

The orbit method and analysis of automorphic forms

by

PAUL D. NELSON

*Eidgenössische Technische Hochschule
Zurich, Switzerland*

AKSHAY VENKATESH

*Institute for Advanced Study
Princeton, NJ, U.S.A.*

Contents

1. Introduction	2
Part I. Microlocal analysis on Lie group representations I: definitions and basic properties	23
2. Operators attached to Schwartz functions	24
3. Operator classes	28
4. Symbol classes	31
5. Operators attached to symbols	36
6. The Kirillov formula	41
Part II. Microlocal analysis on Lie group representations II: proofs and refinements	45
7. Star product asymptotics	45
8. Proofs concerning the operator assignment	56
9. Infinitesimal characters	72
10. Localizing	78
11. Regular coadjoint multi-orbits	83
12. Trace estimates	86
Part III. Gan–Gross–Prasad pairs: geometry and asymptotics	90
13. Basic definitions and invariant theory	91
14. Stability	96
15. Satake parameters and L -functions	103
16. Volume forms	107
17. Measures and integrals	115
18. Relative characters: disintegration	121
19. Relative characters: asymptotics in the stable case	122
Appendix A. Some technicalities related to the Plancherel formula	136

Part IV. Inverse branching	142
20. Overview	142
21. The case of compact groups	144
22. The distinguished archimedean place	144
23. Preliminaries on representations of p -adic groups	152
24. The case of auxiliary p -adic places	157
Part V. Application to the averaged Gan–Gross–Prasad period	167
25. Setting	167
26. Construction of limit states	172
27. Equidistribution of limit states	180
28. Recap and overview of the proof	187
29. Spectral expansion and truncation of the H -period	189
30. The smoothly weighted asymptotic formula	192
31. The normalized asymptotic formula	200
References	204

1. Introduction

The problem of giving asymptotic formulas for moments of large-degree L -functions has proven challenging. In approaching this problem, one encounters difficult analytic questions in the representation theory of reductive Lie groups, involving complicated multi-dimensional oscillatory integrals.

The orbit method (see, e.g., [Ki2], [Vo3]) is a philosophy for, among other purposes, reducing difficult problems in the representation theory of Lie groups to simpler problems in symplectic geometry. It has been widely applied in the algebraic side of that theory.

This paper develops the orbit method in a quantitative analytic form. We combine the tools thus developed with an indirect application of Ratner’s theorem to study moments of automorphic L -functions on higher-rank groups.

1.1. Overview of results

We refer the reader who is not familiar with automorphic forms to §1.3 and onwards for an introduction, in explicit terms, to the main ideas of this paper.

Let $\mathbf{H} \hookrightarrow \mathbf{G}$ be an inclusion of reductive groups over a number field F . Let Π and Σ be cuspidal automorphic representations of \mathbf{G} and \mathbf{H} , respectively. Assuming that $\mathbf{H} \hookrightarrow \mathbf{G}$ is a strong Gelfand pair, and under a temperedness assumption, one may define an “automorphic branching coefficient” $\mathcal{L}(\Pi, \Sigma) \geq 0$ which quantifies how vectors in Π

correlate with Σ . We recall this definition in a simple setting in §1.4 and more fully in §25.4.

We focus on the ‘‘Gan–Gross–Prasad’’ case (§13) in which

$$(\mathbf{G}, \mathbf{H}) \text{ is a form of either } (\mathrm{SO}_n, \mathrm{SO}_{n-1}) \text{ or } (\mathrm{U}_n, \mathrm{U}_{n-1}).$$

The definition of $\mathcal{L}(\Pi, \Sigma)$ then applies, at least for tempered Π and Σ , and one expects $\mathcal{L}(\Pi, \Sigma)$ to be related to special values of L -functions: Ichino–Ikeda [II] and N. Harris [HarN] conjecture the formula

$$\mathcal{L}(\Pi, \Sigma) = 2^{-\beta} \frac{L^{(R)}\left(\frac{1}{2}, -, \Pi \times \Sigma^\vee\right)}{L^{(R)}(1, \mathrm{Ad}, \Pi \times \Sigma^\vee)} \Delta_G^{(R)}, \quad (1.1)$$

whose terms are as follows (see [II] and [HarN] for details):

- R is a fixed set of places outside of which Π and Σ are spherical (and thus $\mathcal{L}(\Pi, \Sigma)$ depends upon R);
- Σ^\vee denotes the contragredient of the unitary representation Σ ; it is isomorphic to the conjugate representation $\bar{\Sigma}$, and we will occasionally use the latter notation;
- $L^{(R)}$ denotes an L -function without Euler factors in R ;
- 2^β is the size of the Arthur component group of $\Pi \boxtimes \Sigma^\vee$ on $\mathbf{G} \times \mathbf{H}$;
- Δ_G is the L -function whose local factor at almost every prime p equals

$$\frac{p^{\dim(G)}}{\#\mathbf{G}(\mathbf{F}_p)}$$

(see [Gr]); e.g., $\Delta_G = \zeta(2)\zeta(4) \dots \zeta(2n)$ for $G = \mathrm{SO}_{2n+1}$; $\Delta_G^{(R)}$ omits factors at R .

The formula (1.1) has been proved in the unitary case, under local assumptions which allow one to use a simple form of the trace formula, by W. Zhang [Z1] (see also [Z2, §2.2] and [BP4]).

Fix one such Π . What are the asymptotic statistics of $\mathcal{L}(\Pi, \Sigma)$, as Σ varies over a large family? For example, what are the moments? Predictions for these may be obtained via (1.1) and random matrix heuristics (§1.2) for families of L -functions. To verify such predictions rigorously has proved an interesting challenge, testing our understanding of families of automorphic forms and L -functions. It has been successfully undertaken in many low-degree cases, where obtaining strong error estimates remains an active area of research (see, e.g., [BFK+], [BM], [KMS], [BFS+] and references).

We aim here to explore some first cases of arbitrarily large degree. Our main result (Theorem 31.11) may be summarized informally as follows.

THEOREM 1.1. *Assume certain local conditions, including the compactness of the quotients $\mathbf{G}(F)\backslash\mathbf{G}(\mathbb{A})$ and $\mathbf{H}(F)\backslash\mathbf{H}(\mathbb{A})$. For each sufficiently small positive real h , let \mathcal{F}_h be the family of all Σ as above which are locally distinguished by Π , have Satake parameters at some fixed archimedean place inside the rescaling $h^{-1}\Omega$ of some nice fixed compact set Ω , and have “fixed level” at the remaining places in R . Then, the branching coefficient $\mathcal{L}(\Pi, \Sigma)$, averaged over $\Sigma \in \mathcal{F}_h$, is asymptotic to $\frac{1}{2}$:*

$$\lim_{h \rightarrow 0} \frac{1}{|\mathcal{F}_h|} \sum_{\Sigma \in \mathcal{F}_h} \mathcal{L}(\Pi, \Sigma) = \frac{1}{2}. \quad (1.2)$$

For “typical” Π and Σ , we expect that (1.1) holds with $\beta=2$ (see §25.5 for further explanation). Our result should thus translate, under (1.1), as follows:

$$\text{The average value of } \frac{L^{(R)}(\frac{1}{2}, -, \Pi \times \Sigma^\vee)}{L^{(R)}(1, \text{Ad}, \Pi \times \Sigma^\vee)} \Delta_G^{(R)} \text{ is 2.} \quad (1.3)$$

We outline in §1.2 why (1.3) agrees with random matrix theory heuristics for orthogonal families of L -functions with positive root number.

One way to normalize the strength of (1.2) is to note (§31.4) that the size of the family \mathcal{F}_h is roughly the fourth power of the analytic conductor of the relevant L -function. By ignoring all but one term and slowly shrinking the family, we obtain a “weakly subconvex” bound

$$\mathcal{L}(\Pi, \Sigma) = o(\text{cond}(\Pi \times \Sigma)^{1/4}) \quad (1.4)$$

(compare with [So], [SoTh]). The hypotheses relevant for (1.4) are that Π is fixed, while Σ traverses a sequence whose archimedean Satake parameters all tend off to ∞ at the same rate.

The new ideas used to obtain (1.2) are based on the orbit method, applied in two ways:

- To determine the asymptotics of complicated oscillatory integrals on higher-rank groups. For instance, Theorem 19.3 gives general and uniform asymptotic expansions of relative characters away from the conductor-dropping locus. This analysis gives a robust supply of analytic test vectors for the local matrix coefficient integrals as in the conjecture of Ichino–Ikeda. We hope these to be of general use in analytic problems involving families of automorphic forms in higher rank.

- To obtain invariant measures towards which we can apply measure-theoretic techniques. Indeed, a major global ingredient for (1.2) is an application of Ratner’s theorem to the case of measures invariant by the centralizer of a regular nilpotent element in $\mathbf{G}(F)$. The estimate (1.2) is ineffective, and the application of Ratner is solely responsible for the ineffectivity. We expect that an effective version of Ratner’s theorem for the

case at hand would lead to a subconvex estimate for $\mathcal{L}(\Pi, \Sigma)$; this perhaps contributes interest to the problem of effectivization.

To implement these, we develop a microlocal calculus for Lie group representations, which may be understood as a quantitative, analytic form of the orbit method and the philosophy of geometric quantization.

These basic ideas do not interact with the arithmetic nature of the setting; in particular, we do not use Hecke operators. However, the problem of averaging and bounding L -functions seems to be the most interesting source of applications at the moment.

1.2. Compatibility with random matrix heuristics

We briefly outline why our result (1.3) should be compatible with the standard heuristics.

Random matrix theory heuristics (see, e.g., [BK] and references, such as [KaSa], [CF], [CFK+], [KeSn1], [KeSn2]) suggest, for a family of L -functions $L(\frac{1}{2}, f)$ parameterized by elements f of some nice enough family \mathcal{F} , that

$$\frac{1}{|\mathcal{F}|} \sum_{\mathcal{F}} L^{(R)}\left(\frac{1}{2}, f\right) \approx (\text{global factor } g) \times \prod_p \langle a_p \rangle,$$

where

- p runs over the finite primes of F outside R ,
- $\langle a_p \rangle$ is the expectation

$$\langle a_p \rangle := \frac{1}{|\mathcal{F}|} \sum_{\mathcal{F}} L_p\left(\frac{1}{2}, f\right)$$

of the central value of the p th Euler factor, and

- the “global factor” g is described by random matrix theory, and given here by the following limit of integrals taken with respect to probability Haar measures:

$$g = \lim_{N \rightarrow \infty} \int_{x \in \text{SO}(2N)} \det(1-x).$$

Indeed, the symmetry type [KaSa] of the family of L -functions implicit in (1.2) is $O(2N)$, i.e., it is an orthogonal family with positive root numbers. One may see this by, for instance, considering the analogous situation with the number field replaced by a function field (see also [ShTe] and [SST]).

Each integral

$$\int_{x \in \text{SO}(2N)} \det(1-x)$$

equals 2, independent of N . Indeed, $\det(1-x)$ is the (super-)trace of $x \in \mathrm{SO}(2N)$ on $\Lambda^*(\mathbb{C}^{2N})$. The average trace of x on $\Lambda^j(\mathbb{C}^{2N})$ computes the dimension of invariants, which is 1 for $j \in \{0, 2N\}$ and 0 otherwise. Thus, $g=2$.

We sketch here why $\langle a_p \rangle = 1$ for every p . We expect the families \mathcal{F}_h considered here to have the (provable) property that the local component σ_p at p of a uniformly random element $\Sigma \in \mathcal{F}_h$ becomes distributed, as $h \rightarrow 0$, according to the Plancherel measure $d\sigma_p$, thus

$$\langle a_p \rangle = \int_{\sigma_p} \underbrace{\frac{L_p(\frac{1}{2}, -, \pi_p \times \sigma_p^\vee)}{L_p(1, \mathrm{Ad}, \pi_p \times \sigma_p^\vee)} \Delta_{G,p}}_{=: I(\pi_p, \sigma_p)} d\sigma_p.$$

Ichino and Ikeda [II, Theorem 1.2] have shown in the orthogonal case, and N. Harris [HarN, Theorem 2.12] in the unitary case, that this integral can be rewritten as a matrix coefficient integral:

$$I(\pi_p, \sigma_p) = \int_{h \in H_p} \langle hv, v \rangle \langle u, hu \rangle dh,$$

where $v \in \pi_p$ and $u \in \sigma_p$ are spherical unit vectors and dh is the Haar measure assigning unit volume to a hyperspecial maximal compact subgroup. Using the Plancherel formula, one can show (§18) that

$$\int_{\sigma_p} I(\pi_p, \sigma_p) d\sigma_p = 1,$$

as required.

We have not discussed yet the quantity β as in (1.1). As mentioned above, we expect for “typical” Π and Σ that $\beta=2$. However, for “atypical” Π (i.e., endoscopic lifts), one can have $\beta > 2$ in the entire family. In that case, the limit of (1.3) is instead a larger power of 2. Correspondingly, the L -function in question factors. It seems to us that our result continues to match with L -function heuristics after appropriately accounting for this factorization and variation of root numbers, but we have not checked all details.

The random matrix theory predictions for higher moments involve somewhat more complicated coefficients than the quantity $g=2$ appearing above. It may be possible to study higher moments by adapting our method to periods of Eisenstein series, and would be interesting to obtain in that way some geometric perspective on those predictions.

1.3. Basic setup

We now simultaneously outline the contents of this paper and sketch the main arguments.

Suppose given a pair of unimodular Lie groups G and H , with $H \leq G$. An example relevant for our main theorem is when

$$G \text{ and } H \text{ are the real points of the split forms of } \mathrm{SO}_{n+1} \text{ and } \mathrm{SO}_n.$$

We assume that representations of G have “multiplicity-free restrictions” to H , as happens in the indicated example (see [SZ] for a precise statement and proof).

Suppose also given a lattice Γ in G for which $\Gamma_H := \Gamma \cap H$ is a lattice in H . We assume that both quotients

$$[G] := \Gamma \backslash G \quad \text{and} \quad [H] := \Gamma_H \backslash H$$

are compact. We equip them with Haar measures.

For the motivating applications to L -functions, we must consider adelic quotients. This entails some additional work at “auxiliary places” (see §24), which we do not discuss further in this overview.

1.4. Branching coefficients: comparing global and local restrictions

Let

$$\pi \subseteq L^2([G])$$

be an irreducible unitary subrepresentation, with the group G acting by right translation. We assume that π is tempered. The branching coefficients of interest arise from comparing the *two* natural ways to restrict π to H :

- (Globally) Take a smooth vector $v \in \pi$. It defines a function on $[G]$. We may restrict it to obtain a function $v|_{[H]}$ on $[H]$. The L^2 -norm of that restriction may be decomposed as

$$\int_{[H]} |v|^2 = \sum_{\text{irreducible } \sigma \subseteq L^2([H])} \|\text{projection of } v|_{[H]} \text{ to } \sigma\|^2. \quad (1.5)$$

- (Locally) Consider π as an abstract unitary representation of G . We may restrict it to obtain an abstract unitary representation $\pi|_H$ of H . We verify in the examples of interest (see §18) that this restriction decomposes as a direct integral

$$\pi|_H = \int_{\sigma} m_{\sigma} \sigma,$$

weighted by multiplicities $m_{\sigma} \in \{0, 1\}$, and taken over tempered irreducible unitary representation σ of H with respect to Plancherel measure. We may define the components $\text{pr}_{\sigma}(v) \in \sigma$ of a smooth vector $v \in \pi$ with respect to this decomposition, and we have

$$\|v\|^2 = \int_{\sigma} \|\text{pr}_{\sigma}(v)\|^2. \quad (1.6)$$

A priori, $\|\text{pr}_{\sigma}(v)\|$ is defined only as a measurable function of σ , but there is a natural way to define it pointwise in the cases of interest.

Let $\sigma \subseteq L^2([H])$ be an irreducible tempered subrepresentation for which $m_\sigma = 1$; we refer to such σ as π -*distinguished*. By the multiplicity-1 property, we may define a proportionality constant $\mathcal{L}(\pi, \sigma) \geq 0$ by requiring that, for every smooth vector $v \in \pi$,

$$\|\text{projection of } v|_{[H]} \text{ to } \sigma\|^2 = \mathcal{L}(\pi, \sigma) \cdot \|\text{pr}_\sigma(v)\|^2. \quad (1.7)$$

We note that $\mathcal{L}(\pi, \sigma)$ depends upon the choices of Haar measure.

1.5. Objective

Let h traverse a sequence of positive reals tending to zero, and let \mathcal{F}_h be a corresponding sequence of families consisting of irreducible tempered representations $\sigma \subseteq L^2([H])$ for which $m_\sigma = 1$.

Assume that each family \mathcal{F}_h arises from some nice subset $\tilde{\mathcal{F}}_h$ of the π -distinguished tempered dual of H as the set of all irreducible $\sigma \subseteq L^2([H])$ for which the class of σ belongs to $\tilde{\mathcal{F}}_h$. Assume that $|\mathcal{F}_h| \rightarrow \infty$ as $h \rightarrow 0$. (In our main theorem, we take for $\tilde{\mathcal{F}}_h$ the set of σ whose infinitesimal character belongs to the rescaling $h^{-1}\Omega$ of some fixed non-empty bounded open set Ω .) Our aim is to determine the asymptotics of the sums

$$\sum_{\sigma \in \mathcal{F}_h} \mathcal{L}(\pi, \sigma).$$

In our motivating examples, these are (in some cases conjecturally) proportional to sums of special values of L -functions.

We drop the subscript h in what follows for notational simplicity.

1.6. Strategy

The rough idea of our proof is to find a vector $v \in \pi$ which simultaneously “picks out the family \mathcal{F} ” in that

$$\|\text{pr}_\sigma(v)\|^2 \approx \begin{cases} 1, & \text{if } \sigma \in \tilde{\mathcal{F}}, \\ 0, & \text{if } \sigma \notin \tilde{\mathcal{F}}, \end{cases} \quad (1.8)$$

and “becomes equidistributed” in that

$$\frac{\int_{[H]} |v|^2}{\text{vol}([H])} \sim \frac{\int_{[G]} |v|^2}{\text{vol}([G])}. \quad (1.9)$$

The vector v will of course depend upon the family $\mathcal{F} = \mathcal{F}_h$, hence upon the asymptotic parameter $h > 0$, and the above estimates are to be understood as holding asymptotically in the $h \rightarrow 0$ limit. Note that (1.8) is a purely local problem of harmonic analysis

in the representation π , whereas (1.9) is a global problem: it relates to the specific way in which π is embedded in $L^2([G])$.

The Weyl law on $[H]$ says that the cardinality of \mathcal{F} is approximately the volume of $[H]$ times the Plancherel measure of $\tilde{\mathcal{F}}$, and thus, by (1.8) and (1.6),

$$|\mathcal{F}| \approx \text{vol}([H]) \|v\|^2. \quad (1.10)$$

Comparing (1.5), (1.7) and (1.8) with (1.10) yields the required asymptotic formula

$$\frac{1}{|\mathcal{F}|} \sum_{\sigma \in \mathcal{F}} \mathcal{L}(\pi, \sigma) \sim \frac{1}{\text{vol}([G])}. \quad (1.11)$$

In applications, the volumes are defined using Tamagawa measure and given then by

$$\text{vol}([H]) = \text{vol}([G]) = 2.$$

Thus, (1.11) leads to (1.2).

Observe, finally, that it suffices to produce *families* of vectors v_i , $i \in I$, for which the analogues of (1.8) and (1.9) hold on average over the index set I .

We note that this basic strategy has appeared (sometimes implicitly) in several antecedent works (see, e.g., [Sar], [BR2], [Ve], [MV]). We would like to note, in particular, the influence of the ideas of Bernstein and Reznikov in exploiting the uniqueness of an invariant functional. The novelty here is that we execute this strategy successfully in arbitrarily large rank.

1.7. Microlocal calculus for Lie group representations

To implement this strategy, we need some way to produce and analyze (families of) vectors in the representation π . Our approach, inspired by microlocal analysis, is to work with vectors implicitly through their symmetry properties under group elements $g \in G$ within a suitable shrinking neighborhood of the identity element. Since we do not expect our readers to have extensive prior familiarity with microlocal analysis, we describe here its content in our setting. The discussion in this section is rather informal, but we hope that the reader will find it helpful in navigating the many technical estimates of the text.

Let \mathfrak{g} denote the Lie algebra of G , and \mathfrak{g}^\wedge the Pontryagin dual of \mathfrak{g} ; we identify \mathfrak{g}^\wedge with the imaginary dual $i\mathfrak{g}^*$, and denote by $e^{x\xi} \in \mathbb{C}^{(1)}$ the image of $(x, \xi) \in \mathfrak{g} \times \mathfrak{g}^\wedge$ under the natural pairing.

The shrinking neighborhood in question depends upon an infinitesimal scaling parameter $\hbar \rightarrow 0$. Informally, we say that a vector $v \in \pi$ is *microlocalized at* $\xi \in \mathfrak{g}^\wedge$ if

$$\pi(\exp(x))v \approx e^{x\xi/\hbar}v \quad (1.12)$$

for all $x \in \mathfrak{g}$ of size $|x| = O(\hbar)$. In practice, we choose \hbar small enough in terms of π that such vectors exist. For instance, if π is generated by a Maass form φ of Laplace eigenvalue λ on a locally symmetric space $\Gamma \backslash G/K$, then we choose \hbar comparable to or smaller than the “wavelength scale” $|\lambda|^{-1/2}$ of φ ; similar considerations apply more generally, with the role of wavelength scale played by the inverse root norm $|\lambda_\pi|^{-1/2}$ of the infinitesimal character of π (see §9.8). We quantify the heuristic (1.12) in several different ways throughout this paper, typically by working with sequences of representations $\pi = \pi_\hbar$ and vectors $v = v_\hbar$ that vary with \hbar and asking that the difference between the left- and right-hand sides of (1.12) decay at some specified rate as $\hbar \rightarrow 0$.

The group elements $g = \exp(x)$ with $|x| = O(\hbar)$ approximately commute as $\hbar \rightarrow 0$, and so the operators $\pi(g)$ may be approximately simultaneously diagonalized; their common approximate eigenvectors are the microlocalized vectors. We might thus hope for an approximate decomposition

$$\pi \approx \bigoplus_{\xi} \mathbb{C}v_{\xi}, \quad (1.13)$$

where ξ traverses a subset of \mathfrak{g}^\wedge and $v_{\xi} \in \pi$ is microlocalized at ξ , and to have an approximate functional calculus

$$\text{Op}_{\hbar}: \{\text{symbols } \mathfrak{g}^\wedge \rightarrow \mathbb{C}\} \longrightarrow \{\text{operators on } \pi\}, \quad (1.14)$$

$$\text{Op}_{\hbar}(a)v_{\xi} \approx a(\xi)v_{\xi}, \quad (1.15)$$

where “symbol” refers to a class of functions with suitable regularity.

To implement such ideas rigorously, we write down an operator assignment (1.14) and verify that it has the properties suggested by the heuristics (1.13) and (1.15). The assignment is similar to the classical Weyl calculus in the theory of pseudodifferential operators, which may be recovered (more-or-less) by specializing our discussion to standard representations of Heisenberg groups. The definition and basic properties of this assignment are philosophically unsurprising. If $a: \mathfrak{g}^\wedge \rightarrow \mathbb{C}$ is the Fourier mode corresponding to a small enough element x in the Lie algebra of G , then $\text{Op}_{\hbar}(a) = \pi(\exp(\hbar x))$. If a is real-valued, then $\text{Op}_{\hbar}(a)$ is self-adjoint. If a is positive, then $\text{Op}_{\hbar}(a)$ is asymptotically positive as $\hbar \rightarrow 0$. The association $a \mapsto \text{Op}_{\hbar}(a)$ is nearly G -equivariant. One has composition formulas relating $\text{Op}_{\hbar}(a)\text{Op}_{\hbar}(b)$ to $\text{Op}_{\hbar}(a \star_{\hbar} b)$ for a suitable star product $a \star_{\hbar} b$, bounds for operator and trace norms, and so on; see Theorems 4.5–5.8, 6.3–8.11 and 12.2.

Our main input concerning π is the Kirillov character formula, due in this case to Rossmann, which asserts roughly (see §12 for a precise statement) that

$$\mathrm{tr}(\mathrm{Op}_h(a)) \approx \int_{\mathfrak{h} \mathcal{O}_\pi} a \, d\omega, \quad (1.16)$$

where $\mathcal{O}_\pi \subseteq \mathfrak{g}^\wedge$ is a G -orbit (or finite union thereof), called the *coadjoint orbit attached to π* . In the metaphor of microlocal analysis, the coadjoint orbit \mathcal{O}_π is the “phase space” underlying π ; it bears the same relationship to π as the cotangent bundle T^*M of a manifold M does to $L^2(M)$. The coadjoint orbit is equipped with a canonical symplectic structure, and $d\omega$ is the associated symplectic volume form.

For each real-valued symbol a , we may informally identify the self-adjoint operator $\mathrm{Op}_h(a)$ with a family of vectors $(v_i)_{i \in I}$ by writing it in the form $\sum_{i \in I} v_i \otimes \bar{v}_i$ for some $v_i \in \pi$; here we have identified operators on π with elements of (a suitable completion of) $\pi \otimes \bar{\pi}$. If the symbol a is taken to be suitably concentrated near a regular point $\xi \in \mathfrak{g}^\wedge$, then the corresponding family is essentially a singleton (i.e., of cardinality $O(h^{-\varepsilon})$), consisting of vectors microlocalized at ξ . By decomposing the constant symbol $a=1$ associated with the identity operator $\mathrm{Op}_h(a)=1$ and appealing to (1.16), we may write any reasonable vector as a linear combination of microlocalized vectors v_ξ , taken over representatives ξ/h for a partition of the coadjoint orbit \mathcal{O}_π into pieces of unit symplectic volume. We may thus regard the microlocalized vectors as giving an approximate basis (1.13), with respect to which the $\mathrm{Op}_h(a)$ act as the approximate multipliers (1.15); in other words,


$$\mathrm{Op}_h(a) \approx \sum_{\xi \in \mathfrak{h} \mathcal{O}_\pi} a(\xi) v_\xi \otimes \bar{v}_\xi.$$

In this way, the Op_h -calculus parameterizes weighted families of microlocalized vectors. These considerations apply uniformly across the various classes of tempered representations of G (principal series, discrete series, etc.) and without reference to any explicit model.

Informally, if the dominant contribution to the decomposition of a vector $v \in \pi$ as a sum $\sum v_\xi$ of microlocalized vectors comes from those ξ belonging to some nice subset of $\mathfrak{h} \mathcal{O}_\pi$, then we refer to that subset as the *microlocal support* of v ; equivalently, it describes where the distribution on \mathfrak{g}^\wedge given by $a \mapsto \langle \mathrm{Op}_h(a)v, v \rangle$ is concentrated. A vector is then microlocalized if its microlocal support is concentrated near a specific point.

These notions from microlocal analysis play a central role throughout the paper, so we illustrate their content in a couple basic examples. Figure 1 depicts coadjoint orbits for certain representations π of the groups

$$G = \mathrm{SO}(3) \quad \text{and} \quad G = \mathrm{SO}(1, 2) \cong \mathrm{PGL}_2(\mathbb{R}),$$

Figure 1. Coadjoint orbits for $(\mathrm{SO}_3, \mathrm{SO}_2)$.

respectively; in the latter case, π belongs to the holomorphic discrete series.⁽¹⁾ Each of these groups contains the compact subgroup $H = \mathrm{SO}(2)$. The circles drawn on the orbits are level sets for the “ z -coordinate” projection $\mathfrak{g}^* \rightarrow \mathfrak{h}^*$ dual to the inclusion $\mathfrak{h} \hookrightarrow \mathfrak{g}$ of Lie algebras. They divide the orbits into strips of equal symplectic volume, say of volume 1, which correspond under the orbit method philosophy to the basis of π given by H -isotypic weight vectors e_n ; the strip should be regarded as the microlocal support of the corresponding weight vector. We may normalize the weights n to be even integers lying in

$$\begin{cases} [-k, k], & \text{if } G = \mathrm{SO}(3), \\ (-\infty, -k-2] \cup [k+2, \infty), & \text{if } G = \mathrm{SO}(1, 2) \cong \mathrm{PGL}_2(\mathbb{R}), \end{cases} \quad (1.17)$$

for some non-negative even integer k .

Let us now take k tending off to ∞ , but simultaneously zoom out our camera by the factor k , so that the above picture of the coadjoint orbit \mathcal{O}_π remains constant. Which weight vectors should we then consider to be “microlocalized”? That is to say, for which vectors does the “zoomed out” microlocal support concentrate within some small distance of a specific point in the picture as $k \rightarrow \infty$? The strength of this notion depends upon the definition of “small distance”, which can sensibly range from the weakest scale $o(1)$ to the Planck scale $O(k^{-1/2})$.

Vectors of highest or lowest weight ($e_{\pm k}$ or $e_{\pm(k+2)}$) are microlocalized, even at the Planck scale, since the corresponding strips are concentrated near the “caps” of the coadjoint orbit (i.e., the regions of extremal z -coordinate). By contrast, “typical” weight vectors—such as the weight-zero vector e_0 in the representation of $\mathrm{SO}(3)$, corresponding to the equatorial strip—are *not* microlocalized, even at the weakest scale. In particular, weight vectors do *not* give an approximate basis of microlocalized vectors as in (1.13). The partition of the coadjoint orbit corresponding to a microlocalized basis would instead

⁽¹⁾ The pictures were created using the online graphing calculator GeoGebra [Hoh].




Figure 2. Rescaling to the nilcone.

have every partition element concentrated near a specific point.

Microlocalized vectors occasionally go in the literature by other names, such as “coherent states”. They are extremely useful for the sort of asymptotic analysis pursued in this paper. Among other desirable properties, their matrix coefficients behave simply near the identity, and are as concentrated as possible; we discuss this phenomenon further in §1.10.

In the body of this paper, we do not often refer explicitly to microlocalized vectors. We instead work with them implicitly through their approximate projectors $\text{Op}_h(a)$. We hope that by phrasing the introduction in terms of microlocalized vectors, it may serve as a useful guide to the ideas behind the arguments executed in the body.

1.8. Measure classification; equidistribution

In the discussion starting in §1.7, we allowed both π and h to vary simultaneously. Let us suppose now that the representation π is held fixed, independent of the scaling parameter h . As $h \rightarrow 0$, the rescaled coadjoint orbit $h\mathcal{O}_\pi$ then tends to a subset of the nilcone $\mathcal{N} \subseteq \mathfrak{g}^\wedge$. Figure 2 depicts this for the coadjoint orbits corresponding to fixed principal series representations of $G = \text{SO}(1, 2) \cong \text{PGL}_2(\mathbb{R})$. The operators $\text{Op}_h(a)$ are negligible unless a is supported near \mathcal{N} .

We thereby obtain in the $h \rightarrow 0$ limit a sequence

$$\{\text{functions } \mathcal{N} \rightarrow \mathbb{R}\} \xrightarrow{\text{Op}_h} \{\text{self-adjoint } T \in \text{End}(\pi)\} \longrightarrow \{\text{measures on } [G]\}, \quad (1.18)$$

where the final map sends $T = \sum_i v_i \otimes \bar{v}_i \in \text{End}(\pi) \cong \pi \otimes \bar{\pi}$ to $\sum_i |v_i|^2(x) dx$, with dx the chosen Haar measure on $[G]$. In §26, we implement this informal discussion rigorously and construct a G -equivariant limit map from functions on \mathcal{N} to measures on $[G]$. We emphasize that this limit map is insensitive to the shape of the unscaled coadjoint orbit \mathcal{O}_π , whose role becomes replaced in the limit by a subset of the nilcone \mathcal{N} .

A key observation is now that the limits of the measures on $[G]$ obtained via (1.18) may be understood using Ratner’s theorem. The application of Ratner is indirect, be-

cause these measures themselves do not acquire any obvious additional invariance; rather, they may be decomposed into measures having unipotent invariance.

Indeed, after suitably rescaling, we may describe the limiting behavior of the sequence (1.18) in terms of a G -invariant measure μ on the product space $\mathcal{N} \times [G]$. Let $\mathcal{N}_{\text{reg}} \subseteq \mathcal{N}$ denote the regular subset; it is the union of the open G -orbits on the nilcone \mathcal{N} , and its complement $\mathcal{N} \setminus \mathcal{N}_{\text{reg}}$ has lower dimension. We assume that π is generic, or equivalently, has maximal Gelfand–Kirillov dimension (cf. §11.4.2); for instance, the principal series satisfy this assumption. Then, the support of μ intersects $\mathcal{N}_{\text{reg}} \times [G]$. By disintegrating the restriction of μ to $\mathcal{N}_{\text{reg}} \times [G]$ over the projection to \mathcal{N}_{reg} , we obtain a non-trivial family of measures μ_ξ on $[G]$ indexed by $\xi \in \mathcal{N}_{\text{reg}}$. Speaking informally, we may regard μ_ξ as the $\hbar \rightarrow 0$ limit of an average of L^2 -masses $|v_\xi|^2(x) dx$ taken over all vectors v_ξ microlocalized at ξ . In any case, each such measure μ_ξ is invariant by the centralizer of the regular nilpotent element ξ . In favorable situations, an application of Ratner’s theorem then forces the μ_ξ and hence μ itself to be uniform.

This last paragraph mimics, in the context of Lie group representations, some of the semiclassical ideas behind the construction of the microlocal lift. We discuss this connection at more length in a sequel to this paper.

The argument just described gives a rich supply of families $(v_i)_{i \in I}$ of vectors $v_i \in \pi$ for which

$$\int_{[G]} |v_i|^2 \Psi \sim \frac{\int_{[G]} |v_i|^2}{\text{vol}([G])} \int_{[G]} \Psi, \quad \text{on average over } i \in I, \quad (1.19)$$

for fixed continuous functions Ψ on $[G]$. Although the characteristic function of $[H] \subseteq [G]$ is not itself continuous, we may deduce an averaged form of (1.9) by applying a similar argument to the derivatives of the functions obtained via (1.18) and appealing to the Sobolev lemma. This approach may be understood as an infinitesimal variant of the “period thickening” technique of Bernstein–Reznikov [BR1].

1.9. Branching and stability

Having indicated how we achieve the objective (1.9), we turn now to the problem of producing (families of) vectors v which pick out the family \mathcal{F} as in (1.8).

This is a *quantitative* version of the branching problem in representation theory: we wish to understand not only how a representation of G restricts to H , but in fact the behavior of individual vectors under the restriction process. Our approach is inspired by the following basic principle of the orbit method: restricting π to H should correspond to

$$\text{disintegrating the coadjoint orbit } \mathcal{O}_\pi \text{ along the projection } \mathfrak{g}^\wedge \longrightarrow \mathfrak{h}^\wedge. \quad (1.20)$$

For example, the distinction of σ by π should be equivalent, at least asymptotically, to the existence of a solution to the equation

$$\xi|_{\mathfrak{h}} = \eta, \quad (1.21)$$

with $(\xi, \eta) \in \mathcal{O}_\pi \times \mathcal{O}_\sigma \subseteq \mathfrak{g}^\wedge \times \mathfrak{h}^\wedge$.

The geometry of the projection (1.20) plays an important role in our argument, so we devote a fair amount of space to studying it in a purely algebraic context (see §13, §14 and §16). Of particular significance is the branch locus of (1.20), i.e., the locus where the map $\mathcal{O}_\pi \rightarrow \mathfrak{h}^\wedge$ induced by (1.20) fails to have surjective differential. Many features of our analysis break down near this branch locus.

We recall, from geometric invariant theory, that $\xi \in \mathfrak{g}^\wedge$ is called *stable* if the following conditions are satisfied (see §14 for details):

- ξ does not lie in the branch locus; equivalently, ξ has trivial \mathfrak{h} -stabilizer;
- ξ has closed \mathbf{H} -orbit, where \mathbf{H} is the algebraic group underlying H .

This notion plays a fundamental role in our paper, and appears to be analytically significant: the complement of the stable locus is where the analytic conductor $C(\pi \times \bar{\sigma})$ of the Rankin–Selberg L -function $L(\pi \times \bar{\sigma}, s)$ drops (see §15.4).

For instance, in the basic examples depicted in Figure 1, with $G \in \{\mathrm{SO}(3), \mathrm{SO}(1, 2)\}$ and $H = \mathrm{SO}(2)$, the coadjoint orbits for H are just singletons $\{\eta\}$, corresponding to 1-dimensional representations, and the equation (1.21) says that ξ should have z -coordinate given by η ; in other words, the solutions to (1.21) are the horizontal slices shown in the diagram.

The stable case, in Figure 1, is when ξ is not at the north or south poles of the sphere. Note that the group $H = \mathrm{SO}(2)$ of rotations fixing the z -axis acts transitively on any circle in \mathcal{O}_π with z -coordinate η , and freely away from the poles. This is a general pattern: the set of solutions to (1.21)—if non-empty—admits a diagonal action by H , which is simply transitive in the stable case (cf. §14.3 and §17).

Figure 3 depicts the coadjoint orbit of a principal series representation π of $G = \mathrm{SO}(1, 2) \cong \mathrm{PGL}_2(\mathbb{R})$, with H being the diagonal subgroup $\mathrm{SO}(1, 1) \cong \mathrm{GL}_1(\mathbb{R})$. One sees again that H acts simply transitively on generic fibers (i.e., away from the “cross”).

1.10. Analysis of matrix coefficient integrals; inverse branching

Having set up the necessary preliminaries regarding the geometry of orbits, we return to the problem of producing (families of) vectors v which pick out the family \mathcal{F} as in (1.8). The solution involves several steps which may be of independent interest:

- (i) We prove in §19 some asymptotic formulas for $\|\mathrm{pr}_\sigma(v)\|^2$, on average over v , when the pair (π, σ) is tempered and the infinitesimal characters of π and σ satisfy a stability




Figure 3. H -orbits on the one-sheeted hyperboloid.

condition (see §14.2 for the exact definition). In more technical terms, we compute asymptotically the Fourier transform of relative characters in a small neighborhood of the identity. The asymptotic formulas readily give a solution (§22) to the inverse problem of producing (families of) v approximating a given function $\sigma \mapsto \|\mathrm{pr}_\sigma(v)\|^2$. The proofs depend heavily upon the operator calculus discussed in §1.7.

(ii) We must estimate the “undesirable” contributions to (1.5) coming from σ which are either non-tempered, non-stable or of “high frequency”. The proofs (see, e.g., §22.3 and §29) apply the operator calculus. We refer to §28 for some discussion of the relevant subtleties.

As illustration, let us indicate how the methods underlying steps (i) and (ii) address the *analytic test vector problem* for a tempered pair (π, σ) in the stable case, i.e., away from conductor-dropping. Informally speaking, that problem is to find

“nice” vectors $v \in \pi$ and $u \in \sigma$ for which the local period $|\langle \mathrm{pr}_\sigma(v), u \rangle|^2$ is “large”.

Solutions to this problem have appeared in period-based approaches to the subconvexity period (see, e.g., [MV, §3.6.1], [NPD2, §2.17.1], [BR2], [Re], [DN], [Bl], [NPD1, Remark 50], [HNS, Theorem 1.2]); the point is that the period formula (1.7) and a “trivial” estimate for the global period $\int_{[H]} v \bar{u}$ often suffice to recover the convexity bound for L -function, suggesting the possibility for improvement in arithmetic settings via Hecke amplification.

The problem reduces to a robust understanding of the asymptotics of the local periods, which may be expressed in terms of matrix coefficient integrals as follows (see §18):

$$|\langle \mathrm{pr}_\sigma(v), u \rangle|^2 = \int_{s \in H} \langle \pi(s)v, v \rangle \langle u, \sigma(s)u \rangle. \quad (1.22)$$

Such integrals had previously been estimated in some low-rank examples (see, e.g., [Wa], [Wo], [BR2], [NPS], [FS2], [FS1]) via explicit calculation.

To analyze (1.22), first decompose as in §1.7 to reduce to the case that v and u are microlocalized at some elements $\xi \in \mathfrak{h}\mathcal{O}_\pi$ and $\eta \in \mathfrak{h}\mathcal{O}_\sigma$ of the rescaled coadjoint orbits. If ξ and η do not satisfy (1.21), at least up to some small error, then the resulting integral is readily seen to be negligible.

We thereby reduce to giving an asymptotic expansion for (1.22) when v and u are microlocalized at a stable solution (ξ, η) to (1.21). One of the main local results of this paper (§19) establishes such expansions after averaging over small families of microlocalized vectors, of cardinality $O(\mathfrak{h}^{-\varepsilon})$ for any fixed $\varepsilon > 0$. Such averages are harmless for the applications pursued here; we note that the basic technique applies also in p -adic settings, where such averages may often be omitted (see [HNS, §3] for illustration in a basic example). Our expansions give in particular a solution to the test vector problem (modulo taking small averages), which appears to be new even in some low-rank cases (e.g., for discrete series representations in the triple product setting $H = \mathrm{GL}_2(\mathbb{R}) \leq G = \mathrm{GL}_2(\mathbb{R})^2$).

We summarize the proof of the local result just described. The stability hypothesis ensures that, if a group element $s \in H$ is not too close to the identity, then it moves any solution to (1.21) a fair bit (cf. §19.7 for details), so the microlocal support of the pair of vectors $\pi(s)v$ and $\sigma(s)v$ is *disjoint* from that of v and u , and so the matrix coefficient integrand in (1.22) is negligible. We may thus truncate the integral (1.22) to the range where s is small. We evaluate the contribution from this range asymptotically (§19.5), after a small average over v and u , by applying the Kirillov formula for π and σ . The application involves some pleasant book-keeping (§16 and §17) concerning disintegration of volume forms (e.g., of $d\omega$ on \mathcal{O}_π under $\mathcal{O}_\pi \rightarrow \mathfrak{h}^\wedge // H$).

The methods described here are more than adequate to produce families of v which pick out the fairly coarse families \mathcal{F} required by our motivating application; see §22. We hope that they will be useful in many other problems involving asymptotic analysis of special functions arising from representations of higher-rank groups. There are several avenues for extending our methods further. For instance, it seems to us an intriguing problem to obtain analogous asymptotic expansions in non-stable cases.

1.11. Comparison with other work

We briefly discuss the relationship of the ideas of this paper to some others:

(i) The general philosophy that equidistribution results lead to mean value theorems for L -functions was advocated in [Ve], but the source of the equidistribution in this paper is quite different from that in [Ve]. The latter dealt exclusively with translates of a given vector by a large group element, which do not pick out sufficiently flexible families to be

useful for the higher-rank applications pursued here.

(ii) Bernstein and Reznikov initiated the systematic application of identities such as (1.7), with a carefully chosen vector, to the estimation of branching coefficients (on average or individually). Their ideas have deeply influenced our basic strategy (§1.6).

(iii) The orbit method [Ki2] is very well known in the representation theory of Lie groups and widely used in the algebraic side of that theory. It is particularly satisfactory in the nilpotent case [Ki1], [ShTs]. However, it does not seem to have been used much in the way of this paper—that is to say, developed quantitatively along the lines of microlocal analysis and used as an analytical tool. One example where it has been applied in such a context is the asymptotic analysis of Wigner $6j$ symbols; see [Rob]. We hope it will be useful in many other contexts like this.

Although we have focused on applications to the averaged Gan–Gross–Prasad period, we hope that the ideas and methods of this paper, especially those related to the orbit method, should be helpful in a broader class of problems involving analysis of automorphic forms in higher rank.

1.12. Reading suggestions

The reader might wish to peruse Part I and then to skip directly to Part V, referring back to earlier results as needed.

1.13. Acknowledgments

We would like to thank Joseph Bernstein, Philippe Michel, and Andre Reznikov for many inspiring discussions, Emmanuel Kowalski for helpful feedback on an earlier draft, Davide Matasci for some corrections and the anonymous referee for a thorough reading and many helpful suggestions that have improved this paper.

During the completion of this paper, A. V. was an Infosys member at the Institute for Advanced Study. He would like to thank the IAS for providing wonderful working conditions. Over the course of writing the paper he received support from the Packard foundation and the National Science Foundation.

Parts of this work were carried out during visits of P. N. to Stanford in 2013–2016, at a 2014 summer school at the Institut des Hautes Études Scientifiques, while in residence in Spring 2017 at the Mathematical Sciences Research Institute (supported there by the National Science Foundation under Grant No. DMS-1440140), and during a couple short-term visits in Spring 2018 to the Institute for Advanced Study. He thanks all of these institutions for their support and excellent working conditions.

1.14. General notation and conventions

We collect these here together for the reader's convenience. We have attempted to include reminders where appropriate throughout the text.

1.14.1. Asymptotic notation

We write $X=O(Y)$ or $X\ll Y$ or $Y\gg X$ to denote that $|X|\leq c|Y|$ for some quantity c which is “constant” or “fixed”, with precise meanings clarified in each situation. We write $X\asymp Y$ to denote that $X\ll Y\ll X$. “Implied constants” are allowed to depend freely upon any local fields, groups, norms, or bases under consideration.

1.14.2. The asymptotic parameter h

Throughout this paper, the symbol h will denote a parameter in $(0, 1]$. Whenever we use this notation, there will always be an implicit subset $\mathcal{H}\subseteq(0, 1]$ in which the parameter h takes values; this subset will usually (but not always) be a discrete set with 0 as its only limit point, reflecting, e.g., the wavelength scale of a family of Laplace eigenfunctions.

We always require the implied constants in any asymptotic notation to be independent of h . We use notation such as $O(h^\infty)$ to denote a quantity of the form $O(h^N)$ for any fixed N .

By an “ h -dependent element” s of some set S , we mean an element that depends, perhaps implicitly, upon the parameter $h\in\mathcal{H}$, thus $s=s(h)$. We might understand s more formally as an assignment $\mathcal{H}\rightarrow S$, $h\mapsto s(h)$. For instance, an h -dependent positive real $c=c(h)\in\mathbb{R}_+^\times$ is a map $\mathcal{H}\rightarrow\mathbb{R}_+^\times$. The parameter h itself may be understood as an h -dependent positive real, corresponding to the identity map. The terminology applies even if the set itself varies with h , thus $S=S(h)$; an h -dependent element of S is then an element of the Cartesian product $\prod_h S(h)$.

1.14.3. Translating between h -dependent and absolute formulations

The results stated in this paper in terms of h -dependent elements can often be reformulated quantitatively for individual h . Namely, a theorem formulated in terms of an arbitrary h -dependent sequence π_h is formally equivalent to an “absolute” statement valid for all π and all $h\in(0, 1]$. More frequently, we encounter theorems involving an arbitrary h -dependent sequence π_h of representations and an h -dependent sequence f_h of functions which remain bounded with respect to certain norms, where the norms themselves are typically h -dependent. Such theorems may likewise be reformulated as quantitative

bounds valid for all π , f and $h \in (0, 1]$. We illustrate this translation in Remark 6.4. As one sees in this example, the resulting absolute statements are more explicit but less succinct. For this reason, we have usually preferred the h -dependent formulation.

1.14.4. Groups

Let k be a field of characteristic zero. We denote by \bar{k} an algebraic closure of k .

We generally denote by X the set of k -points of a k -variety \mathbf{X} (and vice versa, if \mathbf{X} is clear by context). We identify \mathbf{X} with X when $k = \bar{k}$.

By a *reductive group \mathbf{G} over k* , we will always mean a *connected reductive algebraic k -group \mathbf{G}* . The set G of k -points of \mathbf{G} is then Zariski dense (see [Bo2, Corollary 18.3]).

We will frequently use restriction of scalars to regard reductive groups over \mathbb{C} also as reductive groups over \mathbb{R} .

For an algebraic k -group \mathbf{G} , we denote by G , \mathfrak{g} and \mathfrak{g}^* the respective sets of k -points of \mathbf{G} , of its Lie algebra, and of the k -dual of its Lie algebra. The group G acts on \mathfrak{g} and \mathfrak{g}^* . Recall that an element of one of the latter spaces is *regular* if its orbit has maximal dimension, or equivalently, if its centralizer has minimal dimension. We use a subscripted “reg” to denote subsets of regular elements.

Suppose now that \mathbf{G} is an algebraic group over an archimedean local field F , thus $F = \mathbb{R}$ or $F = \mathbb{C}$. We may then regard G as a *real* Lie group; complex Lie groups do not play a role in this paper.

For a Lie group G , we write \mathfrak{g} for its Lie algebra and \mathfrak{g}^\wedge for the Pontryagin dual of \mathfrak{g} . We identify \mathfrak{g}^\wedge with $i\mathfrak{g}^*$ (see §2.1 for details).

1.14.5. Unitary representations

Let \mathbf{G} be a reductive group over a local field k of characteristic zero. As above, we denote by G the group of k -points of \mathbf{G} . By definition, a unitary representation π of G is a Hilbert space V equipped with a strongly continuous homomorphism π from G to the group of unitary operators on V . In Part I, Part II and Appendix A, we commit the standard abuse of notation by writing π for the Hilbert space V ; we then denote by π^∞ the subspace of smooth vectors (where in the non-archimedean case, “smooth” means “open stabilizer”). In §18, §19, Part IV and Part V, it will be convenient instead to write π for the subspace of smooth vectors in V . It will occasionally be useful also to confuse “unitary” with “unitarizable”. The reader will be reminded locally of these conventions.

1.14.6. Topologies on vector spaces

When given a vector space V defined as the subset of some larger space \tilde{V} on which some $[0, +\infty]$ -valued seminorms take finite values, we define $I(V)$ to be the set of restrictions to V of those seminorms, and equip V with its *evident topology*: that for which an open base at $v_0 \in V$ is given by the sets $\{v \in V : \|v - v_0\|_i < \varepsilon \text{ for all } i \in M\}$, where ε runs over the positive reals and M over the finite subsets of $I(V)$. In all examples we consider, the seminorms will separate points, so V will be a Hausdorff locally convex space.

Example 1.2. This discussion applies to $V = C^\infty(U)$, for U an open subset of \mathbb{R}^n , taking for $I(V)$ the set of pairs (K, D) , where $K \subseteq U$ is compact and D is a differential operator (with smooth coefficients, say) on \mathbb{R}^n , with

$$\|f\|_{(K,D)} := \sup_{x \in K} |Df(x)|.$$

It applies also to the Sobolev spaces π^s , $s \in \mathbb{Z} \cup \{\infty\}$, defined below (§3.2) for a unitary representation π of a Lie group.

Unwinding the definitions, a linear map $T: V \rightarrow W$ between two spaces so equipped is *continuous* if for each $j \in I(W)$ there is a finite subset $M \subseteq I(V)$ and a scalar $C \geq 0$ such that

$$\|Tv\|_j \leq C \sum_{i \in M} \|v\|_i \quad \text{for all } v \in V, \quad (1.23)$$

while a family of such linear maps $T_\alpha: V \rightarrow W$ is *equicontinuous* if, for each j , we may choose M and C uniformly with respect to the family's indexing parameter α . We might also describe the latter situation by saying that T_α is *continuous, uniformly in α* .

We can make an analogous definition *even if the spaces themselves vary*, provided that the indexing sets for their seminorms admit natural identifications. Thus, suppose $T_\alpha: V_\alpha \rightarrow W_\alpha$ is a family of linear maps between spaces as above, and suppose also given identifications $I(V_\alpha) = I(V_\beta)$ and $I(W_\alpha) = I(W_\beta)$ for all α and β . We may then speak as above of T_α being *continuous, uniformly in α* .

Example 1.3. Fix a Lie group G and an element x of its Lie algebra. As π varies over the unitary representations of G , the family of operators $\pi^\infty \rightarrow \pi^\infty$ defined by x is continuous, uniformly in π .

Let V be a vector space arising as the increasing union of a sequence of topological vector subspaces V_m as above, with continuous inclusions.

Example 1.4. $C_c^\infty(U)$, for U as in Example 1.2, can be described in this way, as can the space $\pi^{-\infty} = \bigcup_{s \in \mathbb{Z}} \pi^s$ of distributional vectors defined below (§3.2).

By the *evident topology* on V , we then mean the “locally convex inductive limit” topology, which may be characterized in either of the following ways:

- (i) it is the finest for which the inclusions $V_m \hookrightarrow V$ are continuous;
- (ii) for every locally convex space T , the continuous maps $V \rightarrow T$ are precisely those that restrict to continuous maps $V_m \rightarrow T$ for each m .

Given another space W as above, we say that a family of maps $T_\alpha: V \rightarrow W$ is *continuous, uniformly in α* if for each m , the family of restrictions $T_\alpha: V_m \rightarrow W$ has the property explained above.

We consider several examples (§4.5, §5.4 and §12.1) of vector spaces V consisting of “ h -dependent vectors in a varying family of vector spaces, equipped with the evident topology”. More formally, suppose given:

- h -dependent vector spaces $V(h)$;
- an indexing set \mathcal{I} for h -dependent seminorms on $V(h)$; thus, for each $i \in \mathcal{I}$, we are given an h -dependent seminorm $i(h): V(h) \rightarrow [0, \infty]$.

Let V denote the subspace of the Cartesian product $\prod_h V(h)$ consisting of h -dependent vectors $v=v(h)$ for which the seminorm

$$\|v\|_{V,i} := \sup_h \|v(h)\|_{i(h)} \tag{1.24}$$

is finite for each $i \in \mathcal{I}$. We then topologize V by means of these seminorms.

Still in the situation just described, suppose given an h -dependent positive real $c>0$. We denote by cV the image of $V \subseteq \prod_h V(h)$ under the bijection $v \mapsto cv = c(h)v(h)$, equipped with the seminorms and topology transported from V via this bijection. Thus, an h -dependent vector $v=v(h) \in V(h)$ belongs to cV if every seminorm

$$\|v\|_{cV,i} := \sup_{h \in (0,1]} \|c(h)^{-1}v(h)\|_{i(h)}$$

is finite. For instance, we may speak of the space $h^{10}V$; it is the image of V under multiplication by h^{10} . We denote by $h^\infty V$ the intersection $\bigcap_\eta h^\eta V$, topologized in the evident way. In practice, $h^\infty V$ consists of h -dependent vectors which are “negligible” in the $h \rightarrow 0$ limit.

Example 1.5. Let V be a finite-dimensional real vector space. We topologize the Schwartz space $\mathcal{S}(V)$, as usual, by means of the seminorms $f \mapsto \|Df\|_{L^\infty}$ indexed by the polynomial-coefficient differential operators D on V . Per the above conventions, $h^\eta \mathcal{S}(V)$ is the space of h -dependent Schwartz functions $f=f(h)$ whose seminorms satisfy, for each D as above,

$$\sup_h h^{-\eta} \|Df(h)\|_{L^\infty} < \infty.$$

We apply this notation even when $\eta=0$; elements of $h^0 \mathcal{S}(V)$ are then h -dependent Schwartz functions whose seminorms are bounded uniformly with respect to h .

1.14.7. Miscellaneous

For an element ξ of a normed space, we often write

$$\langle \xi \rangle := (1 + |\xi|^2)^{1/2}.$$

This quantifies the size of ξ , but is never smaller than 1.

When we have equipped some space X (e.g., a group G as above) with a “standard” measure μ (e.g., a Haar measure), we often write $\int_{x \in X} f(x)$ as shorthand for the integral $\int_{x \in X} f(x) d\mu(x)$ taken with respect to the standard measure.

We extend addition to a binary operation $+$ on the extended real line $\mathbb{R} \cup \{\pm\infty\}$, given by

$$\infty + (-\infty) = (-\infty) + \infty := -\infty,$$

and in other cases in the obvious way; this extension is used starting in §4.5.

Part I. Microlocal analysis on Lie group representations I: definitions and basic properties

Let G be a unimodular Lie group over \mathbb{R} , with Lie algebra \mathfrak{g} . Let π be a unitary representation of G . This part (Part I) and its sequel (Part II) will study the basic quantitative properties of an assignment Op from functions on the dual of \mathfrak{g} to operators on π . We will describe the contents of both parts here; we suggest that the reader peruse Part I and consult Part II as needed.

§2–§5 and their sequels §7–§8 concern aspects of this assignment which apply to any unitary representation π , such as $\pi = L^2(G)$.⁽²⁾ Their results may be equivalently formulated in terms of the convolution structure on $L^1(G)$, and more generally on spaces of distributions supported near and singular only at the identity element. In particular, all estimates in these sections are uniform in π .

In §9, we record some preliminaries concerning the relationship between representations of real reductive groups and their infinitesimal characters. These are relevant for us because our main result concerns averages over automorphic representations having infinitesimal character in a prescribed region.

§6 and its sequel §12 establish finer properties of Op for G reductive and π irreducible and tempered. The relevant consequence of these assumptions is the Kirillov-type formula.

⁽²⁾ It may be possible to reduce from the general case to this particular one and then to appeal to the pseudodifferential calculus as in [T, Proposition 1.1], but doing so does not seem to yield an overall simplification, so we have opted instead for a direct and self-contained treatment.

We note that §19 in Part III applies the results of Part I to determine (under certain assumptions) the “asymptotic decomposition” of Op under restriction to certain subgroups.

The contents of Part I may be understood as generalizations of standard results in the theory of pseudodifferential operators (see [Hör], [Bea1] and references), which corresponds roughly to the case in which G is 2-step nilpotent. We note some minor differences:

(i) The exponential map is neither injective nor surjective for the groups G of interest to us, so we work within a fixed small enough neighborhood of the identity element of G .

(ii) The Baker–Campbell–Hausdorff formula is much simpler when G is 2-step nilpotent; for the G of interest, it contains arbitrarily nested commutators. It was not obvious to us at the outset whether this would present an obstruction.

(iii) The appropriate way to generalize the standard operator classes in the theory of pseudodifferential operators was not obvious to us; Definition 3.1, given in §3, took some work to identify.

We note that the \star -product implicit in our operator calculus on the Lie algebra has been previously considered by Rieffel [Ri], who observed that it gives a deformation quantization of the Poisson structure. As Rieffel observes in [Ri, p. 658]: “At the heuristic level, our results also mesh nicely with Kirillov’s orbit method for the theory of group representations, and perhaps the connection can be made stronger”. The results of this paper concerning the operator calculi may be seen as steps in this direction. We mention also the work of Cahen and B. Harris (see for instance [Cah], [HarB]).

Notation

We denote by $\pi^\infty \leq \pi$ the (dense) subspace of smooth vectors, by \mathfrak{U} the universal enveloping algebra of $\mathfrak{g}_\mathbb{C}$, and by $\pi: \mathfrak{U} \rightarrow \text{End}(\pi^\infty)$ the induced map. For $u \in \mathfrak{U}$, we occasionally write simply u instead of $\pi(u)$ when it is clear from context that u is acting on π .

2. Operators attached to Schwartz functions

We define here the basic construct of the microlocal calculus on representations: an assignment from Schwartz functions on the dual of the Lie algebra to operators. See (1.14) and surrounding discussion for motivation. We extend and refine this assignment in later sections.

2.1. Measures, Fourier transforms, etc.

We fix an open neighborhood \mathcal{G} of the origin in the Lie algebra \mathfrak{g} , taken sufficiently small. In particular, $\exp:\mathcal{G}\rightarrow G$ is an analytic isomorphism onto its image.

We choose Haar measures dg on G and dx on \mathfrak{g} satisfying the following compatibility: for $x\in\mathcal{G}$ and $g=\exp(x)$, we have $dg=j(x)dx$, where the analytic function $j:\mathcal{G}\rightarrow\mathbb{R}_{>0}$ satisfies $j(0)=1$.

Since \mathcal{G} is small, the image $j(\mathcal{G})$ is a precompact subset of $\mathbb{R}_{>0}$. We assume also that $-x\in\mathcal{G}$ whenever $x\in\mathcal{G}$.

We use the letters $\{x, y, z\}$ for elements of \mathfrak{g} and $\{\xi, \eta, \zeta\}$ for elements of its imaginary dual $i\mathfrak{g}^*:=\text{Hom}_{\mathbb{R}}(\mathfrak{g}, i\mathbb{R})$. We denote the natural pairing by $x\xi:=\xi x:=\xi(x)\in i\mathbb{R}$.

Recall that $\mathfrak{g}^\wedge=\text{Hom}(\mathfrak{g}, \mathbb{C}^{(1)})$ denotes the Pontryagin dual. We identify $i\mathfrak{g}^*$ with \mathfrak{g}^\wedge via the canonical isomorphism $\xi\mapsto[x\mapsto e^{x\xi}]$; we will find it clearer to work with \mathfrak{g}^\wedge for analytic purposes, $i\mathfrak{g}^*$ for algebraic purposes. In particular, we identify \mathfrak{g} with the space of $i\mathbb{R}$ -valued linear functions on \mathfrak{g}^\wedge .

We write $\mathcal{S}(\dots)$ for the Schwartz space. We equip \mathfrak{g}^\wedge with the Haar measure $d\xi$ for which the Fourier transforms $\mathcal{S}(\mathfrak{g})\ni\phi\mapsto\phi^\wedge\in\mathcal{S}(\mathfrak{g}^\wedge)$ and $\mathcal{S}(\mathfrak{g}^\wedge)\ni a\mapsto a^\vee\in\mathcal{S}(\mathfrak{g})$ defined by

$$\phi^\wedge(\xi):=\int_{x\in\mathfrak{g}}\phi(x)e^{x\xi}dx\quad\text{and}\quad a^\vee(x):=\int_{\xi\in\mathfrak{g}^\wedge}a(\xi)e^{-x\xi}d\xi$$

are mutually inverse.

