succinctly described as follows. Let γ_t denote the (Gaussian) heat kernel on $\mathbb{R}^k$. This function has an analytic continuation to all $z \in \mathbb{C}^k$, which I denote $\gamma_{\mathbb{C}}(t, z)$. The Segal–Bargmann transform is defined on $L^2(\mathbb{R}^k, \gamma_t)$ by the integral formula

$$(B_t f)(z) = \int_{\mathbb{R}^k} \gamma_{\mathbb{C}}(t, z - x) f(x)\, dx, \qquad z \in \mathbb{C}^k. \tag{3}$$

The theorem is that B_t is a unitary map from $L^2(\mathbb{R}^k, \gamma_t)$ onto the space $\mathcal{H}L^2(\mathbb{C}^k, \mu_t)$ of holomorphic functions in L^2 of the (variance rescaled) heat kernel measure μ_t on $\mathbb{C}^k$. It is sometimes called the coherent state transform: coherent states are minimum uncertainty states (i.e. those giving equality in Schrödinger's strong form of the Heisenberg uncertainty principle). In finite dimensions, they are Gaussian wave-packets of a certain form, and the complexified heat kernel density $\gamma_{\mathbb{C}}(t, z - x)$ is a coherent state for each fixed $z \in \mathbb{C}^k$.

In 1994, Brian Hall introduced a generalization of the Segal–Bargmann transform, replacing $\mathbb{R}^k$ with a compact Lie group K as the state space. Later, Bruce Driver generalized this to Lie groups of compact type (those whose Lie algebra possesses an Ad-invariant inner product), and then together they generalized the transform further, realizing that the time parameter of the heat kernel in the domain need not be the same as the one used in the transform. The transform has the same form as (3): the heat kernel ρ_t on K has an analytic continuation $\rho_{\mathbb{C}}(t, z)$ to the complexification $K_{\mathbb{C}}$ of K. (For example, if $K = \mathbb{U}_N$, then $K_{\mathbb{C}} = \mathbb{GL}_N$. In the compact type case, the Lie algebra of $K_{\mathbb{C}}$ is equal to $\mathrm{Lie}(K) \otimes_{\mathbb{R}} \mathbb{C}$.) Then the transform itself is defined by

$$(B_{s,t} f)(z) = \int_K \rho_{\mathbb{C}}(t, zx^{-1}) f(x)\, dx, \qquad z \in K_{\mathbb{C}} \tag{4}$$

and converges whenever $0 < t < 2s$. It is a unitary map from $L^2(K, \rho_s)$ onto $\mathcal{H}L^2(K_{\mathbb{C}}, \mu_{s,t})$ for an appropriate heat kernel measure $\mu_{s,t}$ on $K_{\mathbb{C}}$. Driver and Hall then used this transform in the context of Yang–Mills theory.

In [9], together with Driver and Hall, I explored the large-N limit of the transform $B_{s,t}$ of (4), where the group K is taken to be $\mathbb{U}_N$. Considering the matrix-valued version $\mathbf{B}^N_{s,t}$ of the transform (acting entry-wise on $\mathbb{M}_N$-valued functions), we identify a family of L^2-functions we call *trace polynomials* that are invariant under the action of $\mathbf{B}^N_{s,t}$. These trace polynomials are dense in the space of all Ad-equivariant functions from $\mathbb{U}_N$ to $\mathbb{M}_N$. If f is a trigonometric polynomial on the circle, then viewing f as a matrix-valued function on $\mathbb{U}_N$, $\mathbf{B}^n_{s,t} f$ is not typically a trigonometric polynomial. However, due to concentration of measure results that we prove (related to those I used in [13, 14]), as $N \to \infty$ this function $\mathbf{B}^N_{s,t} f$ *does* converge, in an appropriate sense, to a trigonometric polynomial. This allows us to identify the large-N limit as a transform $\mathscr{G}_{s,t}$ on trigonometric polynomials. It then extends to a unitary map from $L^2(\nu_s)$ (where ν_s is the large-N limit of the empirical spectral distribution of the unitary Brownian motion U^N_s, as described in Section 1.1) onto a reproducing kernel Hilbert space of analytic functions $\mathscr{A}_{s,t}$ on a certain precompact domain in $\mathbb{C}$. We use PDE methods to identify this transform as a generalization of the one $\mathscr{G}_{t,t}$ that Philippe Biane introduced from a free probability viewpoint in 1997. This paper appeared in JFA.