We let G act on \mathfrak{g} by the adjoint action $g\cdot x:=\text{Ad}(g)x$, on $\mathfrak{g}^\wedge=i\mathfrak{g}^*$ by the coadjoint action $x(g\cdot\xi):=(g^{-1}\cdot x)\xi$, and on functions f on either space by $g\cdot f:=f(g^{-1}\cdot-)$.

2.2. The basic operator map

For a neighborhood \mathcal{G} of the origin in \mathfrak{g} as above, we let $\mathcal{X}(\mathcal{G})$ denote the set of all ‘‘cut-offs’’ $\chi\in C_c^\infty(\mathcal{G})$ with the following properties:

- (i) χ is even: $\chi(x)=\chi(-x)$ for all x ;
- (ii) $\chi(x)\in[0, 1]$ for all x ; in particular, χ is real-valued;
- (iii) $\chi(x)=1$ for all x in some neighborhood of the origin.

For \mathcal{G} as above, $\chi\in\mathcal{X}(\mathcal{G})$ and $a\in\mathcal{S}(\mathfrak{g}^\wedge)$, we define the bounded operator

$$\text{Op}(a, \chi : \pi) \in \text{End}(\pi)$$

by the formula

$$\text{Op}(a, \chi : \pi)v := \int_{x\in\mathfrak{g}} \chi(x)a^\vee(x)\pi(\exp(x))v dx. \quad (2.1)$$

We abbreviate $\text{Op}(a, \chi):=\text{Op}(a, \chi : \pi)$ when π is clear from context. Starting in §5.4, we further abbreviate $\text{Op}(a):=\text{Op}(a, \chi)$.

Recall that h denotes a parameter in $(0, 1]$. For a as above, we define the rescaled function $a_h(\xi) := a(h\xi)$ and the correspondingly rescaled operator

$$\text{Op}_h(a, \chi) := \text{Op}(a_h, \chi) = \int_{x \in \mathfrak{g}} \chi(hx) a^\vee(x) \pi(\exp(hx)) dx. \quad (2.2)$$

As informal motivation for this definition, recall from (1.15) the desiderata that

$$\text{Op}_h(a)v \approx a(\xi)v$$

for $v \in \pi$ microlocalized at $\xi \in \mathfrak{g}^\wedge$, i.e., for which

$$\pi(\exp(hx))v \approx e^{x\xi}v$$

whenever $|x| \ll 1$. Indeed, since $a^\vee(x)$ is small unless $|x| \ll 1$, in which case $\chi(hx) = 1$, we see from (2.2) that

$$\text{Op}_h(a, \chi)v \approx \left(\int_{x \in \mathfrak{g}} a^\vee(x) e^{x\xi} dx \right) v = a(\xi)v.$$

2.3. Adjoints

Our assumptions on χ imply that the adjoint $\text{Op}(a, \chi)^*$ of $\text{Op}(a, \chi)$, regarded as a bounded operator on π , is given by $\text{Op}(\bar{a}, \chi)$.

2.4. Equivariance

Let $a \in \mathcal{S}(\mathfrak{g}^\wedge)$, $g \in G$ and $\chi \in \mathcal{X}(\mathcal{G})$.

Assume that $\text{Ad}(g)\text{supp}(\chi) \subseteq \mathcal{G}$. Then, $g \cdot \chi \in \mathcal{X}(\mathcal{G})$, and we verify readily that

$$\pi(g)\text{Op}(a, \chi)\pi(g)^{-1} = \text{Op}(g \cdot a, g \cdot \chi).$$

2.5. Composition: preliminary discussion

For $f \in C_c^\infty(G)$, we define $\pi(f) \in \text{End}(\pi)$ by

$$\pi(f)v := \int_{g \in G} f(g)\pi(g)v dg.$$

Then,

$$\pi(f_1)\pi(f_2) = \pi(f_1 * f_2), \quad (2.3)$$

where $*$ denotes convolution. This fact unwinds to a composition formula for $\text{Op}(a, \chi)$:

Suppose given a pair of open neighborhoods \mathcal{G} and \mathcal{G}' of the origin as above and cut-offs $\chi \in \mathcal{X}(\mathcal{G})$ and $\chi' \in \mathcal{X}(\mathcal{G}')$ for which the following conditions hold:

- (i) $\exp(\mathcal{G})\exp(\mathcal{G}) \subseteq \exp(\mathcal{G}')$, so that we may define $*$: $\mathcal{G} \times \mathcal{G} \rightarrow \mathcal{G}'$ by

$$x * y := \log(\exp(x)\exp(y)).$$

- (ii) $\chi'(x * y) = 1$ for all $x, y \in \text{supp}(\chi)$.

Let us denote temporarily by $\phi \mapsto f$ the topological isomorphism

$$C_c^\infty(\mathcal{G}) \longrightarrow C_c^\infty(\exp(\mathcal{G}))$$

defined by the rule $f(\exp(x)) := \phi(x)j(x)^{-1}$. Then, $f(g) dg$ is the push-forward of $\phi(x) dx$, and so

$$\pi(f) = \int_{x \in \mathfrak{g}} \phi(x) \pi(\exp(x)) dx.$$

We may define continuous bilinear operators \star on $C_c^\infty(\mathcal{G})$, and then on $\mathcal{S}(\mathfrak{g}^\wedge)$, by requiring that the diagram

$$\begin{array}{ccc} C_c^\infty(\exp(\mathcal{G}))^2 & \xrightarrow{*} & C_c^\infty(\exp(\mathcal{G}')) \\ \cong \downarrow \phi \mapsto f & & f \mapsto \phi \downarrow \cong \\ C_c^\infty(\mathcal{G})^2 & \xrightarrow{*} & C_c^\infty(\mathcal{G}') \\ \uparrow a \mapsto \chi a^\vee & & \phi \mapsto \phi^\wedge \downarrow \\ \mathcal{S}(\mathfrak{g}^\wedge)^2 & \xrightarrow{*} & \mathcal{S}(\mathfrak{g}^\wedge) \end{array}$$

commute. The top arrow is defined thanks to (i). The middle arrow may be described conveniently in terms of the Fourier transform: for $\phi_1, \phi_2 \in C_c^\infty(\mathcal{G})$ and $\zeta \in \mathfrak{g}^\wedge$,

$$\int_{z \in \mathfrak{g}} \phi_1 \star \phi_2(z) e^{z\zeta} dz = \int_{x, y \in \mathfrak{g}} \phi_1(x) \phi_2(y) e^{(x*y)\zeta} dx dy. \quad (2.4)$$

(Indeed, by definition, $\phi_1 \star \phi_2(z) dz$ is obtained from $\phi_1(x) \phi_2(y) dx dy$ by pushing forward to the group, convolving, and then pulling back to the Lie algebra; testing this definition against $z \mapsto e^{z\zeta}$ gives (2.4).) The bottom arrow, which we refer to as the *star product*, is given by $a_1 \star a_2 = (\chi a_1^\vee \star \chi a_2^\vee)^\wedge$. Note that \star depends upon χ .

LEMMA 2.1. For $a_1, a_2 \in \mathcal{S}(\mathfrak{g}^\wedge)$,

$$\text{Op}(a_1, \chi) \text{Op}(a_2, \chi) = \text{Op}(a_1 \star a_2, \chi'). \quad (2.5)$$

Proof. Set $\phi_i := \chi a_i^\vee \in C_c^\infty(\mathcal{G})$, and let $f_i \in C_c^\infty(\exp(\mathcal{G}))$ be as associated above. The identity (2.5) follows from (2.3) upon unwinding the definitions and noting that $\chi' \equiv 1$ on the support of $\chi a_1^\vee \star \chi a_2^\vee$. \square

We also have the rescaled composition formula

$$\text{Op}_h(a, \chi) \text{Op}_h(b, \chi) = \text{Op}(a \star_h b, \chi'),$$

where the rescaled star product $a \star_h b$ is defined by requiring that

$$(a \star_h b)_h = a_h \star b_h.$$

We note that \star is recovered from \star_h by taking $h=1$.

3. Operator classes

The operator map (2.1) has been defined for functions a belonging to the Schwartz space of \mathfrak{g}^\wedge . We want to extend it to other functions, e.g., functions a that are permitted polynomial growth at ∞ . In that case, $\text{Op}(a)$ is no longer a bounded map $\pi \rightarrow \pi$; rather, it behaves more like the (densely defined) operator on π induced by an element of the universal enveloping algebra.

Motivated by this, we define certain classes of densely defined operators on π . These classes will eventually serve as the target of Op , after we extend its definition to various symbol classes.

3.1. The operator Δ and its inverse

Let us fix a basis $\mathcal{B} := \mathcal{B}(\mathfrak{g})$ of \mathfrak{g} , and set

$$\Delta := \Delta_G := 1 - \sum_{x \in \mathcal{B}} x^2 \in \mathfrak{U}.$$

We will often confuse Δ with its image $\pi(\Delta) \in \text{End}(\pi^\infty)$. It has the following properties:

(i) It induces a densely-defined self-adjoint positive operator on π with bounded inverse Δ^{-1} (see [NS]). The operator norm of Δ^{-1} is ≤ 1 .

(ii) For $n \geq 0$, let $D(\Delta^n) \subseteq \pi$ denote the domain of the densely-defined self-adjoint operator extending $\pi(\Delta^n)$. Then, $\pi^\infty = \bigcap_{n \geq 0} D(\Delta^n)$ (see [NE, Corollary 9.3]). Consequently, Δ^{-1} acts on π^∞ .

3.2. Sobolev spaces

For $s \in \mathbb{Z}$, we define an inner product $\langle -, - \rangle_{\pi^s}$ on π^∞ by the rule

$$\langle v_1, v_2 \rangle_{\pi^s} := \langle \Delta^s v_1, v_2 \rangle.$$

We denote by π^s the Hilbert space completion of π^∞ with respect to the associated norm $\|v\|_{\pi^s} := \langle v, v \rangle_{\pi^s}^{1/2}$. These norms increase with s , and so, up to natural identifications,

$$\pi^\infty = \bigcap_s \pi^s \leq \dots \leq \pi^{s+1} \leq \pi^s \leq \pi^{s-1} \leq \dots \leq \pi^{-\infty} := \bigcup_s \pi^s.$$

These spaces come with evident topologies (§1.14.6).

The inner product on π induces a duality between π^s and π^{-s} .

We note that, for each $s \in \mathbb{Z}_{\geq 0}$, there exist $C_s > c_s > 0$, depending upon \mathcal{B} , such that

$$c_s \|v\|_{\pi^s}^2 \leq \sum_{r=0}^s \sum_{x_1, \dots, x_r \in \mathcal{B}} \|x_1 \dots x_r v\|^2 \leq C_s \|v\|_{\pi^s}^2 \quad (3.1)$$

(see, e.g., [NE, proof of Lemma 6.3]).

3.3. Definition of operator classes

By an *operator on π* , we mean simply a linear map $T: \pi^\infty \rightarrow \pi^{-\infty}$. For each $x \in \mathfrak{g}$, the commutator

$$\theta_x(T) := [\pi(x), T]$$

is likewise an operator on π . Indeed, by (3.1), we have for $s \geq 1$ that $\pi(x): \pi^s \rightarrow \pi^{s-1}$. By duality, it follows that $\pi(x): \pi^{1-s} \rightarrow \pi^{-s}$. In particular, $\pi(x)$ acts on both π^∞ and $\pi^{-\infty}$, so both compositions $\pi(x) \circ T$ and $T \circ \pi(x)$ are defined.

The map $x \mapsto \theta_x$ extends to an algebra morphism $\mathfrak{A} \rightarrow \text{End}(\{\text{operators on } \pi\})$, denoted $u \mapsto \theta_u$. For example, for $x_1, \dots, x_n \in \mathfrak{g}$,

$$\theta_{x_1 \dots x_n}(T) = [\pi(x_1), [\pi(x_2), \dots, [\pi(x_n), T]]].$$

Definition 3.1. For $m \in \mathbb{Z}$, we say that an operator T on π has *order* $\leq m$ if, for each $s \in \mathbb{Z}$ and $u \in \mathfrak{A}$, the operator $\theta_u(T)$ induces a bounded map

$$\theta_u(T): \pi^s \longrightarrow \pi^{s-m}.$$

Remark 3.2. When π is a standard representation of a Heisenberg group, this definition is closely related to Beals's characterization [Bea2] of pseudodifferential operators.

We denote by $\Psi^m := \Psi^m(\pi)$ the space of operators on π of order $\leq m$, by $\Psi^{-\infty} := \bigcap_m \Psi^m$ the space of “smoothing operators”, and by $\Psi^\infty := \bigcup_m \Psi^m$ the space of “finite-order operators”. Then,

$$\Psi^{-\infty} \subseteq \dots \subseteq \Psi^{m-1} \subseteq \Psi^m \subseteq \Psi^{m+1} \subseteq \dots \subseteq \Psi^\infty. \quad (3.2)$$

These spaces come with evident topologies (§1.14.6); thus, for $m \in \mathbb{Z}$, the relevant seminorms on Ψ^m are $T \mapsto \|\theta_u(T)\|_{\pi^s \rightarrow \pi^{s-m}}$, taken over $s \in \mathbb{Z}$ and $u \in \mathfrak{U}$. The inclusions (3.2) are continuous.

We note that Ψ^m depends, implicitly, upon the group G which we regard as acting on π .

3.4. Composition

Observe that finite-order operators act on the space of smooth vectors, i.e.,

$$\Psi^\infty \subseteq \text{End}(\pi^\infty).$$

We may thus compose such operators.

LEMMA 3.3. *For $m_1, m_2 \in \mathbb{Z} \cup \{\pm\infty\}$, composition induces continuous maps*

$$\Psi^{m_1} \times \Psi^{m_2} \longrightarrow \Psi^{m_1+m_2},$$

where, as usual, $\infty + (-\infty) := -\infty$.

Proof. For $T_1 \in \Psi^{m_1}$, $T_2 \in \Psi^{m_2}$ and $u \in \mathfrak{U}$, we may write $\theta_u(T_1 T_2)$ as a sum of expressions $\theta_{u_1}(T_1) \theta_{u_2}(T_2)$ with $u_1, u_2 \in \mathfrak{U}$. For $s \in \mathbb{Z}$, the compositions of bounded maps

$$\pi^s \xrightarrow{\theta_{u_2}(T_2)} \pi^{s-m_2} \xrightarrow{\theta_{u_1}(T_1)} \pi^{s-m_1-m_2}$$

are bounded, and so $\theta_u(T_1 T_2)$ induces a bounded map $\pi^s \rightarrow \pi^{s-m_1-m_2}$, as required. \square

3.5. Differential operators

It is easy to see that Ψ^0 contains the identity operator. We record some further examples.

LEMMA 3.4. *Let $m \in \mathbb{Z}$. Then*

- $\pi(\Delta)^m \in \Psi^{2m}$;
- $\pi(x_1 \dots x_m) \in \Psi^m$ if $m \geq 0$ and $x_1, \dots, x_m \in \mathfrak{g}$.

The proof is elementary but somewhat tedious, hence postponed to §8.5.

3.6. Smoothing operators

LEMMA 3.5. *An operator T on π belongs to $\Psi^{-\infty}$ if and only if, for each $N \in \mathbb{Z}_{\geq 0}$, the operator $\Delta^N T \Delta^N$ induces a bounded map $\pi \rightarrow \pi$. The corresponding seminorms*

$$T \mapsto \|\Delta^N T \Delta^N\|_{\pi \rightarrow \pi}$$

describe the topology on $\Psi^{-\infty}$.

Proof. This follows readily from (3.1) and the definitions. \square

LEMMA 3.6. *For $f \in C_c^\infty(G)$, one has $\pi(f) \in \Psi^{-\infty}$, and the induced map*

$$C_c^\infty(G) \longrightarrow \Psi^{-\infty}$$

is continuous.

Proof. Set $f' := \Delta^N * f * \Delta^N$. We note that $\Delta^N \pi(f) \Delta^N = \pi(f')$, $\|\pi(f')\|_{\pi \rightarrow \pi} \leq \|f'\|_{L^1}$, and that the map $f \mapsto f'$ on $C_c^\infty(G)$ is continuous. \square

4. Symbol classes

As promised, we now define various enlargements of the Schwartz space; we will later extend Op to these spaces.

We also study the behavior of the star products \star and \star_h , defined above, on these enlarged spaces. Eventually, under Op, these products will be intertwined with operator composition.

4.1. Multi-index notation

Temporarily, denote by $n := \dim(G)$ the dimension of the underlying Lie group. For convenience, we choose a basis for \mathfrak{g} . This choice defines coordinates $\mathfrak{g} \ni x \mapsto (x_1, \dots, x_n) \in \mathbb{R}^n$ and $\mathfrak{g}^\wedge \ni \xi \mapsto (\xi_1, \dots, \xi_n) \in i\mathbb{R}^n$ for which $x\xi = x_1\xi_1 + \dots + x_n\xi_n$.

For each ‘‘multi-index’’ $\alpha \in \mathbb{Z}_{\geq 0}^n$ we set

$$|\alpha| := \alpha_1 + \dots + \alpha_n, \quad \alpha! := \alpha_1! \dots \alpha_n!, \quad x^\alpha := x_1^{\alpha_1} \dots x_n^{\alpha_n} \in \mathbb{R}, \quad \xi^\alpha := \xi_1^{\alpha_1} \dots \xi_n^{\alpha_n} \in i^{|\alpha|} \mathbb{R}.$$

We define the differential operators ∂^α on $C^\infty(\mathfrak{g}^\wedge)$ and on $C^\infty(\mathfrak{g})$ by requiring that

$$(\partial^\alpha a)^\vee(x) = x^\alpha a^\vee(x) \quad \text{and} \quad (\partial^\alpha \phi)^\wedge(\xi) = (-\xi)^\alpha \phi^\wedge(\xi) \quad (4.1)$$

for $a \in C^\infty(\mathfrak{g}^\wedge)$ and $\phi \in C^\infty(\mathfrak{g})$; the formal Taylor expansions then read

$$a(\zeta + \xi) = \sum_{\alpha} \frac{\xi^\alpha}{\alpha!} \partial^\alpha a(\zeta) \quad \text{and} \quad \phi(z + x) = \sum_{\alpha} \frac{x^\alpha}{\alpha!} \partial^\alpha \phi(z).$$

We fix norms $|\cdot|$ on \mathfrak{g} and \mathfrak{g}^\wedge . Recall that we abbreviate $\langle \xi \rangle := (1 + |\xi|^2)^{1/2}$.

4.2. Formal expansion of the star product

Let $*$, \star and \star_h be as in §2.5. We define the analytic map $\{\cdot, \cdot\}: \mathcal{G} \times \mathcal{G} \rightarrow \mathfrak{g}$ to be the “remainder term” $\{x, y\} := x \star y - x - y$ in the Baker–Campbell–Hausdorff formula. It follows then from (2.4) that, for $a, b \in \mathcal{S}(\mathfrak{g}^\wedge)$,

$$a \star b(\zeta) = \int_{x, y \in \mathfrak{g}} a^\vee(x) b^\vee(y) e^{x\zeta} e^{y\zeta} e^{\{x, y\}\zeta} \chi(x) \chi(y) dx dy. \quad (4.2)$$

The factor $e^{\{x, y\}\zeta} \in \mathbb{C}^{(1)}$ defines an analytic function of $(x, y, \zeta) \in \mathcal{G} \times \mathcal{G} \times \mathfrak{g}^\wedge$, and so admits a power series expansion

$$e^{\{x, y\}\zeta} = \sum_{\alpha, \beta, \gamma} c_{\alpha\beta\gamma} x^\alpha y^\beta \zeta^\gamma. \quad (4.3)$$

The estimate $\{x, y\} = O(|x||y|)$ controls which monomials $x^\alpha y^\beta \zeta^\gamma$ actually appear in this expansion. Using superscript to indicate degree, the terms appearing are 1, then $x^r y^s \zeta^1$, $r, s \geq 1$, then $x^r y^s \zeta^2$, $r, s \geq 2$, and so on. In particular, $x^a y^b \zeta^c$ appears only if $j = a + b - c$ is a non-negative integer, and each j corresponds to finitely many terms.

By grouping the right-hand side of (4.3) in this way, substituting into (4.2) and casually discarding the truncations χ , we arrive at the formal asymptotic expansion

$$a \star b \sim \sum_{j \geq 0} a \star^j b, \quad (4.4)$$

where \star^j is the finite bidifferential operator on $C^\infty(\mathfrak{g}^\wedge)$ defined by

$$a \star^j b(\zeta) := \sum_{\substack{\alpha, \beta, \gamma \\ |\alpha| + |\beta| - |\gamma| = j \\ |\gamma| \leq \min(|\alpha|, |\beta|) \\ \max(|\alpha|, |\beta|) \leq j}} c_{\alpha\beta\gamma} \zeta^\gamma \partial^\alpha a(\zeta) \partial^\beta b(\zeta). \quad (4.5)$$

(We have introduced some redundant summation conditions for emphasis.) The operator \star^j has order $\leq j$ with respect to both variables and is homogeneous of degree j with respect to dilation, i.e.,

$$a_h \star^j b_h = h^j (a \star^j b)_h,$$

thus (4.4) and (4.5) suggest the rescaled formal expansion

$$a \star_h b \sim \sum_{j \geq 0} h^j a \star^j b. \quad (4.6)$$

The leading term is $a \star^0 b = ab$ (pointwise multiplication of functions), while the next term $a \star^1 b$ is a multiple of the Poisson bracket (cf. Gutt [Gu] for explicit formulas for general j).

4.3. Basic symbol classes

Definition 4.1. Let $m \in \mathbb{R}$. We write S^m for the space of smooth functions $a: \mathfrak{g}^\wedge \rightarrow \mathbb{C}$ so that for each multi-index α there exists $C_\alpha \geq 0$ (depending upon a) so that, for all $\xi \in \mathfrak{g}^\wedge$,

$$|\partial^\alpha a(\xi)| \leq C_\alpha \langle \xi \rangle^{m-|\alpha|}, \quad (4.7)$$

where, as usual, $\langle \xi \rangle = (1 + |\xi|^2)^{1/2}$. We extend this definition to $m = -\infty$ by taking intersections and to $m = +\infty$ by taking unions. We refer to elements $a \in S^m$ as *symbols of order* $\leq m$.

We equip S^m with its evident topology (§1.14.6).

Example 4.2. If $m \in \mathbb{Z}_{\geq 0}$, then a polynomial of degree $\leq m$ defines an element of S^m . The space $S^{-\infty}$ coincides with the Schwartz space $\mathcal{S}(\mathfrak{g}^\wedge)$.

For finite m , we may characterize the elements $a \in S^m$ informally as those which oscillate dyadically and are bounded by a multiple of $\langle \xi \rangle^m$. To see why this is the case, take $m=0$ for concreteness. For $R > 1$, the estimate (4.7) says that the rescaled function $\xi \mapsto a(R\xi)$ is smooth, with fixed bounds for all derivatives inside the dyadic annulus $1 < |\xi| < 2$. In fact, it is bounded by C_0 , its derivatives bounded by $\max_{|\alpha|=1} C_\alpha$, its second derivatives bounded by $\max_{|\alpha|=2} C_\alpha$, and so on.

4.4. \hbar -dependent symbol classes

We encourage the reader to skim this section, and all further discussion of \hbar -dependent symbol classes, on a first reading; these classes will be exploited in the *proofs* of our core technical results (Parts II and III), but not in their *applications* (Parts IV and V).

Recall that \hbar denotes a small positive parameter, and that symbols in S^m oscillate on dyadic ranges. We will occasionally need to consider symbols that vary in a controlled manner with \hbar and oscillate on slightly smaller than dyadic ranges.

Recall (§1.14.2) that an \hbar -dependent function $a: \mathfrak{g}^\wedge \rightarrow \mathbb{C}$ is a function which depends, perhaps implicitly, upon \hbar :

$$a(\xi) := a(\xi; \hbar).$$

We still denote, as before, by a_\hbar the rescaled \hbar -dependent function

$$a_\hbar(\xi) := a(\hbar \xi) = a(\hbar \xi; \hbar).$$

Definition 4.3. Let $m \in \mathbb{R}$, $\delta \in [0, 1)$. Let $a: \mathfrak{g}^\wedge \rightarrow \mathbb{C}$ be a smooth \hbar -dependent function. We write

$$a \in S_\delta^m$$

if for each multi-index α there exists $C_\alpha \geq 0$ such that, for all $\xi \in \mathfrak{g}^\wedge$ and $h \in (0, 1]$,

$$|\partial^\alpha a(\xi)| \leq C_\alpha h^{-\delta|\alpha|} \langle \xi \rangle^{m-|\alpha|}. \quad (4.8)$$

For instance, $a \in S_0^m$ means that

- $a(\cdot; h)$ belongs to S^m for each h ;
- the constants C_α defining the membership $a(\cdot; h) \in S^m$ may be taken *independent* of h .

Informally, S_δ^m consists of elements which oscillate at the scale $\xi + O(h^\delta \langle \xi \rangle)$ (see §7.6 for details); it is obtained from S_0^m by “adjoining smoothed characteristic functions of balls of rescaled radius h^δ ”. As explained below, $\delta = \frac{1}{2}$ corresponds to the Planck scale. The most important range for us is when $\delta \in [0, \frac{1}{2})$; taking $\delta = 0$ means we work at dyadic scales, while taking δ close to $\frac{1}{2}$ means we work “just above the Planck scale”.

We write $S_\delta^m(\mathfrak{g}^\wedge)$ when we wish to indicate explicitly which Lie algebra \mathfrak{g} is being considered. We extend the definition to $m = \pm\infty$ as before, and equip these spaces with their evident topologies. We note that, as m and δ increase, the spaces S_δ^m are related by continuous inclusions.

We may combine the notation S_δ^m introduced here with the notation $h^\eta V$ introduced in §1.14.6 to obtain the space

$$h^\eta S_\delta^m,$$

consisting of smooth h -dependent functions $a: \mathfrak{g}^\wedge \rightarrow \mathbb{C}$ satisfying

$$|\partial^\alpha a(\xi)| \leq h^\eta C_\alpha h^{-\delta|\alpha|} \langle \xi \rangle^{m-\alpha}.$$

Example 4.4. Fix $f \in C_c^\infty(\mathbb{R}_{>0})$ (independent of h) and define the h -dependent elements $a, b \in C_c^\infty(\mathfrak{g}^\wedge \setminus \{0\})$ by the formulas

$$a(\xi) := f\left(\frac{|\xi| - 1}{h^\delta}\right) \quad \text{and} \quad b(\xi) := h^\delta f(|\xi|).$$

Then, a belongs to $S_\delta^{-\infty}$ and b belongs to $h^\delta S_0^{-\infty}$, but not in general the other way around.

For a treatment of h -dependent classes in the setting of microlocal analysis on manifolds, we mention [Küs].

4.5. Basic properties

Pointwise multiplication defines continuous maps

$$(- \cdot -): S^{m_1} \times S^{m_2} \longrightarrow S^{m_1+m_2},$$

where, as usual, $\infty + (-\infty) := -\infty$. For any multi-index α , differentiation and monomial-multiplication give continuous maps

$$\begin{aligned}\partial^\alpha: S^m &\longrightarrow S^{m-|\alpha|}, \\ \xi^\alpha: S^m &\longrightarrow S^{m+|\alpha|}.\end{aligned}$$

For the h -dependent classes, we have analogously

$$\begin{aligned}(\cdot \cdot \cdot): S_{\delta_1}^{m_1} \times S_{\delta_2}^{m_2} &\longrightarrow S_{\max(\delta_1, \delta_2)}^{m_1+m_2}, \\ \partial^\alpha: S_\delta^m &\longrightarrow h^{-\delta|\alpha|} S_\delta^{m-|\alpha|}, \\ \xi^\alpha: S_\delta^m &\longrightarrow S_\delta^{m+|\alpha|},\end{aligned}\tag{4.9}$$

with notation as in §4.4.

4.6. Star product asymptotics

Using the properties of \star^j indicated in §4.2, we verify readily that

$$\star^j: S^{m_1} \times S^{m_2} \longrightarrow S^{m_1+m_2-j},\tag{4.10}$$

$$\star^j: S_\delta^{m_1} \times S_\delta^{m_2} \longrightarrow h^{-2\delta j} S_\delta^{m_1+m_2-j}.\tag{4.11}$$

From (4.10), we see that the expansion (4.6) of the star product converges formally with respect to the symbol classes S^m . From (4.11) and its proof, we see that the same holds for S_δ^m if $\delta < \frac{1}{2}$, but not if $\delta > \frac{1}{2}$. Indeed, if the symbols a and b oscillate substantially at scales finer than about $h^{1/2}$, then the summands in the formal expansion $\sum_{j \geq 0} h^j a \star^j b$ do not decay as $h \rightarrow 0$ and $j \rightarrow \infty$. The scale $h^{1/2}$ is thus the natural limit of our calculus, corresponding to the ‘‘Planck scale’’, or in the language of §1.7, to projections onto individual vectors; we discuss the latter point further in Remark 6.6.

These observations may motivate the following result.

THEOREM 4.5. (i) *There is a unique continuous bilinear extension*

$$\star: S^\infty \times S^\infty \longrightarrow S^\infty\tag{4.12}$$

of the star product \star , defined initially on Schwartz spaces as in §2.5. It induces continuous bilinear maps

$$\star: S^{m_1} \times S^{m_2} \longrightarrow S^{m_1+m_2}\tag{4.13}$$

and, for $\delta \in [0, \frac{1}{2})$,

$$\star_h: S_\delta^{m_1} \times S_\delta^{m_2} \longrightarrow S_\delta^{m_1+m_2}.$$

(ii) Fix $J \in \mathbb{Z}_{\geq 0}$. If $a \in S^{m_1}$ and $b \in S^{m_2}$, then

$$a \star b \equiv \sum_{0 \leq j < J} a \star^j b \pmod{S^{m_1+m_2-J}},$$

where the remainder term

$$r := a \star b - \sum_{0 \leq j < J} a \star^j b \in S^{m_1+m_2-J}$$

varies continuously with a and b . In particular,

$$\begin{aligned} a \star_{\hbar} b &\equiv ab \pmod{\hbar S^{m_1+m_2-1}}, \\ a \star_{\hbar} b &\equiv ab + \hbar a \star^1 b \pmod{\hbar^2 S^{m_1+m_2-2}}. \end{aligned}$$

Similarly, for $\delta \in [0, \frac{1}{2})$, $a \in S_{\delta}^{m_1}$ and $b \in S_{\delta}^{m_2}$,

$$a \star_{\hbar} b \equiv \sum_{0 \leq j < J} \hbar^j a \star^j b \pmod{\hbar^{(1-2\delta)J} S_{\delta}^{m_1+m_2-J}},$$

with continuously-varying remainder.

We verify this (in a slightly more general form) in §7 below. The proof is an application of integration by parts and Taylor's theorem to the integral representation (4.2).

Remark 4.6. Our discussion applies with minor modifications to slightly more general symbol classes, e.g., for $m \in \mathbb{R}$, $\rho \in (0, 1]$ and $\delta \in [0, 1)$, to the class $S_{\delta, \rho}^m$ defined by the condition

$$|\partial^{\alpha} a(\xi)| \leq C_{\alpha} \hbar^{-\delta|\alpha|} \langle \xi \rangle^{m-\rho|\alpha|},$$

with the most important range being when $\rho \in (\frac{1}{2}, 1]$ and $\delta \in [0, \frac{1}{2})$. One could also rescale in more general ways than we do here, or work with symbols that are substantially rougher in directions transverse to the foliation of \mathfrak{g}^{\wedge} by coadjoint orbits. We are content here to develop the minimal machinery required for our motivating applications, leaving such extensions to the interested reader.

5. Operators attached to symbols

5.1. Weak definition of the operator map

Let $\chi \in \mathcal{X}(\mathfrak{g})$ (cf. §2.2). By §3.6, the assignment $\text{Op}(\cdot, \chi): S^{-\infty} \rightarrow \Psi^{-\infty}$ is defined and continuous. We calculate that

$$\langle \text{Op}(a, \chi)u, v \rangle = \int_{\xi \in \mathfrak{g}^{\wedge}} a(\xi) \left(\int_{x \in \mathfrak{g}} e^{-x\xi} \chi(x) \langle \pi(\exp(x))u, v \rangle dx \right) d\xi \quad (5.1)$$

for all $u, v \in \pi^\infty$.

We observe that, for any tempered distribution a on \mathfrak{g}^\wedge , the formula (5.1) defines an operator $\text{Op}(a, \chi): \pi^\infty \rightarrow \pi^{-\infty}$, sending smooth vectors to distributional vectors. (To see this, we need only note that the function $\mathfrak{g} \ni x \mapsto \chi(x) \langle \pi(\exp(x))u, v \rangle$ belongs to the Schwartz space.) This observation applies in particular to $a \in S^\infty$. We may similarly extend the definition of the rescaled analogue $\text{Op}_h(a, \chi)$.

5.2. Polynomial symbols

The identification of \mathfrak{g} with the space of $i\mathbb{R}$ -valued linear functions on \mathfrak{g}^\wedge extends to identify $\text{Sym}(\mathfrak{g}_\mathbb{C})$ with the space of polynomial symbols. For instance, if $p = y_1 \dots y_n$ with each $y_i \in \mathfrak{g}$, then $p(\xi) := y_1(\xi) \dots y_n(\xi)$. Which operators arise from such symbols?

LEMMA 5.1. *For $p \in \text{Sym}(\mathfrak{g}_\mathbb{C})$, we have*

$$\text{Op}(p, \chi) = \pi(\text{sym}(p)).$$

Here $\text{sym}: \text{Sym}(\mathfrak{g}_\mathbb{C}) \rightarrow \mathfrak{U}$ denote the symmetrization map, i.e., the linear isomorphism that sends a monomial to the average of its permutations. The proof is given in §8.1.

We assume henceforth when working with Op that the norm $|\cdot|$ is chosen so that $\mathcal{B}(\mathfrak{g})$ is an orthonormal basis. Then, for $p(\xi) := 1 + |\xi|^2 = \langle \xi \rangle^2$, we have $\text{Op}(p) = \Delta$ (cf. §3.1).

5.3. h -dependence

When working with the rescaled operator assignment Op_h , we allow the representation π to be h -dependent, i.e., to vary implicitly with the small parameter $h \in (0, 1]$:

$$\pi = \pi(h).$$

We emphasize that the variation of π with h is arbitrary (without, e.g., continuity or measurability requirements); indeed, in our applications, we consider only those h belonging to some discrete subset of $(0, 1]$.

Definition 5.2. Let π be an h -dependent unitary representation of G . Fix $\delta \in [0, 1)$. We denote by $\Psi_\mathfrak{g}^\delta$ the space of h -dependent operators $T = T(h)$ on $\pi = \pi(h)$ with the property that, for each $u \in \mathfrak{U}$ and $s \in \mathbb{Z}$, there exists $C_{u,s} \geq 0$ (independent of h) so that, for all $h \in (0, 1]$,

$$\|\theta_u^\delta(T)\|_{\pi^s \rightarrow \pi^{s-m}} \leq C_{u,s},$$

where $u \mapsto \theta_u^\delta$ denotes the linear map given, for $u = x_1 \dots x_n$ ($x_1, \dots, x_n \in \mathfrak{g}$), by $\theta_u^\delta := h^{n\delta} \theta_u$, with $\theta_u^0 = \theta_u$ as in §3.3. We extend the definition to $m = \pm\infty$ by taking intersections or unions.

For example, in the (most important) special case $\delta=0$, the space Ψ_0^m consists of \hbar -dependent elements of Ψ^m whose seminorms are uniformly bounded with respect to \hbar . In general, Ψ_δ^m consists of \hbar -dependent elements of Ψ^m whose seminorms vary with \hbar in the indicated manner, depending upon the exponent δ . The results in §3.4 remain valid for Ψ_δ^m (with the same proof), while the results in §3.5 hold for Ψ_0^m , hence for any Ψ_δ^m .

The classes Ψ_δ^m enter into our applications only in a crude way, in the estimation of remainder terms and the proofs of a-priori bounds. For such purposes, it is useful to note that estimates involving Ψ^m yield slightly modified estimates for Ψ_δ^m .

LEMMA 5.3. *Let π and δ be as in the definition above.*

(i) *Fix $m \in \mathbb{Z}$. Let $\nu: \Psi^m \rightarrow \mathbb{R}_{\geq 0}$ be an \hbar -dependent \hbar -uniformly continuous seminorm. Then, for large enough fixed $M \geq 0$ and all $T \in \Psi_\delta^m$, we have $\nu(T) \ll \hbar^{-M}$.*

(ii) *Fix $N \geq 0$. Let $\ell: \Psi^{-N} \rightarrow \mathbb{C}$ be an \hbar -dependent \hbar -uniformly continuous linear map. Then, for some fixed $N' \geq N$, the restriction of ℓ to $\Psi^{-N'}$ satisfies $|\ell(T)| \leq \nu(T)$ for all $T \in \Psi_\delta^{-N'}$ and some continuous seminorm ν on $\Psi_\delta^{-N'}$.*

Proof. (i) By definition, we may find a finite family of pairs (u, s) such that, for each $T \in \Psi^m$, the quantity $\nu(T)$ is bounded by a constant multiple of $\|\theta_u(T)\|_{\pi^s \rightarrow \pi^{s-m}}$ for some such pair. We may assume that $u = x_1 \dots x_n$ with $x_j \in \mathfrak{g}$. The conclusion then holds for any M larger than the maximum value of $n\delta$.

(ii) We may bound $|\ell|$ in terms of finitely many pairs (u, s) as in the proof of (i), with $u = x_1 \dots x_n$. We take for $N' - N$ the maximum value of n and use that each $\pi(x_j) \in \Psi^1$. \square

For an \hbar -dependent positive scalar $c = c(\hbar)$, we denote, as in §1.14.2, by $c\Psi_\delta^m$ the image of Ψ_δ^m under the map $T \mapsto cT = c(\hbar)T(\hbar)$.

We define

$$\hbar^\infty \Psi_\delta^m := \bigcap_{\eta} \hbar^\eta \Psi^m$$

and topologize it as in §1.14.2. This space is independent of δ ; for instance, for finite m , it consists of \hbar -dependent operators T on π that induce continuous maps $T: \pi^s \rightarrow \pi^{s-m}$ having operator norms $O(\hbar^N)$ for all fixed s and N . We often drop the subscript δ and simply write $\hbar^\infty \Psi^m$.

5.4. Variation with respect to the cut-off

Our operator assignment is not particularly sensitive to the choice of cut-off.

LEMMA 5.4. Fix $\chi_1, \chi_2 \in \mathcal{X}(\mathcal{G})$.

(i) For $a \in S^\infty$,

$$\mathrm{Op}(a, \chi_1) \equiv \mathrm{Op}(a, \chi_2) \pmod{\Psi^{-\infty}}, \quad (5.2)$$

with remainder $R := \mathrm{Op}(a, \chi_1) - \mathrm{Op}(a, \chi_2) \in \Psi^{-\infty}$ varying continuously with a .

(ii) More generally, for $a \in S_\delta^\infty$,

$$\mathrm{Op}_h(a, \chi_1) \equiv \mathrm{Op}_h(a, \chi_2) \pmod{h^\infty \Psi^{-\infty}}, \quad (5.3)$$

with continuously-varying remainder.

The latter continuity means explicitly that, for all $M, N \geq 0$, we may write

$$\mathrm{Op}_h(a, \chi_1) - \mathrm{Op}_h(a, \chi_2) = h^N R,$$

with $R \in \Psi^{-M}$, and the induced map $a \mapsto R$ is continuous, uniformly in h .

The proof, given in §8.3, amounts to noting that the Fourier transforms of our symbols are represented away from the origin in \mathfrak{g} by smooth functions of rapid decay. The singularity at the origin is related to the order of the symbol. These observations are the analogue, in our setup, of the fact that kernels of pseudodifferential operators are smooth away from the diagonal.

We henceforth fix χ and χ' as in §2.5, and abbreviate

$$\mathrm{Op}(a) := \mathrm{Op}(a : \pi) := \mathrm{Op}(a, \chi) = \mathrm{Op}(a, \chi : \pi),$$

and similarly for Op_h .

5.5. Equivariance

By combining §2.4 and §5.4, we see that Op is nearly G -equivariant. Indeed, for each fixed group element $g \in G$, we may find a cut-off $\chi_1 \in \mathcal{X}(G)$ such that $g^{-1} \cdot \chi_1 \in \mathcal{X}(G)$; then, with all congruences taken modulo $\Psi^{-\infty}$,

$$\mathrm{Op}(g \cdot a) := \mathrm{Op}(g \cdot a, \chi) \equiv \mathrm{Op}(g \cdot a, \chi_1) = \mathrm{Op}(a, g^{-1} \cdot \chi_1) \equiv \mathrm{Op}(a).$$

This estimate remains valid for g in any fixed compact subset of G modulo the center. A finer assertion holds for the h -dependent classes.

LEMMA 5.5. Fix $\delta \in [0, 1)$ and $\varepsilon > 0$. Let $g \in G$ be an h -dependent group element such that

$$\|\mathrm{Ad}(g)\| \ll h^{-1+\delta+\varepsilon}. \quad (5.4)$$

Then, for $a \in S_\delta^\infty$, the h -dependent symbol $g \cdot a$ satisfies

$$\mathrm{Op}_h(g \cdot a) \equiv \pi(g) \mathrm{Op}_h(a) \pi(g)^{-1} \pmod{h^\infty \Psi^{-\infty}}.$$

The proof is given in §8.4.

5.6. Operator class memberships

The symbol and operator classes have been defined so as to interact nicely.

THEOREM 5.6. *For any $m \in \mathbb{Z} \cup \{\pm\infty\}$, we have*

$$\text{Op}(S^m) \subseteq \Psi^m,$$

and the induced map is continuous. In particular, elements of $\text{Op}(S^\infty)$ act on π^∞ , and so may be composed. Their compositions satisfy

$$\text{Op}(S^{m_1})\text{Op}(S^{m_2}) \subseteq \text{Op}(S^{m_1+m_2}) + \Psi^{-\infty};$$

more precisely,

$$\text{Op}(a)\text{Op}(b) = \text{Op}(a \star b, \chi') \equiv \text{Op}(a \star b) \pmod{\Psi^{-\infty}},$$

with continuously-varying remainder.

The proof is given in §8.7. By combining with the asymptotic expansion of the star product (Theorem 4.5), we obtain the following result.

COROLLARY 5.7. *Fix $m_1, m_2 < \infty$ and $J \in \mathbb{Z}_{\geq 0}$. Then, for $a \in S^{m_1}$ and $b \in S^{m_2}$, we have*

$$\text{Op}(a)\text{Op}(b) \equiv \sum_{0 \leq j < J} \text{Op}(a \star^j b) \pmod{\Psi^{m_1+m_2-J}},$$

with continuously-varying remainder.

We also have the following rescaled analogues.

THEOREM 5.8. *Fix $\delta \in [0, \frac{1}{2})$. For $m \in \mathbb{Z}$, we have*

$$\text{Op}_h(S_\delta^m) \subseteq h^{\min(0, m)} \Psi_\delta^m.$$

For $a, b \in S_\delta^\infty$, we have

$$\text{Op}_h(a)\text{Op}_h(b) = \text{Op}_h(a \star_h b, \chi') \equiv \text{Op}_h(a \star_h b) \pmod{h^\infty \Psi^{-\infty}}.$$

The proof is given in §8.8. The following consequence will be very useful.

COROLLARY 5.9. *Let $m_1, m_2 < \infty$ and $\delta \in [0, \frac{1}{2})$. For $M, N \in \mathbb{Z}_{\geq 0}$ there exists $J \in \mathbb{Z}_{\geq 0}$, depending only upon (m_1, m_2, M, N, δ) , such that, for all $a \in S_\delta^{m_1}$ and $b \in S_\delta^{m_2}$,*

$$\text{Op}_h(a)\text{Op}_h(b) \equiv \sum_{0 \leq j < J} h^j \text{Op}_h(a \star^j b) \pmod{h^N \Psi_\delta^{-M}}. \quad (5.5)$$

Proof. By Theorems 4.5 and 5.8, the two sides of (5.5) agree modulo

$$\mathcal{E} := h^\infty \Psi^\infty + \text{Op}_h(h^{(1-2\delta)J} S_\delta^{m_1+m_2-J}).$$

We choose $M', N' \in \mathbb{Z}_{\geq 0}$, with $m_1+m_2-M' \leq -M$, $(1-2\delta)M' \geq M$ and $(1-2\delta)N' \geq N$. We take $J := M' + N'$. Then, $(1-2\delta)J \geq M+N$ and $m_1+m_2-J \leq -M-N' \leq -M$, so

$$h^{(1-2\delta)J} S_\delta^{m_1+m_2-J} \subseteq h^{M+N} S_\delta^{-M}.$$

By another application of Theorem 5.8, we conclude that $\mathcal{E} \subseteq h^N \Psi_\delta^{-M}$. \square

A very useful generalization, involving proper subgroups G_1 and G_2 of G , will be given in §8.9.

6. The Kirillov formula

Our discussion thus far has been quite general: π was any unitary representation of a unimodular Lie group G . Conversely, the only control we have established over the operators we have constructed on π is through their operator norms. We now consider a more restrictive situation and derive correspondingly stronger control.

Let \mathbf{G} be a reductive algebraic group over an archimedean local field F . (Our discussion applies somewhat more broadly, e.g., to nilpotent groups, but we focus on the case relevant for our applications.) By restriction of scalars, we may suppose $F = \mathbb{R}$. Per our general conventions, G denotes the group of real points of \mathbf{G} .

Let π be a *tempered irreducible* unitary representation of G . We can then apply the character formula for π (§6.2), expressed in Kirillov form in terms of coadjoint orbits (§6.1), to study the traces and trace norms of the operators constructed via our calculus.

6.1. Coadjoint orbits

The survey article [Ki2] is a useful reference for the following discussion. A *coadjoint orbit* \mathcal{O} is an orbit of G on $\mathfrak{g}^\wedge \cong i\mathfrak{g}^*$. Being an orbit of a Lie group, it is a smooth manifold.

Each such orbit carries moreover a canonical G -invariant symplectic structure

$$\sigma := \sigma_{\mathcal{O}},$$

given at each $\xi \in \mathcal{O}$ by the alternating form

$$\sigma_\xi(\text{ad}_x^* \xi, \text{ad}_y^* \xi) := \frac{[x, y] \xi}{i}$$

on the tangent space $T_\xi(\mathcal{O}) = \{\text{ad}_x^* \xi : x \in \mathfrak{g}\}$.

In particular, \mathcal{O} is even-dimensional. We denote by

$$d := d(\mathcal{O}) := \frac{1}{2} \dim_{\mathbb{R}}(\mathcal{O})$$

half the real dimension of \mathcal{O} ; it is an integer with $2d \leq \dim \mathfrak{g}$.

The d -fold wedge product σ^d defines a volume form on \mathcal{O} . We refer to the measure induced by the volume form

$$\omega := \omega_{\mathcal{O}} := \frac{1}{d!} \left(\frac{\sigma}{2\pi} \right)^d$$

as the *normalized symplectic measure* on \mathcal{O} .

If we choose local coordinates x_i and ξ_j on \mathcal{O} , $1 \leq i, j \leq d$, with respect to which

$$\sigma = \sum_j dx_j \wedge d\xi_j,$$

then

$$\omega = \prod_j \left(dx_j \wedge \frac{d\xi_j}{2\pi} \right).$$

Integration with respect to ω defines a measure on \mathfrak{g}^\wedge [Rao]. One verifies readily that these measures enjoy the homogeneity property: if $t \in \mathbb{R}_+^\times$ and $f \in C_c^\infty(\mathfrak{g})$, then

$$\int_{x \in \mathcal{O}} f(x) d\omega_{\mathcal{O}}(x) = t^{-d(\mathcal{O})} \int_{x \in t\mathcal{O}} f(t^{-1}x) d\omega_{t\mathcal{O}}(x). \quad (6.1)$$

A coadjoint orbit is called *regular* if it consists of regular elements, or equivalently, has maximal dimension among all coadjoint orbits.

To each coadjoint orbit \mathcal{O} we may assign an *infinitesimal character* $[\mathcal{O}]$ in the geometric invariant theory (GIT) quotient $[\mathfrak{g}^\wedge]$ of \mathfrak{g}^\wedge by G ; we recall the details below in §9 and §12.

By a *coadjoint multiorbit*, we will mean a finite union of coadjoint orbits sharing the same infinitesimal character. A *regular coadjoint multiorbit* is a coadjoint multiorbit whose elements are regular. The notation and terminology introduced above carries over with minor modifications; for instance, given a non-empty coadjoint multiorbit, we may speak of its infinitesimal character or its normalized symplectic measure. We note that the number of coadjoint orbits having a given infinitesimal character is bounded by a constant depending only upon G (see [Kos2, Theorem 3 and Remark 16] and [Wh, §3]), so a coadjoint multiorbit consists of a uniformly bounded number of orbits.

6.2. The Kirillov character formula

Harish–Chandra showed (see, e.g., [Kn, §X]) that there is a locally L^1 -function $\chi_\pi: G \rightarrow \mathbb{C}$, the *character* of π , with the following property. Having fixed a Haar measure dg on G , we may associate with each $f \in C_c^\infty(G)$ a smooth compactly-supported measure $f(g) dg$ on G and an operator

$$\pi(f) = \int_{g \in G} \pi(g) f(g) dg.$$

Then, $\pi(f)$ is trace class, with trace given by

$$\chi_\pi(f) = \int_{g \in G} \chi_\pi(g) f(g) dg.$$

The facts recalled thus far apply to any irreducible admissible representation. Recall now that π is assumed tempered. A fundamental theorem of Rossmann [Ros1], [Ros2], [Ros4] gives the validity of the Kirillov formula for π , i.e., gives an exact formula for the “Fourier transform” of χ_π in a neighborhood of the identity element. Recall from §2.1 the definition of j .

THEOREM 6.1. *There is a unique non-empty regular coadjoint multiorbit $\mathcal{O}_\pi \subseteq \mathfrak{g}^\wedge$ such that, for all x in some fixed neighborhood of the origin in \mathfrak{g} , we have the identity of distributions*

$$\begin{aligned} \chi_\pi(e^x) \sqrt{j(x)} &= \text{Fourier transform of normalized symplectic measure on } \mathcal{O}_\pi \\ &:= \int_{\xi \in \mathcal{O}_\pi} e^{x\xi} d\omega_{\mathcal{O}_\pi}(\xi). \end{aligned} \quad (6.2)$$

Moreover, the following statements hold:

- (i) *The infinitesimal character of \mathcal{O}_π is the infinitesimal character of π (cf. §9).*
- (ii) *If the infinitesimal character of π is regular, then \mathcal{O}_π is a coadjoint orbit.*

We emphasize the following:

- The statement of (6.2) is independent of choices of Haar measure.
- In general, G may have several orbits with the same infinitesimal character as π .

It is remarkable that only a single one contributes in the case of regular infinitesimal character. (We do not directly use this fact in this paper.)

- We have defined \mathcal{O}_π only for tempered π .

We henceforth denote by d the maximal dimension of any coadjoint orbit, so that for each π as above, every orbit in \mathcal{O}_π is $2d$ -dimensional.

Remark 6.2. In the p -adic case, the Harish–Chandra/Howe local character expansion gives a result of similar spirit but less precise. It describes the Fourier transform of the character, but only in a neighborhood of the identity that depends on the representation π . As a result, it detects only the geometry of the coadjoint orbit “at infinity”.

6.3. Trace estimates

The Kirillov formula implies that

$$\mathrm{tr}(\mathrm{Op}(a)) = \int_{\mathcal{O}_\pi} (j^{-1/2} \chi a^\vee)^\wedge d\omega_{\mathcal{O}_\pi} \quad (6.3)$$

for all $a \in S^{-\infty}(\mathfrak{g}^\wedge)$. By simple estimates (see §8.2 and Appendix A.3), this conclusion extends continuously to $a \in S^{-N}$ for large enough $N \in \mathbb{Z}_{\geq 0}$. By appeal to the homogeneity property (6.1), we see more generally that

$$\mathrm{tr}(\mathrm{Op}_\mathfrak{h}(a)) = \mathfrak{h}^{-d} \int_{\mathfrak{h}\mathcal{O}_\pi} (j_\mathfrak{h}^{-1/2} \chi_\mathfrak{h} a^\vee)^\wedge d\omega_{\mathfrak{h}\mathcal{O}_\pi}, \quad (6.4)$$

where $j_\mathfrak{h}(x) := j(\mathfrak{h}x)$ and $\chi_\mathfrak{h}(x) := \chi(\mathfrak{h}x)$. Using these formulas, we will establish in §12.3 some refined and generalized forms of the following result.

THEOREM 6.3. *Fix N sufficiently large in terms of the reductive Lie group G .*

(i) *Let $a \in S^{-N}$, let π be a tempered irreducible unitary representation of G , and let $\mathfrak{h} \in (0, 1]$. Then,*

$$\mathrm{tr}(\mathrm{Op}_\mathfrak{h}(a)) = \mathfrak{h}^{-d} \left(\int_{\mathfrak{h}\mathcal{O}_\pi} a d\omega_{\mathfrak{h}\mathcal{O}_\pi} + O(\mathfrak{h}) \right). \quad (6.5)$$

The implied constant is independent of π and \mathfrak{h} , and may be taken to depend continuously upon a .

(ii) *Let π be a tempered irreducible unitary representation of G . Any*

$$T \in \Psi^{-N} := \Psi^{-N}(\pi)$$

(see §3.3) defines a trace class operator on π . The trace norm depends continuously upon T , uniformly in π (in the sense of §1.14.6).

(iii) *Let $a \in S_\delta^{-N}$ for some fixed $\delta \in [0, \frac{1}{2})$. Let π be an \mathfrak{h} -dependent tempered irreducible unitary representation of G . Then, the trace norms of $\mathfrak{h}^d \mathrm{Op}_\mathfrak{h}(a)$ are bounded, uniformly in π and \mathfrak{h} , and continuously with respect to a .*

We note that part (i) follows readily from (6.3), a Taylor expansion of $j_\mathfrak{h}$ and $\chi_\mathfrak{h}$, and some a-priori bounds for integrals over coadjoint orbits, while part (ii) follows from the uniform trace class property of Δ^{-N} . Part (iii) is established using the symbol/operator calculi.

Remark 6.4. To illustrate the content of our “ \mathfrak{h} -dependent” notation, we record an equivalent formulation of part (iii) of Theorem 6.3. Let $N \in \mathbb{Z}_{\geq 0}$ be large enough in terms of G . Let $\delta \in [0, \frac{1}{2})$. There exist $C \geq 0$ and $J \in \mathbb{Z}_{\geq 0}$ such that, for each tempered irreducible

unitary representation π of G , each symbol $a \in S^{-N}$ and each scaling parameter $h \in (0, 1]$, the trace norm of the operator $h^d \text{Op}_h(a)$ on π is bounded by $C \sum_{|\alpha| \leq J} \nu_{\alpha, h}(a)$, where

$$\nu_{\alpha, h}(a) := \sup_{\xi \in \mathfrak{g}^\wedge} \frac{|\partial^\alpha a(\xi)|}{h^{-\delta|\alpha|} \langle \xi \rangle^{-N-|\alpha|}}$$

denotes the infimum of all scalars $C_\alpha \geq 0$ for which the specialization to $m := -N$ of the inequality (4.8) holds for all $\xi \in \mathfrak{g}^\wedge$.

Remark 6.5. In our applications of Theorem 6.3, it is important that π and h may vary simultaneously, but our calculus is interesting only if the rescaled orbit $h \mathcal{O}_\pi$ does not “escape to ∞ ” as $h \rightarrow 0$. In many examples of interest, it happens that $(h \mathcal{O}_\pi, d\omega_{h \mathcal{O}_\pi})$ converges to some “limit orbit” $(\mathcal{O}, d\omega)$ (§11.4); studying the limiting behavior of the calculus will be a major concern of the later parts of the paper (§26, §27). A special case relevant (but not sufficient) for our aims is when π is *independent* of h and generic; the limit orbit \mathcal{O} is then contained in the regular subset of the nilcone (§11.4.2).

Remark 6.6. We have noted already (in §4.6) that the $h^{1/2}$ scale is a natural limit to our calculus. Using (6.5), this may now be understood as follows: If the symbol a is a smooth approximation to the characteristic of a ball, with origin some regular element $\xi \in h \mathcal{O}_\pi$ and with radius $C h^{1/2}$, then §2.3 and (5.5) suggest that $\text{Op}_h(a)$ should approximate a self-adjoint idempotent, i.e., an orthogonal projector onto a subspace V of π , consisting of vectors “microlocalized within $C h^{-1/2}$ of $h^{-1} \xi$ ” (cf. §1.7). But (6.5) suggests that $\dim(V) \asymp C^{2d}$, which makes sense only if C is not too small. Conversely, our calculus allows one to work with such symbols provided that $C \gg h^{-\varepsilon}$ for fixed but arbitrarily small $\varepsilon > 0$, hence to construct and manipulate approximate projectors onto subspaces of size $h^{-o(1)}$; in other words, to work with an approximate orthonormal basis of π consisting of vectors microlocalized at elements $\xi \in h \mathcal{O}_\pi$.

Part II. Microlocal analysis on Lie group representations II: proofs and refinements

We now give proofs of results in Part I, as well as certain refinements that will be useful at localized points of the later treatment. The reader might wish to skim or skip Part II on a first reading.

7. Star product asymptotics

The aim of this section is to prove Theorem 4.5 in a generalized form (Theorem 7.4) that will be very convenient in applications.

7.1. Setup

Let \mathfrak{g}_1 and \mathfrak{g}_2 be subalgebras of \mathfrak{g} . We assume that they arise as the Lie algebras of some unimodular Lie subgroups G_1 and G_2 of G . The most important example is when $\mathfrak{g}_1 = \mathfrak{g}_2 = \mathfrak{g}$.

We fix some sufficiently small even precompact open neighborhoods $\mathcal{G} \subset \mathfrak{g}$, $\mathcal{G}_1 \subset \mathfrak{g}_1$ and $\mathcal{G}_2 \subset \mathfrak{g}_2$ of the respective origins, with \mathcal{G}_1 and \mathcal{G}_2 small enough in terms of \mathcal{G} . In particular, the maps $\star, \{\cdot, \cdot\}: \mathcal{G}_1 \times \mathcal{G}_2 \rightarrow \mathcal{G}$ given (as in §2.5 and §4.2) by

$$x \star y := \log(\exp(x)\exp(y)) \quad \text{and} \quad \{x, y\} := x \star y - x - y$$

are defined and analytic.

We assume that $\mathfrak{g}_1 + \mathfrak{g}_2 = \mathfrak{g}$, or equivalently, that the multiplication map $G_1 \times G_2 \rightarrow G$ is submersive near the identity element. The restriction map

$$\mathfrak{g}^\wedge \ni \zeta \mapsto (\zeta_1, \zeta_2) \in \mathfrak{g}_1^\wedge \times \mathfrak{g}_2^\wedge$$

is then injective.

Convolution defines a continuous map

$$C_c^\infty(\exp(\mathcal{G}_1)) \times C_c^\infty(\exp(\mathcal{G}_2)) \longrightarrow C_c^\infty(\exp(\mathcal{G})).$$

By fixing cut-offs $\chi_j \in \mathcal{X}(\mathcal{G}_j)$ as in §2.2, and arguing as in §2.5, we may define a “star product”

$$\star: \mathcal{S}(\mathfrak{g}_1^\wedge) \times \mathcal{S}(\mathfrak{g}_2^\wedge) \longrightarrow \mathcal{S}(\mathfrak{g}^\wedge)$$

by the formula $a \star b := (a^\vee \chi_1 \star b^\vee \chi_2)^\wedge$; it admits the integral representation

$$a \star b(\zeta) := \int_{x,y} a^\vee(x) b^\vee(y) e^{x\zeta_1} e^{y\zeta_2} e^{\{x,y\}\zeta} \chi_1(x) \chi_2(y).$$

Here, the integral is over $(x, y) \in \mathfrak{g}_1 \times \mathfrak{g}_2$ with respect to some fixed Haar measures.

For a unitary representation π of G , we may define $\text{Op}(a, \chi_1)$ and $\text{Op}(b, \chi_2)$ as in §2, acting on π via its restrictions to G_1 and G_2 . We then have

$$\text{Op}(a, \chi_1) \text{Op}(b, \chi_2) = \text{Op}(a \star b, \chi')$$

for any cut-off χ' on \mathfrak{g} taking the value 1 on \mathcal{G} .

We wish to extend the domain of definition of \star and understand its asymptotic behavior. As in §4.2, we may expand into homogeneous components, namely

$$\Omega(x, y, \zeta) := e^{\{x,y\}\zeta} = \sum_{\substack{\alpha, \beta, \gamma \\ |\gamma| \leq \min(|\alpha|, |\beta|)}} c_{\alpha\beta\gamma} x^\alpha y^\beta \zeta^\gamma = \sum_{j \geq 0} \Omega_j(x, y, \zeta), \quad (7.1)$$

where Ω_j denotes the contribution from $|\alpha|+|\beta|-|\gamma|=j$. This again suggests the formal asymptotic expansion $a \star b \sim \sum_{j \geq 0} a \star^j b$, where

$$a \star^j b(\zeta) := \int_{x,y} a^\vee(x) b^\vee(y) e^{x\zeta_1} e^{y\zeta_2} \Omega_j(x, y, \zeta) = \sum_{\substack{\alpha, \beta, \gamma \\ |\gamma| \leq \min(|\alpha|, |\beta|) \\ |\alpha|+|\beta|-|\gamma|=j}} c_{\alpha\beta\gamma} \zeta^\gamma \partial^\alpha a(\zeta) \partial^\beta b(\zeta).$$

7.2. When should the star product map symbols to symbols?

Let $m_1, m_2 \in \mathbb{Z} \cup \{\pm\}$. For $j=1, 2$, we introduce the temporary abbreviation

$$S^{m_j} := S^{m_j}(\mathfrak{g}_j^\wedge),$$

and similarly

$$S_\delta^{m_j} := S_\delta^{m_j}(\mathfrak{g}_j^\wedge).$$

When does \star extend naturally to a continuous map with domain $S^{m_1} \times S^{m_2}$ and codomain one of our symbol classes? We might ask first whether the finite-order differential operators \star^j admit such an extension, starting with the simplest case $j=0$, for which

$$a \star^0 b(\zeta) = a(\zeta_1) b(\zeta_2).$$

Example 7.1. Suppose $\mathfrak{g}_2 \neq \mathfrak{g}$. Take for $a \in S^0(\mathfrak{g}_1^\wedge)$ a constant symbol and for $b \in S^{-\infty}(\mathfrak{g}_2^\wedge)$ a compactly-supported symbol, both non-zero. One can verify then that $a \star^0 b \in C^\infty(\mathfrak{g}^\wedge)$ does not belong to $S^\infty(\mathfrak{g}^\wedge)$.

Definition 7.2. We say that the pair (m_1, m_2) is *admissible* (relative to $\mathfrak{g}_1, \mathfrak{g}_2 \subseteq \mathfrak{g}$) if both of the following implications are satisfied:

- if $\mathfrak{g}_1 \neq \mathfrak{g}$, then $m_2 = -\infty$;
- if $\mathfrak{g}_2 \neq \mathfrak{g}$, then $m_1 = -\infty$.

For instance, if $\mathfrak{g}_1 = \mathfrak{g}_2 = \mathfrak{g}$, then any pair is admissible, while if $\mathfrak{g}_1 \neq \mathfrak{g}$ and $\mathfrak{g}_2 \neq \mathfrak{g}$, then $(-\infty, -\infty)$ is the only admissible pair. One verifies readily that, if (m_1, m_2) is admissible, then

$$\star^j: S^{m_1} \times S^{m_2} \longrightarrow S^{m_1+m_2-j}(\mathfrak{g}^\wedge),$$

and more generally

$$\star^j: S_\delta^{m_1} \times S_\delta^{m_2} \longrightarrow \mathfrak{h}^{-2\delta j} S_\delta^{m_1+m_2-j}(\mathfrak{g}^\wedge),$$

are defined and continuous. In fact, the converse holds as well.

LEMMA 7.3. *For (m_1, m_2) as above, the following are equivalent:*

- (i) $a \star^0 b \in S^\infty(\mathfrak{g})$ for all $(a, b) \in S^{m_1} \times S^{m_2}$;
- (ii) $a \star^j b \in S^\infty(\mathfrak{g})$ for all $(a, b) \in S^{m_1} \times S^{m_2}$ and $j \in \mathbb{Z}_{\geq 0}$;
- (iii) the pair (m_1, m_2) is admissible.

We have included this lemma for motivational purposes only; the proof is left to the reader.

7.3. Main result

THEOREM 7.4. *The star product \star extends uniquely to a compatible family of continuous maps*

$$\star: S^{m_1}(\mathfrak{g}_1^\wedge) \times S^{m_2}(\mathfrak{g}_2^\wedge) \longrightarrow S^{m_1+m_2}(\mathfrak{g}^\wedge),$$

taken over the admissible pairs (m_1, m_2) .

Fix $\delta_1, \delta_2 \in [0, 1)$ with $\delta_1 + \delta_2 < 1$, $J \in \mathbb{Z}_{\geq 0}$ and an admissible pair (m_1, m_2) . Then, for all $a \in S_{\delta_1}^{m_1}(\mathfrak{g}_1^\wedge)$ and $b \in S_{\delta_2}^{m_2}(\mathfrak{g}_2^\wedge)$, we have the asymptotic expansion

$$a \star_{\mathfrak{h}} b \equiv \sum_{0 \leq j < J} \mathfrak{h}^j a \star^j b \pmod{\mathfrak{h}^{(1-\delta_1-\delta_2)J} S_{\max(\delta_1, \delta_2)}^{m_1+m_2-J}(\mathfrak{g}^\wedge)}, \quad (7.2)$$

with continuously-varying remainder.

The proof is given in §7.7, after some preliminaries.

Remark 7.5. In all applications of Theorem 7.4, except that in §19.5, we take

$$\delta_1 = \delta_2 \in [0, \tfrac{1}{2}).$$

7.4. Taylor's theorem

LEMMA 7.6. *Fix $J \in \mathbb{Z}_{\geq 0}$ and a multi-index $\delta \in \mathbb{Z}_{\geq 0}^{\dim(\mathfrak{g})}$. For $(x, y, \zeta) \in \mathcal{G}_1 \times \mathcal{G}_2 \times \mathfrak{g}^\wedge$, abbreviate*

$$\rho := \max(|x|, |y|, |x| |y| |\zeta|).$$

Then,

$$\partial_\zeta^\delta \Omega(x, y, \zeta) - \sum_{0 \leq j < J} \partial_\zeta^\delta \Omega_j(x, y, \zeta) \ll \rho^J, \quad (7.3)$$

where the implied constant may depend upon (J, δ) but not upon (x, y, ζ) .

Proof. Recall (7.1). We note first that, if $|\gamma| \leq \min(|\alpha|, |\beta|)$ and $|\alpha| + |\beta| - |\gamma| = j$, then

$$|x^\alpha y^\beta \zeta^\gamma| \leq |x|^{|\alpha|} |y|^{|\beta|} |\zeta|^{|\gamma|} \leq \rho^j,$$

as follows readily by induction on $|\gamma|$. Thus,

$$\partial_\zeta^\delta x^\alpha y^\beta \zeta^\gamma \ll (|x| |y|)^{|\delta|} \rho^j \ll \rho^j. \quad (7.4)$$

We observe next, by the analyticity of Ω , that there is a constant $R > 0$ so that

$$|c_{\alpha\beta\gamma}| \ll R^{|\alpha|+|\beta|-|\gamma|}.$$

From this and (7.4), we obtain

$$\partial_\zeta^\delta \Omega_j(x, y, \zeta) \ll (1+j)^{O(1)} (R\rho)^j \quad (7.5)$$

for all $j \geq 0$. By (7.5) (applied with $j < J$) and the trivial estimate

$$\partial_\zeta^\delta \Omega(x, y, \zeta) \ll (|x| |y|)^{|\delta|} \ll 1,$$

we deduce the claim (7.3) in the special case that ρ is bounded uniformly from below, say by $\frac{1}{2}R$. In the remaining case $\rho \leq \frac{1}{2}R$, we deduce (7.3) by summing (7.5) over $j \geq J$. \square

7.5. Integration by parts

For $(\xi, \eta, \zeta) \in \mathfrak{g}_1^\wedge \times \mathfrak{g}_2^\wedge \times \mathfrak{g}^\wedge$, we set

$$F(\xi, \eta, \zeta) := \int_{x,y} e^{x\xi + y\eta} \chi_1(x) \chi_2(y) \Omega(x, y, \zeta). \quad (7.6)$$

Recall that \mathcal{G}_1 and \mathcal{G}_2 have been taken sufficiently small, and that $\mathfrak{g} = \mathfrak{g}_1 + \mathfrak{g}_2$. It will be convenient now to normalize the norms $|\cdot|$ on the dual spaces \mathfrak{g}^\wedge , \mathfrak{g}_1^\wedge and \mathfrak{g}_2^\wedge to be euclidean norms with the property that, for $\zeta \in \mathfrak{g}^\wedge$,

$$|\zeta|^2 = |\zeta_1|^2 + |\zeta_2|^2. \quad (7.7)$$

LEMMA 7.7. *Fix $N \in \mathbb{Z}_{\geq 0}$ and a multi-index γ . Set $t := \sqrt{|\xi|^2 + |\eta|^2}$. Then,*

$$t \geq \frac{1}{2} |\zeta| \implies \partial_\zeta^\gamma F(\xi, \eta, \zeta) \ll t^{-N}. \quad (7.8)$$

The implied constant may depend upon (N, γ) , but not upon (ξ, η, ζ) .

The conclusion holds with $\frac{1}{2}$ replaced by any fixed fraction, provided that \mathcal{G}_1 and \mathcal{G}_2 are taken sufficiently small. The basic idea is that the hypothesis $t \geq \frac{1}{2}|\zeta|$ implies that the integral (7.6) has no stationary point.

Proof. We may write

$$\partial_\zeta^\gamma F(\xi, \eta, \zeta) = \int_{x,y} f(x, y) e^{t\phi(x,y)},$$

where $f(x, y) := \{x, y\}^\gamma \chi_1(x) \chi_2(y)$ and $\phi: \mathcal{G}_1 \times \mathcal{G}_2 \rightarrow i\mathbb{R}$ is given by

$$\phi(x, y) := \frac{x\xi + y\eta}{t} + \{x, y\} \frac{\zeta}{t}.$$

Since \mathcal{G}_1 and \mathcal{G}_2 are small and $|\zeta|/t \leq 2$, the total derivative

$$\partial\phi: \mathcal{G}_1 \times \mathcal{G}_2 \longrightarrow i\mathfrak{g}_1^* \times i\mathfrak{g}_2^*$$

approximates the unit vector $(\xi/t, \eta/t)$. In particular, the euclidean norm $|\partial\phi|(x, y)$ of the total derivative is bounded from below by (say) $\frac{1}{2}$ for all $(x, y) \in \mathcal{G}_1 \times \mathcal{G}_2$. Moreover, ϕ lies in a fixed bounded subset of $C^\infty(\Omega)$. The required estimate follows by ‘‘partial integration’’, as summarized by the following lemma. \square

LEMMA 7.8. *Fix $n, N \in \mathbb{Z}_{\geq 0}$ and $\varepsilon > 0$. Let Ω be an open subset of \mathbb{R}^n and $\phi: \Omega \rightarrow i\mathbb{R}$ be smooth. Assume that the total derivative $\partial\phi: \Omega \rightarrow i\mathbb{R}^n$ has euclidean norm $|\partial\phi|: \Omega \rightarrow \mathbb{R}_{\geq 0}$ bounded from below by ε . Then, for all $f \in C_c^\infty(\Omega)$ and $t > 0$,*

$$\int_{\mathbb{R}^n} f e^{t\phi} \ll t^{-N} \sum_{|\alpha| \leq N} \|\partial^\alpha f\|_{L^1},$$

where the implied constant is independent of (f, t) and depends continuously upon

$$\phi \in C^\infty(\Omega).$$

Proof. We may assume that N is even, say $N = 2r$. Let Δ denote the multiple of the standard Laplacian for which $e^\phi = |\partial\phi|^{-2} \Delta(e^\phi)$, and set $D := |\partial\phi|^{-2} \Delta$. Integrating by parts repeatedly, we obtain

$$I = t^{-N} \int f D^r(e^{t\phi}) = t^{-N} \int D^r(f) e^\phi,$$

so that $|I| \leq t^{-N} \int |D^r f|$. Set $b := |\partial\phi|^{-2}$. We may expand

$$D^r f = \sum C(\alpha, \beta^{(1)}, \dots, \beta^{(r)}) (\partial^\alpha f) (\partial^{\beta^{(1)}} b) \dots (\partial^{\beta^{(r)}} b),$$

with the sum taken over multi-indices $\alpha, \beta^{(1)}, \dots, \beta^{(r)}$ satisfying $|\alpha| + |\beta^{(1)}| + \dots + |\beta^{(r)}| = N$. By the quotient rule for derivatives, we have

$$\|\partial^\beta b\|_{L^\infty(\Omega)} \ll 1.$$

The required estimate follows. \square

7.6. Decomposition into localized symbols

Let $m \in \mathbb{Z}$, $\delta \in [0, 1)$ and $\omega \in \mathfrak{g}^\wedge$. Observe that each $a \in S_\delta^m$ varies mildly over the ball

$$U_\omega := \{\xi \in \mathfrak{g}^\wedge : |\xi - \omega| \leq \frac{1}{2} h^\delta \langle \omega \rangle\}.$$

Definition 7.9. We say that the symbol $a \in S_\delta^m$ is *localized at ω* if it is supported on U_ω .

Note that this terminology depends implicitly upon δ .

LEMMA 7.10. *If $a \in S_\delta^m$ is localized at $\omega \in \mathfrak{g}^\wedge$, then*

$$a(\xi) = \langle \omega \rangle^m \phi\left(\frac{\xi - \omega}{h^\delta \langle \omega \rangle}\right),$$

where $\phi \in C_c^\infty(\mathfrak{g}^\wedge)$ depends continuously upon a . In particular,

- The rescaled Fourier transform of a has the form

$$a_h^\vee(x) = \langle \omega \rangle^m e^{-x\omega/h} A^{\dim(\mathfrak{g})} \phi^\vee(Ax), \quad A := h^{\delta-1} \langle \omega \rangle, \quad (7.9)$$

where $\phi^\vee \in \mathcal{S}(\mathfrak{g})$ depends continuously upon a .

- For fixed $n \in \mathbb{Z}_{\geq 0}$,

$$\int_{x \in \mathfrak{g}} |a_h^\vee(x)| |x|^n \ll A^{-n}, \quad (7.10)$$

with continuous dependence upon a .

Proof. Each assertion follows readily from the definition of S_δ^m . \square

It is not difficult to decompose any symbol into localized symbols. To that end, the following partition of unity is convenient.

LEMMA 7.11. *Fix $\delta \in (0, 1]$. There is an h -dependent countable collection $\Omega = \Omega_{\delta, h} \subseteq \mathfrak{g}^\wedge$ of points $\omega \in \mathfrak{g}^\wedge$ with the following properties:*

- the balls U_ω , for $\omega \in \Omega$, cover \mathfrak{g}^\wedge ;
- for $X \geq 1$, we have $\#\{\omega \in \Omega : |\omega| \leq X\} \ll h^{-O(1)} X^{O(1)}$;
- $\sup_{\omega_1 \in \Omega} \#\{\omega_2 \in \Omega : U_{\omega_1} \cap U_{\omega_2} \neq \emptyset\} \ll 1$;
- we have

$$\sum_{\omega \in \Omega} \phi_\omega\left(\frac{\xi - \omega}{h^\delta \langle \omega \rangle}\right) = 1 \text{ for all } \xi \in \mathfrak{g}^\wedge,$$

where ϕ_ω belongs to a fixed bounded subset of $C_c^\infty(\mathfrak{g}^\wedge)$ and is supported on $\{\xi : |\xi| \leq \frac{1}{2}\}$.

Proof. We construct Ω and ϕ_ω , leaving the remaining verifications to the reader. Fix an element $q \in [2, 3]$ which is generic in a sense to be clarified below. Fix a dyadic partition of unity

$$1 = \psi_0(\xi) + \sum_{n \geq 1} \psi_1(q^{-n}\xi),$$

where $\psi_0 \in C_c^\infty(\mathfrak{g}^\wedge)$ and $\psi_1 \in C_c^\infty(\mathfrak{g}^\wedge \setminus \{0\})$. For $n \geq 1$, write $\psi_n(\xi) := \psi_1(q^{-n}\xi)$. Fix a sufficiently dense lattice $L \subseteq \mathfrak{g}^\wedge$ and an additive partition of unity

$$1 = \sum_{\ell \in L} \rho(\xi - \ell),$$

with $\rho \in C_c^\infty(\mathfrak{g}^\wedge)$. Then, for $n \geq 0$,

$$\psi_n(\xi) = \sum_{\ell \in L} \psi_n(\xi) \rho\left(\frac{\xi}{h^\delta q^n} - \ell\right). \quad (7.11)$$

Take for Ω the set consisting of all $\omega = h^\delta q^n \ell$ for which the corresponding summand in (7.11) is non-zero; the genericity assumption on q implies that these elements are pairwise distinct as n varies. Take for ϕ_ω the corresponding summand. \square

LEMMA 7.12. *Each $a \in S_\delta^m$ may be decomposed as $a = \sum_{\omega \in \Omega} \langle \omega \rangle^m a_\omega$, with Ω as above, where $a_\omega \in S_\delta^0$ is localized at ω and depends continuously upon a .*

Proof. Take

$$a_\omega := \langle \omega \rangle^{-m} a(\xi) \phi_\omega\left(\frac{\xi - \omega}{h^\delta \langle \omega \rangle}\right). \quad \square$$

7.7. Proof of Theorem 7.4

The claimed uniqueness follows from the fact that C_c^∞ has dense image in $S^{-\infty}$ and also in S^∞ (note that $C^m \subseteq S^m$, the closure of the image of C_c^∞ , contains $S^{m'}$ whenever $m' < m$). The existence follows, via a limiting procedure, from the continuity established below in the course of the proof of the asymptotic expansion. For the latter, we may assume that $m_1, m_2 < \infty$. It suffices to consider the following cases:

- (a) $m_1 = m_2 = -\infty$;
- (b) $m_1 = -\infty$, $m_2 \in \mathbb{Z}$ and $\mathfrak{g}_1 = \mathfrak{g}$;
- (c) $m_1 \in \mathbb{Z}$, $m_2 = -\infty$ and $\mathfrak{g}_2 = \mathfrak{g}$;
- (d) $m_1 \in \mathbb{Z}$, $m_2 \in \mathbb{Z}$ and $\mathfrak{g}_1 = \mathfrak{g}_2 = \mathfrak{g}$.

Abbreviate $\delta := \max(\delta_1, \delta_2)$. We must verify then that

$$r := a \star_h b - \sum_{0 \leq j < J} h^j a \star^j b$$

belongs to $h^{(1-\delta_1-\delta_2)J} S_\delta^{m_1+m_2-J}(\mathfrak{g}^\wedge)$, i.e., that, for each fixed multi-index $\gamma \in \mathbb{Z}_{\geq 0}^{\dim(G)}$,

$$\partial^\gamma r(\zeta) \ll h^{(1-\delta_1-\delta_2)J-\delta|\gamma|} \langle \zeta \rangle^{m_1+m_2-J-|\gamma|}, \quad (7.12)$$

where the implied constant may depend upon $(m_1, m_2, J, \delta_1, \delta_2)$ and continuously upon a and b , but not upon h, ζ . If either m_1 or m_2 is $-\infty$, then the meaning of (7.12) is that

$$\partial^\gamma r(\zeta) \ll h^{(1-\delta_1-\delta_2)J-\delta|\gamma|} \langle \zeta \rangle^{-N}$$

holds for each fixed N .

In fact, since the spaces $h^{(1-\delta_1-\delta_2)j} S_\delta^{m+n-j}$ decrease as j increases, the terms $h^j a \star^j b$, for fixed $j \geq J$, satisfy the analogue of the estimate (7.12) required by r . The proof of Theorem 7.4 thereby reduces to that of the following assertion: for each fixed $N \in \mathbb{Z}_{\geq 0}$ and multi-index γ , one has for large enough $J \in \mathbb{Z}_{\geq 0}$ that

$$\partial^\gamma r(\zeta) \ll h^N \langle \zeta \rangle^{-N}. \quad (7.13)$$

If $m_k = -\infty$ for some $k=1, 2$, then it will suffice to show that (7.13) holds under the weaker assumption that m_k is any fixed (negative) integer taken sufficiently small in terms of N . In particular, we may assume that $m_1, m_2 \in \mathbb{Z}$.

We may decompose

$$a = \sum_{\omega_1 \in \Omega_1} \langle \omega \rangle^{m_1} a_{\omega_1} \quad \text{and} \quad b = \sum_{\omega_2 \in \Omega_2} \langle \omega \rangle^{m_2} b_{\omega_2}$$

as in §7.6, where $a_{\omega_1} \in S_{\delta_1}^0(\mathfrak{g}_1^\wedge)$ and $b_{\omega_2} \in S_{\delta_2}^0(\mathfrak{g}_2^\wedge)$ are localized at ω_1 and ω_2 , respectively. We may assume N chosen large enough that

$$h^{N/3} \sum_{\omega_j \in \Omega_j} \langle \omega_j \rangle^{m_j - N} \ll 1, \quad j=1, 2,$$

say. The proof of (7.13) thereby reduces to that of the following result.

PROPOSITION 7.13. *Fix $\delta_1, \delta_2 \in [0, 1)$ such that $\delta_1 + \delta_2 < 1$. Fix $N \in \mathbb{Z}_{\geq 0}$. Fix a multi-index $\gamma \in \mathbb{Z}_{\geq 0}^{\dim(G)}$. Fix $J, M \in \mathbb{Z}_{\geq 0}$ sufficiently large in terms of N and γ . Let $a \in S_{\delta_1}^0(\mathfrak{g}_1^\wedge)$ and $b \in S_{\delta_2}^0(\mathfrak{g}_2^\wedge)$ be localized at $\omega_1 \in \mathfrak{g}_1^\wedge$ and $\omega_2 \in \mathfrak{g}_2^\wedge$, respectively. Set $\delta := \max(\delta_1, \delta_2)$ and*

$$r := a \star_h b - \sum_{0 \leq j < J} h^j a \star^j b.$$

We then have the following estimates, in which implied constants may depend upon $(\delta_1, \delta_2, N, \gamma)$ and continuously upon a and b , but not upon $(h, \zeta, \omega_1, \omega_2)$:

- (a) $\partial^\gamma r(\zeta) \ll h^N \langle \zeta \rangle^{-N} \langle \omega_1 \rangle^M \langle \omega_2 \rangle^M$;
- (b) if $\mathfrak{g}_1 = \mathfrak{g}$, then $\partial^\gamma r(\zeta) \ll h^N \langle \zeta \rangle^{-N} \langle \omega_1 \rangle^M \langle \omega_2 \rangle^{-N}$;
- (c) if $\mathfrak{g}_2 = \mathfrak{g}$, then $\partial^\gamma r(\zeta) \ll h^N \langle \zeta \rangle^{-N} \langle \omega_1 \rangle^{-N} \langle \omega_2 \rangle^M$;
- (d) if $\mathfrak{g}_1 = \mathfrak{g}_2 = \mathfrak{g}$, then $\partial^\gamma r(\zeta) \ll h^N \langle \zeta \rangle^{-N} \langle \omega_1 \rangle^{-N} \langle \omega_2 \rangle^{-N}$.

Proof. To simplify the presentation, we focus on the case $\gamma=0$; the general case follows by the same arguments applied with $a, b, \Omega(x, y, \zeta)$ and $F(\xi, \eta, \zeta)$ replaced by some fixed derivatives (with respect to ζ , in the latter two cases); note that our inputs (§7.4 and §7.5) apply to such derivatives.

We define $Q \in \mathbb{R}_{\geq 1}$ in the various cases as follows:

- (a) $Q := h^{-1} \langle \zeta \rangle$;
- (b) $Q := h^{-1} \langle \zeta \rangle \langle \omega_2 \rangle$;
- (c) $Q := h^{-1} \langle \zeta \rangle \langle \omega_1 \rangle$;
- (d) $Q := h^{-1} \langle \zeta \rangle \langle \omega_1 \rangle \langle \omega_2 \rangle$.

We first note that, for fixed j , the element

$$h^j a \star^j b \in h^{(1-\delta_1-\delta_2)j} S_\delta^{-j} \subseteq S_\delta^0$$

is localized at both ω_1 and ω_2 . Thus, $h^j a \star^j b(\zeta) = 0$ unless $\langle \omega_1 \rangle \asymp \langle \omega_2 \rangle \asymp \langle \zeta \rangle$, in which case

$$h^j a \star^j b(\zeta) \ll 1.$$

Next, set $A := h^{-1+\delta_1} \langle \omega_1 \rangle$ and $B := h^{-1+\delta_2} \langle \omega_2 \rangle$, so that, by (7.10), we have

$$\int_{x,y} |a_h^\vee(x) b_h^\vee(y)| |x|^m |y|^n \ll A^{-m} B^{-n} \quad (7.14)$$

for fixed $m, n \geq 0$. By specializing this estimate to the case $m=n=0$, and recalling that $|\Omega(x, y, \zeta)| \ll 1$ and $h^j a \star^j b(\zeta) \ll 1$ for fixed j , we deduce in particular that $r(\zeta) \ll 1$. This gives an adequate estimate for $r(\zeta)$ in the special case $Q \ll 1$. We may and shall thus assume that Q is sufficiently large.

We now fix $\varepsilon > 0$ small in terms of $\delta_1 + \delta_2$, assume that M is chosen large in terms of (N, ε) , and treat the various cases separately:

(a) The required estimate is trivial unless $|\omega_1| \leq Q^\varepsilon$ and $|\omega_2| \leq Q^\varepsilon$, as we henceforth assume.

Suppose that $|\zeta| \geq Q^{2\varepsilon}$. In that case, $a \star^j b(\zeta) = 0$ for all j , so it will suffice to show that $a \star_h b(\zeta) \ll Q^{-N}$. To that end, recall the function $F: \mathfrak{g}_1^\wedge \times \mathfrak{g}_2^\wedge \times \mathfrak{g}^\wedge \rightarrow \mathbb{C}$ defined in §7.5; we have

$$a \star_h b(\zeta) = h^{-2 \dim(\mathfrak{g})} \int_{\xi, \eta} a(\xi) b(\eta) F\left(\frac{\zeta_1 - \xi}{h}, \frac{\zeta_2 - \eta}{h}, \frac{\zeta}{h}\right). \quad (7.15)$$

For ξ and η with $a(\xi) b(\eta) \neq 0$, we have $|\xi| \asymp \langle \omega_1 \rangle \ll Q^\varepsilon$ and $|\eta| \asymp \langle \omega_2 \rangle \ll Q^\varepsilon$, while

$$|\zeta| = \sqrt{|\zeta_1|^2 + |\zeta_2|^2} \gg Q^{2\varepsilon}.$$

Therefore,

$$t := \sqrt{\left| \frac{\zeta_1 - \xi}{h} \right|^2 + \left| \frac{\zeta_2 - \eta}{h} \right|^2} \geq \frac{1}{2} \left| \frac{\zeta}{h} \right| \quad (7.16)$$

and $t \gg Q^{2\varepsilon}$. The required estimate thus follows from §7.5, together with the trivial estimate $O(Q^{O(1)})$ for the L^1 -norms of a and b .

We have reduced to the case that $|\omega_1|, |\omega_2|, |\zeta| \leq Q^{2\varepsilon}$. We then verify readily, using that $\delta_1 + \delta_2 < 1$ and that ε is small enough in terms of $\delta_1 + \delta_2$, that

$$\max\left(\frac{1}{A}, \frac{1}{B}, \frac{|\zeta|}{hAB}\right) \ll Q^{-\varepsilon}. \quad (7.17)$$

Informally, the key point here is that we have reduced to a case in which

$$\langle \omega_1 \rangle \approx \langle \omega_2 \rangle \approx \langle \zeta \rangle, \quad (7.18)$$

so that

$$\frac{|\zeta|}{hAB} \ll h^{1-\delta_1-\delta_2} \frac{\langle \zeta \rangle}{\langle \omega_1 \rangle \langle \omega_2 \rangle} \lesssim h^{1-\delta_1-\delta_2} \langle \zeta \rangle^{-1} \lesssim Q^{-\varepsilon}.$$

We now split $r(\zeta) = r'(\zeta) + r''(\zeta)$, where

$$r'(\zeta) := \int_{x,y} a_h^\vee(x) b_h^\vee(y) \left(\Omega\left(x, y, \frac{\zeta}{h}\right) - \sum_{0 \leq j < J} \Omega_j\left(x, y, \frac{\zeta}{h}\right) \right) \quad (7.19)$$

and

$$r''(\zeta) := \int_{x,y} a_h^\vee(x) b_h^\vee(y) (\chi_1(x) \chi_2(y) - 1) \Omega\left(x, y, \frac{\zeta}{h}\right).$$

Since $\Omega(x, y, \zeta) \ll 1$, we obtain using (7.14) the estimate $r''(\zeta) \ll (AB)^{-n}$ for any fixed n , which is adequate due to (7.17). We estimate $r'(\zeta)$ using §7.4, (7.14) and (7.17), giving the adequate estimate $r'(\zeta) \ll Q^{-\varepsilon J}$.

(b) We may assume that $|\omega_1| \leq Q^\varepsilon$, since the required estimate is otherwise trivial. We may assume that $|\zeta| \leq Q^{2\varepsilon}$, since otherwise (7.16) holds with $t \gg Q^{2\varepsilon}$, and we may conclude as above; in particular, $|\zeta_2| \leq Q^{2\varepsilon}$. We may assume that $|\omega_2| \leq Q^{3\varepsilon}$, since otherwise (7.16) holds with $t \gg Q^{3\varepsilon}$. We then verify (7.17) and conclude as before.

(c) By the same argument, but with the roles of ω_1 and ω_2 reversed.

(d) If for some $k=1, 2$ we have either

- $|\zeta| \leq Q^\varepsilon$ and $|\omega_k| \geq Q^{2\varepsilon}$, or
- $|\zeta| \geq Q^\varepsilon$ and $|\omega_k| \notin [Q^{-\varepsilon^2} |\zeta|, Q^{\varepsilon^2} |\zeta|]$,

then (7.16) holds with $t \gg Q^{\varepsilon^2}$, so we may conclude as above. In the remaining cases, we have either

- $|\zeta| \leq Q^\varepsilon$ and $|\omega_1|, |\omega_2| \leq Q^{2\varepsilon}$, or
- $|\zeta| \geq Q^\varepsilon$ and $Q^{-\varepsilon^2} |\zeta| \leq |\omega_1|, |\omega_2| \leq Q^{\varepsilon^2} |\zeta|$.

In either case, we verify (7.17) and conclude as before. \square

7.8. Asymptotic expansions for certain convolutions

Here we record a miscellaneous estimate to be applied occasionally. Fix $\delta \in [0, 1)$, $m < \infty$, and $\psi \in C_c^\infty(\mathfrak{g})$. For $a \in S_\delta^m$, we may define $b: \mathfrak{g}^\wedge \rightarrow \mathbb{C}$ by requiring that

$$b_h^\vee = \psi a_h^\vee.$$

It is the Fourier transform of a compactly-supported distribution, hence is smooth.

LEMMA 7.14. *We have $b \in S_\delta^m$. Moreover, for each fixed $J \in \mathbb{Z}_{\geq 0}$,*

$$b \equiv \sum_{0 \leq j < J} (-h)^j \sum_{|\alpha|=j} \frac{\partial^\alpha \psi(0)}{\alpha!} \partial^\alpha a \pmod{h^{(1-\delta)J} S_\delta^{m-J}},$$

with remainder depending continuously upon a .

The proof is similar to, but much simpler than, that of Theorem 7.4, hence left to the reader.

8. Proofs concerning the operator assignment

The main aim of this section is to supply the proofs of Theorems 5.6 and 5.8, as well as some of the miscellaneous results stated in §5. We retain their notation and setup. We also establish a generalization (Theorem 8.11) that will be very useful in applications.

8.1. Polynomial symbols: proofs

We now prove Lemma 5.1. We first recall a characterization of the map sym . Each $p \in \text{Sym}(\mathfrak{g}_\mathbb{C})$ defines a translation-invariant differential operator on $C^\infty(\mathfrak{g})$, that we denote by ∂_p : if $p = y_1 \dots y_n$, then

$$\partial_p \phi(x) = \partial_{t_1=0} \dots \partial_{t_n=0} \phi(x + t_1 y_1 + \dots + t_n y_n).$$

On the other hand, each $r \in \mathfrak{U}$ defines a left-invariant differential operator on $C^\infty(G)$, that we denote simply by r : if $r = y_1 \dots y_n$, then

$$r f(g) = \partial_{t_1=0} \dots \partial_{t_n=0} f(g \exp(t_1 y_1) \dots \exp(t_n y_n)).$$

As one verifies readily using Taylor's theorem, the symmetrization map intertwines the two actions near the origin: if $f(\exp(x)) = \phi(x)$, then

$$\partial_p \phi(0) = \text{sym}(p) f(1). \tag{8.1}$$

Now fix $u, v \in \pi^\infty$. After the change of variables $\xi \mapsto -\xi$ in the definition, we must verify that

$$\langle \pi(\text{sym}(p))u, v \rangle = \int_{\xi \in \mathfrak{g}^\wedge} p(-\xi) \left(\int_{x \in \mathcal{G}} e^{x\xi} \langle \pi(\exp(x))u, v \rangle dx \right) d\xi. \quad (8.2)$$

Define $\phi \in C_c^\infty(\mathcal{G})$ and $f \in C_c^\infty(\exp(\mathcal{G}))$ by

$$f(\exp(x)) := \phi(x) := \chi(x) \langle \pi(\exp(x))u, v \rangle.$$

Since $\chi \equiv 1$ in a neighborhood of the origin, we have $rf(1) = \langle \pi(r)u, v \rangle$ for $r \in \mathfrak{U}$. The left-hand side of (8.2) is thus $\text{sym}(p)f(1)$, while the right-hand side is $(p_- \phi^\wedge)^\vee(0) = \partial_p \phi(0)$, where $p_-(\xi) := p(-\xi)$. We conclude by (8.1).

8.2. Smoothness away from the origin

The following simple estimates, suggested in §5.4, will be employed occasionally.

LEMMA 8.1. (i) *For any $a \in S^\infty$, the distributional Fourier transform a^\vee is represented away from the origin by a smooth function.*

(ii) *Fix integers m and N , and a multi-index α with $|\alpha| + m - N \leq -\dim(\mathfrak{g}) - 1$. Then, for $a \in S^m$ and $x \in \mathfrak{g} \setminus \{0\}$, we have*

$$\partial^\alpha a^\vee(x) \ll |x|^{-N},$$

where the implied constant depends continuously upon a . More generally, for $a \in S_\delta^m$,

$$\partial^\alpha a_h^\vee(x) \ll h^{-\dim(\mathfrak{g})} \left| \frac{x}{h^{1-\delta}} \right|^{-N}. \quad (8.3)$$

(iii) *Let $n \in \mathbb{Z}_{\geq 0}$. If $a \in S^m$ with $m \leq -\dim(\mathfrak{g}) - 1 - n$, then a^\vee is represented near the origin by an n -fold differentiable function.*

Proof. We integrate by parts repeatedly in the integral defining a^\vee . □

8.3. Variation with respect to the cut-off: proofs

We now prove Lemma 5.4. It suffices to prove (ii). Define the h -dependent element $f \in C_c^\infty(G)$ by $a_h^\vee \chi_1 = a_h^\vee \chi_2 + f$. Fix $\varepsilon > 0$ small enough that $\chi_1(x) = \chi_2(x)$ whenever $|x| \leq \varepsilon$. Then, $f(x) \neq 0$ only if $|x| > \varepsilon$; in that case, we may apply (8.3) to see that, for any fixed $N \geq 0$, the h -dependent elements $h^{-N} f$ belong to a fixed bounded subset of $C_c^\infty(G)$. As discussed in §3.6, the map $C_c^\infty(G) \rightarrow \Psi^{-\infty}$ is continuous. The conclusion follows.

8.4. Equivariance: proofs

We now prove Lemma 5.5. We may assume that $\varepsilon \in (0, 1)$ is such that $\delta + \frac{1}{2}\varepsilon \in (0, 1)$, say. Recall that $g \in G$ is assumed to satisfy condition (5.4), which we copy here for convenience:

$$\|\mathrm{Ad}(g)\| \ll h^{-1+\delta+\varepsilon}. \quad (5.4)$$

We fix $h_0 > 0$ sufficiently small in terms of ε , δ and the cut-off χ implicit in the definition of Op . We treat separately the cases $h \geq h_0$ and $h < h_0$. In the range $h \geq h_0$, the rescaled symbol a_h lies in a bounded subset S^∞ , and we see from (5.4) that g lies in a fixed compact subset of G modulo the center. The required conclusion thus follows from the same argument as in §5.5. It remains to treat the range $h < h_0$.

Define $b \in S_\delta^\infty$ so that b^\vee is a smooth truncation of a^\vee to $A := \{x: |x| \leq h^{-\delta-\varepsilon/2}\}$. Using §8.2, we see that

$$a \equiv b \pmod{h^\infty S^{-\infty}} \quad \text{and} \quad g \cdot a \equiv g \cdot b \pmod{h^\infty S^{-\infty}}, \quad (8.4)$$

hence

$$\mathrm{Op}_h(a) \equiv \mathrm{Op}_h(b) \pmod{h^\infty \Psi^{-\infty}}, \quad (8.5)$$

$$\mathrm{Op}_h(g \cdot a) \equiv \mathrm{Op}_h(g \cdot b) \pmod{h^\infty \Psi^{-\infty}}. \quad (8.6)$$

We may also verify directly, using the identity

$$\pi(g)[\pi(x), T]\pi(g)^{-1} = [\pi(\mathrm{Ad}(g)x), T]$$

for $x \in \mathfrak{g}$ and its n -fold iterate, that

$$T \in h^\infty \Psi^{-\infty} \implies \pi(g)T\pi(g)^{-1} \in h^\infty \Psi^{-\infty}. \quad (8.7)$$

For $x \in A$, we have

$$|g \cdot x| \leq \|\mathrm{Ad}(g)\| h^{-\delta-\varepsilon/2} \leq h^{-1+\varepsilon/2}.$$

Since $h < h_0$, it follows (having chosen h_0 suitably) that the cut-off χ implicit in the definition of Op_h satisfies $\chi(hx) = \chi(h(g \cdot x)) = 1$, and so the identity

$$\mathrm{Op}_h(g \cdot b) = \pi(g)\mathrm{Op}_h(b)\pi(g)^{-1} \quad (8.8)$$

holds exactly. We conclude by combining (8.5)–(8.8).

8.5. Membership criteria for operator classes

We establish a basic criterion for membership in the operator classes Ψ^m defined in §3.5. This will be applied to establish Lemma 3.4, and then further in §8.7.

PROPOSITION 8.2. *For $m \in \mathbb{Z}$, an operator T on π belongs to Ψ^m if and only if the following holds for each $u \in \mathfrak{U}$:*

$$\text{if } m \geq 0, \text{ then } \sup_{0 \neq v \in \pi^\infty} \frac{\|\theta_u(T)v\|^2}{\langle \Delta^m v, v \rangle} < \infty; \quad (8.9)$$

$$\text{if } m \leq 0, \text{ then } \sup_{0 \neq v \in \pi^\infty} \sup_{x_1, \dots, x_{-m} \in \{1\} \cup \mathcal{B}} \frac{\|\theta_u(T)x_1 \dots x_{-m}v\|}{\|v\|} < \infty. \quad (8.10)$$

Moreover, the seminorms defining Ψ^m may be bounded in terms of such quantities.

The proofs occupy the remainder of §8.5.

We extend $\pi: \mathfrak{U} \rightarrow \text{End}(\pi^\infty)$ to $\pi: \mathfrak{U}[\Delta^{-1}] \rightarrow \text{End}(\pi^\infty)$. (Here $\mathfrak{U}[\Delta^{-1}]$ is the localization of \mathfrak{U} at Δ : it is the universal ring equipped with a morphism from \mathfrak{U} in which Δ becomes invertible. Although this ring is difficult to describe precisely, we will use it in a rather formal fashion.) We define on $\mathfrak{U}[\Delta^{-1}]$ a \mathbb{Z} -filtration by assigning weight 1 to elements of \mathfrak{g} and weight -2 to Δ^{-1} .

Definition 8.3. For $m \in \mathbb{Z}$, we say that $t \in \mathfrak{U}[\Delta^{-1}]$ has order $\leq m$ if it may be expressed as a linear combination of products $w_1 \dots w_n$, $n \in \mathbb{Z}_{\geq 0}$, for which

- (i) for each $i \in \{1..n\}$, either
 - (a) $w_i \in \mathfrak{g}$, or
 - (b) $w_i = \Delta^{-1}$.
- (ii) If cases (a) and (b) occur n_1 and n_2 times, respectively, then $n_1 - 2n_2 \leq m$.

We denote by $\mathfrak{U} \ni u \mapsto \theta_u \in \text{End}(\mathfrak{U}[\Delta^{-1}])$ the algebra morphism extending $\theta_x(t) := [x, t]$ for $x \in \mathfrak{g}$, so that $\theta_{x_1 \dots x_n}(t) = [x_1, \dots, [x_n, t]]$ for $x_1, \dots, x_n \in \mathfrak{g}$. Observe that

$$\theta_x(\Delta^{-1}) = [x, \Delta^{-1}] = -\Delta^{-1}[x, \Delta]\Delta^{-1} \quad \text{for } x \in \mathfrak{g}. \quad (8.11)$$

LEMMA 8.4. *Let $t \in \mathfrak{U}[\Delta^{-1}]$, $u \in \mathfrak{U}$ and $m \in \mathbb{Z}$. If t has order $\leq m$, then also $\theta_u(t)$ has order $\leq m$.*

Proof. By repeated application of (8.11). □

LEMMA 8.5. *If $t \in \mathfrak{U}[\Delta^{-1}]$ has order ≤ 0 , then $\pi(t)$ induces a bounded operator $\pi \rightarrow \pi$.*

Proof. By repeated application of (8.11), we may write t as a linear combination of products of factors of the form $xy\Delta^{-1}$, with $x, y \in \{1\} \cup \mathcal{B}$. To each such factor we apply the case $s=2$ of (3.1), giving $\|\pi(xy\Delta^{-1})v\|^2 \ll \langle \Delta^2 \Delta^{-1}v, \Delta^{-1}v \rangle = \|v\|^2$. □

Proof of Proposition 8.2. The forward implications are straightforward: if $T \in \Psi^m$, then $\theta_u(T): \pi^m \rightarrow \pi^0$ is bounded, so (8.9) holds, while if $T \in \Psi^{-m}$, then $\theta_u(T): \pi^0 \rightarrow \pi^{-m}$ is bounded, so (8.10) follows from (3.1). We turn to the converse implications, which we treat now in a unified manner. For quantities A and B depending upon an element $v \in \pi^\infty$, we write $A \ll B$ to denote that $|A| \leq c|B|$ for some $c \geq 0$ not depending upon v , and write $A \asymp B$ if $A \ll B \ll A$. Let $m \in \mathbb{Z}$. Assume that, for each $u \in \mathfrak{U}$,

$$\text{if } m \geq 0, \text{ then } \|\theta_u(T)v\| \ll \sup_{x_1, \dots, x_m \in \{1\} \cup \mathcal{B}} \|x_1 \dots x_m v\|, \quad (8.12)$$

$$\text{if } m \leq 0, \text{ then } \sup_{x_1, \dots, x_{-m} \in \{1\} \cup \mathcal{B}} \|\theta_u(T)x_1 \dots x_{-m}v\| \ll \|v\|. \quad (8.13)$$

We must show then, for each $u \in \mathfrak{U}$ and $s \in \mathbb{Z}$, that

$$\langle \Delta^{s-m} \theta_u(T)v, \theta_u(T)v \rangle \ll \langle \Delta^s v, v \rangle. \quad (8.14)$$

To that end, choose $k \in \mathbb{Z}_{\geq 0}$ sufficiently large in terms of m and s ; it will suffice to assume that $2k \geq \max(m-s, -s)$. Then,

$$\langle \Delta^s \Delta^k v, \Delta^k v \rangle = \langle \Delta^{s+2k} v, v \rangle \asymp \sup_{x_1, \dots, x_{s+2k} \in \{1\} \cup \mathcal{B}} \|x_1 \dots x_{s+2k} v\|^2,$$

so by the invertibility of Δ , our task (8.14) reduces to showing that

$$\langle \Delta^{s-m} \theta_u(T) \Delta^k v, \theta_u(T) \Delta^k v \rangle \ll \sup_{x_1, \dots, x_{s+2k} \in \{1\} \cup \mathcal{B}} \|x_1 \dots x_{s+2k} v\|^2. \quad (8.15)$$

In the case $s \geq m$, we expand the definitions of Δ^k and Δ^{s-m} and use identities such as $\theta_u(T)\pi(x) = \pi(x)\theta_u(T) - \theta_{xu}(T)$ for $x \in \mathfrak{g}$ to write the left-hand side of (8.15) as a linear combination of expressions

$$\langle \theta_{u'}(T)x_1 \dots x_{s-m+2k}v, \theta_{u''}(T)y_1 \dots y_{s-m+2k}v \rangle, \quad (8.16)$$

where $x_i, y_i \in \{1\} \cup \mathcal{B}$ and $u', u'' \in \mathfrak{U}$. We then apply the Cauchy–Schwarz inequality to each such expression and invoke the assumed bound for T to conclude. We argue similarly in the case $m \geq s$, but only expand Δ^k . We arrive then at expressions of the form

$$\langle A \theta_{u'}(T)x_1 \dots x_{s-m+2k}v, \theta_{u''}(T)y_1 \dots y_{s-m+2k}v \rangle \quad (8.17)$$

with

$$A = z_1 \dots z_{m-s} \Delta^{s-m} w_1 \dots w_{m-s}, \quad (8.18)$$

where $x_i, y_i, z_i, w_i \in \{1\} \cup \mathcal{B}$ and $u', u'' \in \mathfrak{U}$ are as above. By Lemma 8.5, each such operator A is bounded on π^0 . We may thus apply Cauchy–Schwarz to (8.17) and argue as before to conclude. \square

We now establish Lemma 3.4 in a general form.

PROPOSITION 8.6. *Let $m \in \mathbb{Z}$. If $t \in \mathfrak{A}[\Delta^{-1}]$ has order $\leq m$, then $\pi(t)$ has order $\leq m$, i.e., $\pi(t) \in \Psi^m$.*

Proof. By Lemma 3.3, it suffices to consider the following special cases:

- (a) $t \in \mathfrak{g}$ and $m=1$;
- (b) $t = \Delta^{-1}$ and $m=-2$.

We appeal to the criterion of Proposition 8.2. Let $u \in \mathfrak{A}$. In case (a), we have $\theta_u(t) \in \mathfrak{g}$, so the required estimate (8.9) reduces to the case $s=1$ of (3.1). In case (b), we have, by Lemma 8.4, that $\theta_u(t)$ has order ≤ -2 , hence that $\theta_u(t)x_1x_2$ has order ≤ 0 for $x_1, x_2 \in \{1\} \cup \mathcal{B}$, so the required estimate (8.10) reduces to Lemma 8.5. \square

8.6. Operator norm bounds

We record some estimates to be applied below in the proofs of Theorems 5.6 and 5.8.

PROPOSITION 8.7. (i) *If $a \in S^0$, then $\text{Op}(a)$ defines a bounded operator on π , with operator norm bounded continuously in terms of a .*

(ii) *If $a \in S_\delta^0$ for some fixed $\delta \in [0, \frac{1}{2})$, then $\text{Op}_h(a)$ defines an h -dependent bounded operator on π , with operator norm bounded uniformly in h and continuously in terms of a .*

The proof occupies the remainder of this section. It suffices to establish assertion (ii), which recovers assertion (i) upon taking $\delta=0$ and restricting to h -independent symbols.

We begin with some preliminaries. Let \mathcal{N} denote the norm on $\mathcal{S}(\mathfrak{g}^\wedge)$ given by $\mathcal{N}(a) := \|a^\vee\|_{L^1(\mathfrak{g})}$. It is dilation-invariant: $\mathcal{N}(a) = \mathcal{N}(a_h)$.

Denote by $\|\cdot\|$ the operator norm on $\text{End}(\pi)$. We have the following trivial estimate.

LEMMA 8.8. *For $a \in \mathcal{S}(\mathfrak{g}^\wedge)$, we have $\|\text{Op}_h(a)\| \leq \mathcal{N}(a)$.*

We note the following consequence of Lemma 7.10.

LEMMA 8.9. *Let $a \in S_\delta^0$ be localized at some element $\omega \in \mathfrak{g}^\wedge$. Then, $\mathcal{N}(a) \ll 1$; the implied constant may depend upon δ , and continuously upon a , but not upon ω .*

We recall the Cotlar–Stein lemma (see [Hör, Lemma 18.6.5]).

LEMMA 8.10. *Let V_1 and V_2 be Hilbert spaces. Let $T_j: V_1 \rightarrow V_2$ be a sequence of bounded linear operators. Assume that*

$$\sup_j \sum_k \|T_j^* T_k\|^{1/2} \leq C \quad \text{and} \quad \sup_j \sum_k \|T_j T_k^*\|^{1/2} \leq C, \quad (8.19)$$

Then, the series $T := \sum_j T_j$ converges in the Banach space of bounded linear operators from V_1 to V_2 , and has operator norm $\|T\| \leq C$.

We now prove assertion (ii) of the proposition. Let $a \in S_\delta^0$. As in §7.6, we may write $a = \sum_{\omega \in \Omega} a_\omega$, where $a_\omega \in S_\delta^0$ is localized at ω , and depends continuously upon a . Since $\sum_\omega a_\omega$ converges to a distributionally, we have $\text{Op}_h(a) = \sum_\omega \text{Op}_h(a_\omega)$ as maps $\pi^\infty \rightarrow \pi^{-\infty}$. As noted in §2.3, we have $\text{Op}(a_\omega)^* = \text{Op}(\bar{a}_\omega)$, so $\text{Op}(a_{\omega_1})^* \text{Op}(a_{\omega_2}) = \text{Op}(\bar{a}_{\omega_1} \star_h a_{\omega_2}, \chi')$. By Lemmas 8.8 and 8.10, it will thus suffice to show that

$$\sup_{\omega_1 \in \Omega} \sum_{\omega_2 \in \Omega} \mathcal{N}(\bar{a}_{\omega_1} \star_h a_{\omega_2})^{1/2} \ll 1, \quad (8.20)$$

with continuous dependence upon a . To that end, we fix $N \in \mathbb{Z}_{\geq 0}$ large enough, then fix $J \in \mathbb{Z}_{\geq 0}$ large enough in terms of N , and write

$$\bar{a}_{\omega_1} \star_h a_{\omega_2} = \underbrace{\sum_{0 \leq j < J} h^j \bar{a}_{\omega_1} \star^j a_{\omega_2}}_{=: b_{\omega_1, \omega_2}} + r_{\omega_1, \omega_2}.$$

Using Proposition 7.13, we see that

$$\mathcal{N}(r_{\omega_1, \omega_2}) \ll h^N \langle \omega_1 \rangle^{-N} \langle \omega_2 \rangle^{-N}.$$

In particular, r_{ω_1, ω_2} gives an acceptable contribution to (8.20). On the other hand, the symbol b_{ω_1, ω_2} is localized at ω_1 and depends continuously upon a ; moreover, for given ω_1 , we have $b_{\omega_1, \omega_2} = 0$ for ω_2 outside a set of cardinality $O(1)$. By Lemma 8.9, we deduce that b_{ω_1, ω_2} gives an acceptable contribution to (8.20). The proof is complete.

8.7. Proofs of operator class memberships: without rescaling

We now prove Theorem 5.6. We must verify for $m \in \mathbb{Z}$ (hence for $m \in \mathbb{Z} \cup \{\pm\infty\}$) that

$$\text{Op}(S^m) \subseteq \Psi^m, \quad (8.21)$$

with the induced map continuous, and that

$$\text{Op}(a)\text{Op}(b) = \text{Op}(a \star b, \chi') \quad (8.22)$$

for $a, b \in S^\infty$. Recall that (8.21) implies that

$$\text{Op}(S^\infty)\pi^\infty \subseteq \pi^\infty, \quad (8.23)$$

so that the composition in (8.22) makes sense. In fact, we will prove (8.22) and (8.23) simultaneously, and then combine these with §8.6 to deduce (8.21).

We observe first that, for a polynomial symbol p , the corresponding operator $\text{Op}(p)$ acts both on π^∞ and on $\pi^{-\infty}$. This observation is a consequence of §5.5. It applies in particular when $p(\xi) = \langle \xi \rangle^{2N}$, in which case $\text{Op}(p) = \Delta^N$.

We observe next that the composition law (8.22) holds if both a and b are compactly-supported, as is clear from the preliminary discussion in §2.5.

We observe next that the composition law (8.22) holds under the assumption that either a or b is a polynomial. In view of the previous observations, this assumption permits us to define the composition in (8.22), following §5.1.

Suppose first that b is a polynomial. Let $u, v \in \pi^\infty$. We must check that the quantities

$$\langle \text{Op}(a)\text{Op}(b)u, v \rangle = \int_{\xi \in \mathfrak{g}^\wedge} a(\xi) \left(\int_{x \in \mathfrak{g}} e^{-x\xi} \chi(x) \langle \pi(\exp(x))\text{Op}(b)u, v \rangle dx \right) d\xi \quad (8.24)$$

and

$$\langle \text{Op}(a \star b)u, v \rangle = \int_{\xi \in \mathfrak{g}^\wedge} (a \star b)(\xi) \left(\int_{x \in \mathfrak{g}} e^{-x\xi} \chi(x) \langle \pi(\exp(x))u, v \rangle dx \right) d\xi \quad (8.25)$$

are equal. Note that in either expression, the parenthetical integral over x defines a Schwartz function of ξ .

To verify the required equality, we note first from the star product asymptotics that for given b , the tempered distributions a and $a \star b$ depend continuously upon $a \in S^\infty$. Using that C_c^∞ has dense image in S^∞ (cf. the beginning of §7.7), we may thus reduce to the case that a lies in C_c^∞ . Then, $\text{Op}(a)$ and its adjoint $\text{Op}(a)^* = \text{Op}(\bar{a})$ act on π^∞ . We may rewrite (8.24) as

$$\langle \text{Op}(b)u, \text{Op}(a)^*v \rangle = \int_{\xi \in \mathfrak{g}^\wedge} b(\xi) \left(\int_{x \in \mathfrak{g}} e^{-x\xi} \chi(x) \langle \pi(\exp(x))u, \text{Op}(a)^*v \rangle dx \right) d\xi, \quad (8.26)$$

where, since $\text{Op}(a)^*v \in \pi^\infty$, the parenthetical over x again defines a Schwartz function of ξ . Using now that, for given a , the tempered distributions $a \star b$ and b depend continuously upon $b \in S^\infty$, we may reduce the comparison of (8.25) and (8.26) to the case that also $b \in C_c^\infty$. As noted above, the conclusion is known in that case.

If instead a is a polynomial, then we must check that (8.25) and (8.26) are equal. We argue as before, reducing first to the case that b lies in C_c^∞ .

We next verify (8.23). It will suffice to show that

$$\Delta^N \text{Op}(S^\infty) \pi^\infty \subseteq \pi^0$$

for each $N \in \mathbb{Z}_{\geq 0}$. Let us say that $T \in \text{Op}(S^\infty)$ is *good* if $T\pi^\infty \subseteq \pi^0$.

We have

$$\Delta^N \text{Op}(S^\infty) \subseteq \text{Op}(S^\infty) + \Psi^{-\infty}$$

by the special case of (8.22) already established, so it will suffice to show that every element of $\text{Op}(S^\infty)$ is good.

Let $m \in \mathbb{Z}$, $a \in S^m$, $T := \text{Op}(a)$. If $m \leq 0$, so that $T \in \text{Op}(S^m) \subseteq \text{Op}(S^0)$, then $T\pi^0 \subseteq \pi^0$ by Proposition 8.7, so T is good. For the case $m \geq 0$, we show now by induction on m that T is good. Define $p \in S^2$ by $p(\xi) := \langle \xi \rangle^2$. Then, $a/p \in S^{m-2}$. By the case of (8.22) established above and the result of §5.4, we see that

$$\text{Op}(a/p)\text{Op}(p) = \text{Op}(a/p \star p, \chi') = \text{Op}(a/p \star p) + R,$$

where $R \in \Psi^{-\infty}$. By Theorem 4.5, we have $a/p \star p = a + r$ with $r \in S^{m-1}$. Thus,

$$T = \text{Op}(a/p)\Delta - \text{Op}(r) - R.$$

In particular,

$$T \in \text{Op}(S^{m-2})\Delta + \text{Op}(S^{m-1}) + \Psi^{-\infty}. \quad (8.27)$$

By our inductive hypotheses, we conclude that T is good.

We next verify the general case of (8.22). Let $u, v \in \pi^\infty$. Having established (8.23), we know that, in each of (8.24), (8.25) and (8.26), the parenthetical integral over x defines a Schwartz function of ξ . By a continuity argument as above, we may reduce first to the case that a is compactly supported, then to the case that b is compactly supported, in which the conclusion is known.

In summary, we have established (8.23) and the general case of (8.22). We now establish (8.21). (Continuity will be clear from the proof.) We may assume that $m \in \mathbb{Z}$. We appeal to the criterion in §8.5.

By iterated application of (8.22) (with one symbol a polynomial), it will suffice to verify this criterion in the special case $u=1$. Our task is then to show for $T \in \text{Op}(S^m)$ that there exists $C \geq 0$ so that for all v ,

$$\|Tv\|^2 \leq C \langle \Delta^m v, v \rangle, \quad \text{if } m \geq 0, \quad (8.28)$$

$$\sum_{k=0}^{-m} \sup_{x_1, \dots, x_k \in \mathcal{B}(\mathfrak{g})} \|Tx_1 \dots x_k v\| \leq C \|v\|, \quad \text{if } m \leq 0. \quad (8.29)$$

When $m=0$, either assertion follows from §8.6. The case of (8.29) in which $m < 0$ reduces to the case $m=0$ by the composition law (8.22): for $k \leq -m$, we have

$$Tx_1 \dots x_k \in \text{Op}(S^m)\text{Op}(S^1)^k \subseteq \text{Op}(S^{m+k}) + \Psi^{-\infty} \subseteq \text{Op}(S^0) + \Psi^{-\infty}.$$

For $m > 0$, we may then use the decomposition (8.27) to establish (8.28) inductively.

8.8. Proofs of operator class memberships: with rescaling

We now prove Theorem 5.8.

8.8.1. Composition formula

Fix $\delta \in [0, \frac{1}{2})$. From the $h=1$ case treated above and §5.4, we have for $a, b \in S_\delta^\infty$ that

$$\text{Op}_h(a)\text{Op}_h(b) = \text{Op}_h(a \star_h b, \chi') \equiv \text{Op}_h(a \star_h b) \pmod{h^\infty \Psi^{-\infty}}. \quad (8.30)$$

8.8.2. Operator bounds: overview

It remains to verify for each $m \in \mathbb{Z}$ that

$$\text{Op}_h(S_\delta^m) \subseteq h^{\min(0, m)} \Psi_\delta^m. \quad (8.31)$$

We know by Theorem 5.6 that, for each $a \in S_\delta^m$, the operator $\text{Op}_h(a)$ defines an h -dependent element of Ψ^m . Our task is to estimate suitably the variation of its seminorms with respect to h . We carry this out in the remainder of §8.8.

Before proceeding, it may be instructive to give a plausibility argument for why (8.31) is the natural bound to expect. We focus on the following consequence of (8.31): for $a \in S_\delta^m$ and fixed $s \in \mathbb{Z}$, we have

$$\|\text{Op}_h(a)\|_{\pi^s \rightarrow \pi^{s-m}} \ll h^{\min(0, m)}.$$

(Conversely, in §8.8.5, we will apply the composition formula (8.30) to reduce the proof of (8.31) to that of this consequence.) We consider the case that G is the vector space \mathbb{R}^n , regarded as an abelian Lie group, and that $\pi = L^2(\mathbb{R}^n)$ is the regular representation. In this case, we may omit the cut-off χ in our definition of Op . The assignment Op then attaches to each symbol the corresponding Fourier multiplier. The Sobolev norms $\|\cdot\|_{\pi^s}$ may be given in terms of the Fourier transform $v \mapsto v^\wedge$ on $L^2(\mathbb{R}^n)$ by

$$\|v\|_{\pi^s}^2 = \int_{\xi} |v^\wedge(\xi)|^2 \langle \xi \rangle^m d\xi,$$

where as usual $\langle \xi \rangle = (1 + |\xi|^2)^{1/2}$. It follows that, for $b \in S^m$,

$$\|\text{Op}(b)\|_{\pi^s \rightarrow \pi^{s-m}} = \sup_{\xi \in \mathfrak{g}^\wedge} \frac{|b(\xi)|}{\langle \xi \rangle^m}.$$

Recalling the definition $a_h(\xi) = a(h\xi)$ of our rescaling and the estimate $a(\xi) \ll \langle \xi \rangle^m$ contained in the definition of S_m^δ , it follows that

$$\|\text{Op}_h(a)\|_{\pi^s \rightarrow \pi^{s-m}} = \sup_{\xi \in \mathfrak{g}^\wedge} \frac{|a(\xi)|}{\langle \xi/h \rangle^m} \ll \sup_{\xi \in \mathfrak{g}^\wedge} \frac{\langle \xi \rangle^m}{\langle \xi/h \rangle^m}.$$

It is not hard to see that this last supremum is comparable to $h^{\min(0,m)}$ (consider separately the cases in which $\xi=0$ and in which ξ is large).