Very recently, in [20] we introduced a further generalization of the Segal–Bargmann transform that is, in two precise senses, a completion of the theory. We prove that the heat kernel measure ρ_t on K has an analytic continuation in *both space and time* to a function $\rho_{\mathbb{C}}(\tau, z)$ of $(\tau, z) \in \mathbb{C}_+ \times K_{\mathbb{C}}$ (where $\mathbb{C}_+ = \{\tau = t + iu \colon t > 0, u \in \mathbb{R}\}$ is the right half-plane). Using this function in (4) yields a convergent integral, provided that $\tau \in \mathbb{D}(s, s)$ (the disk in $\mathbb{C}_+$ of radius s centered at s), and we prove that the resultant transform $B_{s,\tau}$ is a unitary map from $L^2(K, \rho_s)$ onto $\mathcal{H}L^2(K_{\mathbb{C}}, \mu_{s,\tau})$ for a certain heat kernel measure $\mu_{s,\tau}$. The two senses in which this setup completes Segal–Bargmann theory are as follows. From a mathematical perspective: the heat kernel $\mu_{s,\tau}$ is

induced by a Riemannian metric on the Lie algebra of $K_{\mathbb{C}}$, which is invariant under the adjoint action of K (a fundamental requirement for the unitarity of the transform). We prove that the corresponding (s, τ) family of metrics are the *only* $\mathrm{Ad}(K)$-invariant metrics on $\mathrm{Lie}(K_{\mathbb{C}})$ (when K is simple); so there cannot be a further generalization. From a physics perspective: in the Euclidean setting, as above, the functions $x \mapsto \gamma_{\mathbb{C}}(\tau, z - x)$ are coherent states; allowing τ to be complex, we now span the set of *all* coherent states, and so this complex time Segal–Bargmann transform is the most natural form from the perspective of geometric quantization.

2.3 [18, 21] Quantum Information Theory

I have recently become interested in quantum information theory, partly because my expertise in random matrix theory (and in particular diffusion processes on unitary groups) is fruitful in approaching some problems of current interest in the field. I have thus far written two papers in this area, and have several other active projects on related topics. Following are brief descriptions of my two quantum information theory papers.

One of the most fundamental questions in quantum information theory concerns the ability to send information over a noisy quantum channel. In 2009, Matt Hastings showed (using tools from random matrix theory) that quantum channels have a remarkable counterintuitive property: two apparently non-interacting quantum channels when combined can have *strictly greater* classical capacity than the sum of their individual classical channel capacities. Equivalently: the minimum output entropy of a quantum channel is *not* additive over tensor product. This violated a decade of conjectures, after Peter Shor and others had shown that half a dozen intuitive entropy / information relationships were equivalent to the additivity of minimum output entropy, among them the additivity of entanglement of formation. It is still a very open question of active research to understand the source of the non-additivity, and to quantify the maximal defect. (This is understood for the specific random channels that Hastings used in his proof, but not in general.) Our paper [18] adds fuel to this fire by proving that the minimum output entropy *is locally additive* over tensor product of channels. In other words: a tensor product of local minimizers of output entropy is a local minimum of the tensor product output entropy. Moreover, we prove this not only for the von Neumann entropy functional but for the more general Renyi entropy (which includes the von Neumann entropy as a degenerate case). This means that the non-additivity proved by Hastings is a *fundamentally non-local effect*, which deepens the mystery of its source. The techniques used to prove this derive from the analytic tools of asymptotic random matrix theory: the linearization method, free difference quotients, the Kato–Rellich theorem, and a version of the functional free Itô formula. The paper has now appeared in Letters in Mathematical Physics.

From a quantum computational perspective, entanglement is a useful tool. This is only true up to a point: a little entanglement can provide counterintuitive boosts in computational efficiency (such as in Shor's algorithm for factoring large integers in polynomial time), but too much entanglement introduces noise and decoherence phenomena impeding quantum computation. It is then natural to ask how entangled the *average* state is; or, more usefully, what proportion of random states have non-zero but low entanglement? This question was addressed by Gross, Flammia, and Eisert in 2009. They measured entanglement of a state T with the *geometric measure of entanglement* $E_g(T) = -2\log_2 \|T\|_\infty$, where $\|T\|_\infty = \max\{|\langle T, \otimes_{j=1}^m \mathbf{v}_j\rangle| : \|\mathbf{v}_1\| = \cdots = \|\mathbf{v}_m\| = 1\}$ is the *spectral norm* of the tensor T (generalizing the spectral radius of a matrix, thought of as a $m = 2$-mode tensor). The quantity $E_g(T)$ is 0 iff T is a pure state; the maximum value it can take is m. Gross et. al. showed that, in the (projective) sphere of all m-qubit states, for all but a Haar measure e^{-m^2} set of states, $E_g(T) \geq m - 2\log_2 m - 2$. That is: with high probability, the geometric measure of entanglement of a random state is close to maximal. This, however, ignores the fact that any multi qubit state accessible to a quantum computer (usually modeled as a photonic system) must be a Boson (symmetric) quantum state. In [21], we address the same question for *Bosons*. In this case, the maximum possible value for the geometric measure of entanglement is much lower, $\log_2 m$, opening the door to much wider availability of (computationally) universal states. We prove that, as in the general setting, with high probability, generic

Boson states are close to the theoretical maximum: $E_g(T) \geq \log_2 m - \log_2 \log_2 m - 3$. In this case, the Haar measure of the complement of this set decays more slowly (super-polynomially but not super-exponentially) than in the general non-symmetric setting. This is owing to the much smaller dimension of the space of Bosons (which is linear in m) compared to the general case (where the dimension is exponential in m).

3 Functional Analysis

Approximately one third of my research papers address questions that, while often connected to problems in probability theory or mathematical physics, are pure functional analysis. In all of these cases, the connecting theme is *sharp constants for functional inequalities in "regular" spaces*. Here regular means holomorphic, or subharmonic, or an appropriate noncommutative analog in the operator theory world. Here I will briefly summarize the results in some of those papers.

3.1 [4, 12] Log Sobolev Inequalities for Log Subharmonic Functions

The logarithmic Sobolev inequality (LSI) is a pervasive tool in probability theory, analysis, and geometry. A measure μ (on Euclidean space) satisfies the classical LSI if, for some constant $c > 0$,

$$\mathrm{Ent}_\mu(f^2) \leq c \int |\nabla f|^2 \, d\mu \tag{5}$$

for all $f \geq 0$ sufficiently integrable. Here $\mathrm{Ent}_\mu(\rho) = \int \rho \log \rho \, d\mu$ is the relative entropy of the probability density ρ; in general we apply it to any $L^1(\mu)$ function g by substituting $\rho = g/\|g\|_1$. The Gaussian measure of variance t satisfies the LSI with constant $c = 1/t$, and there is a huge literature discussing which measures satisfy an LSI in general. The existence of an LSI forces the measure μ to have sub-Gaussian tails, with bounds that depend on the LSI constant c in a *dimension independent* manner; this is the reason it is such a useful tool, especially in asymptotic problems.

There is an equivalent formulation of the LSI (called *hypercontractivity*) which has an unusually strong form when the test function f is the modulus of a holomorphic function. (This *strong* hypercontractivity, discovered by Svante Janson, was explored in a noncommutative context in my first paper [1].) In [4], we discovered that this strong hypercontractivity theorem — in the Gaussian case — actually depends only on a weaker regularity property of the modulus of a holomorphic function: such functions are *log subharmonic*. (I.e. nonnegative functions whose logarithms are subharmonic. Such functions are also subharmonic.) We proved the strong hypercontractivity estimates in this much larger cone of log subharmonic functions. This suggested some kind of strong log Sobolev inequality was at play, and we independently proved one: for the Gaussian measure γ_t of variance t on $\mathbb{R}^k$,

$$\mathrm{Ent}_{\gamma_t}(f^2) \leq \frac{1}{2t} \int fEf \, d\gamma_t, \qquad \text{for } f \text{ log subharmonic}, \tag{6}$$

where $Ef(x) = x \cdot \nabla f(x)$ is the *Euler operator*. Note, integrating by parts, that $\int |\nabla f|^2 \, d\gamma_t = \int f(-\Delta + \frac{1}{t}E) f \, d\gamma_t$, and so if f is subharmonic, one would deduce (6) but with the constant $\frac{1}{t}$; the sharp constant in the log subharmonic category is smaller. We then showed (using various central limit-like approximation schema) that several other measures (for example uniform measures on intervals, and symmetric Bernoulli measures) satisfy strong log Sobolev inequalities in the log subharmonic category.