The case of general G is more involved due to the absence of an analogue for general π of the Fourier transform on $L^2(\mathbb{R}^n)$. In the special case $\delta=0$, we will reduce readily to the h -independent estimates established in §8.7, with the aid of the membership criteria in §8.5. We argue similarly in the case of general δ , but using the composition formula (8.30) as a substitute for the commutator-theoretic arguments in §8.5, which are less efficient when $\delta>0$.

8.8.3. The case $\delta=0$

We treat first the simplest (and most important) case $\delta=0$. Let $a \in S_0^m$, so that

$$\partial^\alpha a(h\xi) \ll \langle h\xi \rangle^{m-|\alpha|}.$$

Writing $b(\xi) := a_h(\xi) = a(h\xi)$ for its rescaling, we see by the chain rule that

$$\partial^\alpha b(\xi) = h^{|\alpha|} \partial^\alpha a(h\xi).$$

Using the estimates

$$\langle h\xi \rangle^m \ll h^{\min(0,m)} \langle \xi \rangle^m \quad \text{and} \quad h\langle \xi \rangle \ll \langle h\xi \rangle,$$

it follows that

$$\partial^\alpha b(\xi) \ll h^{\min(0,m)} \langle \xi \rangle^{m-|\alpha|}$$

In other words, $b \in h^{\min(0,m)} S_0^m$. By the continuity of $\text{Op} = \text{Op}_1: S^m \rightarrow \Psi^m$ established in Theorem 5.6, we deduce that

$$\text{Op}_h(a) = \text{Op}(b) \in h^{\min(0,m)} \Psi_0^m,$$

as required.

8.8.4. Reduction to symbols supported away from the origin

Turning to the case of general $\delta \in [0, 1)$, we treat first the subcase in which $a \in S_\delta^m$ is supported on elements ξ of size $O(h^\delta)$. Set $b(\xi) := a(h^\delta \xi)$, so that $a_h = b_{h^{1-\delta}}$. By the chain rule, $\partial^\alpha b(\xi) = h^{|\alpha|} \partial^\alpha a(h^\delta \xi) \ll 1$. Since b is supported on elements of size $O(1)$, we have $b \in S_0^m$ (in fact, $b \in S_0^{-\infty}$). By the construction of b and the case $\delta=0$ treated above (applied with h replaced by $h^{1-\delta}$), we see that

$$\text{Op}_h(a) = \text{Op}_{h^{1-\delta}}(b) \in (h^{1-\delta})^{\min(0,m)} \Psi_0^m \subseteq h^{\min(0,m)} \Psi_\delta^m,$$

as required. By smoothly decomposing a general symbol, it will suffice from now on to treat the case of symbols supported on $\{\xi:|\xi|\geq h^\delta\}$, say.

We henceforth fix $\delta\in[0,1)$ and take $a\in S_\delta^m$ supported on elements ξ satisfying $|\xi|\geq h^\delta$. By smoothly decomposing a inside S_δ^m , we may assume that there exists a basis element $z\in\mathcal{B}\subseteq\mathfrak{g}$ such that $|z(\xi)|\asymp|\xi|$ for all $\xi\in\text{supp}(a)$. We define the h -dependent operator

$$T:=\text{Op}_h(a)$$

and must verify that $T\in\Psi_\delta^m$, i.e., that for each fixed $u\in\mathfrak{U}$ and $s\in\mathbb{Z}$, the operator norm $\|\theta_u^\delta(T)\|_{\pi^s\rightarrow\pi^{s-m}}$ is $O(h^{\min(0,m)})$, uniformly in h (recall from §5.3 the definition of θ_u^δ).

8.8.5. Reduction to the case $u=1$

It will be convenient to reduce to the case $u=1$. By linearity, we may suppose that $u=x_1\dots x_n$ for some $x_1,\dots,x_n\in\mathfrak{g}$. Regarding x_j as a linear function on \mathfrak{g}^\wedge , we have $\pi(x_j)=h^{-1}\text{Op}_h(x_j)$. Using the composition formula (8.30), we see that

$$\theta_u^\delta(T)\equiv\text{Op}_h(h^{-n}\theta_u^\delta(a))\pmod{h^\infty\Psi^{-\infty}},$$

where $\theta_u^\delta(a)$ is defined inductively by

$$\theta_u^\delta(a)=h^\delta\theta_{x_1\dots x_{n-1}}^\delta(x_n\star_h a-a\star_h x_n)\quad\text{for }n\geq 1$$

and $\theta_1^\delta(a):=a$ for $n=0$. The star-product asymptotics (Theorem 7.4) imply that, for each fixed $y\in\mathfrak{g}$ and $b\in S_\delta^m$, both $y\star_h b$ and $b\star_h y$ are congruent to the pointwise product yb modulo $h^{1-\delta}S_\delta^m$. By induction, it follows that

$$h^{-n}\theta_u^\delta(a)\in S_\delta^m.$$

Thus, the required operator norm bounds $\pi^s\rightarrow\pi^{s-m}$ for each element of $\text{Op}_h(S_\delta^m)$ imply the same for their images under θ_u^δ . The claimed reduction to the case $u=1$ follows.

It remains to verify that $T:\pi^s\rightarrow\pi^{s-m}$ has operator norm $O(h^{\min(0,m)})$ for each fixed $s\in\mathbb{Z}$. As in §8.5, we fix $k\in\mathbb{Z}_{\geq 0}$ large enough in terms of s and m , and reduce to verifying that, for each $v\in\pi$,

$$\langle\Delta^{s-m}T\Delta^k v,T\Delta^k\rangle\ll(h^{\min(0,m)}\|v\|_{\pi^{s+2k}})^2.\tag{8.32}$$

8.8.6. The case $s-m\geq 0$

We argue separately according to the sign of $s-m$, supposing first that $s-m\geq 0$.

Recall that the symbol a is assumed supported on elements ξ with $|\xi| \geq h^\delta$. Let us assume for the moment the stronger support condition $|\xi| \geq 1$.

By expanding the definition of Δ^{s-m} and appealing to Cauchy–Schwarz, we reduce to verifying for all $x_1, \dots, x_{s-m} \in \{1\} \cup \mathcal{B}^c \hookrightarrow \mathfrak{U}$ that

$$\|x_1 \dots x_{s-m} T \Delta^k v\| \ll h^{\min(0,m)} \|v\|_{\pi^{s+2k}}.$$

To see this, we first evaluate $x_1 \dots x_{s-m} T \Delta^k$ using the composition formula. Since

$$\pi(x) = h^{-1} \text{Op}_h(x) \quad \text{for } x \in \mathfrak{g},$$

we obtain

$$x_1 \dots x_{s-m} T \Delta^k \equiv h^{-(s-m+2k)} \text{Op}_h(x_1 \star_h \dots \star_h x_{s-m} \star_h a \star_h p^k),$$

where $p(\xi) := h^2 + |\xi|^2$ and \equiv denotes congruence modulo $h^\infty \Psi^{-\infty}$. (Strictly speaking, \star_h is not associative due to the cut-off χ , so we should specify that this iterated star product is evaluated (say) left-to-right. Changing the order of evaluation introduces negligible errors lying in $h^\infty \Psi^{-\infty}$, so we do not belabor this point.)

By Theorem 5.6, we may find a continuous seminorm ν on S^{s+2k} such that, for all $c \in S^{s+2k}$ and $v \in \pi$, we have

$$\|\text{Op}(c)v\| \leq \nu(c) \|v\|_{\pi^{s+2k}}. \quad (8.33)$$

Since the evaluation at c_h of any S^{s+2k} -seminorm is bounded polynomially in h^{-1} , we see that, if $N \in \mathbb{Z}_{\geq 0}$ is fixed large enough (in terms of m, s, k and δ), then

$$h^{-(s-m+2k)} \nu(c_h) \ll h^{\min(0,m)} \quad \text{for all } c \in h^N S_\delta^{s+2k}. \quad (8.34)$$

Fix $J \in \mathbb{Z}_{\geq 0}$ large enough in terms of m, s, k and N . Let b denote the approximation to $x_1 \star_h \dots \star_h x_{s-m} \star_h a \star_h p^k$ obtained by replacing each star product \star_h with the finite part of its asymptotic expansion obtained by summing over $0 \leq j < J$, as in the statement of Theorem 4.5. Then,

$$b \in S_\delta^{s+2k} \quad \text{and} \quad \text{supp}(b) \subseteq \text{supp}(a).$$

Since J was chosen large enough, we have

$$x_1 \star_h \dots \star_h x_{s-m} \star_h a \star_h p^k = b + c, \quad c \in h^N S_\delta^{-N}.$$

We deduce from (8.34) and (8.33) the acceptable bound

$$h^{-(s-m+2k)} \|\text{Op}_h(c)v\| \ll h^{\min(0,m)} \|v\|_{\pi^{s+2k}}.$$

Our task thereby reduces to verifying that

$$\mathfrak{h}^{-(s-m+2k)} \|\mathrm{Op}_h(b)v\| \ll \mathfrak{h}^{\min(0,m)} \|v\|_{\pi^{s+2k}}. \quad (8.35)$$

Recall from §8.8.4 that we are given an element $z \in \mathcal{B}$ such that $|z(\xi)| \asymp |\xi|$ for all ξ in the support of a , hence also for ξ in the support of b . We now “approximately divide b by z^{s+2k} on the right” with respect to the star product; precisely, we construct $q \in S_\delta^0$ for which $b \approx q \star_h z^{s+2k}$ in the sense that

$$b \equiv q \star_h z^{s+2k} \quad \text{mod } \mathfrak{h}^N S_\delta^{-N}$$

for fixed large enough N . To that end, we take $q = \sum_{0 \leq j < M} \mathfrak{h}^j q_j$, where M is fixed large enough in terms of N and the q_j are chosen so that the required approximation holds in a formal sense; explicitly,

$$q_0 := \frac{b}{z^{s+2k}}, \quad q_1 := \frac{-b \star^1 q_0}{z^{s+2k}} \quad \text{and} \quad q_2 := \frac{-b \star^2 q_0 + b \star^1 q_1}{z^{s+2k}},$$

and so on. We see by induction on j that $q_j \in \mathfrak{h}^{-\delta j} S_\delta^{-(s+2k)j}$. It follows from the composition formula (8.30) that q has the indicated properties.

By another application of the composition formula, we see that the left-hand side of (8.35) is given up to negligible error by

$$\mathfrak{h}^{-(s-m+2k)} \|\mathrm{Op}_h(q) \mathrm{Op}_h(z^{s+2k})v\| = \mathfrak{h}^m \|\mathrm{Op}_h(q)z^{s+2k}v\|.$$

Our task thereby reduces to showing that

$$\mathfrak{h}^m \|\mathrm{Op}_h(q)z^{s+2k}v\| \ll \mathfrak{h}^{\min(0,m)} \|v\|_{\pi^{s+2k}}.$$

To see this, we appeal to part (ii) in Proposition 8.7, which tells us that $\mathrm{Op}_h(q)$ defines a bounded operator $\pi \rightarrow \pi$ with operator norm $O(1)$. Thus,

$$\|\mathrm{Op}_h(q)z^{s+2k}v\| \ll \|z^{s+2k}v\|.$$

By appeal to the lower bound in (3.1) for $\|v\|_{\pi^{s+2k}}$ and the obvious inequality

$$\mathfrak{h}^m \leq \mathfrak{h}^{\min(0,m)},$$

we obtain the required estimate.

This completes our treatment under the assumption that a is supported on $|\xi| \geq 1$. By smooth decomposition inside S_δ^m , it suffices now to treat the complementary case in which a is supported on $\mathfrak{h}^\delta \leq |\xi| \leq 2$, say.

We consider first the subcase in which a is supported on $R \leq |\xi| \leq 2R$ for some h -dependent positive real R with $h^\delta \ll R \ll 1$. The general strategy is then as above, but taking into account that the polynomial symbols x_j and p have respective sizes $\ll R$ and $\ll R^2$ on the support of a . The symbol b constructed as above now lies in $R^{s-m+2k} S_\delta^{-\infty}$. By choosing N and J suitably, the remainder term c remains acceptable. The “quotient” q obtained by approximately dividing b on the right by z^{s+2k} now lies in $R^{-m} S_\delta^0$. Since $h \leq h^\delta \ll R \ll 1$, we have $h^m R^{-m} \ll h^{\min(0,m)}$. We may thus conclude as before.

In the general case that a is supported on $h^\delta \leq |\xi| \leq 2$, we take a smooth dyadic decomposition $a = \sum_R a^{(R)}$, where R runs over powers of two satisfying $h^\delta \ll R \ll 1$ and $a^{(R)}$ is as in the previous paragraph. The quotient q obtained as before has the form $\sum_R q^{(R)}$, where $q^{(R)} \in R^{-m} S_\delta^0$. For each R , the number of R' for which the supports of $q^{(R)}$ and $q^{(R')}$ overlap is $O(1)$. Thus, $q \in \max(1, h^{-\delta m}) S_\delta^0$, and we may conclude as in the previous paragraph.

8.8.7. The case $s - m \leq 0$

The argument is similar in structure to that in the case $s - m \geq 0$ considered above: we treat separately the subcases in which a is supported on $|\xi| \geq 1$ or on $h^\delta \leq |\xi| \leq 2$, and dyadically decompose in the latter case. The only difference is that we arrange the composition and division arguments slightly differently.

We begin by using the composition formula to write the left-hand side of (8.32), up to negligible error, as

$$h^{-4k} \langle \Delta^{s-m} \text{Op}_h(b)v, \text{Op}_h(b)v \rangle,$$

where b is a truncation of the asymptotic expansion of $a \star_h p^k$, with $p(\xi) = h^2 + |\xi|^2$. We then approximately divide the symbol b on the left by z^{m-s} and on the right by z^{s+2k} , giving a symbol q for which $b \approx z^{m-s} \star_h q \star_h z^{s+2k}$ in the same sense as before. In the case that a is supported on $|\xi| \geq 1$, the symbol q lies in S_δ^0 , while in the case that a is supported on $|\xi| \asymp R$, we have $q \in R^{-m} S_\delta^0$. We reduce in either case to verifying that

$$h^{2m} \langle \pi(z)^{s-m} \Delta^{s-m} \pi(z)^{s-m} \text{Op}_h(q)v, \text{Op}_h(q)v \rangle \ll (h^{\min(0,m)} \|v\|_{\pi^{s+2k}})^2.$$

To that end, we see, by Proposition 8.6, that

$$\pi(z)^{s-m} \Delta^{s-m} \pi(z)^{s-m}$$

has operator norm $O(1)$. By Cauchy–Schwarz, we reduce further to showing that

$$h^m \|\text{Op}_h(q)v\| \ll h^{\min(0,m)} \|v\|_{\pi^{s+2k}}.$$

We now conclude exactly as in the case $s - m \geq 0$.

8.9. Generalization to proper subspaces

Let $\mathfrak{g}_1, \mathfrak{g}_2 \leq \mathfrak{g}$ and accompanying notation be as in §7.1; in particular, $\mathfrak{g}_1 + \mathfrak{g}_2 = \mathfrak{g}$. We assume, for convenience, that the cut-offs χ_1 and χ_2 have support taken small enough in terms of the cut-off χ using to define the operator map on $\mathcal{S}(\mathfrak{g}^\wedge)$; this assumption matters little in practice (cf. §5.4).

Let $a \in \mathcal{S}(\mathfrak{g}_1^\wedge)$ and $b \in \mathcal{S}(\mathfrak{g}_2^\wedge)$. As mentioned in §7.1, one can then define operators $\text{Op}_h(a)$ and $\text{Op}_h(b)$ on π , preserving π^∞ , and acting via the restrictions of π to the subgroups G_1 and G_2 . These operators typically do not belong to any of the operator classes we have defined, but their *composition* belongs to $\Psi^{-\infty}$; indeed,

$$\text{Op}(a)\text{Op}(b) = \text{Op}(a \star b)$$

with $a \star b \in \mathcal{S}(\mathfrak{g}^\wedge)$. Some analogues of the above results hold in this setting. For instance, we have the following.

THEOREM 8.11. *Fix $m_1, m_2 \in \mathbb{Z}$ with (m_1, m_2) admissible (§7.2). Fix $\delta \in [0, \frac{1}{2})$. Let $a \in S_\delta^{m_1}(\mathfrak{g}_1^\wedge)$ and $b \in S_\delta^{m_2}(\mathfrak{g}_2^\wedge)$. Then, $\text{Op}_h(a)\text{Op}_h(b) = \text{Op}_h(a \star_h b)$. Moreover, for each fixed $M, N \in \mathbb{Z}_{\geq 0}$, there is a fixed $J \in \mathbb{Z}_{\geq 0}$ such that*

$$\text{Op}_h(a)\text{Op}_h(b) \equiv \sum_{0 \leq j < J} h^j \text{Op}_h(a \star^j b) \pmod{h^N \Psi_\delta^{-M}} \quad (8.36)$$

Proof. We apply the general star product asymptotics (Theorem 7.4) and argue as in the proof of Corollary 5.9. \square

8.10. Disjoint supports

We retain the notation of the previous subsection, and record a simple consequence.

LEMMA 8.12. *Let (δ, m_1, m_2, a, b) be as in the hypotheses of Theorem 8.11. Assume that the preimages in \mathfrak{g}^\wedge of the supports of a and b are disjoint. Then, $\text{Op}_h(a)\text{Op}_h(b) \in h^\infty \Psi^{-\infty}$, with continuous dependence upon a and b .*

The precise meaning of “continuous dependence” is that for any seminorm ℓ defining the topology on $h^\infty \Psi^{-\infty}$ and any $N \geq 0$, we have

$$\ell(\text{Op}_h(a)\text{Op}_h(b)) \leq \nu_1(a)\nu_2(b) h^N$$

for some continuous seminorms ν_j on $S_\delta^{m_j}(\mathfrak{g}_j^\wedge)$.

Proof. We apply (8.36) with large J and use that $a \star^j b = 0$ for all $j \in \mathbb{Z}_{\geq 0}$ and

$$h^\infty \Psi^{-\infty} = \bigcap_{M, N \in \mathbb{Z}_{\geq 0}} h^N \Psi_\delta^{-M}. \quad \square$$

9. Infinitesimal characters

9.1. Overview

Let \mathbf{G} be a reductive group over \mathbb{R} . (Our discussion applies, by restriction of scalars, also to groups over \mathbb{C} .) We denote by G the Lie group of real points of \mathbf{G} , by \mathfrak{g} the Lie algebra, by $\mathfrak{g}_{\mathbb{C}} = \mathfrak{g} \otimes_{\mathbb{R}} \mathbb{C}$ the complexification, by $\mathfrak{g}_{\mathbb{C}}^*$ its complex dual, and by $i\mathfrak{g}^*$ the imaginary dual of \mathfrak{g} , which we may identify with the Pontryagin dual \mathfrak{g}^{\wedge} (cf. §2.1). We regard $i\mathfrak{g}^*$ as a real form of the complex vector space $\mathfrak{g}_{\mathbb{C}}^*$.

We denote by

$$[\mathfrak{g}_{\mathbb{C}}^*] = \mathfrak{g}_{\mathbb{C}}^* // \mathbf{G}$$

the GIT quotient, i.e., the spectrum of the ring of G -invariant polynomials on $\mathfrak{g}_{\mathbb{C}}^*$. This variety has a natural real form, denoted $[i\mathfrak{g}^*]$, corresponding to the polynomials taking real values on $i\mathfrak{g}^*$. We will identify $[\mathfrak{g}_{\mathbb{C}}^*]$ with its set of complex points and likewise $[i\mathfrak{g}^*]$ with its set of real points inside $[\mathfrak{g}_{\mathbb{C}}^*]$. As we recall below, these varieties are affine spaces: the inclusion $[i\mathfrak{g}^*] \hookrightarrow [\mathfrak{g}_{\mathbb{C}}^*]$ looks like $\mathbb{R}^n \hookrightarrow \mathbb{C}^n$.

Example 9.1. Suppose $G = \mathrm{GL}_n(\mathbb{R})$. Using the trace pairing, we may identify $\mathfrak{g}_{\mathbb{C}}^*$ with the space of $n \times n$ complex matrices ξ . The map sending ξ to the characteristic polynomial of ξ/i induces isomorphisms

$$\begin{aligned} [\mathfrak{g}_{\mathbb{C}}^*] &\xrightarrow{\sim} \text{monic polynomials } X^n + p_1 X^{n-1} + \dots + p_{n-1} X + p_n \in \mathbb{C}[X], \\ [i\mathfrak{g}^*] &\xrightarrow{\sim} \text{monic polynomials } X^n + p_1 X^{n-1} + \dots + p_{n-1} X + p_n \in \mathbb{R}[X]. \end{aligned}$$

Let \mathfrak{U} denote the universal enveloping algebra of $\mathfrak{g}_{\mathbb{C}}$, and \mathfrak{Z} its center; the latter acts by scalars on any irreducible representation of G . Harish–Chandra defines an *algebra isomorphism*

$$\gamma: \mathfrak{Z} \simeq \{\text{regular functions on } [\mathfrak{g}_{\mathbb{C}}^*]\}, \quad (9.1)$$

to be recalled below. Each irreducible representation π of G thus gives rise to a point $\lambda_{\pi} \in [\mathfrak{g}_{\mathbb{C}}^*]$, which we refer to as the *infinitesimal character* of π .

The first aim of this section is to record some preliminaries concerning the assignment $\pi \mapsto \lambda_{\pi}$. We then aim to prove, using our operator calculus, some basic estimates involving the λ_{π} .

9.2. Basics concerning the quotient

We denote temporarily by R and $R_{\mathbb{C}}$ the rings of regular functions of the varieties $[i\mathfrak{g}^*]$ and $[\mathfrak{g}_{\mathbb{C}}^*]$ defined above. By definition, $R_{\mathbb{C}} \cong \mathrm{Sym}(\mathfrak{g}_{\mathbb{C}})^G$ is the ring of G -invariant polynomials on $\mathfrak{g}_{\mathbb{C}}^*$; its real form

$$R \cong \mathrm{Sym}(i\mathfrak{g})^G \subseteq R_{\mathbb{C}}$$

consists of the polynomials taking real values on $i\mathfrak{g}^*$.

The \mathbb{R} -algebra R admits a finite set p_1, \dots, p_n of algebraically independent homogeneous generators. Indeed, the corresponding assertion for $R_{\mathbb{C}}$ is a theorem of Chevalley [Ch], and remains true for R due to the following fact whose proof we leave to the reader: if $K \subset L$ are fields and $R = \bigoplus_m R_m$ is a graded K -algebra with the property that $R \otimes_K L$ is polynomial on homogeneous generators, then the same is true for R .

Thus, $R = \mathbb{R}[p_1, \dots, p_n]$ and $R_{\mathbb{C}} = \mathbb{C}[p_1, \dots, p_n]$ are polynomial rings.

Recall that $[\mathfrak{g}_{\mathbb{C}}^*]$ denotes the set of complex points of the spectrum of $R_{\mathbb{C}}$ and $[i\mathfrak{g}^*]$ the set of real points of the spectrum of R ; by definition, $[\mathfrak{g}_{\mathbb{C}}^*]$ consists of \mathbb{C} -algebra maps $R_{\mathbb{C}} \rightarrow \mathbb{C}$ and $[i\mathfrak{g}^*]$ consists of \mathbb{R} -algebra maps $R \rightarrow \mathbb{R}$. There is a natural inclusion $[i\mathfrak{g}^*] \hookrightarrow [\mathfrak{g}_{\mathbb{C}}^*]$. Fixing generators as above, we may identify

$$[i\mathfrak{g}^*] \cong \mathbb{R}^n \subseteq [\mathfrak{g}_{\mathbb{C}}^*] \cong \mathbb{C}^n.$$

In particular, we may speak of the euclidean distance between elements of $[\mathfrak{g}_{\mathbb{C}}^*]$. This depends on the choice of generators p_i above; however, on any fixed compact subset, the notions of distance arising from different choices of generators are comparable, i.e., bounded from above and below in terms of one another.

We denote by $\xi \mapsto [\xi]$ the natural maps $\mathfrak{g}_{\mathbb{C}}^* \rightarrow [\mathfrak{g}_{\mathbb{C}}^*]$ and $i\mathfrak{g}^* \rightarrow [i\mathfrak{g}^*]$. The first of these maps is surjective. More precisely, let $\mathfrak{t}_{\mathbb{C}}$ be a Cartan subalgebra of $\mathfrak{g}_{\mathbb{C}}$, with corresponding Weyl group $W_{\mathbb{C}}$. Then, $\mathfrak{g}_{\mathbb{C}}$ splits as $\mathfrak{n}_{\mathbb{C}}^- \oplus \mathfrak{t}_{\mathbb{C}} \oplus \mathfrak{n}_{\mathbb{C}}$, and we may identify $\mathfrak{t}_{\mathbb{C}}^*$ with the subspace of $\mathfrak{g}_{\mathbb{C}}^*$ orthogonal to $\mathfrak{n}_{\mathbb{C}}^- \oplus \mathfrak{n}_{\mathbb{C}}$. By Chevalley's theorem, the natural composition

$$\mathfrak{t}_{\mathbb{C}}^* \longrightarrow \mathfrak{g}_{\mathbb{C}}^* \longrightarrow [\mathfrak{g}_{\mathbb{C}}^*]$$

induces an isomorphism of complex varieties

$$\mathfrak{t}_{\mathbb{C}}^*/W_{\mathbb{C}} \cong [\mathfrak{g}_{\mathbb{C}}^*]. \quad (9.2)$$

The map $i\mathfrak{g}^* \rightarrow [i\mathfrak{g}^*]$ is not in general surjective (e.g., for $G = \mathrm{SO}(3)$), but its image is readily verified to be Zariski dense.

The complex conjugation $\xi \mapsto \bar{\xi}$ on $\mathfrak{g}_{\mathbb{C}}$ descends to an involution on $[\mathfrak{g}_{\mathbb{C}}^*]$ that we continue to denote by $\lambda \mapsto \bar{\lambda}$. Similarly, the scaling action of $t \in \mathbb{C}^*$ on $\mathfrak{g}_{\mathbb{C}}^*$ given by $\xi \mapsto t\xi$ descends to a scaling action $\lambda \mapsto t\lambda$ on $[\mathfrak{g}_{\mathbb{C}}^*]$. The unique fixed point $[0]$ of the scaling action gives an origin on $[\mathfrak{g}_{\mathbb{C}}^*]$.

For $p \in R$ and $\xi \in \mathfrak{g}_{\mathbb{C}}^*$, one has $\overline{p(\xi)} = p(-\bar{\xi})$; it follows readily that

$$[i\mathfrak{g}^*] = \{\lambda \in [\mathfrak{g}_{\mathbb{C}}^*] : \lambda = -\bar{\lambda}\}. \quad (9.3)$$

9.3. The regular set; description by Cartan subalgebras

We recall that an element of $[\mathfrak{g}_{\mathbb{C}}^*]$ is *regular* if it identifies with a regular point of $\mathfrak{t}_{\mathbb{C}}^*/W_{\mathbb{C}}$, i.e., a point having $|W_{\mathbb{C}}|$ distinct preimages in $\mathfrak{t}_{\mathbb{C}}^*$; equivalently, λ is regular if its preimages in $\mathfrak{g}_{\mathbb{C}}^*$ are regular semisimple. We note that $\xi \in \mathfrak{g}_{\mathbb{C}}^*$ is regular whenever $[\xi]$ is regular (cf. §1.14), but not conversely. As usual, we use a subscripted “reg” to denote the subset of regular elements.

The subset $[i\mathfrak{g}^*]_{\text{reg}}$ of $[i\mathfrak{g}^*]$ is dense and open. We recall how to parameterize

$$[i\mathfrak{g}^*]_{\text{reg}} \cap \text{image}(i\mathfrak{g}^*)$$

in terms of (real) Cartan subgroups T of G (compare with, e.g., [Kn, Theorem 5.22]). For each such T , the complexified Lie algebra $\mathfrak{t}_{\mathbb{C}}$ is a Cartan subalgebra of $\mathfrak{g}_{\mathbb{C}}$. As T varies over a finite set of conjugacy representatives, the images of the maps $(it^*)_{\text{reg}} \rightarrow [i\mathfrak{g}^*]_{\text{reg}}$ partition $[i\mathfrak{g}^*]_{\text{reg}} \cap \text{image}(i\mathfrak{g}^*)$.

9.4. Harish–Chandra isomorphism

We recall the construction of the map γ as in (9.1). (see, e.g., [Kn, p. 220] for further details). Fix a Cartan subalgebra $\mathfrak{t}_{\mathbb{C}} \subseteq \mathfrak{g}$ and a corresponding decomposition

$$\mathfrak{g}_{\mathbb{C}} = \mathfrak{n}_{\mathbb{C}}^- \oplus \mathfrak{t}_{\mathbb{C}} \oplus \mathfrak{n}_{\mathbb{C}}.$$

Let \mathcal{H} denote the universal enveloping algebra (equivalently, symmetric algebra) of $\mathfrak{t}_{\mathbb{C}}$. One has the decomposition

$$\mathfrak{U} = (\mathfrak{n}_{\mathbb{C}}^- \mathfrak{U} + \mathfrak{U} \mathfrak{n}_{\mathbb{C}}) \oplus \mathcal{H}. \quad (9.4)$$

Let $\rho \in \mathfrak{t}^*$ denote the half-sum of positive roots and $\sigma: \mathcal{H} \rightarrow \mathcal{H}$ the algebra automorphism extending $\mathfrak{t}_{\mathbb{C}} \ni t \mapsto t - \rho(t)1_{\mathcal{H}}$. Given $z \in \mathcal{Z}$ with component $z_T \in \mathcal{H}$ relative to the decomposition (9.4), Harish–Chandra defines the element $\gamma(z) := \sigma(z_T) \in \mathcal{H}$. This element turns out to be Weyl-invariant, and thus identifies, via (9.2), with a regular function on $[\mathfrak{g}_{\mathbb{C}}^*]$.

9.5. Basics on infinitesimal characters

Let π be a \mathfrak{U} -module on which \mathfrak{J} acts by scalars. For instance, this happens when π comes from an irreducible representation of G . The *infinitesimal character* $\lambda_{\pi} \in [\mathfrak{g}_{\mathbb{C}}^*]$ is then defined by the following property: each $z \in \mathfrak{J}$ acts on π by the scalar $\gamma(z)(\lambda_{\pi})$.

For any \mathfrak{U} -module π , we may define the dual module π^* and the complex conjugate module $\bar{\pi}$.

We write $\pi^+ := \overline{\pi^*}$ for the conjugate dual. We note that, if π comes from a unitary representation of G , then $\pi \cong \pi^+$. If \mathfrak{Z} acts on π by scalars, then it also acts by scalars on the modules π^* , $\bar{\pi}$ and π^+ , whose infinitesimal characters may be described as follows.

LEMMA 9.2. $\lambda_{\pi^*} = -\lambda_\pi$ and $\lambda_{\bar{\pi}} = \bar{\lambda}_\pi$ and $\lambda_{\pi^+} = -\bar{\lambda}_\pi$.

Proof. This is presumably well known, but we were unable to locate a reference. It suffices to prove for each $\lambda \in [\mathfrak{g}_\mathbb{C}^*]$ that the required identities hold for *some* \mathfrak{U} -module π with $\lambda_\pi = \lambda$.

Let $\mathfrak{g}_s \subseteq \mathfrak{g}_\mathbb{C}$ be a split real form, with split Cartan $\mathfrak{t} \leq \mathfrak{g}_s$, and fix a reductive group G_s with Lie algebra \mathfrak{g}_s . Lift λ to a representative $\lambda \in \mathfrak{t}_\mathbb{C}^*$, and form the corresponding normalized principal series representation $I(\lambda)$ of G_s via induction from some Borel containing $\exp(\mathfrak{t})$. Then, $I(\lambda)$ is a \mathfrak{U} -module whose infinitesimal character is the image of λ (by, e.g., [Kn, Proposition 8.22]), while $I(\lambda)$ has dual $I(-\lambda)$ (by, e.g., calculations as in [Kn, p. 170]) and complex conjugate $I(\bar{\lambda})$ (by construction). \square

In particular, by (9.3), we obtain the following.

COROLLARY 9.3. *If π is an irreducible unitary representation of G , then $\lambda_\pi \in [i\mathfrak{g}^*]$.*

9.6. Langlands classification

Recall that a representation of G is *tempered* if it is unitarizable and weakly contained in $L^2(G)$, or equivalently, if its central character is unitary and the matrix coefficients of its K -finite vectors (with K some maximal compact subgroup of G) belong to $L^{2+\varepsilon}$ modulo the center (see [CHH]). The Langlands classification (see [Kn, Theorem 8.54]) asserts that for each irreducible representation⁽³⁾ π of G , there is a unique G -conjugacy class of pairs (P, σ) , consisting of a parabolic subgroup P of G and a representation σ of the Levi quotient M such that the following statements hold:

- π is the unique irreducible quotient of the induced representation $i_P^G \sigma$;
- σ is tempered on the derived group of M ;
- the absolute value of the central character of σ is strictly dominant.

The infinitesimal characters of σ and π coincide with reference to the natural map $[\mathfrak{m}_\mathbb{C}^*] = \mathfrak{m}_\mathbb{C}^* // M_\mathbb{C} \rightarrow [\mathfrak{g}_\mathbb{C}^*]$ (see [Kn, Proposition 8.22]).

9.7. Infinitesimal criterion for temperedness

The infinitesimal characters of non-tempered representations are located near irregular elements; for lack of a reference, we record the proof.

⁽³⁾ in the sense of (\mathfrak{g}, K) -modules, but this includes unitary representations, see [Kn, Corollary 9.2].

LEMMA 9.4. *For each compact subset Ω of $[i\mathfrak{g}^*]_{\text{reg}}$ there exists $h_0 > 0$ such that, for each $h \in (0, h_0)$, every irreducible unitary representation π with $h\lambda_\pi \in \Omega$ is tempered.*

Proof. First, fix a Cartan subalgebra $\mathfrak{t}_{\mathbb{C}}$ of $\mathfrak{g}_{\mathbb{C}}$ and let $\lambda \in h^{-1}\Omega$. Since λ is regular, any preimage $\tilde{\lambda} \in \mathfrak{t}_{\mathbb{C}}$ under (9.2) satisfies $w \cdot \tilde{\lambda} \neq \tilde{\lambda}$ for all non-trivial elements w of the Weyl group for $\mathfrak{t}_{\mathbb{C}}$. Since Ω is compact, it follows that

$$|w \cdot \tilde{\lambda} - \tilde{\lambda}| \geq c h^{-1} \quad (9.5)$$

for some $c > 0$ depending only upon $\mathfrak{t}_{\mathbb{C}}$ and Ω .

Next, let π be an irreducible unitary representation with $h\lambda_\pi \in \Omega$. We may realize π as the Langlands quotient of $i_P^G \sigma$ for some (P, σ) as in §9.6. Let $|\omega_\sigma|$ denote the (dominant) absolute value of the central character of σ . Assume that π is non-tempered. Then, $P \neq G$ and $|\omega_\sigma|$ is strictly dominant, and in particular non-trivial.

Unitarity also imposes a constraint on $|\omega_\sigma|$. Choose a Levi subgroup M of P . Write \mathfrak{m} for the Lie algebra of M , and \mathfrak{a} for the center of \mathfrak{m} . Let ρ_P be the half-sum of positive roots for \mathfrak{a} on the unipotent radical of P . It follows from boundedness of matrix coefficients that $\rho_P |\omega_\sigma|^{-1}$, considered as a character $\mathfrak{a} \rightarrow \mathbb{C}^*$, is bounded above on the dominant cone [Kn, Chapter XVI, §5, Problems 6 and 7]. This, together with the condition that $|\omega_\sigma|$ is dominant, confines $|\omega_\sigma|$ to a compact subset of \mathfrak{a}^* . (Compare with [SVo, Proposition 7.18].)

Now, passing to a smaller parabolic subgroup if necessary, we may assume that π is a quotient of $i_P^G \sigma$, where

- the restriction of σ to the derived group of the Levi subgroup M of P belongs to the discrete series, and
- the absolute central character $|\omega_\sigma|$ is non-trivial and confined to a compact subset of \mathfrak{a}^* .

Let \mathfrak{m}^0 denote the derived subalgebra of \mathfrak{m} . We then have the splitting $\mathfrak{m} = \mathfrak{a} \oplus \mathfrak{m}^0$, which induces a bijection

$$[\mathfrak{m}_{\mathbb{C}}^*] \simeq \mathfrak{a}_{\mathbb{C}}^* \times [\mathfrak{m}_{\mathbb{C}}^{0*}].$$

By our assumptions on σ , its infinitesimal character decomposes as

$$\lambda_\sigma = (\kappa + \mu, \nu),$$

where $\kappa \in \mathfrak{a}^*$, $\mu \in i\mathfrak{a}^*$ and $\nu \in [i\mathfrak{m}^{0*}]$; moreover, κ is non-zero and confined to a compact set.

By the classification of discrete series [Kn, Theorems 9.20 and 12.21], the parameter ν comes from an imaginary parameter for some (compact) Cartan subgroup of M^0 . We may thus find a Cartan subgroup T of G , contained in M and containing its center, such

that λ_π is the image of $\lambda + \kappa$ for some $\lambda \in \mathfrak{t}^*$; here we identify $\kappa \in \mathfrak{a}^*$ with its pull-back to $\mathfrak{t}_\mathbb{C}^*$. The unitarity of π implies that $\lambda_\pi = -\bar{\lambda}_\pi$; since $\bar{\lambda} = -\lambda$ and $\bar{\kappa} = \kappa$, it follows that

$$\lambda + \kappa = w \cdot (\lambda - \kappa)$$

for some w in the Weyl group of $\mathfrak{t}_\mathbb{C}$. Since $\kappa \neq 0$, the element w is non-trivial.

In summary, we have shown that there exists a Cartan subgroup T of G (which we may assume taken in a finite set of conjugacy representatives), a lift $\tilde{\lambda}_\pi$ of λ_π to $\mathfrak{t}_\mathbb{C}^*$, and a non-trivial Weyl group element w for $\mathfrak{t}_\mathbb{C}$ such that

$$w \cdot \tilde{\lambda}_\pi \in \tilde{\lambda}_\pi + C$$

for some compact $C \subseteq \mathfrak{t}_\mathbb{C}^*$. But if h is small enough, this contradicts (9.5). \square

9.8. Norms

Let us introduce a norm $|\cdot|$ on $[\mathfrak{g}_\mathbb{C}]^*$, i.e. a continuous non-negative function such that $|t\lambda| = |t||\lambda|$ for $t \in \mathbb{C}^*$, and that vanishes only at the origin. Any two such choices are bounded above and below in terms of each other:

$$|\cdot|_1 \asymp |\cdot|_2,$$

and so the explicit choice will not matter.

For example, choosing coordinates $\lambda = (p_1, \dots, p_n)$ on $[\mathfrak{g}_\mathbb{C}^*]$ as in §9.2, where p_j has degree $d_j \geq 1$, the function

$$|\lambda| := \max_j |p_j|^{1/d_j}$$

gives a norm. Alternately, identifying $[\mathfrak{g}_\mathbb{C}^*]$ with the quotient of a Cartan subalgebra by the Weyl group, any Weyl-invariant norm on the Cartan subalgebra gives a norm on $[\mathfrak{g}_\mathbb{C}^*]$.

Now, let π be an irreducible unitary representation of G . Recall the positive-definite densely-defined self-adjoint operator $\Delta := \pi(\Delta)$ on π , as defined in §3.1. Then, the norm $|\lambda_\pi|$ of the infinitesimal character λ_π of π gives a reasonable notion of a “norm of π ”.

LEMMA 9.5. *There is an eigenvalue of Δ on π of size $O(1 + |\lambda_\pi|^2)$. The implied constant depends at most upon \mathbf{G} and the choice of norm.*

Proof. The assertion does not depend on the basis \mathcal{B} of \mathfrak{g} used to define Δ . Choose a Cartan decomposition $\mathfrak{g} = \mathfrak{k} \oplus \mathfrak{p}$ and let $\{y_i\}$ and $\{z_j\}$ be orthonormal bases for \mathfrak{k} and \mathfrak{p} , respectively, as defined with respect to the (possibly-negated) Killing form. Take

$$\mathcal{B} := \{y_i\} \cup \{z_j\}$$

such that $\Delta = 2\Delta_K - 1 - \mathcal{C}$ with

$$\Delta_K := 1 - \sum_i y_i^2 \quad \text{and} \quad \mathcal{C} := - \sum_i y_i^2 + \sum_j z_j^2.$$

Recall that \mathcal{C} (the Casimir operator) defines a quadratic element of \mathfrak{Z} whose eigenvalue on π is thus $O(1 + |\lambda_\pi|^2)$. We thereby reduce to verifying that Δ_K has an eigenvalue of size $O(1 + |\lambda_\pi|^2)$.

For this we use Vogan's theory of minimal K -types (see [Vo2] or [Kn, §XV]). Let μ be a minimal K -type of π . Then, Vogan shows that there exists a parabolic subgroup P and a (possibly not unitary) relative discrete series representation σ of the Levi quotient such that:

- π is a subquotient of an induced representation $I_P^G(\sigma)$, and
- μ is a minimal K -type of $I_P^G(\sigma)$.

This permits us to reduce the question to the case that π is induced from a discrete series, and in turn to the case of discrete series. In that case, the desired result follows from Blattner's formula (see [HS] or [Kn, Theorem 12.26 c]). \square

9.9. Harish–Chandra versus symmetrization

The symmetrization map (cf. §5.2) is the linear isomorphism $\text{Sym}(\mathfrak{g}_{\mathbb{C}})^G \xrightarrow{\text{sym}} \mathfrak{Z}$ given by averaging over permutations. We consider the composition

$$\text{Sym}(\mathfrak{g}_{\mathbb{C}})^G \xrightarrow{\text{sym}} \mathfrak{Z} \xrightarrow{\gamma} \text{Sym}(\mathfrak{g}_{\mathbb{C}})^G. \quad (9.6)$$

LEMMA 9.6. *Let $n \in \mathbb{Z}_{\geq 1}$, and let $p \in \text{Sym}(\mathfrak{g}_{\mathbb{C}})^G$ have order $\leq n$. Then,*

$$\gamma \circ \text{sym}(p) - p \quad \text{has order} \leq n - 1. \quad (9.7)$$

Proof. We use that the component $\text{sym}(p)_T \in \mathcal{H}$ of $\text{sym}(p)$ with reference to (9.4) coincides modulo terms of degree $\leq n - 1$ with the restriction of p via $\mathfrak{t}_{\mathbb{C}}^* \hookrightarrow \mathfrak{g}_{\mathbb{C}}^*$. \square

10. Localizing

Here we record some results of the following theme: if a vector v is “microlocalized” at a point $\xi \in \mathfrak{g}^\wedge$ (§1.7), and a symbol a is supported away from that point, then $\text{Op}(a)v$ is negligibly small. The method of proof—to approximately divide one symbol by another—is ubiquitous in microlocal analysis. The main result is Lemma 10.4, which we apply below in a few places (§19.2, §22.3, §22.4 and §29.3) as an a-priori estimate.

In §10.1 and §10.2, π is an \hbar -dependent unitary representation of a unimodular Lie group G ; in §10.3, we specialize further.

10.1. Division

LEMMA 10.1. Fix $0 \leq \delta_1, \delta_2 < \frac{1}{2}$ and $M, N \in \mathbb{Z}_{\geq 0}$. Set $\delta := \max(\delta_1, \delta_2)$. For each $a \in S_{\delta_1}^{-\infty}$ and $q \in S_0^\infty$ for which

$$a(\xi) \neq 0 \implies |q(\xi)| \geq h^{\delta_2},$$

there exist $b, b' \in h^{-\delta_2} S_\delta^{-\infty}$ such that

$$\begin{aligned} \text{Op}_h(a) &\equiv \text{Op}_h(q) \text{Op}_h(b) \pmod{h^N \Psi_\delta^{-M}}, \\ \text{Op}_h(a) &\equiv \text{Op}_h(b') \text{Op}_h(q) \pmod{h^N \Psi_\delta^{-M}}. \end{aligned}$$

Remark 10.2. One can extend this result to $M=N=\infty$ via ‘‘Borel summation’’ as in [Hör, Proposition 18.1.3], but we have no need to do so.

Proof. (The idea of the proof was applied already in §8.8.6.) We construct b ; one may similarly construct b' . Set $b_0 := a/q$. By the quotient rule, we see that $b_0 \in h^{-\delta_2} S_\delta^{-\infty}$. We now inductively construct b_j for $j \geq 1$ so that

$$q \star_h \sum_{j \geq 0} h^j b_j \sim a$$

in the formal sense (i.e., comparing coefficients of powers of h). Explicitly,

$$b_0 := \frac{a}{q}, \quad b_1 := \frac{-q \star^1 b_0}{q}, \quad b_2 := \frac{-q \star^2 b_0 + q \star^1 b_1}{q},$$

and so on. We see by induction that $\text{supp}(b_j) \subseteq \text{supp}(a)$ and that

$$b_j \in h^{-\delta_2 - 2\delta j} S_\delta^{-\infty}. \tag{10.1}$$

We now take

$$b := \sum_{0 \leq j < J} h^j b_j,$$

with $J \in \mathbb{Z}_{\geq 0}$ large but fixed, and appeal to (5.5). \square

10.2. Localizing near the locus of a symbol

Suppose now given an h -dependent h -uniformly continuous map

$$\mathcal{H}: \Psi^{-\infty} \longrightarrow \mathbb{C}$$

and a symbol $p \in S_0^\infty$ with the property that there is an (h -dependent) scalar $\lambda \in \mathbb{C}$ such that either of the following conditions hold:

$$\mathcal{H}(\text{Op}_h(p)T) = \lambda \mathcal{H}(T) \quad \text{for all } T \in \Psi^{-\infty}, \quad (10.2)$$

$$\mathcal{H}(T\text{Op}_h(p)) = \lambda \mathcal{H}(T) \quad \text{for all } T \in \Psi^{-\infty}. \quad (10.3)$$

We will verify that $\mathcal{H}(\text{Op}_h(a))$ is small if a is supported away from the vanishing locus of p . Let us first choose (as we may) an h -uniformly continuous seminorm ν on $\Psi^{-\infty}$ such that $|\mathcal{H}| \leq \nu$.

LEMMA 10.3. *Fix $\delta \in [0, \frac{1}{2})$ and $N \in \mathbb{Z}_{\geq 0}$. Let $a \in S_\delta^{-\infty}$. Then,*

$$\mathcal{H}(\text{Op}_h(a)) \ll h^{-M} \langle \lambda \rangle^{-N}, \quad (10.4)$$

where M is fixed large enough in terms of the seminorm ν . If moreover

$$a(\xi) \neq 0 \implies |p(\xi) - \lambda| \geq h^\delta,$$

then

$$\mathcal{H}(\text{Op}_h(a)) \ll h^N \langle \lambda \rangle^{-N}. \quad (10.5)$$

Proof. We assume (10.2); a similar proof applies under (10.3). By Lemma 5.3, we may choose M so that the restriction of ν to Ψ_δ^{-M} is continuous. We have

$$\begin{aligned} \mathcal{H}(\text{Op}_h(a)) &= \lambda^{-N} \mathcal{H}(\text{Op}_h(p)^N \text{Op}_h(a)), \\ \text{Op}_h(p)^N \text{Op}_h(a) &\subseteq \text{Op}_h(S_\delta^{-\infty}) + h^\infty \Psi^{-\infty}, \\ \text{Op}_h(S_\delta^{-\infty}) &\subseteq \text{Op}_h(S_\delta^{-M}) \subseteq h^{-M} \Psi_\delta^{-M}. \end{aligned}$$

By applying these both with the given value of N and with $N=0$, we obtain (10.4). We turn to (10.5). Fix $\varepsilon > 0$ with $\delta + \varepsilon < \frac{1}{2}$. By applying (10.4) with N sufficiently large, we reduce to the case that $|\lambda| \leq h^{-\varepsilon}$. We must verify then that

$$\mathcal{H}(\text{Op}_h(a)) \ll h^N. \quad (10.6)$$

Set $q := h^\varepsilon(p - \lambda) \in S_0^\infty$. We have $|q| \geq h^{\delta + \varepsilon}$ on the support of a . By §10.1, we may write $\text{Op}_h(a) \equiv \text{Op}_h(q)\text{Op}_h(b)$ and $h^N \Psi_\delta^{-M}$ for some b . Since $\mathcal{H}(\text{Op}_h(q)\text{Op}_h(b)) = 0$ and \mathcal{H} induces a continuous map $\Psi_\delta^{-M} \rightarrow \mathbb{C}$, the required estimate (10.6) follows. \square

10.3. Localizing near an infinitesimal character

We now establish an analogue of §10.2 which will be very useful in applications. Fix an inclusion $\mathbf{H} \hookrightarrow \mathbf{G}$ of reductive groups over \mathbb{R} . Let π and σ be \mathfrak{h} -dependent irreducible unitary representation of G and H , respectively. Let

$$\mathcal{H}: \Psi^{-\infty}(\pi) \longrightarrow \mathbb{C}$$

be an \mathfrak{h} -dependent \mathfrak{h} -uniformly continuous map which factors as an $(H \times H)$ -equivariant \mathfrak{h} -uniformly continuous composition

$$\Psi^{-\infty}(\pi) \longrightarrow \Psi^{-\infty}(\sigma) \longrightarrow \mathbb{C}.$$

We quantify the rescaled frequencies of π and σ via their infinitesimal characters:

$$\langle \mathfrak{h}\lambda_\pi \rangle := (1 + |\mathfrak{h}\lambda_\pi|^2)^{1/2} \quad \text{and} \quad \langle \mathfrak{h}\lambda_\sigma \rangle := (1 + |\mathfrak{h}\lambda_\sigma|^2)^{1/2}.$$

Define

$$\text{Op}: S^m := S^m(\mathfrak{g}^\wedge) \longrightarrow \Psi^m := \Psi^m(\pi),$$

as usual. Let $a \in S^{-\infty}$. We will verify that $\mathcal{H}(\text{Op}_\mathfrak{h}(a))$ is small unless a is supported on elements $\xi \in \mathfrak{g}^\wedge$ for which

$$([\xi], [\xi|_\mathfrak{h}]) \approx (\mathfrak{h}\lambda_\pi, \mathfrak{h}\lambda_\sigma) \in [\mathfrak{g}^\wedge] \times [\mathfrak{h}^\wedge].$$

More precisely, let us choose an \mathfrak{h} -uniformly continuous seminorm ν on $\Psi^{-\infty}(\pi)$ such that $|\mathcal{H}| \leq \nu$. Then, we have the following result.

LEMMA 10.4. *Fix $\delta \in [0, \frac{1}{2})$ and $N \in \mathbb{Z}_{\geq 0}$. Fix $M \geq 0$ large enough in terms of ν . Let $a \in S_\delta^{-\infty}$. Then, $\mathcal{H}(\text{Op}_\mathfrak{h}(a))$ satisfies the “a-priori estimate”*

$$\mathcal{H}(\text{Op}_\mathfrak{h}(a)) \ll \mathfrak{h}^{-M} \langle \mathfrak{h}\lambda_\pi \rangle^{-N} \langle \mathfrak{h}\lambda_\sigma \rangle^{-N}. \quad (10.7)$$

Suppose now that a is supported on the complement of

$$\{\xi \in \mathfrak{g}^\wedge : \text{dist}([\xi], \mathfrak{h}\lambda_\pi) \leq \mathfrak{h}^\delta \text{ and } \text{dist}([\xi|_\mathfrak{h}], \mathfrak{h}\lambda_\sigma) \leq \mathfrak{h}^\delta\}, \quad (10.8)$$

with dist being the euclidean distance function defined by the coordinates fixed in §9.

Then, $\mathcal{H}(\text{Op}_\mathfrak{h}(a))$ is negligible:

$$\mathcal{H}(\text{Op}_\mathfrak{h}(a)) \ll \mathfrak{h}^N \langle \mathfrak{h}\lambda_\pi \rangle^{-N} \langle \mathfrak{h}\lambda_\sigma \rangle^{-N}. \quad (10.9)$$

Proof. We denote as in §9.2 by $R_{\mathbb{C}} \cong \text{Sym}(\mathfrak{g}_{\mathbb{C}})^G$ and $R = \text{Sym}(i\mathfrak{g})^G$ the rings of G -invariant polynomials on $i\mathfrak{g}^*$ taking complex and real values, respectively. Recall (from §5.2, equation (9.1)) that, for $p \in R_{\mathbb{C}}$,

$$\text{Op}(p) \text{ acts by the scalar } (\gamma \circ \text{sym}(p))(\lambda_{\pi}). \quad (10.10)$$

For $p \in R$, we denote by $p' \in R$ the element for which

$$\gamma \circ \text{sym}(p'_h) = p_h. \quad (10.11)$$

Regarding p' as a polynomial symbol on \mathfrak{g}^{\wedge} , we see that

$$\text{Op}_h(p') \text{ acts by the scalar } p_h(\lambda_{\sigma}) = p(h\lambda_{\sigma}). \quad (10.12)$$

On the other hand, if we fix p and regard p' as an h -dependent polynomial symbol, then Lemma 9.6 gives

$$p' = O(1) \quad \text{and} \quad p' = p + O(h); \quad (10.13)$$

more precisely, (10.13) says that $p' \in h^0 S_0^{\infty}$ and $p' - p \in h^1 S_0^{\infty}$.

Fix a set $\{p\} = \{p_G\} \sqcup \{p_H\}$, where

- p_G runs over a system of generators for the ring $\text{Sym}(i\mathfrak{g})^G$, corresponding to the coordinate functions defining the distance function on $[\mathfrak{g}^{\wedge}]$, and
- p_H runs over a similar system for $\text{Sym}(i\mathfrak{h})^H$.

Let the assignment $p \mapsto p'$ be as above, applied either to H or to G . For each such p , our assumptions concerning \mathcal{H} imply that

$$\mathcal{H}(\text{Op}_h(p')T) = p(h\lambda_{\sigma})\mathcal{H}(T) \quad (10.14)$$

for all $T \in \Psi^{-\infty}$. We have $|h\lambda_{\pi}| + |h\lambda_{\sigma}| \asymp \max_p |p(h\lambda_{\sigma})|$, so the first estimate (10.7) follows immediately from (10.14) and §10.2. In verifying (10.9), we may thus suppose that $|h\lambda_{\pi}| + |h\lambda_{\sigma}| \leq h^{-\varepsilon}$ for any fixed $\varepsilon > 0$. By (10.13), we have

$$|p(\xi) - p'(\xi)| \ll h^{1-\varepsilon},$$

so that, for small enough h ,

$$\begin{aligned} \text{dist}([\xi|_h], h\lambda_{\sigma}) > h^{\delta} &\implies \max_p |p(\xi) - p(h\lambda_{\sigma})| \gg h^{\delta} \\ &\implies \max_p |p'(\xi) - p(h\lambda_{\sigma})| > h^{\delta+\varepsilon}, \end{aligned}$$

say, and similarly if $\text{dist}([\xi], h\lambda_{\pi}) > h^{\delta}$. Having fixed ε small enough, we may suppose that $\delta + \varepsilon < \frac{1}{2}$. If a is supported on the complement of (10.8), then we may decompose it into pieces indexed by p supported on the sets

$$\{\xi : |p'(\xi) - p(h\lambda_{\sigma})| > h^{\delta+\varepsilon}\}.$$

We conclude by applying §10.2 to each such piece. \square

11. Regular coadjoint multiorbits

G is the set of real points of a real reductive group \mathbf{G} , with notation as above.

11.1. Notation and terminology

For each infinitesimal character $\lambda \in [\mathfrak{g}^\wedge]$, we may form the fiber $\mathcal{O}^\lambda := \{\xi \in \mathfrak{g}^\wedge : [\xi] = \lambda\}$. For example, if $\lambda = [0]$, then \mathcal{O}^λ is the nilcone \mathcal{N} . As noted already in §6.1, each such fiber \mathcal{O}^λ consists of a uniformly bounded finite number of G -orbits. We recall that a *coadjoint multiorbit* $\mathcal{O} \subseteq \mathfrak{g}^\wedge$ is a G -invariant set contained in \mathcal{O}^λ for some λ , and that \mathcal{O} is *regular* if it consists of regular elements; then,

$$\mathcal{O} \text{ is regular} \iff \mathcal{O} \text{ is relatively open} \iff \mathcal{O} \text{ has maximal dimension,}$$

where “relatively open” is with respect to \mathcal{O}^λ , and “maximal” means with respect to all coadjoint orbits.

Recall that to each coadjoint orbit \mathcal{O} we may attach its normalized symplectic measure $\omega := \omega_{\mathcal{O}}$ on \mathcal{O} . Let $a \in C_c(\mathfrak{g}^\wedge)$. By a result of Rao [Rao], the integral $\int_{\mathcal{O}} a \, d\omega$ converges. Hence, ω may be regarded as a measure on \mathfrak{g}^\wedge .

11.2. Topology

We temporarily denote by \mathfrak{R} the set of regular coadjoint multiorbits $\mathcal{O} \subseteq \mathfrak{g}_{\text{reg}}^\wedge$. We equip \mathfrak{R} with the topology induced from the inclusion

$$\mathcal{O} \longmapsto \omega_{\mathcal{O}} \in \{\text{locally finite Radon measures on } \mathfrak{g}^\wedge\},$$

where we endow the target with the weak-* topology. Thus, a sequence \mathcal{O}_j of regular coadjoint multiorbits converges to \mathcal{O} in \mathfrak{R} if the corresponding symplectic measures $\omega_{\mathcal{O}_j}$ tend to $\omega_{\mathcal{O}}$. We note that the infinitesimal character map $\mathfrak{R} \setminus \{\emptyset\} \rightarrow [\mathfrak{g}^\wedge]$ is continuous, and has finite fibers.

We note that the topology on \mathfrak{R} is Hausdorff: the topology on the target of the above inclusion is Hausdorff, and the map $\mathcal{O} \mapsto \omega_{\mathcal{O}}$ is injective in view of the regularity of \mathcal{O} .

It is a non-trivial fact that the topology on \mathfrak{R} may be described more simply.

THEOREM 11.1. *Let $\mathcal{O}_j \in \mathfrak{R}$ be a sequence of regular coadjoint multiorbits. Set*

$$\mathcal{O} := \{\xi \in \mathfrak{g}_{\text{reg}}^\wedge : \xi_j \rightarrow \xi \text{ for some } \xi_j \in \mathcal{O}_j\}.$$

(i) *If \mathcal{O}_j has a non-empty limit $\mathcal{O}_{\text{lim}} \in \mathfrak{R}$, then $\mathcal{O}_{\text{lim}} = \mathcal{O}$.*

(ii) *If \mathcal{O} is non-empty and contains all regular subsequential limits of the sequence \mathcal{O}_j , then $\mathcal{O} \in \mathfrak{R}$ and $\mathcal{O}_j \rightarrow \mathcal{O}$.*

This follows from arguments of Rossmann [Ros2], [Ros3], who also gives in [Ros3] some characterizations of when \mathcal{O} is non-empty. Since it does not appear to have been stated explicitly in the above form, we outline the proof. We consider first the special case involving full preimages under the map $\mathfrak{g}_{\text{reg}}^\wedge \rightarrow [\mathfrak{g}^\wedge]$. To that end, for $\lambda \in [\mathfrak{g}^\wedge]$, set

$$\omega_\lambda := \omega_{\mathcal{O}_{\text{reg}}}^\lambda,$$

where by convention $\omega_\lambda := 0$ if $\mathcal{O}_{\text{reg}}^\lambda = \emptyset$.

LEMMA 11.2. *The measures ω_λ vary continuously with respect to $\lambda \in [\mathfrak{g}^\wedge]$.*

Proof. Suppose that $\lambda_j \rightarrow \lambda$. We must verify that $\omega_{\lambda_j} \rightarrow \omega_\lambda$. By a diagonalization argument, we may assume that the λ_j are regular. After passing to a subsequence, we may assume that there is a Cartan subalgebra \mathfrak{h} of \mathfrak{g} such that λ_j is the image of $t_j \in \mathfrak{h}_{\text{reg}}$, with the t_j lying in the same connected component $C \subseteq \mathfrak{h}_{\text{reg}}$ and having a limit $t \in \overline{C} \subseteq \mathfrak{h}$. The required conclusion in that case is stated explicitly in the second paragraph of [Ros2, p. 59] and follows from [Ros2, Lemma D] and the arguments of [Ros2, pp. 59–62], which rely in turn upon results of Harish–Chandra. \square

LEMMA 11.3. *Fix $\xi_0 \in \mathfrak{g}_{\text{reg}}^\wedge$, and let $\xi_0 \in U \subseteq \mathfrak{g}_{\text{reg}}^\wedge$ be a sufficiently small open neighborhood. Let \mathcal{O} be any regular coadjoint multiorbit. Clearly, $\omega_{\mathcal{O}}|_U$ is non-zero if and only if $\mathcal{O} \cap U \neq \emptyset$. In that case, $\omega_{\mathcal{O}}|_U = \omega_{G \cdot \xi}|_U$ for any $\xi \in \mathcal{O} \cap U$. The family of measures $\omega_{G \cdot \xi}|_U$ varies continuously with respect to $\xi \in U$.*

Proof. We have for each $\lambda \in [\mathfrak{g}^\wedge]$ that the intersection $U \cap \mathcal{O}^\lambda = U \cap \mathcal{O}_{\text{reg}}^\lambda$ is non-empty precisely when there is some $\xi \in U$ with $[\xi] = \lambda$. Since the differential of the map $\mathfrak{g}_{\text{reg}}^\wedge \rightarrow [\mathfrak{g}^\wedge]$ is surjective at every point [Kos2, Theorem 0.1], we have in that case—for sufficiently small U —that $U \cap \mathcal{O}^\lambda = U \cap (G \cdot \xi)$. The claims then follow from the prior lemma. \square

Proof of Theorem 11.1. Let $\lambda_j = [\mathcal{O}_j] \in [\mathfrak{g}^\wedge]$ denote the infinitesimal character of \mathcal{O}_j .

Suppose first that the \mathcal{O}_j have a non-empty limit $\mathcal{O}_{\text{lim}} \in \mathfrak{R}$, with infinitesimal character $\lambda := [\mathcal{O}_{\text{lim}}]$. Then, $\lambda_j \rightarrow \lambda$; it follows that \mathcal{O} is non-empty. It is therefore a non-empty regular coadjoint multiorbit with infinitesimal character λ . Using Lemma 11.3, we see that

$$(\text{restriction to } \mathfrak{g}_{\text{reg}}^\wedge \text{ of } \omega_{\mathcal{O}_j}) \longrightarrow (\text{restriction to } \mathfrak{g}_{\text{reg}}^\wedge \text{ of } \omega_{\mathcal{O}}). \quad (11.1)$$

Since $\mathcal{O}_j \rightarrow \mathcal{O}_{\text{lim}}$, this forces $\mathcal{O}_{\text{lim}} = \mathcal{O}$.

Conversely, suppose that \mathcal{O} is non-empty and contains the regular subsequential limits of the \mathcal{O}_j . We can find a sequence $\xi_j \in \mathcal{O}_j$ converging to some $\xi \in \mathcal{O}$, so that $\lambda_j = [\xi_j]$ likewise tends to $\lambda := [\xi]$. Since every element of \mathcal{O} arises in this way, we see that \mathcal{O} is a non-empty regular coadjoint multiorbit with infinitesimal character λ . We aim to

verify that $\mathcal{O}_j \rightarrow \mathcal{O}$. Using Lemma 11.3, we see that (11.1) holds. The remaining point is to understand what happens on $\mathfrak{g}^\wedge - \mathfrak{g}_{\text{reg}}^\wedge$. Since $\omega_{\mathcal{O}_j} \leq \omega_{\lambda_j}$, we see from Lemma 11.2 that $\omega_{\mathcal{O}_j}$ admits (after passing to a subsequence) a limit measure ω' which is G -invariant and bounded by ω_λ , hence is a non-negative linear combination of the measures $\omega_{\mathcal{O}'}$ attached to G -orbits $\mathcal{O}' \subseteq \mathcal{O}_{\text{reg}}^\lambda$. By (11.1), we deduce that $\omega' = \omega_{\mathcal{O}}$. Thus, $\mathcal{O}_j \rightarrow \mathcal{O}$. \square

11.3. Bounds for symplectic measures

In what follows, we denote by $\sup_{\mathcal{O}}$ a supremum taken over all non-empty regular coadjoint multiorbits.

LEMMA 11.4. *For each $a \in C_c(\mathfrak{g}^\wedge)$,*

$$\sup_{\mathcal{O}} \left| \int_{\mathcal{O}} a \, d\omega \right| < \infty. \quad (11.2)$$

Proof. It will suffice to show that $\sup_{\lambda \in [\mathfrak{g}^\wedge]} |\omega_\lambda(a)| < \infty$, with ω_λ attached to $\mathcal{O}_{\text{reg}}^\lambda$ as in §11.2. Recall that $|\omega_\lambda(a)| < \infty$. The image in $[\mathfrak{g}^\wedge]$ of the support of a is compact, so we may conclude via the continuity noted in Lemma 11.2. \square

Combining this with the homogeneity property (6.1) of the symplectic measures, we obtain a basic estimate valid uniformly over \mathfrak{g}^* . Recall the discussion of norms from §9.8 and the abbreviation $\langle \lambda \rangle := (1 + |\lambda|^2)^{1/2}$, which applies in particular with $\lambda = [\mathcal{O}]$.

LEMMA 11.5. *For each $\varepsilon > 0$,*

$$\sup_{\mathcal{O}} \langle [\mathcal{O}] \rangle^\varepsilon \int_{\xi \in \mathcal{O}} \langle \xi \rangle^{-d(\mathcal{O}) - \varepsilon} \, d\omega_{\mathcal{O}}(\xi) < \infty.$$

Proof. The contribution from $|\xi| \leq 1$ may be estimated using the prior lemma, noting that this contribution vanishes identically if $\langle [\mathcal{O}] \rangle > 1$. For the remaining contribution, we split the ξ -integral into dyadic shells $A \leq |\xi| \leq 2A$, with $A = 2^n$ for $n \in \mathbb{Z}_{\geq 0}$. By (6.1) and (11.2), the contribution from each such shell is bounded by a constant multiple of $A^{-\varepsilon}$, while the smallest A giving a non-zero contribution has size $\gg \langle [\mathcal{O}] \rangle$. We conclude by summing dyadically over A . \square

11.4. Limit orbits

Let π be a tempered irreducible unitary representation of G , the real points of a reductive group \mathbf{G} over \mathbb{R} . We now allow π to vary with a positive parameter $h \rightarrow 0$ traversing some sequence $\{h\} \subseteq (0, 1)$. We emphasize again that the dependence of π and σ upon h is not assumed to be (e.g.) continuous or measurable.

11.4.1. Definition

Let \mathcal{O} be a regular coadjoint multiorbit. We say that \mathcal{O} is the *limit orbit* of π (or, pedantically speaking, *limit multiorbit*) if $\lim_{\mathfrak{h} \rightarrow 0} \mathfrak{h} \mathcal{O}_\pi = \mathcal{O}$ in the sense of §11.2. We then often abbreviate $\omega_{\mathcal{O}}$ to simply ω .

We note that π admits at most one limit orbit, since the topology on \mathfrak{X} is Hausdorff (see §11.2). We note also that if π admits the non-empty limit orbit \mathcal{O} , then $\mathfrak{h}\lambda_\pi$ converges to the infinitesimal character $\lambda := [\mathcal{O}]$ of the limit orbit.

11.4.2. The \mathfrak{h} -independent case

Recall that π is *generic* if

- G is quasi-split, i.e., contains a Borel subgroup defined over \mathbb{R} , and
- π admits a Whittaker model (with respect to some non-degenerate character of the unipotent radical of that Borel).

We refer to [Vol] for definitions concerning Gelfand–Kirillov dimension.

LEMMA 11.6. *Suppose that π is \mathfrak{h} -independent. Then, π has a limit orbit \mathcal{O} , where \mathcal{O} is contained in the regular subset \mathcal{N}_{reg} of the nilcone. The following are equivalent:*

- (i) \mathcal{O} is non-empty;
- (ii) π has maximal Gelfand–Kirillov dimension;
- (iii) π is generic.

Proof. The initial assertions and the equivalence of (i) and (ii) follow from [BV] and [Ros5, Theorem C and D]. (The initial assertions also follow readily from Theorem 11.1.) The equivalence of (ii) and (iii) follows from [Kos3, Theorem 6.7.2].⁽⁴⁾ \square

Remark 11.7. Condition (i) is what is relevant for the purposes of this paper; we have invoked its equivalence with (ii) and (iii) only to simplify the statement of our result.

12. Trace estimates

We now prove (a sharper form of) Theorem 6.3. We retain the setup of §11.4, and define $\text{Op}: S^m := S^m(\mathfrak{g}^\wedge) \rightarrow \Psi^m := \Psi^m(\pi)$ as usual. Recall that we write

$$\langle \xi \rangle := (1 + |\xi|^2)^{1/2}$$

for an element ξ of a normed space.

⁽⁴⁾ Kostant proves what is required here modulo the possibility of replacing π by another Hilbert space representation π' ; Kostant’s Hilbert space representations do not preserve the inner product, so it is not obvious that π and π' are “the same”. But Casselman [Cas1] shows that the spaces of smooth vectors in the two representations are isomorphic.

12.1. Spaces of operators associated with various norms

For $T: \pi \rightarrow \pi$ and $p=1, 2$ or ∞ , we let $\|T\|_p$ denote respectively the trace, Hilbert–Schmidt or operator norm of T . We let $\mathcal{T}_p := \mathcal{T}_p(\pi)$ denote the space of operators T on π with the property that for each $u \in \mathfrak{U}$, the operator $\theta_u(T)$ (defined as in §3.3) induces a bounded map $\pi \rightarrow \pi$ for which $\|\theta_u(T)\|_p < \infty$. We equip \mathcal{T}_p with its evident topology (§1.14.6). Note (by §8.5) that $\mathcal{T}_\infty = \Psi^0$ is not a new space.

Given an \mathfrak{h} -dependent positive real c , we denote as in §1.14.2 and §5.3 by $c\mathcal{T}_p$ the space of \mathfrak{h} -dependent $T \in \mathcal{T}_p$ for which the seminorms of $c^{-1}T$ are bounded uniformly in \mathfrak{h} . As usual, expressions involving \mathfrak{h}^∞ are to be understood as holding whenever ∞ is replaced by an arbitrary fixed number N .

12.2. Approximate inverses for Δ

Set

$$\Delta_{\mathfrak{h}} := 1 - \mathfrak{h}^2 \sum_{x \in \mathcal{B}(\mathfrak{g})} x^2.$$

By the spectral theory of the self-adjoint operator Δ , the operator $\pi(\Delta_{\mathfrak{h}})$ is invertible, and its inverse has operator norm ≤ 1 . The following lemma allows one to control its inverse via integral operators.

LEMMA 12.1. *For each $N \in \mathbb{Z}_{\geq 0}$ there exist positive reals \mathfrak{h}_0 and C , depending upon G and χ but not upon π , so that the following holds for $0 < \mathfrak{h} \leq \mathfrak{h}_0$.*

Define $b \in S^0$ by $b(\xi) := \langle \xi \rangle^{-N}$. Then,

$$A(\mathfrak{h}) := \Delta_{\mathfrak{h}}^N \text{Op}_{\mathfrak{h}}(b)^2$$

defines an invertible operator on π . Moreover, the operator norms of $A(\mathfrak{h})$ and $A(\mathfrak{h})^{-1}$ are bounded by C .

Proof. By applying the composition formula, we see that the \mathfrak{h} -dependent operator $A(\mathfrak{h})$ belongs to $1 + \mathfrak{h}\Psi^{-1} \subseteq 1 + \mathfrak{h}\Psi^0$. For small enough \mathfrak{h} , it follows by the Neumann lemma that $A(\mathfrak{h})$ is invertible with inverse of operator norm $O(1)$, as required. \square

12.3. Results

Let π be an \mathfrak{h} -dependent irreducible tempered representation of G . We adopt here the convention that implied constants in any asymptotic notation are independent of π and \mathfrak{h} , and must depend continuously upon any symbols under consideration. Recall from §6.1 the definition of $d \in \mathbb{Z}_{\geq 0}$.

THEOREM 12.2. *Let assumptions and conventions be as above. Let $N \in 2\mathbb{Z}_{\geq 0}$ satisfy*

$$N > d.$$

(i) *For $\delta \in [0, 1)$ and $a \in S_\delta^{-N}$, we have*

$$\mathrm{tr}(\mathrm{Op}_h(a)) \ll h^{-d} \langle h\lambda_\pi \rangle^{d-N} \quad (12.1)$$

and

$$h^d \mathrm{tr}(\mathrm{Op}_h(a)) = \int_{\mathcal{O}} a \, d\omega_{\mathcal{O}} + O(h^{1-\delta} \langle h\lambda_\pi \rangle^{d-N-1}). \quad (12.2)$$

In particular, if a is h -independent and π admits the limit orbit (\mathcal{O}, π) , then

$$\lim_{h \rightarrow 0} h^d \mathrm{tr}(\mathrm{Op}_h(a)) = \int_{\mathcal{O}} a \, d\omega. \quad (12.3)$$

For each $j \in \mathbb{Z}_{\geq 0}$, there is a constant coefficient differential operator \mathcal{D}_j on \mathfrak{g}^\wedge of pure degree j , so that for $a \in S_\delta^{-\infty}$ and fixed $J, N' \in \mathbb{Z}_{\geq 0}$,

$$h^d \mathrm{tr}(\mathrm{Op}_h(a)) = \sum_{0 \leq j < J} h^j \int_{\mathcal{O}} \mathcal{D}_j a \, d\omega_{\mathcal{O}} + O(h^{(1-\delta)J} \langle h\lambda_\pi \rangle^{-N'}). \quad (12.4)$$

(ii) $\Delta_h^{-N/2}$ (cf. §12.2) *is trace class, with $\|\Delta_h^{-N/2}\|_1 = \mathrm{tr}(\Delta_h^{-N/2}) \ll h^{-d} \langle h\lambda_\pi \rangle^{d-N}$.*

(iii) *For any operator T on π , the trace norm is majorized as follows:*

$$\|T\|_1 \leq C \langle \lambda_\pi \rangle^{d-N} \|T\|_{\pi^0 \rightarrow \pi^N}, \quad (12.5)$$

where C depends only upon G and N . For $p=1, 2$, we have $\Psi^{-N} \subseteq \mathcal{T}_p$; the normalized map

$$\begin{aligned} \Psi^{-N} &\longrightarrow \mathcal{T}_p, \\ T &\longmapsto \langle \lambda_\pi \rangle^{N-d} T, \end{aligned} \quad (12.6)$$

is continuous, uniformly in π . These conclusions remain valid also for non-tempered π , possibly with a larger value of N .

(iv) *For fixed $x, y \in \mathfrak{U}$ and any $T \in h^0 \Psi^{-\infty}$, we have $\|xTy\|_2 \ll 1$, with continuous dependence upon T .*

(v) *We have*

$$\begin{aligned} \mathrm{Op}_h(S_0^{-N}) &\subseteq h^{-d} \langle h\lambda_\pi \rangle^{d-N} \mathcal{T}_1, \\ \mathrm{Op}_h(S_0^{-N}) &\subseteq h^{-d/2} \langle h\lambda_\pi \rangle^{(d-N)/2} \mathcal{T}_2. \end{aligned}$$

(vi) For $k \in \mathbb{Z}_{\geq 1}$ and $a_1, \dots, a_k \in S_0^{-N}$,

$$\mathrm{Op}_h(a_1) \dots \mathrm{Op}_h(a_k) \in h^{-d} \langle h\lambda_\pi \rangle^{d-N} \mathcal{T}_1, \quad (12.7)$$

$$h^d \mathrm{Op}_h(a_1) \dots \mathrm{Op}_h(a_k) \equiv h^d \mathrm{Op}_h(a_1 \dots a_k) \pmod{h \langle h\lambda_\pi \rangle^{d-N} \mathcal{T}_1}. \quad (12.8)$$

(vii) Fix $\varepsilon > 0$. Let $g \in G$ be an h -dependent element with $\|\mathrm{Ad}(g)\| \leq h^{-1+\varepsilon}$. Let $a \in S_0^\infty$. Then,

$$h^d \mathrm{Op}_h(g \cdot a) \equiv h^d \pi(g) \mathrm{Op}_h(a) \pi(g)^{-1} \pmod{h^\infty \langle h\lambda_\pi \rangle^{d-N} \mathcal{T}_1}. \quad (12.9)$$

Proof. We will frequently apply Theorems 4.5–5.8.

(i) The first assertion reduces to the estimate

$$\sup_{h \in (0,1]} \sup_{\mathcal{O}: d(\mathcal{O})=d} \langle h[\mathcal{O}] \rangle^\varepsilon h^d \int_{\xi \in \mathcal{O}} \langle h\xi \rangle^{-d-\varepsilon} d\omega_{\mathcal{O}}(\xi) < \infty, \quad (12.10)$$

which follows in term from (6.1) and the results in §11.3. The remaining assertions follow by expanding $(j^{-1/2} \chi a^\vee)^\wedge$ using §7.8 and recalling that $j(0)=1$.

(ii) By spectral theory, we may assume that h is sufficiently small. Set

$$b(\xi) := \langle \xi \rangle^{-N/2}.$$

By §12.2, we have

$$\mathrm{tr}(\Delta_h^{-N/2}) \ll \mathrm{tr}(\mathrm{Op}_h(b)^2).$$

By applying the proof of (i) to

$$\mathrm{tr}(\mathrm{Op}_h(b)^2) = \mathrm{tr}(\mathrm{Op}_h(b \star_h b, \chi')),$$

we obtain an adequate estimate for $\mathrm{tr}(\mathrm{Op}_h(b)^2)$.

(iii) Let $T \in \Psi^{-N}$. By (ii) and the inequality

$$\|A\|_1 \leq \|\Delta_h^{-N/2}\|_1 \|\Delta_h^{N/2} A\|_\infty, \quad (12.11)$$

applied with $A := \theta_u(T)$ and $h := 1$, we obtain the estimate (12.5), as well as the required inclusion for $p=1$. We deduce the case $p=2$ via $\|A\|_2 \leq \|A\|_1^{1/2} \|A\|_\infty^{1/2}$.

The necessary input in this argument was the uniform trace class property of $\Delta^{-N/2}$. This is presumably well known, and holds also non-tempered π , as follows, e.g., from the proof of part (i) of Lemma A.3.

(iv) By a similar argument as in (iii).

(v) We must estimate $\|\theta_u(\mathrm{Op}_h(a))\|_p$ for fixed $u \in \mathfrak{U}$ and $p=1, 2$. By differentiating the composition formula, we have

$$\theta_u(\mathrm{Op}_h(a)) \equiv \mathrm{Op}_h(\theta_u(a)) \pmod{h^\infty \Psi^{-\infty}}.$$

By (iii), we thereby reduce to the case $u=1$. Using the identity

$$\|\mathrm{Op}_h(a)\|_2^2 = \mathrm{tr}(\mathrm{Op}_h(a)\mathrm{Op}_h(a)^*) = \mathrm{tr}(\mathrm{Op}_h(a)\mathrm{Op}_h(\bar{a})),$$

the composition formula, and (iii), we reduce further to the case $p=1$, in which it remains to show that $\|\mathrm{Op}_h(a)\|_1 \ll h^{-d} \langle h\lambda_\pi \rangle^{d-N}$. For this we apply (12.11) with $A := \mathrm{Op}_h(a)$ and appeal to (ii) and the consequence $\|\Delta_h^{N/2}\mathrm{Op}_h(a)\|_\infty \ll 1$ of the composition formula.

The remaining results (vi) and (vii) may be proved similarly (for (vii), cf. §8.4). \square

Part III. Gan–Gross–Prasad pairs: geometry and asymptotics

Let k be a field of characteristic zero.

For our purposes, a *Gan–Gross–Prasad pair* (henceforth “GGP pair”) over k is a pair (\mathbf{G}, \mathbf{H}) of algebraic k -groups equipped with

- an inclusion $\mathbf{H} \hookrightarrow \mathbf{G}$ of algebraic k -groups, isomorphic to one of the standard inclusions

$$\mathbf{SO}_n \hookrightarrow \mathbf{SO}_{n+1}, \quad \mathbf{U}_n \hookrightarrow \mathbf{U}_{n+1} \quad \text{or} \quad \mathbf{GL}_n \hookrightarrow \mathbf{GL}_{n+1}, \quad (\text{III.1})$$

and

- an action (called the “standard action”) of \mathbf{G} on some vector space V , i.e., an embedding $\mathbf{G} \hookrightarrow \mathbf{GL}(V)$. See below for complete details.

We make this definition precise in §13.2. The general linear example may be understood as a special case of the unitary example, as we will find very convenient in our proofs. We refer to the first case as the orthogonal case, and the latter two as unitary cases.

The study of such pairs, locally and globally, was initiated (in the special orthogonal case) by Gross and Prasad [GP]. A broader formalism was developed in the paper of Gan, Gross and Prasad [GGP]. The cases of $(\mathrm{SO}_2, \mathrm{SO}_3)$ and $(\mathrm{SO}_3, \mathrm{SO}_4)$ have played an important role in the analytic theory of L -functions for GL_2 . It is therefore very natural to consider the analytic theory of Gan–Gross–Prasad periods in higher rank.

Part III contains several algebraic and analytic preliminaries concerning GGP pairs, many of which may be of independent interest.

The main aim of §13 and §14 is to study in detail certain algebraic properties of the restriction to \mathbf{H} of the adjoint (equivalently, coadjoint) representation of \mathbf{G} . In the language of §1.10, we are studying how H acts on the system of solutions to

$$\xi|_{\mathfrak{h}} = \eta, \quad \xi \in \mathfrak{g}^\wedge \text{ and } \eta \in \mathfrak{h}^\wedge;$$

this is relevant for us because (for k a local field) it models the “asymptotic decomposition” of the restriction to H of a unitary representation of G . The “nice” solutions will turn out to form a smoothly-varying family of \mathbf{H} -torsors. The definition of “nice” is formulated in §14 in terms of the GIT notion of stability and then related in §15 to the absence of “conductor dropping” for the associated Rankin–Selberg L -function. In §16, we study, e.g., how integrals over G -orbits in \mathfrak{g}^\wedge can be disintegrated in terms of integrals over the H -orbits discussed previously. In §17, we apply our results to archimedean GGP pairs. The main output is that we can write the integral over a coadjoint orbit \mathcal{O} for G explicitly in terms of integrals over the H -orbits on \mathcal{O} , with control over how everything varies in families. The main aim of §18 and §19 is then to prove a “quantum analogue” of such integral formulas, involving the asymptotic decomposition of the restriction to H of a tempered irreducible representation of G .

The pictures in §1.7 and §1.10 may usefully illustrate the discussion below.

13. Basic definitions and invariant theory

13.1. Orthogonal groups and unitary groups

As above, let k be a field of characteristic zero. Let k_1 be either a quadratic étale k -algebra or k itself, thus either

$$k_1 = k, \quad k_1/k \text{ is a quadratic field extension,} \quad \text{or} \quad k_1 = k \times k. \quad (13.1)$$

The three cases indicated in (III.1) will arise accordingly. In the third case, k is embedded diagonally in the product.

Let ι denote the involution of k_1 fixing k , and let $(V, \langle - \cdot - \rangle)$ be a k_1 -vector space equipped with a non-degenerate ι -linear symmetric bilinear form. We then denote by $\dim(V)$, $\text{End}(V)$ and $\text{GL}(V)$ the dimension, endomorphisms and automorphisms of V as a k_1 -vector space. We may define the connected automorphism group

$$\mathbf{G} = \mathbf{Aut}(V/k_1, \langle - \cdot - \rangle)^0.$$

It is a k -algebraic group which comes with a standard representation $\mathbf{G} \hookrightarrow \mathbf{GL}(V)$.

We denote as usual by $\mathfrak{g} \hookrightarrow \text{End}(V)$ the Lie algebra of \mathbf{G} . We write

$$[\mathfrak{g}] = \mathfrak{g} // \mathbf{G}$$

for the set of k -points of the GIT quotient (cf. §9.1 for an example). For $x \in \mathfrak{g}$, we write $[x]$ for its image in $[\mathfrak{g}]$.

For $x \in \text{End}(V)$, we denote by $x^* \in \text{End}(V)$ the conjugate-adjoint, defined by means of the rule $\langle u, xv \rangle = \langle x^*u, v \rangle$ for all $u, v \in V$. Then, $x \mapsto x^*$ is ι -linear, and $(xy)^* = y^*x^*$. The Lie algebra \mathfrak{g} consists of those $x \in \text{End}(V)$ that are skew-adjoint: $x^* = -x$.

13.2. Standard inclusions

Retaining the setting of §13.1, fix $e \in V$ for which $k_1 e$ is a free rank-1 k_1 -module on which the form $\langle -, - \rangle$ is non-degenerate, and set $V_H = (k_1 e)^\perp \subseteq V$. One then has the splitting $V = V_H \oplus k_1 e$, which induces an inclusion

$$\mathbf{H} := \mathbf{Aut}(V_H/k_1, \langle -, - \rangle)^0 \hookrightarrow \mathbf{G} := \mathbf{Aut}(V/k_1, \langle \rangle)^0.$$

A GGP pair over k is defined to be such an inclusion, together with the accompanying standard representations.

13.3. Extension to an algebraic closure

The action of G on V is k_1 -linear, and so preserves the k -linear decomposition of V into isotypic subspaces for k_1 . We explicate this decomposition first in the special case that k is algebraically closed, so that either $k_1 = k \times k$ or $k_1 = k$:

(i) (Unitary case) If $k_1 = k \times k$, then we have a decomposition of k -vector spaces

$$V = V^+ \oplus V^-,$$

where k_1 acts on V^\pm via the two projections to k . The form $\langle -, - \rangle$ on $V^+ \times V^-$ is valued in the first factor k of k_1 , and thus induces a k -valued perfect pairing $[-, -]$ between V^+ and V^- ; the Hermitian form is described in terms of $[-, -]$ as

$$\langle v_1^+ + v_1^-, v_2^+ + v_2^- \rangle = ([v_1^+, v_2^-], [v_2^+, v_1^-]) \in k \times k.$$

Moreover, $G \cong \text{GL}(V^+)$ identifies with the set of pairs $(g, {}^t g^{-1})$ in $\text{GL}(V^+) \times \text{GL}(V^-)$. The vector e is given by $e^+ + e^-$, with $e^+ \in V^+$ and $e^- \in V^-$ satisfying $\langle e^+, e^- \rangle \neq 0$.

There is a corresponding splitting $V_H = V_H^+ \oplus V_H^-$.

(ii) (Orthogonal case) If $k_1 = k$, so that $G = \text{SO}(V)$, then we set

$$V^+ := V^- := V.$$

In general, we fix an algebraic closure \bar{k} of k and apply the above considerations with $(V \otimes_k \bar{k}, \bar{k}, k_1 \otimes_k \bar{k})$ playing the role of “ (V, k, k_1) ”. We obtain in this way G -invariant \bar{k} -subspaces V^\pm of $V \otimes_k \bar{k}$.

13.4. Eigenvalues and eigenvectors

13.4.1. Invariants in terms of eigenvalues

We recall the description of the ring of G -invariant polynomial functions on \mathfrak{g} (“invariant functions” for short). Each $x \in \mathfrak{g}$ gives a k_1 -linear endomorphism of V and thus has a characteristic polynomial in $k_1[t]$. Since $x^* = -x$, the n th coefficient of the characteristic polynomial of x belongs to

$$k_1^\pm = \{x \in k_1 : \iota(x) = \pm x\},$$

where \pm depends upon the parity of n . The spaces k_1^\pm are at most 1-dimensional over k ; fixing bases, we obtain invariant functions $\mathfrak{g} \rightarrow k$. These freely generate the ring of invariant functions in all cases except when G is even orthogonal, in which the Pfaffian $\text{pf}(x)$ gives another invariant satisfying $\text{pf}(x)^2 = \det(x)$.

We denote by $\text{ev}(x)$ the multiset of roots in \bar{k} of the characteristic polynomial for the \bar{k} -linear action of $x \in \mathfrak{g}$ on V^+ , except that in the odd orthogonal case, we subtract the “obvious root” zero with multiplicity 1, noting that it always occurs. Thus $\#\text{ev}(x) = n$, where $n = \dim(V)$ in the unitary, linear and even orthogonal cases and $n = \dim(V) - 1$ in the odd orthogonal case. We may descend ev to a function on $[\mathfrak{g}]$. The map

$$[\mathfrak{g}] \ni [x] \longmapsto \begin{cases} \text{the ordered pair } (\text{ev}(x), \text{pf}(x)), & \text{in the even orthogonal case,} \\ \text{ev}(x), & \text{otherwise,} \end{cases}$$

is then well defined and injective, with readily characterized image.

13.4.2. Geometric characterization

We may describe the set underlying the multiset $\text{ev}(x)$ more geometrically as follows.

LEMMA 13.1. *Let $x \in \mathfrak{g}$ and $c \in k$. The following are equivalent:*

- (i) $c \in \text{ev}(x)$;
- (ii) x has an isotropic eigenvector in V^+ with eigenvalue c ;
- (iii) x has an isotropic eigenvector in V^- with eigenvalue $-c$.

Proof. The equivalence of (ii) and (iii) follows from the definitions of V^+ and V^- .

For the equivalence between (i) and (ii), note that any eigenvector with non-zero eigenvalue is necessarily isotropic. The required equivalence is thus clear if $c \neq 0$. Suppose henceforth that $c=0$. In the unitary case, the space V^+ is itself isotropic, so (ii) just says that x has zero as an eigenvalue on V^+ , which is equivalent to (i).

It remains to show that (i) and (ii) are equivalent in the orthogonal case with $c=0$. Let $W \subseteq V^+$ denote the generalized zero-eigenspace for x . Because the generalized eigenspaces with eigenvalues λ and $-\lambda$ are in perfect pairing with one another, we have $\dim(W) \equiv \dim(V^+) \pmod{2}$; moreover, the restriction of $\langle -, - \rangle$ to W is non-degenerate. Recall that $0 \in \text{ev}(x)$ in the even orthogonal case when $\dim(W) \geq 1$ and in the odd orthogonal case when $\dim(W) \geq 2$. We see in either case that if $0 \in \text{ev}(x)$, then $\dim(W) \geq 2$; since W is non-degenerate, it thus contains a non-zero isotropic vector. Thus (i) implies (ii). For the converse implication, note that if x has an isotropic eigenvector with eigenvalue zero, then W contains a non-zero isotropic subspace, whence $\dim(W) \geq 2$; to complete the proof, it is thus enough to check the following fact.

CLAIM. *If $(W, \langle -, - \rangle)$ is a non-degenerate quadratic space of dimension at least 2, and $x \in \text{End}(W)$ is skew-symmetric and nilpotent, then there is a non-zero isotropic vector in the kernel of x .*

This is clear if $x=0$. Otherwise, there is a non-zero vector $v \in \ker(x)$ which belongs to the image of x , say $v=xu$, and then v does the trick:

$$\langle v, v \rangle = -\langle u, xv \rangle = 0.$$

(Said differently, if x in the claim is non-zero, then it arises from the action of $\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ under some homomorphism $\mathfrak{sl}_2 \rightarrow \mathfrak{so}_W$, and we choose v to be a highest weight vector.) \square

13.4.3. Regular elements

We note also a related characterization of regular elements.

LEMMA 13.2. *Let $x \in \mathfrak{g}$ and suppose that every isotropic subspace of V^+ on which x acts by a scalar is at most 1-dimensional. In the orthogonal case, assume moreover that the kernel of x is at most 2-dimensional. Then, x is regular.*

Proof. In the unitary case, this is well-known: the word “isotropic” can be ignored since the form vanishes on V^+ , and an $n \times n$ matrix is regular if and only if each eigenvalue corresponds to a single Jordan block. Accordingly, we focus on the orthogonal case.

Write

$$V^+ = \bigoplus_{\lambda} V_{\lambda}^+ = \left(\bigoplus_{\{\lambda \neq 0\}/\pm} V_{\lambda}^+ \oplus V_{-\lambda}^+ \right) \oplus V_0^+ \quad (13.2)$$

for the decomposition into generalized eigenspaces under x . To prove regularity, we must show that the dimension of the centralizer of x is minimal. Since any element of this centralizer preserves each summand in (13.2), the minimal centralizer dimension for SO_{2n} or SO_{2n+1} is n , and the summands $V_\lambda^+ \oplus V_{-\lambda}^+$ are even-dimensional, we reduce to showing the following:

(i) x induces a regular element of the orthogonal group of $V_\lambda^+ + V_{-\lambda}^+$ for each $\lambda \neq 0$:

Our condition implies that $x - \lambda$ is a regular nilpotent on V_λ^+ , so its centralizer in $\mathrm{GL}(V_\lambda^+)$ has minimal dimension, which implies the above claim.

(ii) x induces a regular element of the orthogonal group of V_0^+ . In what follows we denote this restriction simply by x .

By the Jacobson–Morozov theorem, x is the image of a non-zero nilpotent in \mathfrak{sl}_2 under a homomorphism $\mathfrak{sl}_2 \rightarrow \mathfrak{so}(V_0^+)$. The assumption on $\ker(x)$ implies V_0^+ decomposes into a sum of at most two irreducible representations $U \oplus U'$. If U and U' are each of dimension ≥ 2 , then we see by considering weights that $\ker(x)$ is contained in $\mathrm{image}(x)$; the skew-symmetry of x then implies that $\ker(x)$ is isotropic and of dimension ≥ 2 , contrary to hypothesis. We may thus suppose that $\dim(U') \leq 1$. We may also suppose that $\dim(U) \geq 2$; if not, then x is the zero element of \mathfrak{so}_1 or \mathfrak{so}_2 and is regular.

Now $\dim(U)$ is necessarily odd-dimensional, otherwise \mathfrak{sl}_2 does not preserve an orthogonal form on it. It is now a routine computation with \mathfrak{sl}_2 -representations to compute that the centralizer of x in $\mathfrak{so}(V_0^+)$ has the correct dimension; explicitly, this centralizer is given by

$$\begin{cases} x, x^3, \dots, x^{\dim(U)-2}, & \text{if } U' = \{0\}, \\ x, x^3, \dots, x^{\dim(U)-2}, & \text{if } \dim(U') = 1, \end{cases}$$

where y is a skew-symmetric transformation sending a generator of U' to a non-zero highest weight vector in U and mapping U to U' . \square

13.5. The spherical property

Given a GGP pair (\mathbf{G}, \mathbf{H}) , we will have occasion to consider the diagonal inclusion

$$\mathbf{H} \longrightarrow \mathbf{G} \times \mathbf{H} \tag{13.3}$$

This is a *spherical pair*, in the following well-known sense (see, e.g., [BP3, §6.4, p. 142, equation (1)]).

LEMMA 13.3. \mathbf{H} has an open orbit, with trivial stabilizer, on the flag variety of

$$\mathbf{G} \times \mathbf{H}.$$

14. Stability

Recall that k is a field of characteristic zero.

14.1. Preliminaries

Let \mathbf{H} be a reductive algebraic k -group which acts (algebraically) on an affine k -variety \mathbf{M} . Recall, following Mumford, that an element $x \in \mathbf{M}(k)$ is called **H-stable** if

- (i) the stabilizer of x in \mathbf{H} is finite, and
- (ii) the orbit $\mathbf{H} \cdot x$ is closed.

We summarize some background from geometric invariant theory.

LEMMA 14.1. *Suppose we are in the setting just described with k algebraically closed. Let $M//H$ denote the spectrum of the ring of H -invariant regular functions on M , and $\phi: M \rightarrow M//H$ the canonical map. Then, $M//H$ is an affine variety, and ϕ is surjective. Each H -invariant morphism with domain M factors uniquely through ϕ . If M is irreducible, then so is $M//H$.*

Let $M^s \subseteq M$ denote the subset of **H-stable** elements. Then, M^s and

$$(M//H)^s := \phi(M^s)$$

are open (but possibly empty). If the isotropy group of every point in M^s inside H is trivial, then the induced map $\phi^s: M^s \rightarrow (M//H)^s$ is a principal H -bundle (indeed, it is locally trivial in the étale topology).

For the last statement, we refer in particular to [MFK, Proposition 0.9] and [Mi, p. 120].

14.1.1. Moment map interpretation

If \mathbf{M} is smooth, then the action of \mathbf{H} on \mathbf{M} induces an action on the cotangent bundle $T^*\mathbf{M}$ and, by duality, an H -equivariant moment map

$$\Phi: T^*M \longrightarrow \mathfrak{h}^*.$$

We then verify readily that condition (i) in the above definition is equivalent to

- (i') The moment map Φ induces, by differentiation, a *surjective* map $T_x^*M \rightarrow \mathfrak{h}^*$.

14.2. Characterization for GGP pairs

Let $\mathbf{H} \hookrightarrow \mathbf{G} \hookrightarrow \mathbf{GL}(V)$ be a GGP pair over k . We retain the accompanying notation of §13.2. We denote by $E \in \text{End}(V)$ the orthogonal projection with image V_H . For $x \in \text{End}(V)$, we denote by $x_H \in \text{End}(V_H)$ the restriction of ExE to V_H . We then have a commutative H -equivariant diagram

$$\begin{array}{ccc} \mathfrak{g} & \xrightarrow{\cong} & \mathfrak{g}^* \\ x \mapsto x_H \downarrow & & \downarrow \text{restriction} \\ \mathfrak{h} & \xrightarrow{\cong} & \mathfrak{h}^* \end{array}$$

in which the horizontal isomorphisms are induced by the trace pairing on $\text{End}(V)$. Via this diagram, the definitions and results below admit equivalent formulations in terms of the coadjoint representations. We sometimes also write ξ_H for the restriction of $\xi \in \mathfrak{g}^*$ to \mathfrak{h}^* .

THEOREM 14.2. *Let $x \in \mathfrak{g}$. The following are equivalent:*

- (i) x is \mathbf{H} -stable.
- (ii) $\text{ev}(x) \cap \text{ev}(x_H) = \emptyset$.

Definition 14.3. We say that a pair $(\lambda, \mu) \in [\mathfrak{g}] \times [\mathfrak{h}]$ is *stable* if $\text{ev}(\lambda) \cap \text{ev}(\mu) = \emptyset$.

This condition is equivalent to asking that the multiset sum $\text{ev}(\lambda) + \text{ev}(-\mu)$ not contain zero. We will later (§15) interpret that sum in terms of the Satake parameters of an associated L -function. In this way, the failure of stability will be related to the situation where the conductor of this L -function drops.

We note that the set $\{(\lambda, \mu) : (\lambda, \mu) \text{ is stable}\}$ is dense and open in $[\mathfrak{g}] \times [\mathfrak{h}]$. Moreover, for each $\lambda \in [\mathfrak{g}]$, the set $\{\mu : (\lambda, \mu) \text{ is stable}\}$ is dense and open in $[\mathfrak{h}]$.

We turn to the proof of Theorem 14.2. We may and shall assume that k is algebraically closed, so that either $k_1 = k \times k$ (the *unitary case*) or $k_1 = k$ (the *orthogonal case*). A key ingredient is the following geometric characterization of when x and x_H share an eigenvalue.

LEMMA 14.4. *Let $x \in \mathfrak{g}$ and $c \in k$. The following are equivalent:*

- (i) $c \in \text{ev}(x) \cap \text{ev}(x_H)$;
- (ii) x has either
 - an isotropic eigenvector in V_H^+ with eigenvalue c , or
 - an isotropic eigenvector in V_H^- with eigenvalue $-c$.

To see the significance of the two cases, consider the simple example of $G = \mathbf{GL}_2(k)$: if a 2×2 matrix $(a_{ij})_{i,j} \in \mathfrak{g}$ has an eigenvalue that coincides with the upper-left entry a_{11} , then either a_{12} or a_{21} must vanish, corresponding to the two cases above.

Proof. Recall from §13.4.2 that c belongs to $\text{ev}(x)$ exactly when there is an isotropic vector in V^\pm with eigenvalue $\pm c$. Thus, (ii) implies (i).

Conversely, suppose $c \in \text{ev}(x) \cap \text{ev}(x_H)$. There is then

- an isotropic eigenvector $v \in V_H^+$ for x_H with eigenvalue c , and also
- an isotropic eigenvector $w \in V^-$ for x with eigenvalue $-c$.

We have

$$\langle (x-c)v, w \rangle = -\langle v, (x+c)w \rangle = 0. \quad (14.1)$$

and thus either

- $xv=cv$, so that v is an isotropic eigenvector for x in V_H^+ , or
- $xv \neq cv$, so that e^+ (the V^+ component of e , as in §13.2) is a multiple of $(x-c)v$. By (14.1), it follows that $\langle e^+, w \rangle = 0$, hence w is an isotropic eigenvector for x in V_H^- . \square

Recall the Hilbert–Mumford criterion: x is not \mathbf{H} -stable if and only if there is a non-trivial 1-parameter subgroup $\gamma: \mathbb{G}_m \rightarrow \mathbf{H}$ with respect to which x has only non-negative weights; we say then that γ *witnesses the failure of stability for x* . Equivalently, consider the decomposition of V into weight spaces for γ :

$$V = \bigoplus_{i \in \mathbb{Z}} V_i, \quad \gamma(t) \text{ acts on } V_i \text{ by the scalar } t^i.$$

Then,

$$x \text{ is not stable} \iff \text{there exists } \gamma \neq 1 \text{ so that } xV_i \subseteq \bigoplus_{j \geq i} V_j \text{ for all } i. \quad (14.2)$$

To prove Theorem 14.2, it is enough to show that the following are equivalent:

- (a) x has an isotropic eigenvector in $V_H^+ \cup V_H^-$;
- (b) Some γ as above witnesses the failure of stability for x .

(a) implies (b): We assume that x has an isotropic eigenvector $v_1 \in V_H^+$; the other case is handled analogously. By a standard lemma, we may choose an isotropic vector $v_2 \in V_H^-$ for which $\langle v_1, v_2 \rangle = 1$; we refer to [Se, p.29] in the orthogonal case $k_1 = k$, while in the unitary case $k_1 = k \times k$, we use that the spaces V_H^\pm are themselves totally isotropic and in duality. The subspace $kv_1 \oplus kv_2$ of V_H is then non-degenerate, and so, with

$$W := (kv_1 \oplus kv_2)^\perp,$$

we have

$$V = kv_1 \oplus W \oplus kv_2. \quad (14.3)$$

Write $xv_1 = cv_1$. Then, $\langle (x+c)W, v_1 \rangle = 0$, so $xW \subseteq W + v_1^\perp = v_1 \oplus W$, and thus the matrix of x with respect to (14.3) is upper-triangular. The 1-parameter subgroup γ of H given by $\gamma(t)v_1 = tv_1$, $\gamma(t)|_W = 1$ and $\gamma(t)v_2 = t^{-1}v_2$ then witnesses the failure of stability for x .

(b) implies (a): Suppose γ witnesses the failure of stability for x . Since γ is non-trivial, some V_i with $i \neq 0$ is non-trivial. The spaces $V_{>0} := \bigoplus_{i>0} V_i$ and $V_{<0} := \bigoplus_{i<0} V_i$ are totally isotropic, contained in V_H , and in duality, hence both non-zero. The non-zero x -stable isotropic subspace $V_{>0}$ of V_H thus contains an isotropic eigenvector v for x . In the unitary case, we further split $v = v^+ + v^-$ to get an isotropic eigenvector in $V_H^+ \cup V_H^-$.

14.3. Fibers of $x \mapsto ([x], [x_H])$

We retain the setting of §14.2, and assume that k is algebraically closed.

THEOREM 14.5. *The morphism of varieties*

$$\begin{aligned} \{H\text{-stable } x \in \mathfrak{g}\} &\longrightarrow \{\text{stable } (\lambda, \mu) \in [\mathfrak{g}] \times [\mathfrak{h}]\} \\ x &\longmapsto ([x], [x_H]) \end{aligned} \tag{14.4}$$

defines a principal H -bundle. In particular, for stable $(\lambda, \mu) \in [\mathfrak{g}] \times [\mathfrak{h}]$, the fiber

$$\mathcal{O}^{\lambda, \mu} := \{x \in \mathfrak{g} : [x] = \lambda \text{ and } [x_H] = \mu\}, \tag{14.5}$$

which consists entirely of H -stable elements, by the prior theorem, is an H -torsor.

As before, it suffices to consider the first map (14.4). We may and shall assume that k is algebraically closed. We require several lemmas.

LEMMA 14.6. *The map $\mathfrak{g} \rightarrow [\mathfrak{g}] \times [\mathfrak{h}]$ given by $x \mapsto ([x], [x_H])$ is surjective.*

Proof. Recall from §13.5 that the quotient \mathbf{X} of $\mathbf{G} \times \mathbf{H}$ by the diagonally embedded \mathbf{H} is a spherical variety, i.e., each Borel subgroup of $\mathbf{G} \times \mathbf{H}$ has an open orbit on \mathbf{X} . In particular, fixing such a Borel \mathbf{B} , the associated moment map

$$T^*X \longrightarrow \mathfrak{g}^* \oplus \mathfrak{h}^*$$

has image which surjects onto \mathfrak{b}^* , because \mathbf{B} acts simply transitively on an open subset of \mathbf{X} .

This readily implies that the composition $T^*X \rightarrow [\mathfrak{g}^*] \times [\mathfrak{h}^*]$ is surjective. Indeed, if \mathfrak{t} is the torus quotient of \mathfrak{b} , and $\lambda \in \mathfrak{t}^* \hookrightarrow \mathfrak{b}^*$, all extensions of λ from \mathfrak{b}^* to $\mathfrak{g}^* \oplus \mathfrak{h}^*$ have the same image in $[\mathfrak{g}^*] \times [\mathfrak{h}^*]$; this common image corresponds to the class of λ in \mathfrak{t}^* modulo the Weyl group.

The map $T_{x_0}^*X \rightarrow [\mathfrak{g}^*] \times [\mathfrak{h}^*]$ is thus also surjective, where $x_0 \in X$ belongs to the open orbit. The image of the moment map restricted to the fiber $T_{x_0}^*X$ is thus

$$\text{orthogonal complement of } \text{diag } \mathfrak{h} \text{ inside } \mathfrak{g}^* \oplus \mathfrak{h}^*,$$

which is precisely the set $\{(\xi, -\xi_H) : \xi \in \mathfrak{g}^*\}$. Thus, every element of $[\mathfrak{g}^*] \times [\mathfrak{h}^*]$ is of the form $[\xi] \times [-\xi_H]$; negating the second coordinate gives the result. \square

LEMMA 14.7. *Let $x \in \mathfrak{g}$ be H -stable. Then, x and x_H are regular.*

Proof. Suppose, to the contrary, that either x or x_H is irregular. We divide into cases using §13.4.3 and, in each case, produce an isotropic eigenvector $v \in V_H^+$, contradicting stability by Lemma 14.4:

(a) x has an isotropic eigenspace $W \subset V^+$ of dimension ≥ 2 . Take v to be any non-zero element of $W \cap V_H^+$.

(a)' x_H has an isotropic eigenspace $W \subset V_H^+$ of dimension ≥ 2 . Take v to be any non-zero element of the kernel of $W \xrightarrow{x} V^+ / V_H^+$.

(b) there is a subspace W of V^+ , of dimension ≥ 3 , on which x is identically zero. Take v to be a non-zero isotropic element of $W \cap V_H^+$; it is possible because this space is at least 2-dimensional.

(b)' there is a subspace W of V_H^+ , of dimension ≥ 3 , on which x_H is identically zero. Take v to be a non-zero isotropic element of the kernel of $W \xrightarrow{x} V^+ / V_H^+$; it is possible because this space is at least 2-dimensional. \square

The proof of Theorem 14.5 requires also a further stability characterization.

LEMMA 14.8. *The equivalent conditions (i) and (ii) of Theorem 14.2 are also equivalent to the following:*

(iii) *The $k_1[x]$ -module*

$$k_1[x]e \subset V$$

generated by e (see §13.2) is

- *all of V in the unitary cases;*
- *a non-degenerate subspace of codimension ≤ 1 in the orthogonal cases.*

Remark 14.9. Rallis and Schiffmann [RS, Theorems 6.1 and 17.1] obtained a related⁽⁵⁾ equivalence. See also [BGW] and [Z1].

Proof. Observe that

$$U = (k_1[x]e)^\perp \subset V.$$

is a $k_1[x]$ -stable subspace of V_H which is non-degenerate precisely when $k_1[x]e$ is. In view of Lemma 14.4, it is enough to show that the following are equivalent:

- (a) x has no isotropic eigenvector in $V_H^+ \cup V_H^-$;
- (b) U is trivial in unitary cases, and non-degenerate of dimension ≤ 1 in orthogonal cases.

⁽⁵⁾ We note that in the orthogonal case, our results do not exactly agree with theirs, due to a slight inaccuracy in the latter: the first paragraph of the proof of [RS, Theorem 17.1] suggests that the Lie algebra of the orthogonal group of a non-trivial quadratic space is non-trivial, which fails when the latter is 1-dimensional.

If (b) holds, then U is anisotropic; since any eigenvector in $V_H^+ \cup V_H^-$ belongs to U , it follows that (a) holds.

Conversely, assuming (a), we proceed separately in the unitary and orthogonal cases:

- In the unitary case, the spaces V^\pm are isotropic. Since x is k_1 -linear, our hypothesis implies that the spaces U^\pm contain no eigenvectors for x . It follows that $U^\pm = \{0\}$ and thus $U = \{0\}$.

- In the orthogonal case, there is a maximal isotropic subspace X of U stabilized by $x|_U$. If $X \neq \{0\}$, then $x|_U$ has an eigenvector in X , hence an isotropic eigenvector, contrary to our hypothesis. Thus $X = \{0\}$. Therefore $\dim(U) \leq 1$ and, if U is non-zero, then the quadratic form must be non-zero on it. \square

We finally prove Theorem 14.5. Consider the unique morphism j fitting into the commutative diagram

$$\begin{array}{ccc} \mathfrak{g}^s & & \\ \downarrow & \searrow^{x \mapsto ([x], [x_H])} & \\ \mathfrak{g}^s // H & \xrightarrow{j} & ([\mathfrak{g}] \times [\mathfrak{h}])^s, \end{array}$$

where a superscripted s denotes the subset of stable elements in the sense of either Theorem 14.2 or the definition that follows it.

By Lemma 14.6 (and again Theorem 14.2), j is surjective. We will show in addition, with notation as in (14.5), that if $(\lambda, \mu) \in ([\mathfrak{g}] \times [\mathfrak{h}])^s$, then

$$\text{if } x, y \in \mathcal{O}^{\lambda, \mu}, \text{ there is a unique } h \in H \text{ with } h \cdot x = y.$$

This implies that j is injective; then (by Zariski's main theorem, using characteristic zero) j will be an isomorphism. It also implies (see Lemma 14.1) that the map $\mathfrak{g}^s \rightarrow \mathfrak{g}^s // H$ is a principal H -bundle.

By Lemma 14.7, x and y belong to the unique regular (open) G -orbit \mathcal{O} contained in the fiber \mathcal{O}^λ above λ .

Therefore, there exists $g_0 \in G$ such that $g_0 \cdot y = x$. Set $e' := g_0 e$. From $[x_H] = [y_H]$, we deduce that

$$[x_H] = [x_{H'}], \tag{14.6}$$

where we define $x_{H'} := E' x E'$, with E' the orthogonal projection onto the orthocomplement of e' . Since H is the G -stabilizer of e , it will suffice to show the following.

CLAIM. There is a unique g in the G -stabilizer of x for which $ge' = e$.

Indeed, the element $h := gg_0$ then belongs to H and satisfies $h \cdot y = x$, and any such h arises from some g as in the claim.

We show first that, for all $n \in \mathbb{Z}_{\geq 0}$,

$$\langle x^n e, e \rangle = \langle x^n e', e' \rangle. \quad (14.7)$$

For this, it suffices to show that the formal power series

$$\sum_{n \in \mathbb{Z}_{\geq 0}} t^n \langle x^n e, e \rangle = \langle e, (1-tx)^{-1} e \rangle$$

is unchanged by replacing (x, e) with (x, e') . To that end, it suffices in view of (14.6) to establish the identity

$$\frac{\langle e, (1-tx)^{-1} e \rangle}{\langle e, e \rangle} = \frac{\det(1-tx_H)}{\det(1-tx)}. \quad (14.8)$$

For this, we extend e to an orthogonal k_1 -basis $e=e_1, e_2, \dots, e_n$ of V . The left-hand side of (14.8) is then the upper-left matrix entry of $(1-tx)^{-1}$. The identity (14.8) follows from Cramer's rule, noting that the matrix of $(1-tx_H)$ with respect to the basis e_2, \dots, e_n of V_H is the lower-right $(n-1) \times (n-1)$ submatrix of the matrix of $(1-tx)$ with respect to e_1, \dots, e_n .

Consider the submodules $W := k_1[x]e$ and $W' := k_1[x]e'$ of V . Let U and U' denote the orthogonal complements in V of W and W' , respectively. By Lemma 14.8, the spaces U and U' are non-degenerate, and so $V = W \oplus U = W' \oplus U'$. Since x is skew-adjoint, it preserves these decompositions. By (14.7) there is a unique isometric $k_1[x]$ -equivariant isomorphism $g: W' \rightarrow W$ for which $ge' = e$. In particular, $\dim(W) = \dim(W')$, so that

$$\delta := \dim(U) = \dim(U').$$

If $\delta \neq 0$, then we are in the orthogonal case with $\delta=1$. By Witt's theorem, there is an extension of g to an isometric isomorphism $V \rightarrow V$. There are two such extensions, which may be obtained from one another by composing with the non-trivial element of the orthogonal group of the line U . There is thus a unique extension which belongs to G (the *connected component* of the orthogonal group of V). This extension remains x -equivariant—indeed, $gU = U'$, and x acts on the 1-dimensional spaces U and U' by zero. The proof of the claim is thus complete in this case.

Suppose now that $\delta=0$, so that $g: V \rightarrow V$ is the unique isometric x -equivariant morphism for which $ge' = e$. If we are not in the orthogonal case, then g defines an element of G , and so the proof of the claim is likewise complete. What remains to be checked is that in the orthogonal case, g belongs to G , or equivalently, either that $\det(h)=1$ or $\det(g)=1$. To do this, we must use the one piece of information not used to date, namely, that not merely eigenvalues but also Pfaffians match.

Consider first the case that $\dim(V)$ is odd. Then,

$$\text{pf}(x_H) = \text{pf}(h \cdot y_H) = \det(h)\text{pf}(y_H) = \det(h)\text{pf}(x_H),$$

so if $\text{pf}(x_H) \neq 0$, then $\det(h) = 1$. Otherwise, $\text{pf}(x_H) = 0$. Then, $0 \in \text{ev}(x_H)$, so the kernel of x_H is non-zero, necessarily even-dimensional, and so of dimension ≥ 2 . Thus, the kernel of x contains a non-zero element $v \in V_H$ (not necessarily isotropic). Clearly, $x^n v$ is orthogonal to e for all $n \in \mathbb{Z}_{\geq 0}$, and hence $v \in U$, contrary to our assumption that $\delta = 0$.

If $\dim(V)$ is even, then we may argue as above that $\det(g) = 1$ (since $g \cdot x = x$) unless $\text{pf}(x) = 0$. In the case $\text{pf}(x) = 0$, the kernel of x is again even-dimensional, and thus contains a non-zero element $v \in V_H$; as above, this contradicts the assumption that $\delta = 0$. We thereby deduce as required that $\det(g) = 1$.

This completes the proof of the claim, hence of Theorem 14.5.

15. Satake parameters and L -functions

In this section we show that the notion of stability is closely related to the analytic notion of conductor dropping. This is not used in the proof of our main result but, of course, is helpful in interpreting it.

15.1. Local Langlands and infinitesimal character

Let us first recall the relationship between the Langlands parameters and infinitesimal characters (see, e.g., [ALTV, §6] or [Bo1, §11]). Let \mathbf{G} be a reductive group over \mathbb{R} . Let π be an irreducible admissible representation of G . The local Langlands parameterization attaches to π a conjugacy class of representations of the real Weil group:

$$\phi_\pi: W_{\mathbb{R}} \longrightarrow \underbrace{G^\vee \rtimes \text{Gal}(\mathbb{C}/\mathbb{R})}_{{}^L G}.$$

Here G^\vee is the complex dual group of \mathbf{G} , and the right-hand side defines the Langlands dual ${}^L G$.

Now restriction to $\mathbb{C}^* = W_{\mathbb{C}} \subset W_{\mathbb{R}}$ gives a homomorphism $\phi_\pi^0: \mathbb{C}^* \rightarrow G^\vee$, which we may express uniquely as

$$\phi_\pi^0(e^t) = \exp(t\lambda_\pi + \bar{t}\mu_\pi)$$

for some commuting elements λ_π and μ_π in \mathfrak{g}^\vee , the Lie algebra of G^\vee .

If π is tempered, then the image of ϕ_π is bounded, which imposes the constraint

$$\mu_\pi = -\bar{\lambda}_\pi, \quad \text{thus } \phi_\pi^0(e^t) = \exp(2i\text{Im}(t\lambda_\pi)), \quad (15.1)$$

where we write $t\lambda_\pi = \text{Re}(t\lambda_\pi) + i\text{Im}(t\lambda_\pi)$.

Example 15.1. If $G = \mathrm{GL}_n(\mathbb{R})$ and π is a principal series representation with parameters $i\nu_1, \dots, i\nu_n \in i\mathbb{R}$, then one may choose ϕ_π so that $\lambda_\pi = \mu_\pi = \mathrm{diag}(i\nu_1, \dots, i\nu_n)$. If $G = \mathrm{GL}_2(\mathbb{R})$ and π factors through the discrete series representation of $\mathrm{PGL}_2(\mathbb{R})$ of lowest weight k , then one may arrange that

$$\lambda_\pi = -\mu_\pi = \mathrm{diag}\left(\frac{1}{2}(k-1), \frac{1}{2}(1-k)\right).$$

We observe next that one may identify

$$\mathfrak{g}^\vee // G^\vee \simeq [\mathfrak{g}_\mathbb{C}^*]. \quad (15.2)$$

Indeed, if we fix maximal tori $T^\vee \subset G^\vee$ and $\mathbf{T}_\mathbb{C} \subset \mathbf{G}$, then T^\vee and $\mathbf{T}_\mathbb{C}$ are dual, canonically up to the action of the Weyl group. For complex tori T_1 and T_2 , an identification of T_1 with the dual of T_2 identifies the Lie algebra of T_1 with the complex linear dual of the Lie algebra for T_2 . So, there is a canonical identification

$$\mathbb{C}[(\mathfrak{t}^\vee)^*]^W \simeq \mathbb{C}[\mathfrak{t}_\mathbb{C}]^W$$

which induces (15.2).

LEMMA 15.2. *The identification (15.2) carries λ_π to the infinitesimal character of π .*

The lemma justifies our notational abuse of using the same symbol λ_π above as we had in §9 for the infinitesimal character.

Proof. Recall (§9.6) that the representation π is a summand of the unitarily normalized parabolic induction from a discrete series representation σ on the Levi factor \mathbf{M} of a parabolic subgroup $\mathbf{P} \subset \mathbf{G}$. The Langlands parameters of σ and π are related by means of the natural inclusion of the L -group of \mathbf{M} into the L -group of \mathbf{G} , and the natural map

$$[\mathfrak{m}_\mathbb{C}^*] \longrightarrow [\mathfrak{g}_\mathbb{C}^*]$$

carries the infinitesimal character of σ to that of π . The two maps (of dual groups, and of dual Lie algebras) are compatible with reference to (15.2).

We thereby reduce to the case of a discrete series representation, and then (by the characterization of the local Langlands correspondence for discrete series via infinitesimal characters; see [Bo1, §11.2]) we further reduce to the case of a real torus \mathbf{S} . This follows from [Bo1, §9.3, equation (2)] (note the misprint: the second occurrence of $\sigma \cdot x$ should be $\sigma \cdot \bar{x}$); see also [AV, §6]. \square

15.2. GGP pairs

For simplicity, we restrict ourselves to the case $K=\mathbb{R}$, and leave the straightforward extensions to $K=\mathbb{C}$ to the reader.

Let (\mathbf{G}, \mathbf{H}) now be a GGP pair, in the sense of §13, over $K=\mathbb{R}$; let K_1/K be the associated K -algebra, and let $n = \lceil \frac{1}{2} \dim_K(V) \rceil$, where V is the associated K_1 -vector space. Thus,

$$\begin{aligned} \mathbf{G} &= (\text{form of } \mathrm{SO}_{2n} \text{ or } \mathrm{SO}_{2n+1} \text{ or } \mathrm{GL}_n \text{ over } K), \\ \mathbf{H} &= (\text{form of } \mathrm{SO}_{2n-1} \text{ or } \mathrm{SO}_{2n} \text{ or } \mathrm{GL}_{n-1} \text{ over } K). \end{aligned}$$

15.2.1. Dual Lie algebra

The K -group \mathbf{G} admits a representation by K_1 -linear automorphisms of V . The form

$$x, y \in \mathfrak{g} \longmapsto \mathrm{tr}_V(xy)$$

is actually K -valued and non-degenerate; it identifies $\mathfrak{g} \simeq \mathfrak{g}^*$ (duality of real vector spaces).