In the sequel paper [12], we directly investigated the connection between strong hypercontractivity and our strong LSI. We proved, for a large class of probability measures (characterized by smooth densities with certain weak rigidity under the action of the Euclidean group), that our inequality (6) is precisely equivalent to strong hypercontractivity. This resolved a wrinkle in Gross's pioneering work on the subject. He originally

showed that the general LSI (5) is equivalent to hypercontractivity. In 2001, he also gave a proof of Janson's strong hypercontractivity theorem, *using* (5). But that proof could not be reversed: the LSI had to hold in the full space, not just the holomorphic space. Our main theorem in [12] is *intrinsic*: the strong LSI and strong hypercontractivity are equivalent, with tests functions in the same cone of functions in each case.

3.2 [7, 11] Sharp Constants for Duality in Gaussian Holomorphic L^p Spaces

Let γ_t denote the Gaussian measure of variance t on $\mathbb{C}^n$. The holomorphic L^2 space $\mathcal{H}L^2(\mathbb{C}^n, \gamma_t)$ is the image space of the Segal–Bargmann transform (see (3) above). The corresponding holomorphic L^p spaces $\mathcal{H}L^p(\mathbb{C}^n, \gamma_t)$ are Banach spaces of holomorphic functions. They are not, it turns out, complex interpolation scale: it is known that the usual map $L^{p'}(\mathbb{C}^n, \gamma_t) \to L^p(\mathbb{C}^n, \gamma_t)^*$ given by $f \mapsto \int (\cdot) \bar{f}\, d\gamma_t$, when restricted to $\mathcal{H}L^{p'}$, is not surjective if $p \neq 2$. Instead, the dual space can be identified with a *dilated* version of the space: using the identification $f \mapsto \Lambda(f) = \int (\cdot) \bar{f}\, d\gamma_{1/2}$,

$$\mathcal{H}L^p(\mathbb{C}^n, \gamma_{1/p})^* \cong \mathcal{H}L^{p'}(\mathbb{C}^n, \gamma_{1/p'}).$$

(The dilation means that one should really "attach" the Gaussian density γ_1 to the test function, and look at L^p-spaces taking the p^{th} power of the density as well. This has a nice interpretation, suggesting the considered objects should be holomorphic sections over the base Euclidean space; this is natural from the perspective of geometric quantization, where the Segal–Bargmann transform plays a role.) This was proven by Janson, Peetre, and Rochberg; they showed that the two spaces are equal (under Λ) with equivalent norms, and their estimates showed a maximum distortion of 2^n.

The subject of my two joint papers [7, 11] with Will Gryc is the computation of the sharp constants comparing the dual norm on $\mathcal{H}L^p(\mathbb{C}^n, \gamma_{1/p})$ to the norm on $\mathcal{H}L^{p'}(\mathbb{C}^n, \gamma_{1/p'})$. The lower bound can be computed directly in terms of the norm of the orthogonal projection operator $P\colon L^2(\mathbb{C}^n, \gamma_{1/2}) \to \mathcal{H}L^2(\mathbb{C}^n, \gamma_{1/2})$, considered (as an integral operator) as a transformation $L^p(\mathbb{C}^n, \gamma_{1/p}) \to \mathcal{H}L^p(\mathbb{C}^n, \gamma_{1/p})$. Using Lieb's important result that Gaussian integral kernels have only Gaussian maximizers, and careful application of Minkowski's inequality for integrals, we were able to reduce the computation of the norm of P to an optimization problem in 6 real dimensions, which (through very involved calculations) we solved to compute the exact sharp constant. The result is that

$$\|f\|_{\mathcal{H}L^{p'}(\mathbb{C}^n, \gamma_{1/p'})} \leq \|\Lambda(f)\|_{\mathcal{H}L^p(\mathbb{C}^n, \gamma_{1/p})^*} \leq C(p)^n \|f\|_{\mathcal{H}L^{p'}(\mathbb{C}^n, \gamma_{1/p'})} \tag{7}$$

where $C(p) = \frac{1}{2} p^{1/p} (p')^{1/p'}$. Moreover, the inequality on the left is sharp.