We may assign to each $x \in \mathfrak{g}$ a multiset $\mathrm{ev}(x)$ of complex numbers—namely, the multiset of eigenvalues of x in the standard representation, where we remove zero with multiplicity 1 in the odd orthogonal case. By means of the identification above, we may also make sense of $\mathrm{ev}(x)$ for $x \in \mathfrak{g}^*$; this is a set of size $2n$, $2n$ and n in the three cases above. Similarly, $\mathrm{ev}(y)$ for $y \in \mathfrak{h}^*$ is a set of size $2n-2$, $2n$ and $n-1$ in the three cases above.

15.3. Rankin–Selberg L -function for GGP pairs

The tensor product of standard representations on $H^\vee \times G^\vee$ extends to a homomorphism

$$L(H^\vee \times G^\vee) \longrightarrow \mathrm{Sp}_{2n-2} \times \mathrm{O}_{2n} \text{ or } \mathrm{O}_{2n} \times \mathrm{Sp}_{2n} \text{ or } (\mathrm{GL}_{n-1} \times \mathrm{GL}_n(\mathbb{C})) \rtimes \{\pm 1\}.$$

We define the Rankin–Selberg representations ρ to be the natural representations of the right-hand side of dimensions $(2n-2) \cdot 2n$, $(2n)^2$ and $2n(n-1)$, respectively; in the last case we induce the standard representation of $(\mathrm{GL}_{n-1} \times \mathrm{GL}_n(\mathbb{C}))$ to the disconnected group.

LEMMA 15.3. *The correspondence (15.2)*

$$A^\vee \times B^\vee \in \mathfrak{g}^\vee // G^\vee \times \mathfrak{h}^\vee // H^\vee \longleftrightarrow A' \times B' \in [\mathfrak{g}_{\mathbb{C}}^*] \times [\mathfrak{h}_{\mathbb{C}}^*]$$

has the property that

$$\begin{aligned} & \text{eigenvalues of } \rho(A^\vee \times B^\vee) \\ &= (\text{ev}(A') + \text{ev}(B')) \times \begin{cases} \{1, -1\}, & \text{in the unitary case,} \\ \{1\}, & \text{in the orthogonal cases.} \end{cases} \end{aligned} \quad (15.3)$$

On the right-hand side we interpret the sum of two multisets as all pairwise sums of elements from the individual multisets.

Proof. Let Ω_G be the multiset of weights arising from the standard representation for $\mathfrak{g}_{\mathbb{C}}$, where we remove all zero weights. Define Ω_H similarly and let $\Omega = \Omega_G \times \Omega_H$; this Ω is a Weyl-invariant multiset inside $\mathfrak{t}_{G,\mathbb{C}}^* \oplus \mathfrak{t}_{H,\mathbb{C}}^*$. (We add subscripts to clarify whether dealing with a torus for G or a torus for H .)

Let Ω^\vee be the multiset of weights for the Rankin–Selberg representation of $\mathfrak{g}^\vee \oplus \mathfrak{h}^\vee$, as just described above; it is a Weyl-invariant multiset in $(\mathfrak{t}_G^\vee)^* \oplus (\mathfrak{t}_H^\vee)^*$.

Therefore, what we must show is that Ω corresponds to Ω^\vee under (15.2) and the trace duality, i.e., under the sequence of identifications

$$\mathfrak{t}_{G,\mathbb{C}}^* \xrightarrow{\text{trace pairing}} \mathfrak{t}_{G,\mathbb{C}} \longrightarrow (\mathfrak{t}_G^\vee)^*, \quad (15.4)$$

and its analogue for H . The final identification is well defined only up to the Weyl group (this ambiguity makes no difference for comparing Weyl-invariant multisets).

There is a standard basis for roots for SO_{2n} , SO_{2n+1} and GL_n , labelled as

$$\{\pm e_i \pm e_j\}_{[n]}, \quad \{\pm e_i \pm e_j, \pm e_j\}_{[n]} \quad \text{and} \quad \{e_i - e_j\}_{[n]},$$

respectively, where we use $[n]$ as a shorthand for $1 \leq i, j \leq n$, and we always omit zero roots. Label similarly the roots for H as

$$\{\pm f_i \pm f_j, \pm f_i\}_{[n-1]}, \quad \{\pm f_i \pm f_j\}_{[n]} \quad [\text{and}] \quad \{f_i - f_j\}_{[n-1]},$$

in the three respective cases. The composition (15.4) carries e_i to e_i^\vee (the corresponding standard basis for the space of cocharacters); with these identifications we readily compute

$$\Omega = \{\pm e_i \pm f_j\} \quad \text{and} \quad \Omega^\vee = \{\pm e_i^\vee \pm f_j^\vee\}. \quad \square$$

15.4. Satake parameters, conductor drop, and stability

The set appearing in (15.3) almost determines the Rankin–Selberg L -function.

Any irreducible admissible representation of the real Weil group is at most 2-dimensional. If it has bounded image then its restriction to \mathbb{C}^* is of the form

$$z \mapsto |z|^{it} \quad \text{or} \quad z \mapsto z^{n/2+it} \bar{z}^{-n/2+it} \oplus z^{-n/2+it} \bar{z}^{n/2+it}, \quad t \in \mathbb{R} \text{ and } n \in \mathbb{Z}.$$

For a complex number $z = x + iy$, write $z^+ = |x| + iy$. The associated L -factor is given respectively by

$$\Gamma_{\mathbb{R}}(s + \varepsilon + it) \quad \text{or} \quad \prod_{i=1}^2 \Gamma_{\mathbb{R}}\left(s + \left(\frac{n}{2} + it\right)^+ + \varepsilon_i\right),$$

where $\varepsilon_i \in \{0, 1\}$ and where $\Gamma_{\mathbb{R}}(s) = \pi^{-s/2} \Gamma(s/2)$ is the real Γ function.

From this, it readily follows that if π is a tempered representation of G with $\lambda_{\pi} = A^{\vee}$ (equivalently, in the notation of the prior lemma, with infinitesimal character A'), and similarly σ a tempered representation of H with $\lambda_{\sigma} = B^{\vee}$ (equivalently, with infinitesimal character B'), then we have

$$L_{\mathbb{R}}(\pi \times \sigma, \rho, s) = \prod_i \Gamma_{\mathbb{R}}(s + \lambda_i^+ + \varepsilon_i),$$

where the λ_i range through the multiset appearing in (15.3). For this reason, we will refer to the multiset on the left-hand side of (15.3) as the *multiset of Satake parameters for the Rankin–Selberg L -function*.

We may now reinterpret Theorem 14.2. Here, we denote by σ^{\vee} the contragredient of σ ; its Satake parameters are the negatives of those of σ .

LEMMA 15.4. *Let π and σ be irreducible representations of G and H , respectively, with infinitesimal characters $\lambda \in [\mathfrak{g}_{\mathbb{C}}^*]$ and $\mu \in [\mathfrak{h}_{\mathbb{C}}^*]$.*

The following conditions are equivalent:

- (a) (λ, μ) is stable;
- (b) no Satake parameter for the local L -factor $L(\pi \times \sigma^{\vee}, \rho, s)$ is equal to zero.

The significance of this reinterpretation is that (b) is related to an important analytic phenomenon—dropping of the analytical conductor. It would be interesting to see if this relation between stability and conductor drop extends to other integral representations.

16. Volume forms

Let k be a field of characteristic zero. A *volume form* on a smooth k -variety \mathbf{Y} is simply an everywhere non-vanishing global section of the bundle of top-degree algebraic differential forms. When $k = \mathbb{R}$, volume forms give rise to measures (§17.1).

The purpose of this section is to describe the various volume forms that exist on a Lie algebra, its dual, and its coadjoint orbits, as well as the relationships between these forms that arise in the context of a GGP pair. As we explain in §17.5 and §19, the results obtained here model the asymptotic representation theory of G and H .

We note that, in order to evaluate the constant in the main theorem of this paper, we really need the exact relationships between these volume forms (rather than, say, their relationship up to an unspecified proportionality constant).

Some special cases of the foregoing results have been established for certain compact groups G and H (see, e.g., [Ba, Proposition 4.2] and [Ol, Proposition 3.1]).

Throughout this section, we work over an algebraically closed field k of characteristic zero. In the following section, we deduce results over $k=\mathbb{R}$.

16.1. Fibral volume forms

Recall that a short exact sequence $X \rightarrow Y \rightarrow Z$ of vector spaces induces a natural isomorphism $\det(Y^*) \cong \det(X^*) \otimes \det(Z^*)$, where \det denotes the top exterior power.

More generally, given a smooth morphism of varieties, a volume form on source and target induces a volume form on each fiber. To be precise, let $f: Y \rightarrow Z$ be a morphism of smooth irreducible varieties, and fix a regular value $z \in Z$, i.e., for each $y \in f^{-1}(z)$, the induced map $T_y Y \rightarrow T_z Z$ of tangent spaces is surjective. Then, the fiber $X = f^{-1}(z)$ is smooth, and we have a sequence of maps

$$\begin{array}{ccc} X & \xrightarrow{\text{inclusion}} & Y \\ & & \downarrow f \\ & & Z. \end{array}$$

We obtain for each $x \in X$ a short exact sequence $T_x X \rightarrow T_x Y \rightarrow T_z Z$, and hence an identification

$$\det(T_x^* Y) \cong \det(T_x^* X) \otimes \det(T_z^* Z). \quad (16.1)$$

Let β and γ be nowhere vanishing volume forms on Y and Z , respectively. There is then a unique volume form α on X so that $\beta = \alpha \otimes \gamma$ under (16.1) at each point of X . We refer to α as the *fibral volume form* with respect to β and γ , and express this symbolically by $\alpha = \beta/\gamma$. We shall also say, in this situation, that the sequence $X \rightarrow Y \rightarrow Z$ is compatible with the volume forms.

16.2. Haar forms

Let V be a finite-dimensional vector space over k . A *Haar form* β on V is defined to be a translation-invariant volume form. It induces a *dual form* β^* on the dual space V^* . This terminology will be applied most frequently when $V = \mathfrak{h}^*$ for a reductive k -group \mathbf{H} .

Example 16.1. Suppose $\mathbf{H} = \mathbf{GL}_2(k)$. We may identify \mathfrak{h}^* with the space of 2×2 matrices. The Haar forms on \mathfrak{h}^* are the non-zero multiples of $d\xi_{11} \wedge d\xi_{12} \wedge d\xi_{21} \wedge d\xi_{22}$.

16.3. Symplectic volume forms

Let \mathbf{G} be a reductive k -group. Recall (from §1.14) that regular elements of \mathfrak{g}^* are those whose stabilizer has minimal dimension, and that we denote subsets of regular elements by a subscripted reg , as in $\mathfrak{g}_{\text{reg}}^*$. We have the following result.

LEMMA 16.2. *Every fiber \mathcal{O}^λ of $\mathfrak{g} \rightarrow [\mathfrak{g}]$ contains a unique open orbit $\mathcal{O}_{\text{reg}}^\lambda$. This orbit consists precisely of the regular elements of that fiber. An element $\xi \in \mathfrak{g}^*$ is regular if and only if the linear map $\mathfrak{g}^* = T_\xi \mathfrak{g}^* \rightarrow T_{[\xi]}[\mathfrak{g}^*]$ obtained by differentiating the projection $\mathfrak{g} \rightarrow [\mathfrak{g}]$ is surjective.*

Proof. See [Kos2, Theorem 0.1 and Theorem 3]. □

The orbit $\mathcal{O}_{\text{reg}}^\lambda$ carries a canonical G -invariant symplectic form σ , and so also a G -invariant symplectic volume form $\sigma^d/d!$, by the algebraic version of the discussion in §6.1: the symplectic pairing on $T_\xi(\mathcal{O}_{\text{reg}}^\lambda) = \{\text{ad}_x^* \xi : x \in \mathfrak{g}\}$ is given by

$$(\text{ad}_x^* \zeta, \text{ad}_y^* \zeta) \longmapsto \langle \zeta, [x, y] \rangle. \quad (16.2)$$

The same discussion applies to any coadjoint orbit, but we require here only the regular case.

16.4. Affine volume forms

Let \mathbf{H} be a reductive k -group. The quotient $[\mathfrak{h}^*]$ is an affine space (cf. §9.2).

Choose an isomorphism $[\mathfrak{h}^*] \rightarrow \mathbb{A}^r$, or equivalently, generators p_1, \dots, p_r for the ring of G -invariant regular functions on \mathfrak{h}^* . The volume form $dp_1 \wedge \dots \wedge dp_r$ on $[\mathfrak{h}^*]$ is independent, up to scaling, of the choice of generators p_i ; indeed, any two choices will differ by an invertible element of $\mathbb{C}[p_1, \dots, p_r]$. An *affine volume form* on $[\mathfrak{h}^*]$ is then defined to be a non-zero multiple of $dp_1 \wedge \dots \wedge dp_r$.

Example 16.3. Suppose $\mathbf{H}=\mathbf{GL}_r(k)$. We may identify \mathfrak{h}^* with the space of $r \times r$ matrices. By sending a matrix to its characteristic polynomial, we obtain an isomorphism

$$[\mathfrak{h}^*] \xrightarrow{\sim} \text{monic polynomials } x^r + \sum_1^r a_i x^{r-i}.$$

An affine volume form is then given by $da_1 \wedge \dots \wedge da_r$.

LEMMA 16.4. *For each Haar form β_H on \mathfrak{h}^* there is a unique affine form γ_H on $[\mathfrak{h}^*]$ so that for each $\mu \in [\mathfrak{h}^*]$, the fibral volume form for the sequence*

$$\mathcal{O}_{\text{reg}}^\mu \longrightarrow \mathfrak{h}^* \longrightarrow [\mathfrak{h}^*]$$

is the symplectic volume form.

Proof. See [Ros2, Lem C]. □

Later it will be useful to have an explicit formula available for the form on $[\mathfrak{h}]$. Choose as usual a Chevalley basis $H_i, X_\alpha, X_{-\alpha}$ for \mathfrak{h} , where i ranges over simple roots and α over all positive roots; in particular $\alpha_i(H_i)=2$ and $[X_{\alpha_i}, X_{-\alpha_i}]=H_i$. Wedging these together gives a volume form on \mathfrak{h}^* . Write \mathfrak{t} for the Cartan subalgebra spanned by the H_i . Now for $\mu \in \mathfrak{t}^*$ regular the natural projection gives an identification $\mathfrak{t}^* \simeq T_\mu[\mathfrak{h}^*]$. A short computation shows that the pull-back to \mathfrak{h}^* of the affine volume form is given by

$$\prod_{\alpha > 0} \langle \mu, \alpha^\vee \rangle \cdot \bigwedge_i H_i; \tag{16.3}$$

note that the Weyl group acts by the sign character on both factors of (16.3), so the product is invariant.

Definition 16.5. Given a Haar form β_H , the *normalized affine form* γ_H on $[\mathfrak{h}^\wedge]$ is the one associated by the lemma.

Remark 16.6. When working over a local field, the normalized affine form is closely related to the scaling limit of the Plancherel measure; see §17.5.

16.5. Orbital volume forms for a GGP pair

Let $\mathbf{H}^C \rightarrow \mathbf{G}$ be a GGP pair over k . Recall, from Theorem 14.5, that for each stable element $(\lambda, \mu) \in [\mathfrak{g}^*] \times [\mathfrak{h}^*]$, the corresponding fiber $\mathcal{O}^{\lambda, \mu}$ of the map

$$\mathfrak{g}^* \times \mathfrak{h}^* \longrightarrow [\mathfrak{g}^*] \times [\mathfrak{h}^*]$$

is an H -torsor. Fixing a basepoint $\xi \in \mathcal{O}^{\lambda, \mu}$, the orbit map gives an identification

$$H \xrightarrow{\sim} \mathcal{O}^{\lambda, \mu}.$$

Fix a Haar form β_H on \mathfrak{h}^* . We define the *orbital volume form* α to be the volume form on $\mathcal{O}^{\lambda, \mu}$ transferred, via the orbit map, from the volume form on \mathfrak{h} dual to β_H .

THEOREM 16.7. *Fix Haar volume forms β_G and β_H on \mathfrak{g}^* and \mathfrak{h}^* , respectively. Equip $[\mathfrak{g}^*]$ and $[\mathfrak{h}^*]$ with the corresponding normalized affine volume forms γ_G and γ_H , respectively. Let (λ, μ) be stable. Equip $\mathcal{O}_{\text{reg}}^\lambda$ and $\mathcal{O}_{\text{reg}}^\mu$ with their symplectic volume forms. Let α denote the orbital volume form on $\mathcal{O}^{\lambda, \mu}$. Then, in each of the following three sequences, either α or $-\alpha$ is the fibral volume form:*

$$\begin{array}{ccccc}
 & & \mathcal{O}_{\text{reg}}^\lambda \times \mathcal{O}_{\text{reg}}^{-\mu} & \xrightarrow{(\xi, \eta) \mapsto \xi|_{\mathfrak{h} + \eta}} & \mathfrak{h}^* \\
 & \nearrow^{\xi \mapsto (\xi, -\xi|_{\mathfrak{h}})} & & & \\
 \mathcal{O}^{\lambda, \mu} & \longrightarrow & \mathcal{O}_{\text{reg}}^\lambda & \longrightarrow & [\mathfrak{h}^*] \\
 & \searrow & & & \\
 & & \mathfrak{g}_{\text{reg}}^* & \longrightarrow & [\mathfrak{g}^*] \times [\mathfrak{h}^*]
 \end{array} \tag{16.4}$$

Note that, for the top sequence, $0 \in \mathfrak{h}^*$ is a regular value, due to the “trivial stabilizer” consequence of stability; for the middle sequence, $\mu \in [\mathfrak{h}^*]$ is a regular value because “stable implies regular”; for the bottom sequence, $(\lambda, \mu) \in [\mathfrak{g}^*] \times [\mathfrak{h}^*]$ is a regular value because of the “principle bundle” consequence of stability (see §14.3).

The proof for the upper exact sequence is given in §16.5.1, and for the bottom sequence in §16.5.2. The claims for the two lower sequences are readily seen to be equivalent, so this will conclude the proof.

16.5.1. Proof for the top sequence

We examine the top sequence of Theorem 16.7. We equip $\mathcal{O}_{\text{reg}}^\lambda \times \mathcal{O}_{\text{reg}}^{-\mu}$ with its symplectic volume form Ω , and must show that $\Omega/\beta_H = \pm\alpha$. The differential at $\tau \in \mathcal{O}_{\lambda, \mu}$ of the sequence in question fits into the commutative diagram

$$\begin{array}{ccc}
 T_\tau(\mathcal{O}^{\lambda, \mu}) & \longrightarrow & T_\zeta(\mathcal{O}_{\text{reg}}^\lambda \times \mathcal{O}_{\text{reg}}^{-\mu}) \xrightarrow{\text{res}} \mathfrak{h}^* \\
 \uparrow \cong & \nearrow \circ_\zeta \circ \delta & \\
 \mathfrak{h} & &
 \end{array} \tag{16.5}$$

in which $\zeta := (\tau, -\tau|_{\mathfrak{h}})$, o_τ and o_ζ denote differentials of orbit maps, and $\delta: \mathfrak{h} \rightarrow \mathfrak{g} \oplus \mathfrak{h}$ and $\text{res}: \mathfrak{g}^* \oplus \mathfrak{h}^* \rightarrow \mathfrak{h}^*$ are given by $\delta(x) := (x, x)$ and $\text{res}(\xi, \eta) := \xi|_{\mathfrak{h}} + \eta$.

We claim that this is a *Lagrangian fibration*, i.e., that $o_\zeta \circ \delta(\mathfrak{h})$ is Lagrangian and that the duality between \mathfrak{h} and \mathfrak{h}^* is induced by the symplectic structure on the middle term.

We may then conclude via the definition of symplectic volume forms.

To verify the Lagrangian fibration property, it is enough to check, for each $x \in \mathfrak{h}$ and $\eta \in T_\zeta(\mathcal{O}_{\text{reg}}^\lambda \times \mathcal{O}_{\text{reg}}^{-\mu}) \subset \mathfrak{g}^* \times \mathfrak{h}^*$, that $\sigma(o_\zeta \circ \delta(x), \eta) = \langle x, \text{res}(\eta) \rangle$, where σ denotes the symplectic pairing on $T_\zeta(\mathcal{O}_{\text{reg}}^\lambda \times \mathcal{O}_{\text{reg}}^{-\mu})$. Indeed, we may write $\eta = o_\zeta(y)$ for some $y \in \mathfrak{g} \oplus \mathfrak{h}$. By the definition of the symplectic pairing (see (16.2)), we then have

$$\sigma(o_\zeta \circ \delta(x), \eta) = \langle \zeta, [\delta(x), y] \rangle = \langle \delta(x), o_\zeta(y) \rangle = \langle \delta(x), \eta \rangle = \langle x, \text{res}(\eta) \rangle.$$

This concludes the proof.

16.5.2. Proof for the bottom sequence

It will be convenient to deduce the assertion from one that is more explicitly phrased in the context of spherical varieties. The proof that follows borrows ideas that are well known in that context (resolving the cotangent bundle, degenerating to the boundary).

Thus, let \mathbf{M} be any reductive group over k with Lie algebra \mathfrak{m} . Suppose that \mathfrak{s} is a spherical Lie subalgebra of \mathfrak{m} , that is to say, that there is a Borel subgroup \mathbf{B} of \mathbf{M} whose Lie algebra \mathfrak{b} is complementary to \mathfrak{s} ; in particular, \mathfrak{s} is of dimension $\dim(M/B)$. Let $\zeta \in \mathfrak{s}^\perp \subset \mathfrak{m}^*$. We denote by $T_{[\zeta]}$ the tangent space to $[\mathfrak{m}^*]$ at $[\zeta]$, and assume that the sequence

$$\mathfrak{s} \xrightarrow{x \mapsto \text{ad}_x^* \zeta} \mathfrak{s}^\perp \xrightarrow{\text{project}} T_{[\zeta]}. \quad (16.6)$$

is short exact; here the final map is

$$\mathfrak{s}^\perp \hookrightarrow \mathfrak{m}^* \cong T_\zeta \mathfrak{m}^* \longrightarrow T_{[\zeta]}.$$

Note that the composition of the sequence (16.6) is always zero.

CLAIM. *This sequence is compatible, up to signs, with volume forms, where*

- *We fix Haar forms on \mathfrak{s} and \mathfrak{m} , and give \mathfrak{s}^\perp the induced form via*

$$\det(\mathfrak{s}^\perp) \simeq \det(\mathfrak{m})^* \otimes \det(\mathfrak{s}).$$

- *The given Haar form on \mathfrak{m} defines a normalized affine form on $[\mathfrak{m}^*]$, and hence a volume form on $T_{[\zeta]}$.*

Before proving the claim, let us see how it implies the desired result, namely, that the bottom sequence of the theorem is compatible with volume forms. We apply the claim with $(\mathbf{M}, \mathbf{S}) := (\mathbf{G} \times \mathbf{H}, \text{diag } \mathbf{H})$. Then, $\mathfrak{s} = \text{diag } \mathfrak{h} \hookrightarrow \mathfrak{g} \times \mathfrak{h} = \mathfrak{m}$, and we have an isomorphism

$$\begin{aligned} \iota: \mathfrak{s}^\perp &\xrightarrow{\sim} \mathfrak{g}^*, \\ (\xi, -\xi|_{\mathfrak{h}}) &\longmapsto \xi, \end{aligned}$$

compatible with the adjoint actions of $S \cong H$.

Let (λ, μ) be stable. Fix a basepoint $\xi \in \mathcal{O}^{\lambda, \mu} \subset \mathfrak{g}^*$. Then, the sequence (16.6) with $\zeta := \iota^{-1}(\xi)$ is isomorphic to the differential of bottom sequence of the theorem, namely:

$$\begin{array}{ccccc} \mathfrak{s} & \xrightarrow{x \mapsto \text{ad}_x^*(\xi, -\xi_H)} & \mathfrak{s}^\perp & \xrightarrow{[\cdot]} & T_{(\lambda, -\mu)}([\mathfrak{g}^*] \times [\mathfrak{h}^*]) \\ \cong \downarrow & x \mapsto \text{ad}_x^* \xi & \downarrow \cong & & \downarrow \cong \\ T_\xi(\mathcal{O}^{\lambda, \mu}) & \longrightarrow & T_\xi(\mathfrak{g}_{\text{stab}}^*) & \xrightarrow{\xi \mapsto ([\xi], [\xi_H])} & T_{(\lambda, \mu)}([\mathfrak{g}^*] \times [\mathfrak{h}^*]). \end{array}$$

Here we identify $T_\xi(\mathfrak{g}_{\text{stab}}^*) = \mathfrak{g}^*$, with the subscripted stab denoting the subset of \mathbf{H} -stable elements.

The rightmost vertical arrow preserves volume forms, up to sign. Since the claim shows that the upper sequence is compatible with volume forms, the lower sequence is also compatible with the volume forms that are transferred from the upper sequence. The desired result follows.

Proof of the claim. Consider the set of $\zeta \in \mathfrak{s}^\perp$ for which there exists a Borel subalgebra $\mathfrak{b} \leq \mathfrak{m}$ such that

- (i) $\mathfrak{m} = \mathfrak{b} \oplus \mathfrak{s}$;
- (ii) ζ is the composition of the projection $\mathfrak{b} \oplus \mathfrak{s} \rightarrow \mathfrak{b}/[\mathfrak{b}, \mathfrak{b}]$ with a regular character $\bar{\zeta}$ of $\mathfrak{b}/[\mathfrak{b}, \mathfrak{b}]$. Here “regular” means “regular when identified with a character of the torus quotient of \mathfrak{b} ”.

We claim that it is sufficient to prove the claim for such ζ . Indeed, condition (ii) implies that ζ is regular semisimple when considered as an element of \mathfrak{m}^* . For such a ζ , the set of Borel subgroups \mathfrak{b} that satisfy the polarization condition $\zeta([\mathfrak{b}, \mathfrak{b}]) = 0$ is actually finite. It follows from this that

- For ζ arising as in (i) and (ii), the stabilizer of ζ in \mathfrak{m} is contained in \mathfrak{b} , and hence the stabilizer in \mathfrak{s} is trivial. Therefore, the first map of sequence (16.6) is injective.
- For ζ arising as in (i) and (ii), the second map of (16.6) is surjective (and so, counting dimensions, sequence (16.6) is exact). Indeed, writing \mathfrak{t} for the torus quotient

of \mathfrak{b} , there is a natural map $P: \mathfrak{t}^* \rightarrow [\mathfrak{m}^*]$, and the second map of (16.6) sends ζ to the image of $\bar{\zeta}$ under that map. But P is submersive at the regular element $\bar{\zeta}$.

- The dimension of the set of ζ that arise as in (i) and (ii) equals

$$\dim(M/B) + \dim(T) = \dim(\mathfrak{s}^\perp).$$

Let Y be the set of all $\zeta \in \mathfrak{s}^\perp$ for which (16.6) is short exact. Then, Y is a Zariski-open subset of \mathfrak{s}^\perp , and it contains the constructible set X of ζ arising as in (i) and (ii). As the dimension of X coincides with the dimension of Y , X is Zariski-dense inside Y . It follows that it is sufficient to prove the desired assertion for $\zeta \in X$.

We will do this by degeneration. Let L be the variety of pairs

$$(\mathfrak{s}', \zeta'),$$

where \mathfrak{s}' is a subalgebra of \mathfrak{m} of the same dimension as \mathfrak{s} , and such that (16.6), with \mathfrak{s}' replacing \mathfrak{s} , is short exact. Comparing the volume forms on the various factors of (16.6) describes an M -invariant regular function $f: L \rightarrow \mathbb{G}_m$, where $f=1$ at any point (\mathfrak{s}', ζ') where the sequence is compatible with volume forms. Fixing $\zeta, \bar{\zeta}$ and \mathfrak{b} as in (i) and (ii), we will show that $f(\mathfrak{s}, \zeta)=1$.

Let B be the associated Borel subgroup, and fix a maximal torus $T \subset B$ with Lie algebra \mathfrak{t} . We get a splitting $\mathfrak{m} = \bar{\mathfrak{n}} \oplus \mathfrak{t} \oplus \mathfrak{n}$, where $\mathfrak{n} = [\mathfrak{b}, \mathfrak{b}]$. Choose a regular 1-parameter subgroup $\gamma: \mathbb{G}_m \rightarrow T$ with $\langle \gamma, \alpha \rangle < 0$ for every positive root α , so that $\gamma(t)$ contracts \mathfrak{n} as $t \rightarrow \infty$.

Let ζ_t denote the restriction of ζ to \mathfrak{t} ; we regard ζ_t as an element of \mathfrak{m}^* by extending trivially on $\mathfrak{n} \oplus \bar{\mathfrak{n}}$. Then, it is easy to see that $(\bar{\mathfrak{n}}, \zeta_t) \in L$; moreover,

$$\lim_{t \rightarrow \infty} \text{Ad} \gamma(t) \cdot (\mathfrak{s}, \zeta) = (\bar{\mathfrak{n}}, \zeta_t). \quad (16.7)$$

Indeed, by the assumption $\mathfrak{s} \cap \mathfrak{b} = \{0\}$ and a dimension computation, we see that \mathfrak{s} projects onto $\bar{\mathfrak{n}}$ with respect to the splitting $\mathfrak{m} = \bar{\mathfrak{n}} \oplus \mathfrak{b}$. It follows readily from this that

$$\lim_t \text{Ad} \gamma(t) \mathfrak{s} = \bar{\mathfrak{n}}.$$

Furthermore, the character $\text{Ad}(\gamma(t))\zeta$ is trivial on the subspace $\text{Ad}(\gamma(t))\mathfrak{s} \oplus \mathfrak{n}$, which converges in turn to $\bar{\mathfrak{n}} \oplus \mathfrak{n}$. This implies that $\text{Ad}(\gamma(t))\zeta$ converges in \mathfrak{m}^* to the character ζ_t , extended trivially on $\bar{\mathfrak{n}} \oplus \mathfrak{n}$, and so concludes the proof of (16.7).

Since f is T -invariant and regular, it follows that $f(\mathfrak{s}, \zeta) = f(\bar{\mathfrak{n}}, \zeta_t)$. It remains to show that $f(\bar{\mathfrak{n}}, \zeta_t) = 1$. This is a routine computation with (16.3). \square

17. Measures and integrals

We now apply the preceding considerations to define and compare measures on spaces associated with a GGP pair over an archimedean local field.

17.1. Real varieties

We denote by $X=\mathbf{X}(\mathbb{R})$ the set of real points of a smooth real algebraic variety \mathbf{X} . For each Haar measure λ on \mathbb{R} there is an assignment (see [We, §2.2])

$$\{\mathbb{R}\text{-rational top-degree differential forms } \omega \text{ on } \mathbf{X}\} \longrightarrow \{\text{measures } |\omega| \text{ on } X\},$$

which is functorial under pullback by étale maps, compatible with products, satisfies $|f\omega|=|f|\cdot|\omega|$, and is normalized by $|dx|=\lambda$ when \mathbf{X} is the affine line \mathbb{A}^1 .

We henceforth take for λ the measure

$$\frac{\text{Lebesgue}}{\sqrt{2\pi}}.$$

On the affine space \mathbb{A}^n , we then have

$$|dx_1 \wedge \dots \wedge dx_n| = \frac{\text{Lebesgue}}{(2\pi)^{n/2}}. \quad (17.1)$$

We have normalized λ to be Fourier self-dual for the character $\psi(x):=e^{ix}$. This normalization has the following consequence:

Let V be a real vector space. Using ψ , we may identify the real dual V^* with the Pontryagin dual $V^\wedge:=\text{Hom}(V, \mathbb{C}^{(1)})$. Then, dual algebraic volume forms on V and V^* correspond to (Fourier-)dual measures on V and V^\wedge .

17.2. Groups

Let \mathbf{G} be a reductive group over \mathbb{R} . Recall (from §1.14) that we denote by \mathfrak{g} the real Lie algebra, by \mathfrak{g}^* its real dual, and by \mathfrak{g}^\wedge the Pontryagin dual $\text{Hom}(\mathfrak{g}, \mathbb{C}^{(1)})$. Sending $\xi \in i\mathfrak{g}^*$ to $x \mapsto e^{x\xi} \in \mathbb{C}^{(1)}$ gives an identification $i\mathfrak{g}^* \simeq \mathfrak{g}^\wedge$.

We suppose given a Haar measure dg on the Lie group G . There is then a compatible Haar measure dx on \mathfrak{g} , normalized as in §2.1 by requiring that $U \subseteq \mathfrak{g}$ and $\exp(U) \subseteq G$ have similar volumes when U is a small neighborhood of the origin. We obtain also a Fourier-dual measure $d\xi$ on \mathfrak{g}^\wedge . We choose an \mathbb{R} -rational Haar form β on the real vector space \mathfrak{g}^\wedge by requiring that $|\beta|=d\xi$; it normalizes a dual form β^\vee on \mathfrak{g} . Explicitly, if we choose coordinates $x=(x_1, \dots, x_n)$ and $\xi=(\xi_1, \dots, \xi_n)$ as in §4.1 so that

$x\xi = \sum_j x_j \xi_j$ and $dx = dx_1 \dots dx_n$, then $d\xi = (2\pi)^{-n} d\xi_1 \dots d\xi_n$, $\beta = \pm(2\pi)^{-n/2} d\xi_1 \wedge \dots \wedge d\xi_n$ and $\beta^\vee = \pm(2\pi)^{n/2} dx_1 \wedge \dots \wedge dx_n$.

Recall that, for $\lambda \in [\mathfrak{g}^\wedge]$, we set

$$\mathcal{O}^\lambda = \{\xi \in \mathfrak{g}^\wedge : [\xi] = \lambda\}.$$

On the regular subset $\mathcal{O}_{\text{reg}}^\lambda$ we have both an algebraic symplectic volume form

$$\omega_{\text{alg}} = \frac{1}{d!} \sigma^d$$

as in §16.3 and a normalized symplectic measure

$$\omega = \frac{1}{d!} \left(\frac{\sigma}{2\pi} \right)^d$$

as in §6.1. We verify readily that the algebraic form is \mathbb{R} -rational, and satisfies $|\omega_{\text{alg}}| = \omega$.

By the recipe in §16.4, β induces a normalized affine volume form γ on $[\mathfrak{g}^\wedge]$. (We use that \mathfrak{g}^\wedge and $[\mathfrak{g}^\wedge]$ are real forms of $\mathfrak{g}_{\mathbb{C}}$ and $[\mathfrak{g}_{\mathbb{C}}]$.) We verify readily that γ is \mathbb{R} -rational, hence induces a *normalized affine measure* $|\gamma|$ on $[\mathfrak{g}^\wedge]$. By the construction and compatibilities noted previously, we then have

$$\int_{\mathfrak{g}^\wedge} a = \int_{\lambda \in [\mathfrak{g}^\wedge]} \left(\int_{\mathcal{O}_\lambda^{\text{reg}}} a d\omega \right) \quad (17.2)$$

for each $a \in C_c(\mathfrak{g}^\wedge)$.

17.3. GGP pairs

Let (\mathbf{G}, \mathbf{H}) be a GGP pair over an archimedean local field. By restriction of scalars, we may regard \mathbf{G} and \mathbf{H} as reductive groups over \mathbb{R} ; the discussion and notation in §17.2 thus applies.

17.3.1. Stability consequences

For $\lambda \in [\mathfrak{g}^\wedge]$ and $\mu \in [\mathfrak{h}^\wedge]$, we set

$$\mathcal{O}^{\lambda, \mu} := \{\xi \in \mathfrak{g}^\wedge : [\xi] = \lambda \text{ and } [\xi|_{\mathfrak{h}}] = \mu\}.$$

As before, a subscripted “stab” denotes the subset of \mathbf{H} -stable elements.

THEOREM 17.1. *The map $\mathfrak{g}_{\text{stab}}^\wedge \rightarrow \{\text{stable}(\lambda, \mu)\}$ is a principal H -bundle over its image, with fibers $\mathcal{O}^{\lambda, \mu}$. In particular, if $(\lambda, \mu) \in [\mathfrak{g}^\wedge] \times [\mathfrak{h}^\wedge]$ is stable, then either*

- $\mathcal{O}^{\lambda, \mu} = \emptyset$, or
- $\mathcal{O}^{\lambda, \mu}$ is an H -torsor, i.e., a closed H -invariant subset of \mathfrak{g}^\wedge on which H acts simply transitively; moreover, $\mathcal{O}^{\lambda, \mu}$ consists of \mathbf{H} -stable regular elements.

Proof. This follows readily from the corresponding properties (§14.3) established over the algebraic closure \mathbb{C} . \square

More generally, for each regular coadjoint multiorbit $\mathcal{O} \subseteq \mathfrak{g}^\wedge$ and $\mu \in [\mathfrak{h}^\wedge]$, we set

$$\mathcal{O}(\mu) := \{\xi \in \mathcal{O} : [\xi|_{\mathfrak{h}}] = \mu\}.$$

For example, $\mathcal{O}^{\lambda, \mu} = \mathcal{O}^\lambda(\mu)$. Then,

$$\mathcal{O}_{\text{stab}} \longrightarrow [\mathfrak{h}^\wedge] \cap \text{image}(\mathcal{O}_{\text{stab}})$$

is a principal H -bundle, with fibers $\mathcal{O}(\mu)$.

Let π and σ be tempered irreducible unitary representations of G and H , respectively. We set

$$\mathcal{O}_{\pi, \sigma} := \{\xi \in \mathcal{O}_\pi : \xi|_{\mathfrak{h}} \in \mathcal{O}_\sigma\}.$$

We note that

$$\mathcal{O}_{\pi, \sigma} \subseteq \mathcal{O}_\pi(\lambda_\sigma) \subseteq \mathcal{O}^{\lambda_\pi, \lambda_\sigma}.$$

The pictures in §1.7 and §1.10 give some examples to which this notation apply. Theorem 17.1 implies that if $(\lambda_\pi, \lambda_\sigma)$ is stable and $\mathcal{O}_{\pi, \sigma}$ is non-empty, then $\mathcal{O}_{\pi, \sigma}$ is an H -torsor consisting of \mathbf{H} -stable elements, and hence

$$\mathcal{O}_{\pi, \sigma} = \mathcal{O}_\pi(\lambda_\sigma) = \mathcal{O}^{\lambda_\pi, \lambda_\sigma}. \quad (17.3)$$

17.3.2. Integral transforms and identities

We assume given a Haar measure on H . As in §17.2, this choice defines a measure $[\mathfrak{h}^\wedge]$. By the discussion in §16.5 and §17.1, we obtain also—for stable $(\lambda, \mu) \in [\mathfrak{g}^\wedge] \times [\mathfrak{h}^\wedge]$ —a measure on $\mathcal{O}^{\lambda, \mu}$, which we will see below induces a measure on \mathfrak{g}^\wedge . More generally, we adopt the convention that an integral over $\mathcal{O}^{\lambda, \mu}$ is defined to be zero unless (λ, μ) is stable; in that case, the measure is given explicitly by the push-forward of the Haar measure from H , i.e.,

$$\int_{\mathcal{O}^{\lambda, \mu}} a := \int_{s \in H} a(s \cdot \xi_{\lambda, \mu})$$

for any basepoint $\xi_{\lambda,\mu} \in \mathcal{O}^{\lambda,\mu}$. We note (cf. §14.2) that, for given λ , the pair (λ, μ) is stable for μ outside a measure-zero subset.

For a regular coadjoint multiorbit $\mathcal{O} \subseteq [\mathfrak{g}^\wedge]$ and $\mu \in [\mathfrak{h}^\wedge]$, the set $\mathcal{O}(\mu)$ is either empty or of the form $\mathcal{O}^{\lambda,\mu}$ with $\lambda = [\mathcal{O}]$. Thus integration over $\mathcal{O}(\mu)$ is defined.

THEOREM 17.2. *Integration defines a continuous map*

$$\begin{aligned} \{\text{stable } (\lambda, \mu) \in [\mathfrak{g}^\wedge] \times [\mathfrak{h}^\wedge]\} \times \mathcal{S}(\mathfrak{g}^\wedge) &\longrightarrow \mathbb{C}, \\ (\lambda, \mu, a) &\longmapsto \int_{\mathcal{O}^{\lambda,\mu}} a. \end{aligned}$$

We have

$$\int_{\mathcal{O}_{\text{reg}}^\lambda} a \, d\omega = \int_{\mu \in [\mathfrak{h}^\wedge]} \int_{\mathcal{O}^{\lambda,\mu}} a. \quad (17.4)$$

More generally, for any regular coadjoint multiorbit $\mathcal{O} \subseteq \mathfrak{g}^\wedge$,

$$\int_{\mathcal{O}} a \, d\omega = \int_{\mu \in [\mathfrak{h}^\wedge]} \int_{\mathcal{O}(\mu)} a. \quad (17.5)$$

Proof. The convergence follows from the inequalities recorded below in §17.4, the continuity from Theorem 14.5, and the integral formulas from Theorem 16.7. \square

This result will be applied in §22.3.

17.4. A Lojasiewicz-type inequality

We pause to record a technical lemma justifying the convergence and continuity of the integral transforms defined above.

We fix a norm $|\cdot|$ on \mathfrak{g}^\wedge and a faithful finite-dimensional representation of \mathbf{H} . We may use the latter to define an algebraic norm $|\cdot|$ on H : denoting the faithful representation by R , we set

$$|h| := \max(\|R(h)\|, \|R(h^{-1})\|),$$

where $\|\cdot\|$ denotes the operator norm on the space of R .

LEMMA 17.3. *Let $\xi \in \mathfrak{g}^\wedge$ be \mathbf{H} -stable. There are then positive reals c_1 and c_2 and a (topological) neighborhood U of ξ such that, for all $\eta \in U$ and $s \in H$,*

$$|s \cdot \eta| > c_1 |s|^{c_2}. \quad (17.6)$$

Proof. The required estimate is in the spirit of the Lojasiewicz inequality, but for lack of a convenient reference it is simplest for us to argue directly.

Let Z denote the set of \mathbf{H} -stable elements of \mathfrak{g}^\wedge . By Theorem 14.2, the set Z may be described explicitly as the non-vanishing locus of a certain resultant $\rho: \mathfrak{g}^\wedge \rightarrow \mathbb{R}$. We may thus regard Z as the set of real points of the real affine variety $\mathbf{Z} := \text{Spec} \mathbb{R}[\mathfrak{g}^\wedge, 1/\rho]$.

The group \mathbf{H} acts on \mathbf{Z} . The orbit map $(h, z) \mapsto (hz, z)$, regarded as an algebraic map

$$\mathbf{H} \times \mathbf{Z} \longrightarrow \mathbf{Z} \times \mathbf{Z},$$

is a *closed immersion*. Indeed, this property may be checked on complex points, and Theorem 14.5 says that $\mathbf{Z}(\mathbb{C}) \rightarrow \mathbf{Z}(\mathbb{C}) // \mathbf{H}(\mathbb{C})$ is a principal $\mathbf{H}(\mathbb{C})$ -bundle. In particular, the induced map on coordinate rings is surjective.

Choose generators f_1, \dots, f_k for the \mathbb{R} -algebra of regular functions on \mathbf{Z} . The closed immersion property noted above implies that each regular function $P(h)$ on \mathbf{H} may be written as a polynomial in $f_i(z)$ and $f_j(h \cdot z)$. Writing $\|z\| = \max_i |f_i(z)|$, we may thus find positive constants C and K such that, for all $h \in H$ and $z \in Z$, we have

$$P(h) \leq C(\|z\| + \|h \cdot z\|)^K,$$

or equivalently,

$$\|h \cdot z\| \geq C^{-1/K} P(h)^{1/K} - \|z\|.$$

The norm $|s|$ defined for $s \in H$ by a finite-dimensional faithful representation is comparable to $\max_{i \in I} |P_i(s)|$ for some finite collection $(P_i(h))_{i \in I}$ of regular functions on \mathbf{Z} . Restricting η to a compact subset of Z , we obtain

$$|s \cdot \eta| \geq c_1 |s|^{c_2} - c_3 \tag{17.7}$$

for suitable constants $c_1, c_2, c_3 > 0$ (depending only upon the given compact set). When $|s|$ is large enough, the required estimate follows from (17.7). In the remaining range, s and η both lie in fixed compact sets and $s \cdot \eta \neq 0$, giving the adequate estimate $|s \cdot \eta| \gg 1 \gg |s|$. \square

17.5. The scaling limit of Plancherel measure

It is instructive to note that the normalized affine volume measure on $[\mathfrak{h}^\wedge]$ is closely related to the Plancherel measure μ on \widehat{H} (cf. §A.3). We do not use this comparison directly, and so will be brief and sketchy. For an open set $U \subset [\mathfrak{h}^\wedge]$, let

$$\widetilde{U} := \{\sigma \in \widehat{H}_{\text{temp}} : \lambda_\sigma \in U\}$$

denote the set of isomorphism classes of tempered irreducible unitary representations having infinitesimal character in U . We assume for simplicity that we are working over a complex group, so that, for each $\sigma \in \widehat{H}_{\text{temp}}$, we have $\mathcal{O}_\sigma = \mathcal{O}_{\text{reg}}^\mu$ with $\mu := \lambda_\sigma$.

LEMMA 17.4. *Fix a non-empty bounded open subset U of $[\mathfrak{h}^\wedge]$ whose boundary has measure zero (with respect to any measure in the class of smooth measures). Then,*

$$\lim_{\mathfrak{h} \rightarrow 0} \frac{\text{Plancherel measure of } \widetilde{\mathfrak{h}^{-1}U}}{\text{normalized affine measure of } \mathfrak{h}^{-1}U} = 1.$$

While one can prove this simply by examining the explicit form of Plancherel measure, it would then be tedious to check carefully the normalization of constants. We sketch a different argument, presumably well known, which makes the normalization clear. Define $\text{Op}_{\mathfrak{h}}: C_c^\infty(\mathfrak{h}^\wedge) \rightarrow \text{End}(\sigma)$ as usual. Fix $a \in C_c^\infty(\mathfrak{h}^\wedge)$, with dilates $a_{\mathfrak{h}}(\xi) := a(\mathfrak{h}\xi)$ as usual. By the definition of normalized affine measure,

$$\int_{\mu \in [\mathfrak{h}^\wedge]} \left(\int_{\mathcal{O}^\mu} a_{\mathfrak{h}} \right) = \int_{\mathfrak{h}^\wedge} a_{\mathfrak{h}}.$$

Recall that $\text{Op}_{\mathfrak{h}}(a) = \sigma(f)$ for some $f \in C_c^\infty(H)$ supported near 1 and given by

$$f(\exp(x)) := a_{\mathfrak{h}}^\vee(x) \chi(x) j(x)^{-1};$$

since $\chi(0) = j(0) = 1$, we have in particular

$$\int_{\mathfrak{h}^\wedge} a_{\mathfrak{h}} = f(1).$$

By the Plancherel formula (§A.3),

$$f(1) = \int_{\sigma \in \widehat{H}_{\text{temp}}} \text{tr}(\sigma(f)) d\mu(\sigma).$$

By the Kirillov formula (see §12.3),

$$\text{tr}(\sigma(f)) = \text{tr}(\text{Op}_{\mathfrak{h}}(a)) \sim \int_{\mathcal{O}_\sigma} a_{\mathfrak{h}} d\omega_{\mathcal{O}_\sigma}.$$

(Here and henceforth the precise meaning of \sim may be inferred from the precise statements in §12.3.) It follows that

$$\int_{\mu \in [\mathfrak{h}^\wedge]} \left(\int_{\mathcal{O}^\mu} a_{\mathfrak{h}} \right) \sim \int_{\sigma \in \widehat{H}_{\text{temp}}} \left(\int_{\mathcal{O}_\sigma} a_{\mathfrak{h}} d\omega_{\mathcal{O}_\sigma} \right) d\mu(\sigma). \quad (17.8)$$

We now choose a so that $\int_{\mathcal{O}^\mu} a_{\mathfrak{h}}$ approximates the characteristic function $1_U(\lambda)$. The left-hand side of (17.8) then approximates the affine measure of $\mathfrak{h}^{-1}U$, while the right-hand side approximates the Plancherel measure of the set \widetilde{U} of representations σ for which $\lambda_\sigma \in \mathfrak{h}^{-1}U$.

18. Relative characters: disintegration

Let (\mathbf{G}, \mathbf{H}) be a GGP pair over a local field F .

We equip G and H with some Haar measures. To simplify notation, we do not display these Haar measures in our integration notation. Thus $\int_{s \in H} f(s)$ denotes the integral of $f \in L^1(H)$.

By the proofs of [II, Proposition 1.1] and [HarN, §2], the corresponding Harish-Chandra functions (cf. §A.2) satisfy

$$\int_H \Xi_G|_H \cdot \Xi_H < \infty. \quad (18.1)$$

Let π and σ be tempered irreducible unitary representations of G and H , respectively. More precisely, we denote in this subsection by π and σ the spaces of smooth vectors in the underlying Hilbert spaces. Choose an orthonormal basis $\mathcal{B}(\sigma)$ consisting of isotypic vectors for the action of some fixed maximal compact subgroup of H . Similarly, choose an orthonormal basis $\mathcal{B}(\pi)$ for π .

LEMMA 18.1. (i) For $v_1, v_2 \in \pi$, the formula

$$\mathcal{H}_\sigma(v_1 \otimes \bar{v}_2) := \sum_{u \in \mathcal{B}(\sigma)} \int_{s \in H} \langle sv_1, v_2 \rangle \langle u, su \rangle \quad (18.2)$$

converges and defines an H -invariant Hermitian form

$$\mathcal{H}_\sigma: \pi \otimes \bar{\pi} \longrightarrow \mathbb{C}.$$

(ii) For $T \in \pi \otimes \bar{\pi}$, one has

$$\mathrm{tr}(T) = \int_{\sigma \in \hat{H}_{\mathrm{temp}}} \mathcal{H}_\sigma(T), \quad (18.3)$$

where $\mathrm{tr}: \pi \otimes \bar{\pi} \rightarrow \mathbb{C}$ is the linear map for which $\mathrm{tr}(v_1 \otimes \bar{v}_2) := \langle v_1, v_2 \rangle$ and the integral is taken with respect to the Plancherel measure on \hat{H}_{temp} dual to the chosen Haar measure on H . In particular,

$$\langle v_1, v_2 \rangle = \int_{\sigma \in \hat{H}_{\mathrm{temp}}} \sum_{u \in \mathcal{B}(\sigma)} \int_{s \in H} \langle sv_1, v_2 \rangle \langle u, su \rangle. \quad (18.4)$$

(iii) Suppose that F is archimedean, so that the definitions of Part I apply, and let $N \geq 0$ be sufficiently large in terms of G . Then, \mathcal{H}_σ extends to a map

$$\mathcal{H}_\sigma: \Psi^{-N}(\pi) \longrightarrow \mathbb{C}$$

which is continuous, uniformly in π and σ , and given by

$$\mathcal{H}_\sigma(T) = \sum_{u \in \mathcal{B}(\sigma)} \int_{s \in H} \mathrm{tr}(sT) \langle u, su \rangle = \sum_{\substack{v \in \mathcal{B}(\pi) \\ u \in \mathcal{B}(\sigma)}} \int_{s \in H} \langle sTv, v \rangle \langle u, su \rangle. \quad (18.5)$$

Proof. The details of the proof are technical and not particularly interesting, so we have relegated most of them to Appendix A. Let $v_1, v_2 \in \pi$, and define $f: H \rightarrow \mathbb{C}$ by $f(s) := \langle sv_1, v_2 \rangle$. Then

$$f(1) = \langle v_1, v_2 \rangle, \quad \mathcal{H}_\sigma(v_1 \otimes \bar{v}_2) = \text{tr}(\sigma(f)). \quad (18.6)$$

Observe now that assertions (i) and (ii) are formal consequences of the Plancherel formula. That formula does not directly apply, because f is typically not compactly-supported, but an approximation argument gives what is needed. We postpone the details to §A.6. \square

Remark 18.2. It is expected, and known for F non-archimedean (see [Sve]), that \mathcal{H}_σ satisfies the positivity condition

$$\mathcal{H}_\sigma(v \otimes \bar{v}) \geq 0 \quad \text{for all } v \in \pi. \quad (18.7)$$

Assume this. Since $\dim \text{Hom}_H(\pi, \sigma) \leq 1$, we may then write

$$\int_{s \in H} \langle sv, v \rangle \langle u, us \rangle = |\langle \ell_\sigma(v), u \rangle|^2 \quad (18.8)$$

for some H -invariant functional $\ell_\sigma: \pi \rightarrow \sigma$, determined up to phase. One has

$$\mathcal{H}_\sigma(v) = \|\ell_\sigma(v)\|^2.$$

The continuity of ℓ_σ follows from that of \mathcal{H}_σ .

19. Relative characters: asymptotics in the stable case

We assume that (\mathbf{G}, \mathbf{H}) is a GGP pair over an *archimedean* local field, and retain the notation and conventions of §17 and §18. In particular, by restriction of scalars, we may regard \mathbf{G} and \mathbf{H} as real reductive groups, and the corresponding point sets G and H as real Lie groups with Lie algebras \mathfrak{g} and \mathfrak{h} .

19.1. Motivation

Let $\pi \in \widehat{G}_{\text{temp}}$ and $\sigma \in \widehat{H}_{\text{temp}}$ be tempered irreducible unitary representations. We allow π and σ to vary arbitrarily with a scale parameter $\mathfrak{h} \rightarrow 0$, which we normalize so that the rescaled infinitesimal characters $\mathfrak{h}\lambda_\pi$ and $\mathfrak{h}\lambda_\sigma$ remain bounded.

Fix $a \in S^{-\infty}(\mathfrak{g}^\wedge)$, and set

$$\text{Op}_\mathfrak{h}(a) := \text{Op}_\mathfrak{h}(a : \pi) \in \Psi^{-\infty}(\pi).$$

We wish to understand the $\mathfrak{h} \rightarrow 0$ asymptotics of the evaluation $\mathcal{H}_\sigma(\mathrm{Op}_\mathfrak{h}(a))$ of the relative character \mathcal{H}_σ defined in §18.

To see what to expect, observe the following:

$$\int_{\sigma \in \widehat{H}_{\mathrm{temp}}} \mathcal{H}_\sigma(\mathrm{Op}_\mathfrak{h}(a)) \stackrel{(18.3)}{=} \mathrm{tr}(\mathrm{Op}_\mathfrak{h}(a)) \stackrel{(12.2)}{\approx} \int_{\mathcal{O}_\pi} a_\mathfrak{h} d\omega \stackrel{(17.5)}{=} \int_{\mu \in [\mathfrak{h}^\wedge]} \int_{\mathcal{O}_\pi(\mu)} a_\mathfrak{h}.$$

Consideration of the action of the universal enveloping algebra of H suggests that the above sequence localizes to individual σ , i.e., that

$$\mathcal{H}_\sigma(\mathrm{Op}_\mathfrak{h}(a)) \approx \int_{\mathcal{O}_{\pi,\sigma}} a_\mathfrak{h}, \quad (19.1)$$

at least if $(\mathfrak{h}\lambda_\pi, \mathfrak{h}\lambda_\sigma)$ stays away from the boundary of the stable locus.

Strictly speaking, the most one can deduce immediately from such reasoning is an equality like (19.1), but summed over all σ of fixed infinitesimal character. However, there is another way to motivate (19.1), again ignoring issues of rigor. If we were to pretend that the exponential map were an isomorphism with trivial Jacobian and to ignore the cut-off in $\mathrm{Op}_\mathfrak{h}(a)$, we would obtain

$$\mathcal{H}_\sigma(\mathrm{Op}_\mathfrak{h}(a)) = \int_{h \in H} \mathrm{tr}(\pi(h)\mathrm{Op}_\mathfrak{h}(a)) \overline{\chi_\sigma(h)} \quad (19.2)$$

$$\approx \int_{x \in \mathfrak{g}} a_\mathfrak{h}^\vee(x) \int_{y \in \mathfrak{h}} \chi_\pi(e^{y+x}) \chi_\sigma(e^{-y}), \quad (19.3)$$

which leads formally to (19.1) via the Kirillov formula.

The expectation (19.1) belongs to the general philosophy of the orbit method, whereby restricting a representation of G to the subgroup H corresponds to disintegrating its coadjoint orbit along the projection $\mathfrak{g}^\wedge \rightarrow \mathfrak{h}^\wedge$ (cf. §1.9). The main result in §19 (Theorem 19.3 below) confirms this expectation in a sharper form. Our proof will be along the lines of the second argument discussed above; we will chop the H -integral up into ranges depending on how far the group element h is from the identity.

19.2. A-priori estimates

For orientation and later applications, we record some crude bounds. We abbreviate $S_\delta^m := S_\delta^m(\mathfrak{g}^\wedge)$, and write a for an element of $S_\delta^{-\infty}$ for some fixed $0 \leq \delta < \frac{1}{2}$.

LEMMA 19.1. *We have the very weak bound*

$$\mathcal{H}_\sigma(\mathrm{Op}_\mathfrak{h}(a)) \ll \mathfrak{h}^{-N}. \quad (19.4)$$

Proof. Recall from §18 that for N chosen sufficiently large (relative to G), the map $\mathcal{H}_\sigma: \Psi^{-N}(\pi) \rightarrow \mathbb{C}$ is continuous, uniformly in π and σ . By enlarging N suitably and appealing to Lemma 5.3, we deduce that \mathcal{H}_σ induces a continuous map $\mathcal{H}_\sigma: \Psi_\delta^{-N}(\pi) \rightarrow \mathbb{C}$, uniformly in the same sense. Since $a \in S_\delta^{-N}$, we know by Theorem 5.8 that $\text{Op}_h(a) \in \mathfrak{h}^{-N} \Psi_\delta^{-N}$. The required estimate follows. \square

The technique employed in this last proof for passing from Ψ^{-N} to Ψ_δ^{-N} (after possibly enlarging N) will be applied in the remainder of §19 without explicit mention.

Recall from §9 that we identify $[\mathfrak{g}^\wedge]$ and $[\mathfrak{h}^\wedge]$ with euclidean spaces, equipped with distance functions.

LEMMA 19.2. *Assume that a is supported in some fixed subset $U \subseteq \mathfrak{g}^\wedge$ (independent of h , but otherwise arbitrary). Then, the very strong bound*

$$\mathcal{H}_\sigma(\text{Op}_h(a)) \ll h^N \langle h\lambda_\pi \rangle^{-N} \langle h\lambda_\sigma \rangle^{-N} \quad (19.5)$$

holds unless

$$(h\lambda_\pi, h\lambda_\sigma) \text{ is within distance } o_{h \rightarrow 0}(1) \text{ of the image of } U \text{ in } [\mathfrak{g}^\wedge] \times [\mathfrak{h}^\wedge]. \quad (19.6)$$

Proof. We observe first that \mathcal{H}_σ factors as an $(H \times H)$ -equivariant sequence

$$\Psi^{-\infty}(\pi) \longrightarrow \Psi^{-\infty}(\sigma) \xrightarrow{\text{tr}} \mathbb{C},$$

where the first arrow sends T to the operator on σ given by

$$\int_{s \in H} \text{tr}(sT)\sigma(s^{-1}).$$

This sequence is continuous, uniformly in π and σ , by the same argument as in the proof of part (iii) of Lemma 18.1. These observations allow us to deduce the required implication (in sharper form) from the results in §10.3. \square

19.3. Main result

For convenience, we recall some notation and conventions from §17.

- We write $\mathcal{O}_{\pi,\sigma}$ for the intersection of \mathcal{O}_π with the preimage of \mathcal{O}_σ under the projection $\mathfrak{g}^\wedge \rightarrow \mathfrak{h}^\wedge$.

- Integration over the set $\mathcal{O}_{\pi,\sigma}$, or over its rescaling $h\mathcal{O}_{\pi,\sigma}$, is defined to be zero unless that set is non-empty and \mathbf{H} -stable; in that case, it is an H -torsor, i.e., a closed subset of \mathfrak{g}^\wedge on which H acts simply transitively, and we equip it with the transport of the Haar measure from H .

The simplest case $\delta=0$ of the following result is the relevant one for our applications, but we will pass to the general case $\delta>0$ in the course of the proof.

THEOREM 19.3. *Fix a compact subset $U \subseteq \mathfrak{g}^\wedge$ consisting of \mathbf{H} -stable elements. Let h traverse a sequence of positive reals tending to zero. Fix $0 \leq \delta < \frac{1}{2}$, and let $a \in S_\delta^{-\infty}(\mathfrak{g}^\wedge)$ with $\text{supp}(a) \subseteq U$. Let π and σ be h -dependent tempered irreducible unitary representations of G and H . Then,*

$$\mathcal{H}_\sigma(\text{Op}_h(a)) = \int_{h\mathcal{O}_{\pi,\sigma}} a + O(h^{1-2\delta}), \tag{19.7}$$

More precisely, there are differential operators \mathcal{D}_j on \mathfrak{g}^\wedge with the following properties:

- $\mathcal{D}_0 a = a$;
- \mathcal{D}_j has order $\leq 2j$ and has homogeneous degree j : $\mathcal{D}_j(a_h) = h^j(\mathcal{D}_j a)_h$;
- for each fixed $J \in \mathbb{Z}_{\geq 0}$,

$$\mathcal{H}_\sigma(\text{Op}_h(a)) = \sum_{0 \leq j < J} h^j \int_{h\mathcal{O}_{\pi,\sigma}} \mathcal{D}_j a + O(h^{(1-2\delta)J}). \tag{19.8}$$

We may take the implied constant in (19.8) to be $C\nu(a)$, where

- $C \geq 0$ is a scalar depending at most upon U and δ , and
- ν is a continuous seminorm on $S_\delta^{-\infty}(\mathfrak{g}^\wedge)$ depending at most upon J .