Our second paper [11] investigated the inequality on the right side of (7). We proved that Hölder's inequality in the spaces $\mathcal{H}L^p(\mathbb{C}^n, \gamma_{1/p})$ is always a strict inequality, which suggests that the right bound in (7) is *not* sharp for $p \neq 2$. We proved that the sharp constant is between $C(p)^n$ and $C(p)^{n/2}$, and showed that this lower bound is achieved on monomials, and (holomorphic) quadratic exponentials, strongly suggesting that this is the sharp bound. The techniques we developed yielded, as an interesting corollary, the sharpest-known bounds on the Taylor coefficients of a function $f \in \mathcal{H}L^p(\mathbb{C}, \gamma_{1/p})$ in terms of $\|f\|_p$.

3.3 [2, 3, 5] Sharp Constants for $\mathscr{R}$-Diagonal Operators

$\mathscr{R}$-diagonal operators appeared in the context of free probability in the late 1990s. They can be defined through random matrix models: let V be a convex potential function on $\mathbb{R}_+$ which tends to ∞. Then the random matrix with probability density

$$Z \mapsto c_N e^{-N \mathrm{Tr}[V(ZZ^*)]}$$

(for a normalization coefficient c_N) is an $\mathscr{R}$-diagonal model: it has a large-N limit (noncommutative) distribution which we call $\mathscr{R}$-diagonal. Examples include Haar unitaries, and the Ginibre ensemble (all i.i.d. complex Gaussian entries) described in Section 1, whose large-N limit is a *circular operator*. They form a rich class of non-self-adjoint operators (all non-normal except the Haar unitary), closed under freely independent sums and powers. They have an especially nice connection to free entropy, and they played a central role in Haagerup's groundbreaking progress towards solving the invariant subspace conjecture.

Thought of as acting on the free group $\mathbb{F}_k$ by convolution (i.e. the left-regular representation L), the generators of the group can be identified as freely independent Haar unitaries. There is an important classical inequality, know as Haagerup's inequality, on the free group: if $f \in \ell^2(\mathbb{F}_k)$ is supported on the sphere of reduced words of fixed length n, then the operator norm $L(f)$ is controlled by the ℓ^2-norm of f:

$$\|L(f)\|_{\mathrm{op}} \leq (n+1)\|f\|_2. \tag{8}$$

In [2], we proved a strong "holomorphic" version of (8): taking f supported on length-n words in the generators but not their inverses, the sharp constant reduces to $O(n^{1/2})$. (This can be thought of as a "holomorphic" setting since $\Lambda(u^{-1}) = \Lambda(u)^*$ for each generator; hence, such f is a noncommutative polynomial in the u-variables and not their "complex conjugates" u^*.) In fact, we proved a general version of this strong Haagerup inequality for $\mathscr{R}$-diagonal operators. If $a_1, \ldots, a_k$ are $\mathscr{R}$-diagonal and freely independent, then there is a universal constant $c = c(a_1, \ldots, a_k)$ so that, for any homogeneous degree-n noncommutative polynomial p in k variables,

$$\|p(a_1, \ldots, a_k)\|_{\mathrm{op}} \leq c\sqrt{n+1}\|p(a_1, \ldots, a_k)\|_2. \tag{9}$$

(The latter norm $\|\cdot\|_2$ is the Hilbert–Schmidt norm, and it plays a special role for $\mathscr{R}$-diagonal operators: it coincides with the spectral radius.) Our strong Haagerup inequality was quite a surprise to the community; several generalizations and applications of it (to the world of operator spaces, and to quantum groups) have appeared since.

Following [2], I explored the consequences of (9) to Markov semigroups in the $\mathscr{R}$-diagonal context. Then, in [5], joint with Haagerup and Speicher, we completely characterized the sharp asymptotic behavior of $\mathscr{R}$-diagonal Resolvents. If a is $\mathscr{R}$-diagonal, normalized so that the spectral radius $\|a\|_2 = 1$, we proved the universal asymptotic

$$\|(\lambda - a)^{-1}\|_{\mathrm{op}} \sim \sqrt{\frac{27}{32}}\sqrt{\|a\|_4^4 - 1}(\lambda - 1)^{-3/2}, \qquad \text{as } \lambda \downarrow 1.$$

We gave two proofs: one combinatorial using the diagrammatic approach to free cumulant calculations, and one complex analytic using the subordination theory for the Cauchy transform. Note that any normal operator of spectral radius 1 has resolvent of $O((\lambda-1)^{-1})$; it is quite appealing that this family of non-normal operators also has a universal blow-up rate, with constant determined by the fourth moment of the operator.