We note that the maps

$$a \mapsto \int_{h\mathcal{O}_{\pi,\sigma}} a,$$

whose domain we take to be the class of symbols a arising in Theorem 19.3, are h -uniformly continuous. Indeed, by the support assumption on a , the integral on the right-hand side vanishes identically unless the pair $(h\lambda_\pi, h\lambda_\sigma)$ belongs to a fixed compact subset of the set of stable pairs; the claim thus follows from the discussion in §17.3.2. Since \mathcal{D}_j has order $\leq 2j$, we deduce in particular that

$$h^j \int_{h\mathcal{O}_{\pi,\sigma}} \mathcal{D}_j a \ll h^{(1-2\delta)j}, \tag{19.9}$$

which explains why (19.8) remains consistent as J varies.

Theorem 19.3 applies readily to $\mathcal{H}_\sigma(\text{Op}_h(a_1) \dots \text{Op}_h(a_k))$ for fixed k and a_1, \dots, a_k as in the hypothesis: just expand $\text{Op}_h(a_1) \dots \text{Op}_h(a_k)$ using the composition formula (5.5), then apply the uniform continuity of $\mathcal{H}_\sigma: \Psi^{-N}(\pi) \rightarrow \mathbb{C}$ to the remainder and the main formula (19.8) to the other terms. Taking the resulting estimate to leading order gives in particular that

$$\mathcal{H}_\sigma(\text{Op}_h(a_1) \dots \text{Op}_h(a_j)) = \int_{h\mathcal{O}_{\pi,\sigma}} a_1 \dots a_k + O(h^{1-2\delta}). \tag{19.10}$$

Recall from §17.3.1 the definition of the notation (e.g.) $\mathcal{O}(\mathfrak{h}\lambda_\sigma)$. Using Theorem 17.2, it follows also that if π admits a limit orbit (\mathcal{O}, ω) , then

$$\mathcal{H}_\sigma(\mathrm{Op}_\mathfrak{h}(a)) \simeq \begin{cases} \int_{\mathcal{O}(\mathfrak{h}\lambda_\sigma)} a, & \text{if } \mathcal{O}_{\pi, \sigma} \neq \emptyset, \\ 0, & \text{otherwise,} \end{cases} \quad (19.11)$$

where $\mathrm{pr}: \mathfrak{g}^\wedge \rightarrow \mathfrak{h}^\wedge$ is the natural projection and $A \simeq B$ means $A = B + o_\mathfrak{h}(1)$. One has also the analogue of (19.11) for multiple symbols, as in (19.10).

For the proof of Theorem 19.3, we may assume that (19.6) is satisfied, since otherwise the integrals $\int_{\mathfrak{h}\mathcal{O}_{\pi, \sigma}} \mathcal{D}_j a$ are eventually identically zero, and so the claim (19.8) follows from the a-priori estimate (19.5).

19.4. Reduction to symbols on the product

We perform here an important technical reduction for the proof of Theorem 19.3. We introduce the notation

$$\mathbf{M} := \mathbf{G} \times \mathbf{H}, \quad \mathbf{S} := \text{diagonal embedding of } \mathbf{H} \text{ in } \mathbf{M},$$

and set

$$\tau := \pi \boxtimes \bar{\sigma},$$

so that τ is a tempered irreducible representation of the reductive group M ; conversely, every such τ arises in this way. We equip S with the transport of Haar from H . Recall from (18.5) that

$$\mathcal{H}_\sigma(\mathrm{Op}_\mathfrak{h}(a)) = \sum_{\substack{v \in \mathcal{B}(\pi) \\ u \in \mathcal{B}(\sigma)}} \int_{s \in H} \langle s \mathrm{Op}_\mathfrak{h}(a)v, v \rangle \langle u, su \rangle.$$

We may rewrite the integrand as

$$\langle s \mathrm{Op}_\mathfrak{h}(a)v, v \rangle \langle u, su \rangle = \langle s \mathrm{Op}_\mathfrak{h}(a)(v \otimes \bar{u}), v \otimes \bar{u} \rangle,$$

where here s acts on τ diagonally while $\mathrm{Op}_\mathfrak{h}(a)$ acts via the restriction of τ to G . Thus

$$\mathcal{H}_\sigma(\mathrm{Op}_\mathfrak{h}(a)) = \sum_{v \in \mathcal{B}(\tau)} \int_{s \in S} \langle s \mathrm{Op}_\mathfrak{h}(a)v, v \rangle.$$

We now exploit the S -invariance to “fatten up” $\mathrm{Op}_\mathfrak{h}(a)$; this will have the effect of replacing the symbol a on \mathfrak{g}^\wedge by a symbol on \mathfrak{m}^\wedge supported close to

$$\mathfrak{s}^\perp := \{\xi \in \mathfrak{m}^\wedge : \xi|_{\mathfrak{s}} = 0\}.$$

To that end, fix $b \in C_c^\infty(\mathfrak{s}^\wedge)$ supported in a small neighborhood of the origin and identically 1 in a smaller neighborhood. We may then form $\text{Op}_h(b)$, which acts on τ via its restriction to H . By invariance of Haar measure, we have

$$\int_{s \in S} \langle s \text{Op}_h(b) \text{Op}_h(a) v, v \rangle = c_0 \int_{s \in S} \langle s \text{Op}_h(a) v, v \rangle,$$

where

$$c_0 := \int_{\mathfrak{s}} \chi b_h^\vee = 1 + O(h^\infty).$$

Combining this with the a-priori bound $\mathcal{H}_\sigma(\text{Op}_h(a)) \ll h^{-O(1)}$ (see (19.4)), we obtain

$$\mathcal{H}_\sigma(\text{Op}_h(a)) = \sum_{v \in \mathcal{B}(\tau)} \int_{s \in S} \langle s \text{Op}_h(b) \text{Op}_h(a) v, v \rangle + O(h^\infty). \quad (19.12)$$

The composition $\text{Op}_h(b) \text{Op}_h(a)$ of operators on τ is a bit subtle because the symbols a and b are defined using different Lie algebras. We may nevertheless compose them using (8.36); what's crucial here is that a and b both have order $-\infty$, and \mathfrak{m} is spanned by \mathfrak{s} and \mathfrak{g} . We obtain in this way—for any fixed $N_1, N_2 \geq 0$, and large enough fixed $J \geq 0$ —an expansion

$$\text{Op}_h(b) \text{Op}_h(a) \equiv \text{Op}_h(a') \pmod{h^{N_1} \Psi^{-N_2}(\tau)} \quad (19.13)$$

where

$$a' := \sum_{0 \leq j < J} h^j \text{Op}_h(b \star^j a) \in C_c^\infty(\mathfrak{m}^\wedge).$$

Arguing as in §18—using now that $\int_S \Xi_M < \infty$ —we see that the formula

$$\mathcal{H}(T) = \int_{s \in S} \text{tr}(sT) = \sum_{v \in \mathcal{B}(\tau)} \int_{s \in S} \langle sT v, v \rangle \quad (19.14)$$

defined initially by the first equality for smooth finite-rank tensors $T \in \tau \otimes \bar{\tau}$, extends continuously to

$$\mathcal{H}: \Psi^{-N}(\tau) \longrightarrow \mathbb{C},$$

uniformly in π and σ , for N sufficiently large but fixed. These observations give an adequate estimate for the contribution to (19.12) from the remainder term implicit in (19.13). Thus $\mathcal{H}_\sigma(\text{Op}_h(a))$ is given up to acceptable error by $\mathcal{H}(\text{Op}_h(a'))$.

Recall that a was assumed supported on a fixed compact collection of \mathbf{H} -stable elements of \mathfrak{g}^\wedge . We claim that if the support of b is chosen small enough, then a' will be supported on a compact collection of \mathbf{S} -stable elements of \mathfrak{m}^\wedge . Indeed, if the support of b is small, then the symbol a' will be supported close to \mathfrak{s}^\perp . But we have an identification

$$\mathfrak{s}^\perp = \{(\xi, -\xi|_{\mathfrak{h}}) : \xi \in \mathfrak{g}^\wedge\} \cong \mathfrak{g}^\wedge, \quad (19.15)$$

which intertwines the coadjoint actions of \mathbf{S} and \mathbf{H} , hence identifies \mathbf{S} -stable elements of \mathfrak{s}^\perp with \mathbf{H} -stable elements of \mathfrak{g}^\wedge . Since the \mathbf{S} -stable locus in \mathfrak{m}^\wedge is open, the claim follows.

The coadjoint multiorbit of τ is given by

$$\mathcal{O}_\tau = \mathcal{O}_\pi \times \mathcal{O}_\sigma = \{(\xi, -\eta) : \xi \in \mathcal{O}_\pi \text{ and } \eta \in \mathcal{O}_\sigma\},$$

so the identification (19.15) induces

$$\mathcal{O}_\tau \cap \mathfrak{s}^\perp \cong \mathcal{O}_{\pi, \sigma},$$

intertwining S and H . In particular, $\mathcal{O}_\tau \cap \mathfrak{s}^\perp$ is an S -torsor; we equip it with the transport of Haar from S , which is then compatible with the above identification. We similarly equip $\mathfrak{h} \mathcal{O}_\tau \cap \mathfrak{s}^\perp$.

Since $b \equiv 1$ in a neighborhood of the origin in \mathfrak{s}^\wedge , we have $\mathbf{a}'(\xi) = a(\xi|_{\mathfrak{g}})$ for all $\xi \in \mathfrak{m}^\wedge$ close to \mathfrak{s}^\perp . In particular,

$$\int_{\mathfrak{h} \mathcal{O}_\tau \cap \mathfrak{s}^\perp} a' = \int_{\mathfrak{h} \mathcal{O}_{\pi, \sigma}} a.$$

More generally, any homogeneous differential operator \mathcal{D}' on \mathfrak{m}^\wedge induces a homogeneous differential operator \mathcal{D} on \mathfrak{g}^\wedge , of the same homogeneity degree and of no larger order, so that $\mathcal{D}'a'(\xi) = \mathcal{D}a(\xi|_{\mathfrak{g}})$ for $\xi \in \mathfrak{s}^\perp$; in particular

$$\int_{\mathfrak{h} \mathcal{O}_\tau \cap \mathfrak{s}^\perp} \mathcal{D}'a' = \int_{\mathfrak{h} \mathcal{O}_{\pi, \sigma}} \mathcal{D}a.$$

The proof of Theorem 19.3 thereby reduces to that of the following (in which we have relabeled (a', τ) to (a, π)).

THEOREM 19.4. *Fix a compact subset $U \subseteq \mathfrak{m}^\wedge$ consisting of \mathbf{S} -stable elements. Let \mathfrak{h} traverse a sequence of positive reals tending to zero. Fix $0 \leq \delta < \frac{1}{2}$, and let $a \in S_\delta^{-\infty}(\mathfrak{m}^\wedge)$ with $\text{supp}(a) \subseteq U$. Let π be an \mathfrak{h} -dependent tempered irreducible unitary representation of M , and $\text{Op}: S^m(\mathfrak{m}^\wedge) \rightarrow \Psi^m(\pi)$ as usual.*

There are differential operators \mathcal{D}_j on \mathfrak{m}^\wedge , satisfying properties analogous to those enunciated in the statement of Theorem 19.3, so that, for each fixed $J \geq 0$,

$$\mathcal{H}(\text{Op}_{\mathfrak{h}}(a)) = \sum_{0 \leq j < J} \int_{\mathfrak{h} \mathcal{O}_\pi \cap \mathfrak{s}^\perp} \mathcal{D}_j a + O(\mathfrak{h}^{(1-2\delta)J}). \quad (19.16)$$

The proof of Theorem 19.4 occupies the remainder of this section. The discussion in §1.10, phrased in terms of microlocalized vectors, might serve as a useful guide to the following arguments.

For the same reasons as explained in §19.2, we may reduce to the case that $h\lambda_\pi$ is within $o(1)$ of the image of U in $[\mathfrak{m}^\wedge]$, so that $a \mapsto \int_{\mathfrak{h}} \mathcal{O}_\pi \cap \mathfrak{h}^\perp$ is h -uniformly continuous, and

$$h^j \int_{\mathfrak{h}} \mathcal{D}_j a \ll h^{(1-2\delta)j}. \tag{19.17}$$

We may assume that $0 \leq \delta < \frac{1}{2}$ is sufficiently large: the problem becomes more general as δ increases, and the asymptotic expansion (19.16) for a given δ (taken with J sufficiently large) implies it for all smaller values.

Recall (from (19.14)) that

$$\mathcal{H}(\text{Op}_h(a)) = \int_{s \in H} \text{tr}(\pi(s)\text{Op}_h(a)).$$

We will analyze below the contribution to the latter from various ranges of $\|s-1\|$, where $\|\cdot\|$ denotes the operator norm on $\text{End}(\mathfrak{m}^\wedge)$. We note that S contains no non-trivial central elements of M , so that the coadjoint representation of M , restricted to S , is a faithful representation; thus $\|s-1\|$ may be regarded as quantifying the distance from s to the identity element 1 of S .

19.5. Small elements give the expected main term

Let j_S be attached as in §2.1 to the group S .

LEMMA 19.5. *Fix $\Theta \in C_c^\infty(\mathfrak{s})$, with $\Theta \equiv 1$ near 0. Fix $J \in \mathbb{Z}_{\geq 0}$ and $\frac{1}{2} > \delta' > \delta$. Then,*

$$\int_{y \in \mathfrak{s}} \Theta\left(\frac{y}{h^{\delta'}}\right) j_S(y) \text{tr}(\pi(\exp(y))\text{Op}_h(a)) = \sum_{0 \leq j < J} h^j \int_{\mathfrak{h}} \mathcal{D}_j a + O(h^{(1-2\delta)J}), \tag{19.18}$$

with \mathcal{D}_j as in the statement of Theorem 19.4.

The basic idea of the proof is as follows. The left-hand side involves traces of group elements close to the identity, which may be evaluated with the Kirillov formula. The conclusion then follows essentially as in the formal sketch (19.2).

In practice, we implement the proof by using the multiplication law for symbols; the $\Theta(y/h^{\delta'})$ factor is accounted for, in the proof below, by the symbol s . The simplest case for the argument is when $\delta=0$. In that case, both a and the symbol s belong to the “easy” symbol class S^0 defined in §4.3, although their supports are on quite different scales.

Proof. Note first of all that it is permissible to prove the statement with $h^{(1-2\delta)J}$ replaced by $h^{J'}$ so long as $J' \rightarrow \infty$ as $J \rightarrow \infty$; one then applies the modified statement with

a larger value of J to obtain the version above, noting that (by (19.17)) the contribution of $\mathcal{D}_j a$ has size $O(h^{(1-2\delta)j})$.

We will first establish the modified form of (19.18) obtained by omitting the factor $j_S(y)$ from the integrand; we will later explain why including this factor does not affect the required conclusion. With this modification, we may write the left-hand side as

$$\int_{y \in \mathfrak{s}} \Theta\left(\frac{y}{h^{\delta'}}\right) \operatorname{tr}(\pi(\exp(y)) \operatorname{Op}_h(a)) = \operatorname{tr}(\operatorname{Op}_h(b) \operatorname{Op}_h(a)), \quad (19.19)$$

where $b \in \mathcal{S}(\mathfrak{s}^\wedge)$ is defined by requiring that $b_h^\wedge(y) = \Theta(y/h^{\delta'})$, i.e., that

$$b(h\eta) = h^{\delta' \dim(\mathfrak{s})} \Theta^\vee(h^{\delta'} \eta).$$

Using our assumptions on Θ , we check readily that

$$b \in h^{\delta'(\dim \mathfrak{s})} S_{1-\delta'}^{-\infty}(\mathfrak{s}^\wedge) \quad (19.20)$$

and

$$\int_{\eta \in \mathfrak{s}^\wedge} \eta^\alpha b(\eta) = 1_{\alpha=0}, \quad (19.21)$$

for all multi-indices α , where 1_X denotes the indicator function for the condition X . The property (19.21) remains valid up to an additive error $O(h^\infty)$ if we replace \mathfrak{s}^\wedge by a ball of fixed radius about the origin (or indeed, by a ball of radius $O(h^{1-\delta''})$, $\delta'' > \delta'$). We should thus think of b as a very strong approximation to the delta function on \mathfrak{s}^\wedge , with thickness at scale $h^{1-\delta'}$.

Although b is defined using a smaller Lie algebra than a , we can compose $\operatorname{Op}_h(a)$ and $\operatorname{Op}_h(b)$ using Theorem 7.4. Due to the crucial assumption $\delta' > \delta$ (which may be rewritten $\delta + (1-\delta') < 1$), we obtain in this way an asymptotic expansion

$$\operatorname{Op}_h(b) \operatorname{Op}_h(a) = \sum_{0 \leq j < J} h^j \operatorname{Op}_h(b \star^j a) + \operatorname{Op}_h(r),$$

involving star products

$$h^j b \star^j a \in h^{(\delta'-\delta)j} S_{\delta'}^{-\infty}(\mathfrak{m}^\wedge)$$

and remainder $r \in h^{(\delta'-\delta)J} S_{\delta'}^{-\infty}(\mathfrak{m}^\wedge)$. By the consequence (12.1) of the Kirillov formula, we obtain a satisfactory estimate for $\operatorname{tr}(\operatorname{Op}_h(r))$ provided that J is taken sufficiently large. To the remaining terms we apply the Kirillov formula expanded in terms of differential operators (see (12.4)). We obtain an asymptotic expansion for the right-hand side of (19.19) as a linear combination taken over small $j_1, j_2 \geq 0$ of integrals

$$h^{j_1+j_2} \int_{\zeta \in h \mathcal{O}_\pi} \partial^{\alpha'}(\zeta^\gamma \cdot \partial^\alpha a(\zeta)) \cdot \partial^\beta b(\zeta) \quad (19.22)$$

involving multi-indices satisfying $|\alpha|+|\beta|-|\gamma|=j_1$, $|\alpha|, |\beta|\leq j_1$ and $|\alpha'|=j_2$. By the product rule applied to $\partial^{\alpha'}$, followed by partial integration, we may rewrite the above as

$$h^j \int_{h\mathcal{O}_\pi} b\mathcal{D}_j a,$$

where $j=j_1+j_2$ and \mathcal{D}_j has the form indicated in the statement of the theorem.

The map $h\mathcal{O}_\pi \rightarrow \mathfrak{s}^\wedge$ has full rank in a neighborhood of ω , so we may fix a small open neighborhood $N_\omega \subset h\mathcal{O}_\pi$ of ω and local coordinates

$$N_\omega \ni \xi = (\xi_1, \xi_2) \in \mathfrak{s}^\wedge \times (h\mathcal{O}_\pi \cap \mathfrak{s}^\perp), \quad \xi_1 \approx 0, \quad \xi_2 \approx \omega,$$

so that

- the coordinate ξ_1 defines the projection to \mathfrak{s}^\wedge , and
- $\xi = (0, \xi)$ for $\xi \in h\mathcal{O}_\pi \cap \mathfrak{s}^\perp$.

The integral of a function f on $h\mathcal{O}_\pi$ supported on N_ω may be expressed in such coordinates as

$$\int_{h\mathcal{O}_\pi} f = \int_{\xi_1 \in \mathfrak{s}^\wedge} \int_{\xi_2 \in h\mathcal{O}_\pi \cap \mathfrak{s}^\perp} f(\xi_1, \xi_2) w(\xi_1, \xi_2) d\xi_1 d\xi_2,$$

where $d\xi_1$ denotes the given Haar measure, $d\xi_2$ the transport of Haar from S , and w is a smooth Jacobian factor. We have $b(\xi)=b(\xi_1)$, and a is supported in N_ω for small enough h , so

$$\int_{h\mathcal{O}_\pi} b\mathcal{D}_j a = \int_{(\xi_1, \xi_2) \in N_\omega} b(\xi_1) w(\xi_1, \xi_2) \mathcal{D}_j a(\xi_1, \xi_2) d\xi_1 d\xi_2.$$

Using Taylor's theorem, we may write

$$w(\xi_1, \xi_2) \mathcal{D}_j a(\xi_1, \xi_2) = w(0, \xi_2) \mathcal{D}_j a(0, \xi_2) + \sum_{1 \leq |\alpha| < A} c_\alpha(\xi_2) \xi_1^\alpha + O(|\xi_1|^A h^{-\delta A}) \quad (19.23)$$

for any fixed A . The compatibility of the top sequence in (16.4) tells us that $w(0, \xi_2)=1$, so the contribution to $\int_{h\mathcal{O}_\pi} b\mathcal{D}_j a$ from the first term on the right-hand side of (19.23) is

$$\begin{aligned} \int_{(\xi_1, \xi_2) \in N_\omega} b(\xi_1) \mathcal{D}_j a(0, \xi_2) d\xi_1 d\xi_2 &= \left(\int_{\xi_1} b(\xi_1) d\xi_1 \right) \left(\int_{\xi_2} \mathcal{D}_j a(0, \xi_2) d\xi_2 \right) + O(h^\infty) \\ &= \int_{h\mathcal{O}_\pi \cap \mathfrak{s}^\perp} \mathcal{D}_j a + O(h^\infty). \end{aligned}$$

The remaining Taylor monomials contribute $O(h^\infty)$, by (19.21) and the remark thereafter. The contribution from the remainder term is dominated by

$$h^{-\delta A} \int_{\xi_1} |\xi_1|^A |b(\xi_1)| d\xi_1 \ll h^{\delta' \dim(\mathfrak{s}) + (1-\delta-\delta')A},$$

due to (19.20) and the definition of b ; informally, $b(\eta)$ is concentrated on $|\eta| \ll h^{1-\delta'}$. Since $\delta, \delta' < \frac{1}{2}$, we have $1-\delta-\delta' > 0$, so this last estimate is adequate for A sufficiently large.

This completes the proof of the modified assertion obtained by omitting j_S . To incorporate that factor, we define the symbol c by requiring that $c_h^\vee(y) = \Theta(y/h^{\delta'}) j_S(y)$, and then follow the previous argument up to (19.22), leading us to consider

$$h^{j_1+j_2} \int_{\zeta \in \mathfrak{h} \mathcal{O}_\pi} \partial^{\alpha'}(\zeta^\gamma \cdot \partial^\alpha a(\zeta) \cdot \partial^\beta c(\zeta)) \quad (19.24)$$

We apply §7.8 to obtain an asymptotic expansion for c given up to acceptable error by a sum over finitely many multi-indices α of the quantities

$$\frac{\partial^\alpha j_S(0)}{\alpha!} (-h)^{|\alpha|} \partial^\alpha b(\zeta);$$

inserting these into (19.24) yields terms of the form (19.22), which we treat as before. \square

19.6. Huge elements contribute negligibly

LEMMA 19.6. *For each fixed $N \geq 0$ there is a fixed $N' \geq 0$ such that*

$$\int_{s \in S: \|s-1\| \geq h^{-N'}} \mathrm{tr}(\pi(s) \mathrm{Op}_h(a)) \ll h^N. \quad (19.25)$$

Proof. We note that, due to the finitude $\int_S \Xi_M < \infty$ and a Lojasiewicz-type inequality, the estimate

$$\int_{s \in S: \|s-1\| \geq X} \Xi_M(s) \ll X^{-\eta}$$

holds for some fixed $\eta > 0$. By the matrix coefficient bounds for tempered representations (see §A) and the trace norm estimate (12.5), we deduce that the left-hand side of (19.25) is dominated for some fixed $L \geq 0$ (independent of N') by

$$\|\Delta^L \mathrm{Op}_h(a) \Delta^L\|_1 \int_{s \in S: \|s-1\| \geq h^{-N'}} \Xi(s) \ll h^{-4L+\eta N'}.$$

Taking N' large enough gives an adequate estimate. \square

19.7. Medium-sized elements contribute negligibly

This section contains the most delicate arguments.

LEMMA 19.7. *Let $\Omega \subset \mathfrak{m}^\wedge$ be a compact collection of \mathbf{S} -stable elements. For each $\omega \in \Omega$ and $s \in S$ there exists $u \in \mathfrak{m}$ with $|u|=1$ so that*

$$|s \cdot u| \asymp |(s \cdot u)\omega| \gg \|s\|^\varepsilon,$$

where the implied constants and the positive quantity ε depend only upon Ω . (Recall from §2.1 that the natural pairing between \mathfrak{m} and \mathfrak{m}^\wedge is denoted by juxtaposition.)

Proof. Fix a maximal compact subgroup K of S and maximal split Cartan subalgebra \mathfrak{a} of S such that $S = K \exp(\mathfrak{a})K$. We may reduce readily to verifying that the conclusion holds when $s \in \exp(\mathfrak{a})$, say $s = \exp(rz)$ with $r > 0$, $z \in \mathfrak{a}$ and $|z|=1$. Consider the weight decompositions for the adjoint and coadjoint actions of z :

$$\mathfrak{m} = \bigoplus_{t \in \mathbb{R}} \mathfrak{m}_t \quad \text{and} \quad \mathfrak{m}^\wedge = \bigoplus_{t \in \mathbb{R}} \mathfrak{m}_t^\wedge.$$

Each $\omega \in \Omega$ is \mathbf{S} -stable, and so has both positive and negative weights (e.g., by applying Hilbert–Mumford to a dense set of 1-parameter subgroups, or by noting that the S -orbit of ω must be topologically closed). By compactness, each $\omega \in \Omega$ has projection onto $\bigoplus_{t \leq -\varepsilon} \mathfrak{m}_t^\wedge$ of norm $\gg 1$. We may choose a weight vector $u \in \bigoplus_{t \geq \varepsilon} \mathfrak{m}_t$ with $|u|=1$ and $|u\omega| \asymp 1$. Then, $s \cdot u$ is a multiple of u with $|s \cdot u| \gg \|s\|^\varepsilon$. The required estimates follow for ω . The same choice of u works for all ω' in a small neighborhood of ω , so we may conclude by the compactness of Ω . \square

LEMMA 19.8. *Suppose $s \in S$ satisfies $\|s-1\| \geq h^{1/2-\eta}$ for some fixed $\eta > 0$. Then,*

$$\mathrm{tr}(\pi(s)\mathrm{Op}_h(a)) \ll h^\infty.$$

The informal idea is to write $a = \sum_i a_i$, where each a_i has very small microlocal support; the trace of $\pi(s)\mathrm{Op}_h(a_i)$ is then small, because the stability condition implies that $s \cdot \mathrm{supp}(a_i)$ and $\mathrm{supp}(a_i)$ are disjoint. It is worth noting that the result is essentially optimal: if $\|s-1\| \asymp h^{1/2}$, that is, if s is just a bit closer to the identity than the scale prescribed by the lemma, then s does not move Planck-scale balls significantly.

Proof. The problem becomes more general as δ increases, so we may and shall assume that

$$\delta \geq \frac{1}{2} - \frac{1}{2}\eta. \tag{19.26}$$

We fix $\varepsilon > 0$ sufficiently small in terms of δ , η and U .

By decomposing a into $h^{-O(1)}$ many pieces, we may assume that it is supported on a ball

$$B(\omega, h^\delta) := \{\xi \in \mathfrak{m}^\wedge : |\xi - \omega| \leq h^\delta\}$$

centered at some $\omega \in U$. We may choose a compactly-supported “envelope” $\psi \in S_\delta^{-\infty}$ with

- $0 \leq \psi \leq 1$,
- $\psi \equiv 1$ on $B(\omega, 2h^\delta)$, and
- $\psi \equiv 0$ on $B(\omega, 3h^\delta)$.

We may write $\text{tr}(\pi(s)\text{Op}_h(a)) = E_1 + E_2$, where

$$E_1 := \text{tr}(\pi(s)\text{Op}_h(a)\text{Op}_h(1-\psi)) \quad \text{and} \quad E_2 := \text{tr}(\text{Op}_h(\psi)\pi(s)\text{Op}_h(a)).$$

Since a and $1-\psi$ have disjoint supports, we see (by §8.10 and §12.3) that $E_1 \ll h^\infty$.

We turn now to E_2 . The idea is that the translation by s of the support of the symbol a is disjoint from the support of ψ . This idea can be implemented rigorously using the operator calculus when $\|s\|$ is not too large. Indeed, suppose first that $\|s-1\| \leq h^{-\varepsilon}$, which means, by (19.26), that

$$h^{\delta-\eta/2} \leq \|s-1\| \leq h^{-\varepsilon}.$$

The upper bound implies in particular that $\|s\| \ll h^{-\varepsilon} \leq h^{-1+\delta+\varepsilon}$, so the hypothesis (5.4) in §5.5 is satisfied; by the conclusion of that section, the operator norm of

$$\pi(s)\text{Op}_h(\psi)\pi(s)^{-1} - \text{Op}_h(s \cdot \psi)$$

is negligible, so we reduce to showing that

$$\|\text{Op}_h(s \cdot \psi)\text{Op}_h(a)\|_1$$

is negligible. As s distorts Lie algebra elements by at most $\|s\| \ll h^{-\varepsilon}$, we have $s \cdot \psi \in S_{\delta+\varepsilon}^{-\infty}$, and may assume that $\delta+\varepsilon < \frac{1}{2}$. It will thus suffice to verify that $s \cdot \psi$ and a have disjoint supports. To that end, we need only verify for $\xi \in B(\omega, 3h^\delta)$ that

$$\|s-1\| \geq h^{\delta-\eta/2} \implies |s \cdot \xi - \xi| \gg h^{\delta-\varepsilon}, \quad (19.27)$$

say. Note that the union of the sets $B(\omega, 3h^\delta)$, as ω varies over U and h over sufficiently small positive reals, is contained in a fixed compact collection of stable elements. The estimate (19.27) follows in the range $\|s-1\| \geq h^\varepsilon$ from Hilbert–Mumford, as in the proof of Lemma 19.7, and in the remaining range $h^\varepsilon \geq \|s-1\| \geq h^{\delta-\eta/2}$ from Lie algebra considerations, using that ξ has trivial \mathfrak{g} -centralizer.

It remains to handle the case that $\|s-1\| \geq h^{-\varepsilon}$, and hence $\|s\| \gg h^{-\varepsilon}$. Direct application of the symbol calculus does not work as well here, because $s \cdot \psi$ is overly distorted. We instead construct convolution operators along well-chosen lines inside \mathfrak{m} , corresponding to 1-parameter subgroups in G , so that there is no issue of distortion.

Fix a Fourier transform between \mathbb{R} and its Pontryagin dual $\mathbb{R}^\wedge = i\mathbb{R}$, and fix $\Theta \in C_c^\infty(\mathbb{R}^\wedge)$ and $\chi \in C_c^\infty(\mathbb{R})$, each identically 1 in neighborhoods of the respective origins. Choose $u \in \mathfrak{m}$ as in Lemma 19.7, so that $|u|=1$ and $|s \cdot u| \gg h^{-\varepsilon^2}$ and

$$\frac{s \cdot u}{|s \cdot u|} \omega \asymp 1. \quad (19.28)$$

Set $r := h^{1-\varepsilon^3}/|s \cdot u|$, so that $r \ll h^{1+\varepsilon^2-\varepsilon^3}$, and

$$\begin{aligned} \mathcal{C}_1 &:= \int_{t \in \mathbb{R}} \frac{\Theta^\vee(t/r)}{r} \chi\left(\frac{h^{\varepsilon^3} t}{r}\right) \pi(\exp(tu)), \\ \mathcal{C}_2 &:= \pi(s) \mathcal{C}_1 \pi(s)^{-1} = \int_{t \in \mathbb{R}} \frac{\Theta^\vee(t/h^{1-\varepsilon^3})}{h^{1-\varepsilon^3}} \chi\left(\frac{t}{h^{1-2\varepsilon^3}}\right) \pi\left(\exp\left(t \frac{s \cdot u}{|s \cdot u|}\right)\right). \end{aligned}$$

Informally, we should think of \mathcal{C}_1 as a convolution operator in the u direction utilizing a bump function of width substantially smaller than h , whereas \mathcal{C}_2 is a convolution operator in the $s \cdot u$ direction utilizing a bump function of width slightly greater than h .

It will suffice to verify that (with $\|\cdot\|_\infty$ the operator norm)

$$\|\mathcal{C}_1 \text{Op}_h(a) - \text{Op}_h(a)\|_\infty \ll h^\infty, \quad (19.29)$$

$$\|\text{Op}_h(\psi) \mathcal{C}_2\|_\infty \ll h^\infty, \quad (19.30)$$

because then, writing \equiv to denote agreement up to $O(h^\infty)$ and applying §8.10 and §12.3, we have

$$E_2 \equiv \text{tr}(\text{Op}_h(\psi) \pi(s) \mathcal{C}_1 \text{Op}_h(a)) = \text{tr}(\text{Op}_h(\psi) \mathcal{C}_2 \pi(s) \text{Op}_h(a)) \equiv 0.$$

To establish (19.29) and (19.30), let $l_1, l_2 \subseteq \mathfrak{m}$ denote the lines spanned by u and $s \cdot u$, respectively, and observe that we may write $\mathcal{C}_j = \text{Op}_h(b_j)$, with $b_1 \in S_0^0(l_1)$ and $b_2 \in S_{\varepsilon^3}^{-\infty}(l_2)$ satisfying

- $b_1(\xi) = 1 + O(h^\infty)$ for $|\xi| \leq h^{-\varepsilon^4}$, and
- $b_2(\xi) = O(h^\infty)$ for $|\xi| \geq h^{\varepsilon^4}$,

and similarly for derivatives. In particular, b_1 is approximately 1 on the image of the support of a , while (by (19.28)) b_2 is approximately 0 on the image of the support of ψ . The required estimates (19.29) and (19.30) follow from the asymptotic expansion (8.36). \square

19.8. Completion of the proof

We note for $X \geq 1$ that

$$\text{vol}(\{s \in S : \|s-1\| \leq X\}) \ll X^{O(1)} \quad (19.31)$$

(cf. [Wil, Lemma 2.A.2.4]). The definition (19.14) and the results in §19.6 and §19.7 show that, for any fixed $\eta > 0$,

$$\mathcal{H}(\mathrm{Op}_h(a)) = \int_{s \in S: \|s-1\| \geq h^{1/2-\eta}} \mathrm{tr}(\pi(s)\mathrm{Op}_h(a)) + O(h^\infty).$$

We now write $s = \exp(y)$, pull the integral back to the Lie algebra, and combine §19.7 and §19.5 to derive the required asymptotic expansion.

Appendix A. Some technicalities related to the Plancherel formula

The aim of this appendix, which the reader is encouraged to skip, is to supply the unsurprising details required by the proofs in §18.

Let F be a local field, either archimedean or non-archimedean. Let \mathbf{G} be a reductive group over F .

We denote as usual by $\widehat{G}_{\mathrm{temp}} \subseteq \widehat{G}$ the tempered dual of G , thus each $\pi \in \widehat{G}_{\mathrm{temp}}$ is a tempered irreducible unitary representation of G .

We always choose a Haar measure dg on G and a maximal compact subgroup $K := K_G$ of G . For a unitary representation π of G , we denote by $\mathcal{B}(\pi)$ an orthonormal basis consisting of K -isotypic vectors. For $f \in L^1(G)$ we define

$$\pi(f) := \int_{g \in G} \pi(f) f(g) dg,$$

as usual.

When F is archimedean, we retain the notation of Part I ($\mathfrak{U}, \Delta, \dots$), applied to the real Lie group underlying G .

A.1. Uniform bounds for K -types

Assume that F is archimedean. The proof of [Kn, Lemma 10.4] shows that there is an element κ of the universal enveloping algebra of K with the following properties:

- (i) κ acts on each irreducible representation τ of K by a scalar κ_τ ;⁽⁶⁾
- (ii) $\dim(\tau) \leq \kappa_\tau^{1/2}$;
- (iii) $\sum_{\tau \in \widehat{K}} \kappa_\tau^{-1}$ is finite.

(Explicitly, one may take

$$\kappa = -c \sum_{x \in \mathcal{B}(\mathrm{Lie}(K))} x^2$$

for large enough $c > 0$.)

⁽⁶⁾ We use the notation κ_τ for what Knapp denotes $d_\lambda^2(1 + \|\lambda|_{Z_\mathfrak{t}}\|^2)$.

LEMMA A.1. *Let π be an irreducible admissible representation of G .*

- (i) *Let $v \in \pi$ be τ -isotypic. Then, $\dim(Kv)^{1/2} \|v\| \leq \|\kappa v\|$.*
- (ii) *$\pi(\kappa)$ is positive and invertible. $\mathrm{tr}(\pi(\kappa)^{-2}) \leq C$, where C depends only upon G .*

Proof. (i) By [Kn, Theorem 8.1], we have $n_\tau := \mathrm{Hom}_K(\tau, \pi) \leq \dim(\tau)$, so that

$$\dim(Kv) \leq n_\tau \dim(\tau) \leq \dim(\tau)^2.$$

The conclusion follows from the enunciated properties of κ .

- (ii) $\mathrm{tr}(\pi(\kappa)^{-2}) \leq \sum_{\tau \in \hat{K}} n_\tau \dim(\tau) \kappa_\tau^{-2} \leq \sum_{\tau \in \hat{K}} \kappa_\tau^{-1} < \infty$. □

A.2. Bounds for matrix coefficients

Assume that π is tempered.

A.2.1. By [CHH], there is a function $\Xi := \Xi_G : G \rightarrow \mathbb{R}_{>0}$ (depending also upon K), called the *Harish–Chandra spherical function*, with the following property: for any $\pi \in \widehat{G}_{\mathrm{temp}}$ and any K -finite $u, v \in \pi$, one has

$$|\langle gu, v \rangle| \leq \Xi(g) (\dim Ku)^{1/2} (\dim Kv)^{1/2} \|u\| \|v\| \tag{A.1}$$

for all $g \in G$. The function Ξ descends to G/Z , where Z denotes the center of G , and tends to zero at infinity on G/Z .

A.2.2. Assume now that F is archimedean. We may then readily translate the bound (A.1) in terms of the Sobolev norms defined in §3.2.

LEMMA A.2. *For $\pi \in \widehat{G}_{\mathrm{temp}}$, $g \in G$ and $u, v \in \pi^s$,*

$$|\langle gu, v \rangle| \leq c \Xi(g) \|u\|_{\pi^s} \|v\|_{\pi^s}, \tag{A.2}$$

where $c \geq 0$ and $s \in \mathbb{Z}_{\geq 0}$ depend only upon G .

Proof. Let κ be as in §A.1.

Let $v \in \pi$ be K -finite; write its isotypic decomposition as $v = \sum_\tau v_\tau$. By part (i) of Lemma A.1,

$$\sum_\tau \dim(Kv_\tau)^{1/2} \|v_\tau\| \leq \sum_\tau \kappa_\tau^{-1/2} \|\kappa v_\tau\| \leq \left(\sum_\tau \kappa_\tau^{-1} \right)^{1/2} \|\kappa v\|. \tag{A.3}$$

To prove (A.2), we may assume by continuity that the vectors u and v are K -finite. We decompose such vectors into their K -isotypic components, apply (A.1) to the inner product arising from each pair of components, and then apply (A.3), giving

$$|\langle gu, v \rangle| \leq \left(\sum \kappa_\tau^{-1} \right) \Xi(g) \|\kappa u\| \|\kappa v\|.$$

We conclude by appeal to the simple consequence (3.1) of the definition of $\|\cdot\|_{\pi^s}$. \square

A.3. Plancherel formula

Let π be an irreducible unitary representation of G . We denote by χ_π its distributional character, as in §6.2.

The Plancherel formula asserts that, for $f \in C_c^\infty(G)$, we have the identity

$$f(1) = \int_{\pi \in \widehat{G}_{\text{temp}}} \chi_\pi(f), \tag{A.4}$$

with the latter integral taken with respect to a certain measure on $\widehat{G}_{\text{temp}}$, called the *Plancherel measure* dual to dg . For $n \in \mathbb{Z}_{\geq 0}$ large enough in terms of G , the formula extends by continuity to the class of n -fold differentiable compactly-supported functions.

A.4. Some crude growth bounds

A.4.1. Assume that F is non-archimedean. Let U be a compact open subgroup of G . For any admissible representation π of G (e.g., any irreducible unitary representation), the space π^U of U -fixed vectors is finite-dimensional. By applying the Plancherel formula to the normalized characteristic function of U , we see moreover that

$$\int_{\pi \in \widehat{G}_{\text{temp}}} \dim(\pi^U) < \infty.$$

A.4.2. Assume that F is archimedean.

LEMMA A.3. *There is $N \in \mathbb{Z}_{\geq 0}$, depending only upon G , such that*

- (i) $\sup_{\pi \in \widehat{G}} \text{tr}(\pi(\Delta^{-N})) < \infty$;
- (ii) $\int_{\pi \in \widehat{G}_{\text{temp}}} \text{tr}(\pi(\Delta^{-N})) < \infty$.

This is likely well known; we record a proof for completeness.

Proof. Here we require implied constants to be uniform in π .

(i) Let $\kappa \in \mathfrak{U}$ be as in §A.1. Assume that N exceeds twice the degree of κ . By Lemma 8.5, the operator $A := \pi(\kappa)^2 \pi(\Delta^{-N})$ then has uniformly bounded operator norm. Since $\pi(\kappa)^{-2}$ is positive, it follows that

$$\mathrm{tr}(\pi(\Delta^{-N})) \ll \mathrm{tr}(\pi(\kappa)^{-2}).$$

We conclude by part (ii) of Lemma A.1.

(ii) By spectral theory, it suffices to establish the modified conclusion obtained by replacing Δ with the rescaled variant Δ_h (cf. §12.2) for some fixed $h \in (0, 1]$. Set

$$b(\xi) := \langle \xi \rangle^{-N}.$$

Let $\pi \in \widehat{G}_{\mathrm{temp}}$. Then, $T_\pi := \mathrm{Op}_h(b; \pi)^2$ is positive definite. By §12.2, we have

$$\mathrm{tr}(\pi(\Delta_h^{-2N})) \asymp \mathrm{tr}(T_\pi)$$

for h small enough. Let $n \in \mathbb{Z}_{\geq 0}$ large enough that the Plancherel formula holds for n -fold differentiable functions $f \in C_c(G)$, and assume that N is large enough in terms of n . By the composition formula, combined with §8.2, we then have $T_\pi = \pi(f)$, where $f \in C_c(G)$ is n -fold differentiable. Thus,

$$\int_{\pi \in \widehat{G}_{\mathrm{temp}}} \mathrm{tr}(\pi(\Delta_h^{-2N})) \ll \int_{\pi \in \widehat{G}_{\mathrm{temp}}} \mathrm{tr}(\pi(f)) = f(1) < \infty. \quad \square$$

A.5. Plancherel formula, II

We record an extension of the Plancherel formula (A.4) to a larger space of functions that we denote by $\mathcal{F} := \mathcal{F}_G$; in brief, it consists of functions whose $G \times G$ derivatives lie in $L^1(G, \Xi)$.

- In the non-archimedean case, we define \mathcal{F}^U , for each compact open subgroup U of G , to be the space of bi- U -invariant functions $f: G \rightarrow \mathbb{C}$ satisfying

$$\int_{g \in G} \Xi(g) |f(g)| dg < \infty, \quad (\text{A.5})$$

equipped with the evident topology (§1.14.6); we then set $\mathcal{F} := \cup \mathcal{F}^U$, equipped with the direct limit topology.

- In the archimedean case (so that G is regarded as a real Lie group), we take for \mathcal{F} the space of smooth functions $f: G \rightarrow \mathbb{C}$ each of whose $(\times G)$ -derivatives (i.e., allowing applications of both left- and right-invariant differential operators) satisfies the analogue of (A.5); we equip \mathcal{F} with its evident topology (§1.14.6).

In either case, observe that $C_c^\infty(G)$ is dense in \mathcal{F} and (by the Sobolev lemma) that point evaluations on \mathcal{F} are continuous. Let $\pi \in \widehat{G}_{\text{temp}}$. For $f \in \mathcal{F}$, we wish to define and study an operator $\pi(f)$ on the space π^∞ of smooth vectors in π . It is natural to ask that this operator satisfies

$$\langle \pi(f)u, v \rangle = \int_{g \in G} f(g) \langle gu, v \rangle dg \quad (\text{A.6})$$

for $u, v \in \pi^\infty$; note that the right-hand side of (A.6) converges absolutely, due to (A.2) and (A.5).

LEMMA A.4. *Let $f \in \mathcal{F}$ and $\pi \in \widehat{G}_{\text{temp}}$.*

- (i) *There is a unique continuous linear map $\pi(f): \pi^\infty \rightarrow \pi^\infty$ for which (A.6) holds.*
- (ii) *The map $f \mapsto \pi(f)$ is continuous for the trace norm $\|\cdot\|_1$ on the target.*
- (iii) *The map*

$$f \mapsto \int_{\pi \in \widehat{G}_{\text{temp}}} \|\pi(f)\|_1$$

is finite-valued and continuous.

- (iv) *The Plancherel formula (A.4) remains valid.*

Proof. We may initially define $\pi(f)u$ as the anti-linear functional on π^∞ for which (A.6) holds. We aim then to verify that $\pi(f)u$ is represented by a smooth vector and that the resulting map has the required properties.

In the non-archimedean case, f is $U \times U$ -invariant for some open subgroup U of G . We verify readily that the functional $\pi(f)u$ is then U -invariant, hence represented by a unique element of the finite-dimensional space π^U . The remaining assertions may be verified by simpler analogues of the arguments to follow (using §A.4.1 instead of §A.4.2). We turn henceforth to the details of archimedean case.

- (i) Choose $N \in \mathbb{Z}_{\geq 0}$ large enough that we have the bound (cf. §A.2)

$$\langle gu, v \rangle \ll \Xi(g) \|\Delta^N u\| \|\Delta^N v\|$$

for $u, v \in \pi^\infty$. Since

$$\int_{g \in G} f(g) \langle gu, v \rangle dg = \int_{g \in G} (\Delta^{2N} * f * \Delta^{2N})(g) \langle g \Delta^{-2N} u, \Delta^{-2N} v \rangle dg,$$

we then have

$$|\langle \pi(f)u, v \rangle| \ll \nu(f) \|\Delta^{-N} u\| \|\Delta^{-N} v\| \quad (\text{A.7})$$

for some continuous norm

$$\nu(f) := \int_{g \in G} \Xi(g) |\Delta^{2N} * f * \Delta^{2N}(g)|$$

on \mathcal{F} . By summing over v in an orthonormal basis and appealing to §A.4.2, we deduce that $\pi(f)u$ is represented by an element of the Hilbert space π^0 . By a similar argument applied to $\Delta^n \pi(f)u$ for each $n \in \mathbb{Z}_{\geq 0}$, we deduce that $\pi(f)u \in \pi^\infty$ and that the induced map $\pi(f): \pi^\infty \rightarrow \pi^\infty$ is continuous.

(ii) The trace norm of $\pi(f)$ is bounded by

$$\sum_{u,v \in \mathcal{B}(\pi)} |\langle \pi(f)u, v \rangle|,$$

so we conclude by summing (A.7) and appealing to §A.4.2.

(iii) We argue similarly, using now also part (ii) of Lemma A.3.

(iv) We appeal to continuity and the density of $C_c^\infty(G)$ in \mathcal{F} . □

A.6. Proof of Lemma 18.1

By §A.2 and (18.1), the function f defined in (18.6) belongs to the space \mathcal{F}_H from §A.5. Assertions (i) and (ii) thus follow from §A.5.

To establish (iii), we first show that the map

$$\begin{aligned} \Phi: \pi \otimes \bar{\pi} &\longrightarrow \mathcal{F}_H, \\ v_1 \otimes \bar{v}_2 &\longmapsto f, \end{aligned}$$

extends continuously to

$$\Phi: \Psi^{-\infty}(\pi) \longrightarrow \mathcal{F}_H.$$

To that end, observe that for each $t_1, t_2 \in \mathfrak{U}$, we may write

$$t_1 * \Phi(v_1 \otimes \bar{v}_2) * t_2 = \Phi(t_1 v_1 \otimes \overline{t_2^\iota v_2}),$$

with ι being the standard involution on \mathfrak{U} . Thus for each continuous seminorm ν on \mathcal{F}_H there are $C \geq 0$ and $N \in \mathbb{Z}_{\geq 0}$ so that

$$\nu(\Phi(v_1 \otimes \bar{v}_2)) \leq C \|\Delta^N v_1\| \|\Delta^N v_2\|.$$

By summing over orthonormal bases and appealing to §A.4.2, we deduce that

$$\nu(\Phi(T)) \ll \|\Delta^N T \Delta^N\|_2$$

for all $T \in \pi \otimes \bar{\pi}$, where $\|\cdot\|_2$ denotes the Hilbert–Schmidt norm. By Theorem 12.2 (iv), we deduce that Φ extends continuously to $\Psi^{-\infty}(\pi)$ with the required uniformity. By the definition of the topology on $\Psi^{-\infty}(\pi)$, we may pass to $\Psi^{-N}(\pi)$ for some fixed N .

The same argument gives, for $T \in \Psi^{-N}(\pi)$ and with notation as in (18.5), that each of the quantities

$$\sum_{u \in \mathcal{B}(\sigma)} \left| \int_{s \in H} \operatorname{tr}(sT) \langle u, su \rangle \right| \quad \text{and} \quad \sum_{\substack{v \in \mathcal{B}(\pi) \\ u \in \mathcal{B}(\sigma)}} \left| \int_{s \in H} \langle sTv, v \rangle \langle u, su \rangle \right|$$

is finite and depends continuously upon T . The required identity (18.5) thus follows by continuous extension from the finite-rank case.

Part IV. Inverse branching

20. Overview

Let (\mathbf{G}, \mathbf{H}) be a GGP pair over a local field F of characteristic zero. Fix a tempered irreducible representation π of G . More precisely, we abuse notation in what follows, as in §18, by working implicitly with underlying spaces of smooth vectors.

Let \widehat{H} denote the unitary dual of H , $\widehat{H}_{\text{temp}} \subseteq \widehat{H}$ the tempered dual, and $\widehat{H}_{\text{temp}}^\pi$ the π -distinguished subset, i.e.,

$$\widehat{H}_{\text{temp}}^\pi := \{ \sigma \in \widehat{H}_{\text{temp}} : \text{there is a non-zero } H\text{-equivariant map } \ell_\sigma : \pi \rightarrow \sigma \}.$$

For each $\sigma \in \widehat{H}_{\text{temp}}^\pi$, the discussion in §18 gives us a map $\mathcal{H}_\sigma : \pi \otimes \bar{\pi} \rightarrow \mathbb{C}$. It is known at least for F non-archimedean (see [BP3, Theorem 5] in the unitary case and [Wll, Proposition 5.7] in the special orthogonal case) that \mathcal{H}_σ is non-zero precisely when

$$\sigma \in \widehat{H}_{\text{temp}}^\pi.$$

Recall that \mathcal{H}_σ is expected to satisfy the positivity condition (18.7), and that this expectation is a theorem in the non-archimedean case [Sve].

If σ is tempered and \mathcal{H}_σ is non-zero and satisfies the expected positivity condition, then we may write

$$\mathcal{H}_\sigma(v_1 \otimes \bar{v}_2) = \sum_{u \in \mathcal{B}(\sigma)} \int_{s \in H} \langle sv_1, v_2 \rangle \langle u, su \rangle = \langle \ell_\sigma(v_1), \ell_\sigma(v_2) \rangle \quad (20.1)$$

for some ℓ_σ as above, uniquely defined up to a scalar of magnitude 1. If \mathcal{H}_σ vanishes, then we take $\ell_\sigma = 0$.

In the unexpected case that the positivity condition (18.7) is violated (necessarily for F archimedean), we define ℓ_σ by requiring that (20.1) hold, up to some scalar of magnitude 1.

We crudely extend the definition of ℓ_σ to non-tempered $\sigma \in \widehat{H}$ by choosing an H -equivariant map $\ell_\sigma: \pi \rightarrow \sigma$, possibly zero but non-zero if possible (and in that case, unique up to a scalar), and requiring that (20.1) hold. Thus \mathcal{H}_σ , for non-tempered σ , is defined only up to multiplication by a positive real.

We have defined an association

$$\begin{aligned} \pi \otimes \bar{\pi} &\longrightarrow \{\text{functions } \widehat{H} \rightarrow \mathbb{C}\}, \\ T &\longmapsto [\sigma \mapsto \mathcal{H}_\sigma(T)]. \end{aligned}$$

Let us pause to speak informally about the relevance of this association to our aims. Recall, from §1.7, that we may think of self-adjoint elements $\sum_j v_j \otimes \bar{v}_j \in \pi \otimes \bar{\pi}$ as weighted families of vectors in π , and hence the above association as an assignment

$$\left\{ \begin{array}{l} \text{weighted families} \\ \text{of vectors } v \text{ in } \pi \end{array} \right\} \longrightarrow \left\{ \begin{array}{l} \text{weighted families of} \\ \text{representations } \sigma \text{ of } H \end{array} \right\}.$$

To implement the basic strategy of this paper (cf. §1.6 and §1.10), we would like to know that we can approximate any reasonable family of representations in this way, while retaining some control over the family of vectors achieving the approximation. This is the “inverse branching problem” alluded to in the title; by comparison, the classical *branching problem* concerns how a representation of a group decomposes upon restriction to a subgroup, or perhaps how individual vectors decompose.

In the global setting (Part V), the pairs (π, σ) as above will arise as the local components of a pair (Π, Σ) of automorphic forms over a number field, taken unramified outside some fixed set R of places containing the archimedean places. We will single out an individual archimedean place $\mathfrak{q} \in R$ as the “interesting” one, assume that the relevant groups are compact at all other archimedean places, and aim to study families with “increasing frequency at \mathfrak{q} ” and “fixed level at \mathfrak{p} ” for all $\mathfrak{p} \in R \setminus \{\mathfrak{q}\}$, with some fairly flexible definition of “fixed level”. Motivated by this aim, we consider here in Part IV the “inverse branching problem” indicated above in the following aspects:

- For *varying* families of representations σ , taken over a suitable scaling limit, and with F an archimedean local field (§22).
- For *fixed* families of σ in either of the following cases:
 - (the trivial case in which) H is compact (§21);
 - F is non-archimedean (§24), after some general preliminaries (§23).

An important subtlety is that the families of interest to us will *not* in general be “microlocally separated” from the complementary series, e.g., via their infinitesimal character. We must nevertheless exclude the latter from our final formula, due to the absence

of a general conjecture along the lines of Ichino–Ikeda in the non-tempered case. These considerations motivate the estimate (22.11) and are responsible for the main difficulties in §24.

21. The case of compact groups

Suppose that H is compact. Then, $\widehat{H}_{\text{temp}}^\pi$ is a discrete countable set. The Hermitian forms \mathcal{H}_σ describe the canonical decomposition

$$\pi|_H \cong \bigoplus_{\sigma \in \widehat{H}_{\text{temp}}^\pi} \sigma.$$

Thus, for any finitely-supported function $k: \widehat{H}_{\text{temp}}^\pi \rightarrow \mathbb{C}$, there exists $T \in \pi \otimes \bar{\pi}$ such that

$$\mathcal{H}_\sigma(T) = k(\sigma)$$

for all $\sigma \in \widehat{H}_{\text{temp}}^\pi$. If k is valued in the non-negative reals, then we may take T to be positive definite.

22. The distinguished archimedean place

We assume here that F is archimedean. By restriction of scalars, we may regard \mathbf{G} and \mathbf{H} as real reductive groups.

22.1. Setup

We allow the tempered irreducible representation π of G to vary with a positive parameter $\mathfrak{h} \rightarrow 0$. We assume that π has a limit orbit (see §11.4)

$$(\mathcal{O}, \omega) = \lim_{\mathfrak{h} \rightarrow 0} (\mathfrak{h} \mathcal{O}_\pi, \omega_{\mathfrak{h} \mathcal{O}_\pi}).$$

As in §17, we write $\mathcal{O}_{\text{stab}} \subseteq \mathcal{O}$ for the subset of \mathbf{H} -stable elements. We recall that, for each $\mu \in [\mathfrak{h}^\wedge] \cap \text{image}(\mathcal{O}_{\text{stab}})$, the preimage $\mathcal{O}(\mu)$ of $\{\mu\}$ in \mathcal{O} is an H -torsor. The map $\mathcal{O}_{\text{stab}} \rightarrow [\mathfrak{h}^\wedge] \cap \text{image}(\mathcal{O}_{\text{stab}})$ is a principal H -bundle, with fibers $\mathcal{O}(\mu)$.

We assume given a Haar measure on H ; as explained in §17, this choice defines measures on \mathfrak{h} , \mathfrak{h}^\wedge , $[\mathfrak{h}^\wedge]$, and on the sets $\mathcal{O}(\mu)$ as above.

22.2. Orbit-distinction

Let $\sigma \in \widehat{H}_{\text{temp}}$. Recall (from §20) that σ is *distinguished* by π if $\text{Hom}_H(\pi, \sigma) \neq 0$. We say that σ is *orbit-distinguished* by π if $\mathcal{O}_{\pi, \sigma}$ —the intersection of \mathcal{O}_π with the preimage of \mathcal{O}_σ —is non-empty.

Remark 22.1. Our asymptotic expansion of relative characters (Theorem 19.3) implies that

- (i) if $h\lambda_\sigma$ belongs to a fixed compact subset E of $[\mathfrak{h}^\wedge] \cap \text{image}(\mathcal{O}_{\text{stab}})$, and
- (ii) if $h > 0$ is small enough in terms of E ,

then orbit-distinction implies distinction. One expects also the converse implication, that distinction implies orbit-distinction under the stated hypotheses. This would follow from the following:

- Strong multiplicity 1 for archimedean L -packets. This is addressed in unitary cases by the preprint [BP3] and in orthogonal cases by the recent preprint [L].
- That distinction implies non-vanishing of the matrix coefficient integral, known in p -adic cases (cf. §20) and in the unitary archimedean case [BP3, Theorem 5], and likely provable in the orthogonal archimedean case.

In any event, orbit-distinction seems easier to check than distinction, so we are content to formulate our main results in terms of the former notion.

22.3. Main result

Let

$$k \in C_c^\infty([\mathfrak{h}^\wedge] \cap \text{image}(\mathcal{O}_{\text{stab}})). \tag{22.1}$$

For each $h > 0$ we define a function $k_h: \widehat{H} \rightarrow \mathbb{C}$ by setting

$$k_h(\sigma) := k(h\lambda_\sigma)$$

if σ is tempered and $\mathcal{O}_{\pi, \sigma} \neq \emptyset$; otherwise, we set $k_h(\sigma) := 0$.

We may find

$$\text{precompact open subsets } U \subset [\mathfrak{h}^\wedge] \text{ and } V \subset \mathfrak{g}^\wedge, \tag{22.2}$$

with

$$\bar{V} \text{ consisting of } \mathbf{H}\text{-stable elements,} \tag{22.3}$$

such that

$$\text{supp}(k) \subseteq \text{image}(V), \quad \text{image}(\bar{V}) \subseteq U \quad \text{and} \quad \bar{U} \subseteq \text{image}(\mathcal{O}_{\text{stab}}). \tag{22.4}$$

Since $\mathcal{O}_{\text{stab}}$ is a submanifold of \mathfrak{g}^\wedge and the map $\mathcal{O}_{\text{stab}} \rightarrow [\mathfrak{h}^\wedge]$ is a principal H -bundle over its image, we may readily find $a \in C_c^\infty(V)$ such that, for each $\mu \in [\mathfrak{h}^\wedge]$,

$$\int_{\mathcal{O}(\mu)} a = k(\mu). \tag{22.5}$$

From (22.5) and the asymptotic formulas in §19.3 it follows that

$$\mathcal{H}_\sigma(\text{Op}_\hbar(a)) = k_\hbar(\sigma) + o_{\hbar \rightarrow 0}(1) \quad \text{for all } \sigma \in \widehat{H}_{\text{temp}} \text{ with } \hbar \lambda_\sigma \in U. \tag{22.6}$$

If k is real-valued, then we may arrange that a is real-valued. In the language of §1.10 and §20, we have achieved our goal of producing a weighted family of vectors—that obtained by writing $\text{Op}_\hbar(a) = \sum_j v_j \otimes \bar{v}_j$ —that picks off the weighted family of representations described by k_\hbar . We note in passing also that, by (17.5), we have

$$\int_{\mathcal{O}} a \, d\omega = \int_{\mu \in [\mathfrak{h}^\wedge]} \int_{\mathcal{O}(\mu)} a = \int_{[\mathfrak{h}^\wedge]} k, \tag{22.7}$$

with integration over $[\mathfrak{h}^\wedge]$ defined by the normalized affine measure.

We aim now to elaborate upon this observation in somewhat technical ways that will turn out to be convenient for our global applications. It will be useful to work with “positive-definite families”, such as those attached to $\text{Op}_\hbar(a)^2$ for real-valued a , and to bound the error in (22.6) in terms of another such family.