References

[1] **Kemp, T.**: *Hypercontractivity in noncommutative holomorphic spaces.* Commun. Math. Phys. **259** no. 3, 615-637 (2005)

[2] **Kemp, T.**; Speicher, R.: *Strong Haagerup inequalities for free $\mathscr{R}$-diagonal elements.* J. Funct. Anal. **251** no. 1, 141-173 (2007)

[3] **Kemp, T.**: *$\mathscr{R}$-diagonal dilation semigroups.* Math. Z. **264** no. 1, 111-136 (2010)

[4] Graczyk, P.; **Kemp, T.**; Loeb, J.J.: *Hypercontractivity for log-subharmonic functions.* J. Funct. Anal. **258** no. 6, 1785-1805 (2010)

[5] Haagerup, U.; **Kemp, T.**; Speicher, R.: *Resolvents of $\mathscr{R}$-diagonal operators.* Trans. Amer. Math. Soc. **362** no. 11, 6029-6064 (2010)

[6] **Kemp, T.**; Mahlburg, K.; Rattan, A.; Smyth, C.: *Enumeration of noncrossing pairings on bit strings.* J. Combin. Theory Ser. A **118**, no. 1, 129-151 (2011)

[7] Gryc, W.; **Kemp, T.**: *Duality in Segal-Bargmann spaces.* J. Funct. Anal. **261**, 1591-1623 (2011)

[8] **Kemp, T.**; Nourdin, I.; Peccati, G.; Speicher, R.: *Wigner chaos and the fourth moment.* Ann. Prob. **40**, 1577–1635 (2012)

[9] Driver, B.K.; Hall, B.C.; **Kemp, T.**: *The large-N limit of the Segal–Bargmann transform on $\mathbb{U}_N$.* J. Funct. Anal. **265** no. 11, 2585-2644 (2013)

[10] Collins, B.; **Kemp, T.**: *Liberating projections.* J. Funct. Anal. **266** no. 4, 1988-2052 (2014)

[11] Gryc, W.; **Kemp, T.**: *On sharp constants for dual Segal–Bargmann L^p spaces.* J. Math. Anal. Appl. **424**, no. 2, 1198-1222 (2015)

[12] Graczyk, P.; **Kemp, T.**; Leob, J.: *Strong logarithmic Sobolev inequalities for log-subharmonic functions.* Canad. J. Math. 67 (6), 1384-1410 (2015)

[13] **Kemp, T.**: *Heat kernel empirical laws on $\mathbb{U}_N$ and $\mathbb{GL}_N$.* To appear in J. Theor. Probab.

[14] **Kemp, T.**: *The large-N limits of Brownian motions on $\mathbb{GL}_N$.* To appear in Int. Math. Res. Not. IMRN.

[15] Cébron, G.; **Kemp, T.**: *Fluctuations of Brownian motions on $\mathbb{GL}_N$.* To appear in AIHP Probab. Stat.

[16] Driver, B.; Hall, B.; **Kemp, T.**: *Three proofs of the Makeenko–Migdal equation for Yang-Mills theory in the plane.* To appear in CMP.

[17] Driver, B.; Gabriel, F.; Hall, B.; **Kemp, T.**: *The Makeenko–Migdal equation for Yang-Mills theory on compact surfaces.* To appear in CMP.

[18] Gour, G.; **Kemp, T.**: *The minimum Renyi entropy output of a quantum channel is locally additive.* To appear in Letters in Math. Phys.

[19] Collins, B.; Dahlqvist, A.; **Kemp, T.**: *The hard edge of unitary Brownian motion.* To appear in Probab. Theory Related Fields.

[20] Driver, B.; Hall, B.; **Kemp, T.**: *The complex time Segal–Bargmann transform.* Preprint.

[21] Friedland, K.; **Kemp, T.**: *Most Boson quantum states are almost maximally entangled.* Preprint.

[22] **Kemp, T.**; Zimmermann, D.: *Random matrices with log-range correlations, and log-Sobolev inequalities.* Preprint.