We will also need to say something about non-tempered σ . In that case, we have only thus far (cf. §20) normalized \mathcal{H}_σ up to a scalar. It will be convenient now to impose the following more precise normalization, again motivated by global considerations (cf. §29): we suppose given an \hbar -dependent family of maps \mathcal{H}_σ that factors as a composition

$$\Psi^{-\infty}(\pi) \longrightarrow \Psi^{-\infty}(\sigma) \xrightarrow{\text{tr}} \mathbb{C} \tag{22.8}$$

with the first arrow \hbar -uniformly continuous. In practice, this is a fairly weak requirement. We note that the analogous continuity holds in the tempered case by the discussion of §19.2.

THEOREM 22.2. *Let k , U and V be as in (22.1) and (22.2), satisfying the assumptions (22.3) and (22.4). Assume that $k \geq 0$. Then, for each $\varepsilon > 0$ and $N \in \mathbb{Z}_{\geq 0}$, there exist non-negative $a, a_1, a_2, a_{\text{nt}} \in C_c^\infty(V)$ with the following properties:*

- (i) $\int_{\mathcal{O}} a_1^2 \, d\omega$ is bounded by a constant depending only upon k and V , while $\int_{\mathcal{O}} a_2^2 \, d\omega$ and $\int_{\mathcal{O}} a_{\text{nt}}^2 \, d\omega$ are bounded by ε .
- (ii) $|\int_{[\mathfrak{h}^\wedge]} k - \int_{\mathcal{O}} a^2 \, d\omega| \leq \varepsilon$.

(iii) Assume that $\mathfrak{h} > 0$ is sufficiently small.

• Let σ be a tempered irreducible unitary representation of H for which $\mathfrak{h}\lambda_\sigma \in U$. If $\mathcal{O}_{\pi,\sigma}$ is non-empty, then

$$|k_{\mathfrak{h}}(\sigma)| \leq |\mathcal{H}_\sigma(\mathrm{Op}_{\mathfrak{h}}(a_1)^2)| \quad (22.9)$$

and

$$|k_{\mathfrak{h}}(\sigma) - \mathcal{H}_\sigma(\mathrm{Op}_{\mathfrak{h}}(a)^2)| \leq |\mathcal{H}_\sigma(\mathrm{Op}_{\mathfrak{h}}(a_2)^2)|. \quad (22.10)$$

• Let σ be a non-tempered irreducible unitary representation of H such that

$$\mathfrak{h}\lambda_\sigma \in U.$$

Then,

$$\mathcal{H}_\sigma(\mathrm{Op}_{\mathfrak{h}}(a)^2) = \mathcal{H}_\sigma(\mathrm{Op}_{\mathfrak{h}}(a_{\mathrm{nt}})^2) + O(\mathfrak{h}^N). \quad (22.11)$$

The implied constant may depend upon (N, k, a, ε) , but not upon $(\pi, \sigma, \mathfrak{h})$.

Remark 22.3. We note that, since $\mathrm{Op}_{\mathfrak{h}}(a_j)^2$ is positive-definite, the absolute values on the right-hand side of (22.9) and (22.10) should not be necessary; in any event, they will disappear from our analysis when we pass to the global setting, in which the product of local Hermitian forms as above is manifestly positive.

Proof. The informal idea for (22.9) and (22.10) is as in the arguments leading to (22.5) and (22.6): for instance, to get (22.10), we can just choose a such that

$$\int_{\mathcal{O}(\mu)} a^2 d\omega \approx k(\mu),$$

with the difference thus majorized by $\int_{\mathcal{O}(\mu)} a_1^2 d\omega$ for some small a_1 ; the asymptotics for $\mathcal{H}_\sigma(\dots)$ then give the required estimates.

The informal idea for (22.11) is that the infinitesimal characters of non-tempered representations are close to irregular elements, which form a set of measure zero. We may thus construct a_{nt} from a by shrinking its support to be concentrated near the inverse images of irregular elements.

Turning to details, choose a_1 as indicated, depending only upon k and V , so that $\int_{\mathcal{O}(\mu)} a_1^2 \geq k(\mu) + 1$ for $\mu \in \mathrm{supp}(k)$. Fix an open subset $U_0 \subset [\mathfrak{h}^\wedge]$, with $\mathrm{supp}(k) \subseteq U_0$ and $\bar{U}_0 \subseteq \mathrm{image}(V)$. Choose $\varepsilon_1 > 0$ small enough in terms of k and ε , then choose a_2 as indicated so that $\int_{\mathcal{O}(\mu)} a_2^2 \geq 2\varepsilon_1$ for $\mu \in U_0$ and $\int_{\mathcal{O}} a_2^2 d\omega \leq \varepsilon$. Choose a as indicated so that $|k(\mu) - \int_{\mathcal{O}(\mu)} a^2| \leq \varepsilon_1$ for $\mu \in U$.

The set

$$W := \{\xi \in \mathfrak{h}^\wedge \cap \mathrm{image}(\bar{V}) : [\xi] \notin [\mathfrak{h}^\wedge]_{\mathrm{reg}}\}$$

is compact and has measure zero. We may thus find a_{nt} as indicated, with $\int_{\mathcal{O}} a_{\text{nt}}^2 d\omega \leq \varepsilon$, so that $a = a_{\text{nt}}$ in a small neighborhood of W .

Assertion (i) is clear by construction. Assertion (ii) follows as in (22.7) if ε_1 is sufficiently small. Turning to assertion (iii), let $\sigma \in \widehat{H}$ with $h\lambda_\sigma \in U$.

Suppose first that σ is tempered and $\mathcal{O}_{\pi,\sigma} \neq \emptyset$. Then, $k_h(\sigma) = k(h\lambda_\sigma)$, while the asymptotic formulas in §19.3 give

$$\mathcal{H}_\sigma(\text{Op}_h(a)^2) = \int_{\mathcal{O}(h\lambda_\sigma)} a^2 + o_{h \rightarrow 0}(1),$$

and similarly for a_1 and a_2 . Thus (22.9) and (22.10) hold for h small enough in terms of ε_1 .

Suppose next that σ is non-tempered.

By the composition formula (8.36), we may write

$$\text{Op}_h(a)^2 = \text{Op}_h(a_{\text{nt}})^2 + \text{Op}_h(c) + \mathcal{E},$$

where the following statements hold:

- The h -dependent element $c \in C_c^\infty(V)$ is bounded with respect to h and vanishes identically on a small neighborhood of W . Since non-tempered representations have infinitesimal characters close to irregular elements (cf. §9.7), it follows that $c = 0$ on

$$\{\xi \in \mathfrak{h}^\wedge : \text{dist}([\xi], h\lambda_\sigma) \leq \varepsilon_2\}$$

for some small but fixed $\varepsilon_2 > 0$. By §10.3, we deduce that $\mathcal{H}_\sigma(\text{Op}_h(c)) \ll h^N$.

- $\mathcal{E} \in h^{N'} \Psi^{-N'}$, where $N' \rightarrow \infty$ as $J \rightarrow \infty$, so that, by the assumed uniform continuity of (22.8), we have $\mathcal{H}_\sigma(\mathcal{E}) \ll h^N$.

The required estimate (22.11) follows. □

22.4. Auxiliary estimates relevant for Weyl’s law

Recall that our main result concerns the average of an L -function over a family. We record here, for completeness, a technical estimate relevant for computing the *cardinality* of that family (cf. §31.4 for its application).

Let \mathcal{H} denote the Hecke algebra of smooth compactly-supported complex measures on H . Since we have fixed a Haar measure dh on H , we may identify \mathcal{H} with $C_c^\infty(H)$; in particular, we may define the evaluation $f(1)$ at the identity element $1 \in H$ of any $f \in \mathcal{H}$.

LEMMA 22.4. *Let k and U be as in §22.3, with $k \geq 0$. Fix $\varepsilon > 0$ and $N \in \mathbb{Z}_{\geq 0}$, and let $h > 0$ be sufficiently small. There are positive-definite elements $f, f_1 \in \mathcal{H}$, supported on $1 + o_{h \rightarrow 0}(1)$, such that*

$$\left| \int_{[\mathfrak{h}^\wedge]} k - h^d f(1) \right| \leq \varepsilon \tag{22.12}$$

and

$$f_1(1) \leq \varepsilon h^{-d}, \quad (22.13)$$

and, for each $\sigma \in \widehat{H}$,

$$|k_h(\sigma) - \chi_\sigma(f)| \leq \chi_\sigma(f_1) + O(h^N \langle h \lambda_\sigma \rangle^{-N}), \quad (22.14)$$

where $\chi_\sigma: \mathcal{H} \rightarrow \mathbb{C}$ denotes the character.

Before giving the proof, we record the basic idea. If we argue formally—ignoring convergence, truncations, etc.—then, for each $a \in S^{-\infty}(\mathfrak{g}^\wedge)$, the function $f: H \rightarrow \mathbb{C}$ defined by

$$NNf(s) := \text{tr}(\pi(s^{-1})\text{Op}_h(a)),$$

where $\text{Op}_h(a) := \text{Op}_h(a; \pi)$, satisfies

$$\chi_\sigma(f) = \sum_{\substack{v \in \mathcal{B}(\pi) \\ u \in \mathcal{B}(\sigma)}} \int_{s \in H} \langle s \text{Op}_h(a)v, v \rangle \langle u, su \rangle = \mathcal{H}_\sigma(\text{Op}_h(a)),$$

and

$$f(1) = \text{tr}(\text{Op}_h(a)).$$

To make f positive definite, we can argue instead with $\text{Op}_h(a)^2$. The lemma should thus be plausible in view of the analogous passage from (22.5)–(22.7) to Theorem 22.2. The subtlety is that f as defined above is not compactly-supported, so does not belong to \mathcal{H} as we have defined it. To make matters worse, the integral defining $\chi_\sigma(f)$ need not converge when σ is non-tempered. To get around these issues, we truncate f , taking care to do so in a manner that preserves positive definiteness. For σ tempered and stable relative to π , we have seen already that the main contribution to the integral over $s \in H$ defining $\mathcal{H}_\sigma(\text{Op}_h(a))$ comes from s fairly small, so the truncation has negligible impact in that case. We can use the operator calculus and the same argument as in Theorem 22.2 to control the contributions from the remaining σ .

Proof. For $b \in S^{-\infty}(\mathfrak{h}^\wedge)$, let us denote by $\widetilde{\text{Op}}_h(b) \in \mathcal{H}$ the element implicit, for a unitary representation σ of H , in the definition of $\text{Op}_h(b; \sigma) = \sigma(\widetilde{\text{Op}}_h(b))$.

Choose V as in §22.3, choose $\varepsilon_1 > 0$ small enough in terms of k and ε , assume that $h > 0$ is small enough in terms of ε_1 , and choose $a, a_2 \in C_c^\infty(V)$ as in the proof of Theorem 22.2. Define $f_0 \in \mathcal{H}$ by the formula

$$f_0 := \int_{s_1, s_2 \in H} \chi(s_1) \chi(s_2) \langle \pi(s_1) \text{Op}_h(a), \pi(s_2) \text{Op}_h(a) \rangle \delta_{s_2^{-1}s_1},$$

where

- $\chi \in C_c^\infty(H)$ denotes a suitable normalized cut-off: supported near the identity, non-negative-valued, invariant under inversion, constant near the identity, and satisfying $\int_H |\chi|^2 = 1$;
- $\text{Op}_h(a) := \text{Op}_h(a; \pi)$;
- $\delta_{s_2^{-1}s_1}$ is the δ -mass; its integral as above will define a smooth measure;
- we employ the Hilbert–Schmidt inner product on $\text{End}(\pi)$.

Then, f_0 is positive definite. Choose an open $W \subset \mathfrak{h}^\wedge$ such that $W \supseteq \text{image}(\bar{V})$ and $U \supseteq \text{image}(\bar{W})$. Choose $b \in C_c^\infty(\mathfrak{h}^\wedge)$ supported in the preimage of U and with $b=1$ on W . We note, by (8.36), that

$$\text{Op}_h(b; \pi) \text{Op}_h(a) \equiv \text{Op}_h(a) \pmod{\mathfrak{h}^\infty \Psi^{-\infty}(\pi)}.$$

Set

$$f := \widetilde{\text{Op}}_h(b) f_0 \widetilde{\text{Op}}_h(b).$$

(Here and below the only relevant product structure on \mathcal{H} is given by convolution.) In the same way that f was defined in terms of a , let f_+ be defined in terms of a_2 . Arguing as in the proof of Theorem 22.2, let b_{nt} be obtained from b by smoothly truncating to a sufficiently small neighborhood of those elements in the support of b whose image in $[\mathfrak{h}^\wedge]$ is irregular, and set $f_{\text{nt}} := \widetilde{\text{Op}}_h(b_{\text{nt}}) f_0 \widetilde{\text{Op}}_h(b_{\text{nt}})$. Finally, set

$$f_1 := 2\varepsilon_1 f_+ + f_{\text{nt}}.$$

Then, f and f_1 are positive definite. By slowly shrinking the support of χ , we may arrange that they are supported on $1 + o_{\mathfrak{h} \rightarrow 0}(1)$, or indeed on $1 + O(\mathfrak{h}^\eta)$ for any fixed $\eta \in (0, 1)$; this has no effect on the arguments to follow.

We now verify that these constructions lead to the required estimates. We start with (22.12). By unwinding the definitions, we see that

$$f_0(1) = \left(\int_H \chi^2 \right) \text{tr}(\text{Op}_h(a)^2) = \mathfrak{h}^{-d} \left(\int_{\mathcal{O}} a^2 d\omega + o_{\mathfrak{h} \rightarrow 0}(1) \right).$$

As in the proof of Theorem 22.2, it follows that

$$f_0(1) = \mathfrak{h}^{-d} \left(\int_{[\mathfrak{h}^\wedge]} k + o_{\varepsilon_1 \rightarrow 0}(1) \right)$$

for small enough \mathfrak{h} , which gives the modified form of (22.12) obtained by replacing f with f_0 . To obtain the required assertion concerning f , we calculate first that

$$f(1) = \int_{s \in H} \text{tr}(T_s \text{Op}_h(a)^2 T_s),$$

where

$$T_s := \int_{t \in H} \chi(st) \widetilde{\text{Op}}_{\mathfrak{h}}(b)(t) \pi(t). \quad (22.15)$$

Note that $T_s = 0$, unless s is small, and the integrand in (22.15) vanishes unless t is small. By trivially estimating the L^1 -norm of $\widetilde{\text{Op}}_{\mathfrak{h}}(b)$, we see that the operator norm of T_s is $O(1)$ and the trace norm of $\text{Op}_{\mathfrak{h}}(a)^2$ is $O(\mathfrak{h}^{-d})$. Thus, to compute $f(1)$ to accuracy $o(\mathfrak{h}^{-d})$, it suffices to do so after replacing T_s by any modification T'_s differing in operator norm by $o(1)$. To that end, let us pull the integral (22.15) back to the Lie algebra, writing $t = e^y$ with $y \in \mathfrak{h}$. The integrand is concentrated on $|y| = O(\mathfrak{h})$, so we may truncate it to $|y| \leq \mathfrak{h}^{1-\eta}$ for some fixed $\eta > 0$ and then Taylor expand

$$\chi(se^y) = \chi(s) + O(\mathfrak{h}^{1-\eta}).$$

The modification

$$T'_s := \chi(s) \int_{t \in H} \widetilde{\text{Op}}_{\mathfrak{h}}(b)(t) \pi(t) = \chi(s) \text{Op}_{\mathfrak{h}}(b : \pi)$$

is thus acceptable for our purposes, and we obtain

$$f(1) = \left(\int_H \chi^2 \right) \text{tr}(\text{Op}_{\mathfrak{h}}(b) \text{Op}_{\mathfrak{h}}(a)^2 \text{Op}_{\mathfrak{h}}(b)) + o_{\mathfrak{h} \rightarrow 0}(\mathfrak{h}^{-d}), \quad (22.16)$$

say. We appeal now to the composition formula (8.36) to replace $\text{Op}_{\mathfrak{h}}(b) \text{Op}_{\mathfrak{h}}(a)^2 \text{Op}_{\mathfrak{h}}(b)$ with $\text{Op}_{\mathfrak{h}}(a)^2$, and argue as before. This completes the verification of (22.12).

The same arguments applied to f_1 lead to the estimate (22.13).

We turn finally to (22.14). Thus, σ be a unitary representation of H . We consider first the case that $\mathfrak{h}\lambda_{\sigma} \notin U$. Then, $k_{\mathfrak{h}}(\sigma) = 0$. On the other hand, the results in §10.3—applied with (G, π) playing the duplicate role of “ (H, σ) ”, and using the continuity of (12.6) as an a-priori estimate to clean up remainders—give that the trace norm of $\text{Op}_{\mathfrak{h}}(b : \sigma)$ is $O(\mathfrak{h}^N \langle \mathfrak{h}\lambda_{\sigma} \rangle^{-N})$ for any fixed N . As the operator norms of $\sigma(f_0)$, $\text{Op}_{\mathfrak{h}}(b : \sigma)$ and $\text{Op}_{\mathfrak{h}}(b_{\text{nt}} : \sigma)$ are readily bounded by $\mathfrak{h}^{-O(1)}$, the claim (22.14) follows in this case.

It remains to consider the case that $\mathfrak{h}\lambda_{\sigma} \in U$. If σ is tempered, then we may verify—using arguments similar to those leading to (22.16)—that

$$\chi_{\sigma}(f) = \mathcal{H}_{\sigma}(\text{Op}_{\mathfrak{h}}(a)^2) + O(\mathfrak{h}),$$

and similarly for $\chi_{\sigma}(f_1)$. If σ is non-tempered, then, as before, $\text{Op}_{\mathfrak{h}}(b - b_{\text{nt}} : \sigma)$ has trace norm $O(\mathfrak{h}^{\infty})$. In either case, we may conclude as in the proof of Theorem 22.2. \square

23. Preliminaries on representations of p -adic groups

In this section, we recall how tempered representations of a p -adic reductive group fit into families indexed by quotients of certain tori, and explain the relationship of this picture to the Bernstein center.

The considerations of this section apply to any reductive group \mathbf{H} over a non-archimedean local field F of characteristic zero; the group \mathbf{G} plays no role.

23.1. Standard parabolic and Levi subgroups

We fix a minimal parabolic subgroup P_0 of H and Levi subgroup $M_0 < P_0$. A *standard* parabolic subgroup is one that contains P_0 . A *standard* Levi subgroup M is one that contains M_0 and arises as a Levi component of a standard parabolic $P = MN$; thus $P = MP_0$. Each parabolic or Levi subgroup is conjugate to a standard one, so there is little loss of generality in restricting to the latter.

For each standard Levi subgroup M , we have an induction functor $i_M^G = \text{Ind}_P^G$ from smooth representations of M to smooth representations of G , normalized to take unitary representations to unitary representations; here and henceforth “unitary” and “unitarizable” are used interchangeably.

What matters most for our purposes is the set of irreducible subquotients of $i_M^G \tau$. This set is independent of the H -conjugacy class of (M, τ) , and makes sense for any Levi subgroup M , not necessarily standard. If moreover τ is unitary, then so is $i_M^G \tau$, hence subquotients of $i_M^G \tau$ are the same as submodules.

23.2. Good compact open subgroups

Let J be a compact open subgroup of H . Recall that J admits an *Iwahori factorization* with respect to a parabolic subgroup P if there is a Levi decomposition $P = MN$, with associated opposite parabolic subgroup $P^- = MN^-$, so that

$$J = J_{N^-} J_M J_N$$

with

$$J_{N^-} := J \cap N^-, \quad J_M := J \cap M \quad \text{and} \quad J_N := J \cap N.$$

We may and shall fix a maximal compact subgroup K of G for which $KP_0 = H$ (see, e.g., [HR, Corollary 9.12]).

Let us call a compact open subgroup J of H *good* (relative to the choice of K and P_0) if

- (i) J is a normal subgroup of K , and
- (ii) J admits an Iwahori factorization with respect to each standard parabolic subgroup.

LEMMA 23.1. *For suitable choice of K as above, there are good compact open subgroups $(J_n)_{n \geq 0}$ of H that form a neighborhood basis of the identity.*

Proof. We apply [MP, Proposition 4.2], with x a special vertex. \square

23.3. Invariant vectors and induction

The following result is standard.

LEMMA 23.2. *Let J be a good compact open subgroup of H , let $P=MN$ be a parabolic subgroup of H , and τ a supercuspidal representation of M . The following are equivalent:*

- (a) τ admits a non-zero J_M -fixed vector;
- (b1) some subquotient of $i_P^H \tau$ admits a non-zero J -fixed vector;
- (b2) $i_P^H \tau$ admits a non-zero J -fixed vector;
- (c) every subquotient of $i_P^H \tau$ admits a non-zero J -fixed vector.

Proof. The dimension of the space of J -fixed vectors may be expressed as the trace of an averaging operator, and so may be computed in terms of a composition series. Conditions (b1) and (b2) are thus equivalent; henceforth we refer to them together as (b). Obviously (c) implies (b1).

We now show that (a) and (b) are equivalent. We may assume that P and M are standard. The J -fixed vectors in $i_M^H \tau$ are described by pairs (x, v) , where $x \in P \backslash H/J$ and $v \in \tau$ satisfy

$$\delta_P^{1/2} \tau(g)v = v \quad \text{for all } g \in P \cap xJx^{-1}, \quad (23.1)$$

where as usual τ acts via the projection $P \twoheadrightarrow M$. Since $H=PK$, we may assume that $x \in K$; then $xJx^{-1}=J$, so (23.1) just says that v is fixed by the projection to M of $P \cap J = J_M J_N$, i.e., by J_M . Thus (a) and (b) are equivalent.

It remains to see that (a) implies (c). Suppose thus that $\tau^{J_M} \neq 0$. Let π be any subquotient of $i_P^H \tau$. In fact, we may assume that π is a submodule, by a standard argument: there exists a parabolic $Q=NU$ such that the Jacquet module π_U is supercuspidal, up to twist; by the computations of Bernstein and Zelevinsky [BZ, §2.13] there is an element of g carrying (M, τ) to a constituent of (N, π_U) . We may therefore suppose that $N=M$ and

that π_U contains a conjugate $w\tau$ of τ by the normalizer of M ; thus $\pi \hookrightarrow i_{MU}^G(w\tau)$. Since w has a representative belonging to K , it is equivalent whether $w\tau$ or τ has a J_M -fixed vector.

By Frobenius reciprocity, the inclusion $\pi \rightarrow i_P^G \tau$ gives rise to a non-zero map $\pi_N \rightarrow \tau$. Since τ is irreducible, this map is surjective. Taking J_M -invariants gives a surjective map $(\pi_N)^{J_M} \rightarrow \tau^{J_M}$. But a basic theorem [Cas2, Theorem 3.3.3] asserts that the map $\pi^J \rightarrow (\pi_N)^{J_M}$ is surjective. Thus, our assumption $\tau^{J_M} \neq 0$ implies $\pi^J \neq 0$, as required. \square

23.4. Classifications

23.4.1. Terminology

Let π be an irreducible representation of H . Recall that π is *tempered* if it is unitary and its matrix coefficients lie in $L^{2+\varepsilon}$ modulo the center. Recall that π is *square-integrable* if its central character is unitary and its matrix-coefficients are square-integrable modulo the center; in particular, π is unitary and tempered. Recall that π is *supercuspidal* if its matrix coefficients are compactly-supported modulo the center; then π is unitary if and only if its central character is unitary, in which case it is square-integrable. In particular, any supercuspidal representation has an unramified twist which is unitary.

23.4.2. Bernstein–Zelevinsky; infinitesimal characters

By the results in [BZ, §2], we have the following.

LEMMA 23.3. *For each irreducible representation π of H there is a unique H -conjugacy class $[(M, \tau)]$ of pairs (M, τ) , where M is a Levi subgroup and σ is a supercuspidal representation of M , such that π is a subquotient of $i_M^G \tau$.*

The *infinitesimal character* of an irreducible representation π of H is the class

$$\lambda_\pi := [(M, \tau)]$$

arising in the lemma. By an *infinitesimal character* for H we will mean any such class

$$[(M, \tau)].$$

23.4.3. Langlands, from square-integrable to tempered

By [Wld, Proposition III.4.1]. we have the following result.

LEMMA 23.4. *For each tempered irreducible representation π of H there is a unique H -conjugacy class $[(M, \sigma)]$ of pairs (M, σ) , where M is a Levi subgroup and σ is a square-integrable representation of M , such that π is a subquotient (equivalently, submodule) of $i_M^G \sigma$.*

23.5. Bernstein components

In this and the following subsections we recall some facts from [Ber] (cf. [BDK, §2] for a summary).

For a Levi subgroup M of H , let \mathfrak{X}_M denote the group of unramified characters of M , i.e., homomorphisms $\chi: M \rightarrow \mathbb{C}^\times$ that are trivial on the subgroup M^0 on which all algebraic characters have valuation zero. The group \mathfrak{X}_M is a complex torus, i.e., isomorphic to $(\mathbb{C}^\times)^r$, while the subgroup \mathfrak{X}_M^0 of unitary characters identifies with the compact subtorus $(\mathbb{C}^{(1)})^r$.

For each supercuspidal representation τ of M , the set

$$\Theta = \{(M, \tau \otimes \chi) : \chi \in \mathfrak{X}_M\}$$

is a *Bernstein component*, or simply a *component* for short, of the set of infinitesimal characters. Each component may be identified with the quotient of \mathfrak{X}_M by a finite group (see for instance the discussion in [Hai, §3.3.1]).

23.6. Bernstein center

Varying (M, τ) , the set of infinitesimal characters identifies with a disjoint union of finite quotients of complex tori, giving it the structure of a complex algebraic variety, typically with infinitely many components. The *Bernstein center* is the algebra $\mathfrak{Z}(H)$ of regular functions on the variety of infinitesimal characters for H ; it is the direct product over the set of components Θ of the algebra of regular functions on Θ , and we have

$$\text{spectrum of } \mathfrak{Z}(H) = \bigsqcup \Theta,$$

the union taken over all Bernstein components. By [Ber] (cf. [BDK, §2.2]), we have the following result.

LEMMA 23.5. *There is a natural action of $\mathfrak{Z}(H)$ on the category of representations of H : for each $z \in \mathfrak{Z}(H)$ and each representation π of H , there is an associated H -equivariant endomorphism $z: \pi \rightarrow \pi$, such that for each H -equivariant morphism of representations $j: \pi \rightarrow \pi'$, we have $z \circ j = j \circ z$. If π is irreducible, then $z: \pi \rightarrow \pi$ is scalar multiplication by $z(\lambda_\pi)$.*

To apply this in practice, we let \mathcal{H} denote the Hecke algebra of locally constant compactly-supported measures on H , under convolution, regarded as a representation of H under the action defined by left translation. For a vector v in a representation π of H , the action map $\mathcal{H} \rightarrow \pi$ given by $f \mapsto f*v$ is then H -equivariant, so, for each $z \in \mathfrak{Z}(H)$, we have $(z \cdot f)*v = z \cdot (f*v)$. Let J be a compact open subgroup of H that fixes v , and take $f := e_J$, the corresponding averaging operator. Then, $z \cdot v = h_z * v$, where $h_z := z \cdot e_J$ is a central element of the bi- J -invariant subalgebra $\mathcal{H}_J \subseteq \mathcal{H}$. In particular, if π is irreducible, then h_z acts on π^J by the scalar $z(\lambda_\pi)$.

23.7. Components arising from Langlands classification of the tempered dual

23.7.1. By §23.4.3, there is a natural map

$$\mathfrak{l}: \widehat{H}_{\text{temp}} \longrightarrow \left\{ [(M, \tau)] : \begin{array}{l} \tau \text{ is a square-integrable representation} \\ \text{of the Levi subgroup } M \text{ of } H \end{array} \right\} \quad (23.2)$$

assigning to σ the H -conjugacy class $[(M, \tau)]$ of pairs as indicated for which $\sigma \hookrightarrow i_M^H \tau$.

As in §23.5, we may fix M and vary τ in a family of unramified unitary twists $\{\tau \otimes \chi\}_{\chi \in \mathfrak{X}_M^0}$ to write the right-hand side of (23.2) as a disjoint union of subsets \mathcal{D} parameterized by the compact tori \mathfrak{X}_M^0 and identified with quotients \mathfrak{X}_M^0/Γ for some finite subgroups Γ of $\mathfrak{X}_M^0 \rtimes N(M)/M$. To differentiate from the Bernstein components (§23.5), we refer to these subsets \mathcal{D} as *l-components*; we are unaware of any standard terminology. We rewrite (23.2) as

$$\mathfrak{l}: \widehat{H}_{\text{temp}} \longrightarrow \bigsqcup \mathcal{D}.$$

23.7.2. We note in passing, for the sake of orientation, that, by generic irreducibility [Cas2, Theorem 6.6.1], each \mathfrak{l} -component \mathcal{D} contains a non-empty Zariski open subset U such that the map $\mathfrak{l}^{-1}(U) \rightarrow U$ is injective.

23.7.3. If $\mathfrak{l}(\sigma) = [(M, \tau)]$, then the infinitesimal character λ_σ is the image of λ_τ in the space of infinitesimal characters for H under the evident map from infinitesimal characters for M .

In particular,

$$\mathfrak{l}(\sigma_1) = \mathfrak{l}(\sigma_2) \implies \lambda_{\sigma_1} = \lambda_{\sigma_2}.$$

For each \mathfrak{l} -component \mathcal{D} there is thus a (unique) Bernstein component Θ for which $\sigma \mapsto \lambda_\sigma$ descends to a map

$$\mathcal{D} \longrightarrow \Theta.$$

This map is continuous; indeed, as noted in §23.5, both spaces are locally (for the analytic topology) identified with character tori for Levi subgroups, and locally the map is given by a homomorphism of these character tori arising from an inclusion of Levi subgroups.

Two distinct \mathfrak{l} -components \mathcal{D}_1 and \mathcal{D}_2 may map to the same Bernstein component Θ , and their images may overlap. For instance, there do exist (for general H) non-isomorphic square-integrable representations of H having the same infinitesimal character.

23.8. Finiteness

The Bernstein components or \mathfrak{l} -components form countable sets. More precisely, it follows from §23.3 that for any good compact open subgroup J of H , a representation σ of H with $\mathfrak{l}(\sigma)=[(M, \tau)]$ has a non-zero J -fixed vector if and only if the representation τ of the Levi M , taking M standard without loss of generality, has a non-zero J_M -fixed vector. By [BDK, §2.3] and the Plancherel formula (or see [Wld, Theorem VIII.1.2]), only finitely many Bernstein components or \mathfrak{l} -components contain some $[(M, \tau)]$ with this property.

24. The case of auxiliary p -adic places

We consider now the non-archimedean case of the setup of §20, thus (\mathbf{G}, \mathbf{H}) is a GGP pair over a non-archimedean local field F of characteristic zero, and π is a smooth tempered irreducible unitary representation of G . We introduce the abbreviation

$$\Omega := \widehat{H}_{\text{temp}}^{\pi}$$

for the π -distinguished tempered dual of H . The notation and terminology of §23 will be employed freely.

24.1. The structure of Ω

Recall from §23.7 the map $\mathfrak{l}: \widehat{H}_{\text{temp}} \rightarrow \bigsqcup \mathcal{D}$ arising from the Langlands classification.

LEMMA 24.1. (Strong multiplicity 1) *The induced map*

$$\mathfrak{l}: \Omega \longrightarrow \bigsqcup \mathcal{D}$$

is injective.

Proof. Indeed, strong multiplicity 1 [MW], [BP1] implies that each L -packet of tempered representations of H contains at most one π -distinguished element. Since each fiber of \mathfrak{l} is contained in a single L -packet, the conclusion follows.

We note that the results in [MW] and [BP1] are formulated as conditional on certain expected properties of L -packets for classical groups. It is not straightforward for us to extract these properties from the literature, so we observe also that the required conclusion—multiplicity 1 for the full induction of a square-integrable representation—can be verified directly (see [BP2, §14.2 and §14.3]). \square

For $t = [(M, \tau)] \in \mathfrak{l}(\Omega)$, write $\sigma_t := \mathfrak{l}^{-1}(t) \in \Omega$. Then, $i_M^H \tau$ decomposes as a finite direct sum of tempered irreducible representations σ of H ; one of these summands is the given π -distinguished representation σ_t , while strong multiplicity one implies that the remaining summands are not π -distinguished. It follows that for any $v_1, v_2 \in \pi$,

$$\mathcal{H}_{\sigma_t}(v_1 \otimes \bar{v}_2) = \sum_{u \in \mathcal{B}(i_M^H \tau)} \int_{h \in H} \langle hv_1, v_2 \rangle \langle u, hu \rangle. \quad (24.1)$$

Here the sum over u is really a finite sum, since v_1 and v_2 are smooth (under H). The right-hand side of (24.1) is manifestly continuous as t varies within a given \mathfrak{l} -component \mathcal{D} , and so defines a continuously-varying family of Hermitian forms on $\pi \otimes \bar{\pi}$.

LEMMA 24.2. *The image of Ω under \mathfrak{l} is a union of \mathfrak{l} -components.*

Proof. We need only verify that $\mathfrak{l}(\Omega)$ is both closed and open: it is closed, by the upper semicontinuity of multiplicity [FLO, Lemma D.1], and open, by the continuity of $t \mapsto \mathcal{H}_{\sigma_t}$ noted above. \square

Thus \mathfrak{l} identifies Ω with a disjoint union of (typically infinitely many) finite quotients of compact tori, hence equips Ω with a natural topology with respect to which the Hermitian forms \mathcal{H}_{σ} vary continuously. We may also speak of the space $C_c(\Omega)$ of compactly-supported continuous functions. We henceforth refer to Ω and $\mathfrak{l}(\Omega)$ interchangeably.

It seems conceivable to us that the topology just defined on Ω coincides with the topology induced by the Fell topology, but we have not verified this. In any event, the more explicit topology just defined is the relevant one for our purposes.

The map

$$\Omega \longrightarrow \bigsqcup \Theta = \text{spectrum of the Bernstein center} \quad (24.2)$$

is continuous, for the topology just defined on Ω . This continuity follows from the fact the map from \mathfrak{l} -components to Bernstein components, discussed in §23.7.3, is continuous.

24.2. Main results

Let J be a good (see §23.2) compact open subgroup of H . We denote by \widehat{H}^J , $\widehat{H}_{\text{temp}}^J$ and Ω^J the spaces corresponding to representations of H having a non-zero J -fixed vector.

The space Ω^J is a finite union of \mathfrak{l} -components, and as J traverses a neighborhood basis of the identity, we have

$$\Omega = \bigcup_J \Omega^J \quad \text{and} \quad C_c(\Omega) = \bigcup_J C(\Omega^J).$$

The following notion is motivated by our global applications, and related to the class of functions that appear in Sauvageot’s “principe de densité” [Sau].⁽⁷⁾ We henceforth adopt the convention that $k(\sigma) := 0$ for $k \in C_c(\Omega)$ and $\sigma \in \widehat{H} \setminus \Omega$.

Definition 24.3. We say that an \mathfrak{l} -component \mathcal{D} of Ω is *allowable* if there is a good compact open subgroup J as above, with $\mathcal{D} \subseteq \Omega^J$ such that every non-negative

$$k \in C(\mathcal{D}) \subseteq C_c(\Omega)$$

may be approximated in the following senses:

(i) For each $\varepsilon > 0$ there are positive-definite $T, T_+ \in \pi^J \otimes \overline{\pi^J}$, with $\text{tr}(T_+) \leq \varepsilon$, such that, for each $\sigma \in \widehat{H}$,

$$|k(\sigma) - \mathcal{H}_\sigma(T)| \leq \mathcal{H}_\sigma(T_+) \tag{24.3}$$

(We emphasize that, per the general conventions of Part IV, all tensors considered here such as T and T_+ are *smooth*.)

(ii) For each $\varepsilon > 0$ there are $\phi, \phi_+ \in \mathcal{H}_J$, with ϕ_+ positive definite and $\phi_+(1) \leq \varepsilon$, so that, for each $\sigma \in \widehat{H}$,

$$|k(\sigma) - \text{tr}(\sigma(\phi))| \leq \text{tr}(\sigma(\phi_+)). \tag{24.4}$$

We say that a function $k: \Omega \rightarrow \mathbb{C}$ is *allowable* if its support lies in a finite union of allowable \mathfrak{l} -components.

Remark 24.4. Allowability (applied with $\varepsilon = 1$, say) implies that there is a positive-definite $T_+ \in \pi^J \otimes \overline{\pi^J}$ such that

$$|k(\sigma)| \leq \mathcal{H}_\sigma(T_+) \tag{24.5}$$

for all $\sigma \in \widehat{H}$. Similarly there is a positive-definite element ϕ_+ of the bi- J -invariant Hecke algebra $\mathcal{H}_J \cong C_c^\infty(J \backslash H / J)$ such that

$$|k(\sigma)| \leq \text{tr}(\sigma(\phi_+)) \tag{24.6}$$

for all $\sigma \in \widehat{H}$. The upper bounds (24.5) and (24.6) will be useful in applications involving products of several groups such as G .

⁽⁷⁾ There were some points in the original paper [Sau] that we do not understand: specifically the usage of Lemma 2.1 on p. 181. Lemma 2.1 assumes that the algebra in question separates points. This is related to the distinction between \mathfrak{l} -components and Bernstein components.

Also note that the left-hand side of either (24.3) or (24.4) vanishes identically unless σ belongs to \widehat{H}^J , whose Plancherel measure is finite. By the Plancherel formula (cf. §18, §A.3), it follows that, for T and ϕ as in the conclusion, we have

$$\mathrm{tr}(T) + o_{\varepsilon \rightarrow 0}(1) = \int_{\widehat{H}_{\mathrm{temp}}} k = \phi(1) + o_{\varepsilon \rightarrow 0}(1). \quad (24.7)$$

24.3. Cuspidal type components are allowable

There is a class of Bernstein components which are particularly straightforward to analyze.

We say that an \mathfrak{l} -component is *of cuspidal type* if the inducing data is not merely discrete series, but supercuspidal. For example, the \mathfrak{l} -component of unramified principal series is of cuspidal type.

THEOREM 24.5. *Any \mathfrak{l} -component of cuspidal type is allowable.*

It seems reasonable to expect that all \mathfrak{l} -components are allowable.

24.4. Outline of the proof

In this subsection, we give an overview of the argument. Details of the steps are given in the following subsections. The subtlety is in controlling the contribution of the non-tempered spectrum, which meets the tempered spectrum in different ways (for example, the complementary series can approach both the Steinberg representation and a tempered principal series).

In the argument that follows, integrals and volumes are always computed with respect to Plancherel measure.

We must verify conditions (i) and (ii) of Definition 24.3. We first prove (i) and then deduce (ii).

For the proof of (i), it suffices to obtain the required approximation (24.3) for each continuous function

$$k: \mathcal{D} \longrightarrow [0, 1],$$

where \mathcal{D} is an \mathfrak{l} -component of cuspidal type. Recall that we extend k by zero to \widehat{H} .

Step 1. Let Θ denote the union of Bernstein components with J -fixed vectors, where J is chosen small enough so that \mathcal{D} maps into Θ . The complex variety Θ has a natural real form (see §24.5.2 for details); we write $\mathbb{R}[\Theta]$ for the set of regular functions on that

real form, and regard it as a subring of the Bernstein center (§23.6). Given $z \in \mathbb{R}[\Theta]$ and

$$T = \sum_i u_i \otimes \bar{v}_i \in \pi^J \otimes \bar{\pi}^J,$$

we set

$$zTz := \sum_i zu_i \otimes z\bar{v}_i \in \pi^J \otimes \bar{\pi}^J.$$

Then, for each $\sigma \in \widehat{H}$, we have $z(\lambda_\sigma) \in \mathbb{R}$ (§23.6) and

$$\mathcal{H}_\sigma(zTz) = z(\lambda_\sigma)^2 \mathcal{H}_\sigma(T). \quad (24.8)$$

In particular, $\mathcal{H}_\sigma(zTz) = 0$ unless $\lambda_\sigma \in \Theta$.

Step 2. There are finitely many π -distinguished \mathfrak{l} -components besides \mathcal{D} that map into Θ ; let \mathcal{D}' be their union. Since the \mathfrak{l} -component \mathcal{D} is of cuspidal type, we have that

$$\mathcal{D} \longrightarrow \Theta \quad \text{is injective and its image is disjoint from the image of } \mathcal{D}'. \quad (24.9)$$

Step 3. We may find a positive-definite T_0 such that $\mathcal{H}_\sigma(T_0) \geq 1$ for $\sigma \in \mathcal{D} \cup \mathcal{D}'$ (see Lemma 24.6).

Step 4. Fix $\varepsilon < 1$.

Then, by an application of Stone–Weierstrass and the assumed disjointness (see §24.5.3 for details), we may find $z_0 \in \mathbb{R}[\Theta]$ which approximates $\sqrt{k(\sigma)/\mathcal{H}_\sigma(T_0)}$ on \mathcal{D} and approximates zero on \mathcal{D}' :

$$|k(\sigma) - z_0(\lambda_\sigma)^2 \mathcal{H}_\sigma(T_0)| < \varepsilon, \quad \sigma \in \mathcal{D} \cup \mathcal{D}'. \quad (24.10)$$

Set $T := z_0 T_0 z_0$. Then, (24.8) and (24.10) imply that

$$|k(\sigma) - \mathcal{H}_\sigma(T)| < \varepsilon, \quad \sigma \in \widehat{H}_{\text{temp}}^J, \quad (24.11)$$

since the left-hand side vanishes for $\sigma \in \widehat{H}_{\text{temp}}^J \setminus (\mathcal{D} \cup \mathcal{D}')$. Using our assumptions on k and ε , we get

$$\int_{\mathcal{D}'} |\mathcal{H}_\sigma(T)| \ll \varepsilon \quad \text{and} \quad \sup_{\mathcal{D}} |\mathcal{H}_\sigma(T)| \ll 1. \quad (24.12)$$

Here we adopt the convention that implied constants may depend upon (\mathcal{D}, J, k) (hence possibly upon $(\mathcal{D}', \Theta, T_0)$), but must be independent of ε . We also have abused notation and written integrals over \mathcal{D} and \mathcal{D}' , where the domain of the integral is, more precisely, $\mathfrak{l}^{-1}\mathcal{D}$ and $\mathfrak{l}^{-1}\mathcal{D}'$.

Step 5. Now take $T_+ = zTz + \varepsilon T_0$, where $z \in \mathbb{R}[\Theta]$ has the following properties, again achieved by Stone-Weierstrass:

- $z(\lambda_\sigma) \geq 1$ for each non-tempered $\sigma \in \widehat{H}$;
- $z(\lambda_\sigma) \in [-2, 2]$ for each $\sigma \in \widehat{H}$;
- $z(\lambda_\sigma)$ is small on average over $\sigma \in \mathcal{D}$: the integral of $z(\lambda_\sigma)^2$, taken with respect to Plancherel measure, is bounded by ε .

See §24.5.3 for the construction. Note that the first and third requirements “pull in opposite directions” because the intersection

$$\overline{\text{infinitesimal characters of non-tempered representations}} \cap \lambda(\mathcal{D}) \quad (24.13)$$

need not be empty; we can nevertheless simultaneously satisfy them because the measure of the set (24.13) is zero. Note also that we are again using our assumption that \mathcal{D} is of cuspidal type—the same intersection, but with \mathcal{D} replaced by \mathcal{D}' , does not necessarily have measure zero. For example, if $G = \mathrm{PGL}_2(F)$ and \mathcal{D} is the \mathfrak{l} -component of cuspidal type consisting of the unitary principal series representations, then the other \mathfrak{l} -component in \mathcal{D}' is the singleton consisting of the Steinberg representation, whose infinitesimal character is a limit of infinitesimal characters of complementary series representations.

We claim now that

$$|k(\sigma) - \mathcal{H}_\sigma(T)| \leq \mathcal{H}_\sigma(T_+) \quad \text{for all } \sigma \in \widehat{H} \quad (24.14)$$

and

$$\mathrm{tr}(T_+) \ll \varepsilon. \quad (24.15)$$

Assuming the claim, we may replace ε by a smaller constant as needed to obtain the desired pair (T, T_+) . We verify (24.14) separately in the following cases:

- for $\sigma \in \mathcal{D} \cup \mathcal{D}'$, we have $|k(\sigma) - \mathcal{H}_\sigma(T)| < \varepsilon \mathcal{H}_\sigma(T_0) \leq \mathcal{H}_\sigma(T_+)$;
- for $\sigma \in \widehat{H}_{\mathrm{temp}}$ but $\sigma \notin \mathcal{D} \cup \mathcal{D}'$, both sides are zero;
- for non-tempered σ , we have $\mathcal{H}_\sigma(T) \geq 0$, $z(\lambda_\sigma) \geq 1$ and $k(\sigma) = 0$, and hence

$$|k(\sigma) - \mathcal{H}_\sigma(T)| = \mathcal{H}_\sigma(T) \leq \mathcal{H}_\sigma(zTz) \leq \mathcal{H}_\sigma(T_+).$$

We verify (24.15) using that

$$\mathrm{tr}(T_+) = \varepsilon \mathrm{tr}(T_0) + \int_{\sigma} z(\lambda_\sigma)^2 \mathcal{H}_\sigma(T),$$

and the following estimates, deduced using (24.12) and the construction of z :

$$\int_{\sigma \in \mathcal{D}} z(\lambda_\sigma)^2 \mathcal{H}_\sigma(T) \leq \left(\max_{\sigma \in \mathcal{D}} \mathcal{H}_\sigma(T) \right) \int_{\sigma \in \mathcal{D}} z(\lambda_\sigma)^2 \ll \varepsilon$$

and

$$\int_{\sigma \in \mathcal{D}'} z(\lambda_\sigma)^2 \mathcal{H}_\sigma(T) \ll \varepsilon.$$

Step 6. We now prove (ii). One could argue in parallel with the prior argument; however, for (ii), one encounters issues of reducibility that do not occur in (i)—in the context of (i) such issues are effectively eliminated by strong multiplicity 1. We have therefore found it more convenient, although perhaps slightly unnatural, to deduce (ii) from (i).

Given k , take $T \in \pi^J \otimes \pi^{\overline{J}}$ as in (i). Define the bi- J -invariant function $\phi: H \rightarrow \mathbb{C}$ by $\phi_0(h) := \text{tr}(\pi(h)T)$. Write $e_J \in \mathcal{H}_J$ for the normalized characteristic function of J . The dimension of σ^J is uniformly bounded for $\sigma \in \widehat{H}$; let M be an upper bound for this dimension.

In a formal sense, we have $\mathcal{H}_\sigma(T) = \text{tr}(\sigma(\phi_0))$ (see §18). The basic idea of the argument is to use this formal identity to construct functions from the T and T_+ previously constructed. This is not entirely straightforward, because ϕ_0 is not compactly supported.

However, at least if σ is tempered, the integral defining the operator $\sigma(\phi_0)$ converges (see §A.5), and defines a non-negative operator because $T \geq 0$. (To verify the non-negativity, it suffices to treat the case when T is of rank 1, and then it comes from the positivity of the matrix coefficient integral, see [Sve]).

Take

$$\phi_1 = (\text{bi-}J\text{-invariant) truncation of } \phi_0 + \text{small multiple of } e_J,$$

with a large enough truncation; by truncating symmetrically, we arrange that

$$\phi_1(x^{-1}) = \overline{\phi_1(x)}$$

for all $x \in H$. Using the absolute convergence of the matrix coefficient integral defining \mathcal{H}_σ , we obtain the following:

- $|\text{tr}(\sigma(\phi_1)) - \mathcal{H}_\sigma(T)| \leq \varepsilon$ for all $\sigma \in \widehat{H}_{\text{temp}}^J$;
- $\sigma(\phi_1)$ is positive definite for each such σ ;
- $\sigma(\phi_1)$ is zero if σ does not have a J -fixed vector, i.e., $\lambda_\sigma \notin \Theta$.

For $\sigma \in \widehat{H}_{\text{temp}}^J$, we have

$$|\text{tr}(\sigma(\phi_1)) - k(\sigma)| \leq |\text{tr}(\sigma(\phi_1)) - \mathcal{H}_\sigma(T)| + |\mathcal{H}_\sigma(T) - k(\sigma)| \leq 2\varepsilon, \quad (24.16)$$

using (24.11) at the second step.

We now construct a function ϕ_3 which controls ϕ_1 on the non-tempered set, in a sense to be made precise. Fix $\varepsilon' > 0$, and choose z as in Step 5, but now with

$$\int_{\sigma \in \mathcal{D}} z(\lambda_\sigma)^2 < \varepsilon'.$$

Put $\phi_3 = \varepsilon^{-1}(z \cdot \phi_1) * (z \cdot \phi_1)^\vee$, where as usual $f^\vee(g) = \overline{f(g^{-1})}$. Then, ϕ_3 is positive definite and symmetric.

We claim that, for any non-tempered $\sigma \in \widehat{H}$,

$$|\mathrm{tr}(\sigma(\phi_1))| \leq \mathrm{tr}(\sigma(\phi_3)) + M\varepsilon. \tag{24.17}$$

Indeed, choose an orthonormal basis for σ^J consisting of eigenvectors for $\sigma(\phi_1)$. The basis has cardinality at most M . Since $z \cdot \phi_1$ acts self-adjointly on σ^J , its action coincides with that of $(z \cdot \phi_1)^\vee$. Thus, if v belongs to the chosen basis and has eigenvalue c under $\sigma(\phi_1)$, then $\sigma(\phi_3)v = c'v$, where $c' := z(\lambda_\sigma)^2 \varepsilon^{-1} |c|^2 \geq \varepsilon^{-1} |c|^2$. Thus $c' \geq |c|$ whenever $|c| \geq \varepsilon$. We obtain (24.17) by summing over v , considering separately the cases $|c| \geq \varepsilon$ and $|c| \leq \varepsilon$.

Moreover,

$$\phi_3(1) = \int_{\sigma \in \widehat{H}_{\mathrm{temp}}^J} \varepsilon^{-1} z(\lambda_\sigma)^2 \mathrm{tr}(\sigma(\phi_1 * \phi_1)). \tag{24.18}$$

We bound the integrand on the right-hand side as follows:

- Suppose that $\sigma \in \Omega \cap \Gamma^{-1}(\mathcal{D})$, i.e., σ has \mathfrak{l} -parameter in \mathcal{D} and is distinguished. The trace of $\sigma(\phi_1 * \phi_1)$ is bounded by $M \|\phi_1\|_{L^1}^2$. The contribution of such σ to the integral above is therefore

$$\ll \frac{\varepsilon'}{\varepsilon} M \|\phi_1\|_{L^1}^2.$$

- If σ is not as just described, then $k(\sigma) = 0$, so (24.16) implies that $\mathrm{tr}(\sigma(\phi_1)) \leq 2\varepsilon$. But positivity of $\sigma(\phi_1)$ implies that $\mathrm{tr}(\sigma(\phi_1 * \phi_1)) \leq \mathrm{tr}(\sigma(\phi_1))^2$. The integrand on the right-hand side of (24.18) is thus bounded by 16ε .

Taken together, we get

$$\phi_3(1) \ll \frac{\varepsilon'}{\varepsilon} M \|\phi_1\|_{L^1}^2 + \varepsilon. \tag{24.19}$$

Now, ϕ_1 depends on ε . However, choosing first ε and then reducing ε' as appropriate, equations (24.17) and (24.19) show that $(\phi_1, \phi_3 + 2M\varepsilon e_J)$ give the desired pair of functions (up to a final rescaling of ε).

24.5. Proofs for steps 3, 4 and 5

24.5.1. Uniform distinction

LEMMA 24.6. *Let J be a good compact open subgroup of H . There is a positive-definite (smooth) tensor $T \in \pi^J \otimes \overline{\pi^J}$ such that $\mathcal{H}_\sigma(T) \geq 1$ for all $\sigma \in \Omega^J$.*

Proof. Each $\sigma \in \Omega^J$ is π -distinguished and contains non-zero J -invariant vectors, so we may find $x \in \pi^J$ with $\mathcal{H}_\sigma(x \otimes \bar{x}) \geq 2$. By continuity—using the formula (24.1) for \mathcal{H}_σ in terms of matrix coefficients—we then have $\mathcal{H}_{\sigma'}(x \otimes \bar{x}) \geq 1$ for all $\sigma' \in \Omega^J$ in some

neighborhood of σ . Thus, by the compactness of Ω^J , we may find a finite collection of vectors $x_1, \dots, x_n \in \pi^J$ and corresponding finite-rank tensor

$$T = \sum_j x_j \otimes \bar{x}_j$$

such that $\mathcal{H}_\sigma(T_0) \geq 1$ for all $\sigma \in \mathcal{D}$. \square

24.5.2. The real form of a Bernstein component

Let Θ be any Bernstein component. We denote by Θ^{unit} the image in Θ of \widehat{H} , i.e., the set of infinitesimal characters in Θ arising from some *unitary* representation. Since unitary representations are isomorphic to their conjugate-dual, Θ^{unit} is pointwise fixed by the anti-holomorphic involution

$$\Theta \ni [(M, \tau)] \mapsto [(M, \tau^+)],$$

where as usual τ^+ denotes conjugate dual. That involution defines a real form of Θ whose real points contain Θ^{unit} . We henceforth abuse notation slightly by writing $\mathbb{R}[\Theta]$ for the set of regular functions on that real form, with real coefficients; any such function is real-valued on Θ^{unit} .

We write $\Theta^0 \subseteq \Theta^{\text{unit}}$ for the subset consisting of $[(M, \tau)]$, with τ unitary (i.e., $\tau \cong \tau^+$) and set $\Theta^{\text{nt}} := \Theta^{\text{unit}} - \Theta^0$. The set Θ^0 is in general a finite quotient of a compact torus; we equip it with the push-forward of an arbitrary Haar measure on the latter.

Example 24.7. Suppose $G = \text{PGL}_2(F)$ and that $\Theta = \mathfrak{X}/W$ is the principal series component as considered above, so that we may identify $\mathfrak{X} \cong \mathbb{C}^\times$ by sending χ to its value α on a uniformizer and Θ with the quotient of \mathbb{C}^\times by the equivalence relation \sim defined by the inversion map $\alpha \mapsto \alpha^{-1}$. We then have the following identifications (here q denotes the cardinality of the residue field of F):

- $\{\text{real points of } \Theta\} \cong \mathbb{C}^{(1)} \cup \mathbb{R}^\times / \sim$;
- $\Theta^{\text{unit}} \cong \mathbb{C}^{(1)} \cup [q^{-1/2}, q^{1/2}] / \sim$;
- $\Theta^0 \cong \mathbb{C}^{(1)} / \sim$;
- $\Theta^{\text{nt}} \cong (q^{-1/2}, 1) \cup (1, q^{1/2}) / \sim$.

We may identify $\mathbb{C}[\Theta]$ with the ring of Laurent polynomials in α that are invariant under $\alpha \mapsto \alpha^{-1}$, and $\mathbb{R}[\Theta]$ with the subring consisting of those having real coefficients.

LEMMA 24.8. *Then, the closure $\overline{\Theta^{\text{nt}}}$ of Θ^{nt} intersects Θ^0 in a set of measure zero.*

Proof. Let $t = [(M, \tau)] \in \Theta^{\text{unit}}$. Then, there is a unitary representation $\sigma \in \widehat{H}$ with infinitesimal character $\lambda_\sigma = t$. Since σ is isomorphic to its own conjugate dual σ^+ , we

have $\lambda_\sigma = \lambda_{\sigma^+}$, so that $\tau^+ \cong w\tau$ for some $w \in N(M)/M$. If w is trivial, then τ has unitary central character; thus σ is tempered and thus $t \in \Theta^0$.

It follows that each $t \in \overline{\Theta^{\text{nt}}} \cap \Theta^0$ is contained in the set

$$\{[(M, \tau)] \in \Theta^0 : w\tau \cong \tau\}$$

for some $1 \neq w \in N(M)/M$. Each of these sets has measure zero. \square

24.5.3. Applications of Stone–Weierstrass

LEMMA 24.9. *Let \mathcal{D} , \mathcal{D}' and Θ be as in §24.4. For any $k \in C_c(\mathcal{D})$ and $\varepsilon > 0$ there is a regular function $f \in \mathbb{R}[\Theta]$ such that*

- $|k - f| < \varepsilon$ on \mathcal{D} ;
- $|f| < \varepsilon$ on \mathcal{D}' .

Proof. This follows from a variant of Stone–Weierstrass. We spell out some details.

Since \mathcal{D} is compact, infinitesimal character induces a homeomorphism between it and its image $\lambda(\mathcal{D})$ in Θ . Therefore, the continuous function k is pulled back from a continuous function (also denoted k) on $\lambda(\mathcal{D})$.

Moreover, $\lambda(\mathcal{D})$ is disjoint from $\lambda(\mathcal{D}')$ by assumption, and each of these sets is closed. By Tietze’s extension theorem, we may find a continuous function on Θ which induces k on $\lambda(\mathcal{D})$, and is zero on $\lambda(\mathcal{D}')$. The union of these sets is contained in a compact subset of the real points of Θ , and then we apply Stone–Weierstrass as usual. \square

LEMMA 24.10. *For each $\varepsilon > 0$ there is a regular function $f \in \mathbb{R}[\Theta]$ with the following properties:*

- f is valued in $[-2, 2]$ on Θ^{unit} ;
- $f \geq 1$ on Θ^{nt} ;
- $\int_{\Theta^0} f^2 \leq \varepsilon$.

For instance, in the above $\text{PGL}_2(F)$ example, the conclusion of the lemma holds with

$$f := \frac{\alpha^{-n} + \alpha^{-n+2} + \dots + \alpha^n}{n}$$

for $n \in \mathbb{Z}_{\geq 1}$ taken large enough in terms of ε . Note also that, in the third part, we take the measure on Θ^0 to be that induced from Haar measure, but that implies a similar statement for Plancherel measure, which is a continuous multiple of the Haar measure (see, e.g., [Wld]).

Proof. Write $\overline{\Theta^{\text{unit}}}$ for the closure of Θ^{unit} . Then, $\overline{\Theta^{\text{unit}}}$ is compact. By Stone–Weierstrass, $\mathbb{R}[\Theta]$ is dense in the space of continuous real-valued functions ϕ on $\overline{\Theta^{\text{unit}}}$. Choose $\varepsilon_1 > 0$ sufficiently small in terms of ε . It then suffices to find such a ϕ for which

- $|\phi| \leq 2 - \varepsilon_1$ on $\overline{\Theta^{\text{unit}}}$,
- $\phi \geq (1 + \varepsilon_1)$ on Θ^{nt} , and
- $\int_{\Theta^0} \phi^2 \leq \varepsilon_1$,

because then we may find $f \in \mathbb{R}[\Theta]$ with $\|f - \phi\| \leq \varepsilon_1$ on $\overline{\Theta^{\text{unit}}}$, and this satisfies the required conditions.

The existence of ϕ follows from Urysohn’s lemma, using Lemma 24.8. \square

Part V. Application to the averaged Gan–Gross–Prasad period

We aim now to formulate and prove our main result (Theorem 31.11, stated at the very end).

25. Setting

25.1. Basic setup

Let F be a number field; denote by \mathbb{Z}_F its ring of integers (we are using the letter \mathcal{O} for a coadjoint orbit) and by \mathbb{A} its adèle ring.

Let (\mathbf{G}, \mathbf{H}) be a GGP pair over F , in the sense of §13. We denote by \mathfrak{p} a typical place of F (possibly archimedean!) and by $F_{\mathfrak{p}}$ the corresponding completion. When \mathfrak{p} is non-archimedean, we denote by $\mathbb{Z}_{\mathfrak{p}} \subseteq F_{\mathfrak{p}}$ the ring of integers. We set $G_{\mathfrak{p}} := \mathbf{G}(F_{\mathfrak{p}})$ and $H_{\mathfrak{p}} := \mathbf{H}(F_{\mathfrak{p}})$.

We fix a finite set R of places of F which is sufficiently large in the following senses:

- R contains every archimedean place.
- The groups \mathbf{G} and \mathbf{H} admit smooth models over $\mathbb{Z}_F[1/R]$, which we continue to denote by \mathbf{G} and \mathbf{H} . This implies that for each $\mathfrak{p} \notin R$, the subgroups

$$K_{\mathfrak{p}} := \mathbf{G}(\mathbb{Z}_{\mathfrak{p}}) \leq G_{\mathfrak{p}} \quad \text{and} \quad J_{\mathfrak{p}} := \mathbf{H}(\mathbb{Z}_{\mathfrak{p}}) \leq H_{\mathfrak{p}}$$

are hyperspecial maximal compact subgroups.

- The inclusion $\mathbf{H} \hookrightarrow \mathbf{G}$ extends to a closed immersion of the smooth models over $\mathbb{Z}_F[1/R]$, so that $K_{\mathfrak{p}}$ contains $J_{\mathfrak{p}}$.

- Set

$$G_R := \prod_{\mathfrak{p} \in R} G_{\mathfrak{p}}, \quad K := \prod_{\mathfrak{p} \notin R} K_{\mathfrak{p}}, \quad H_R := \prod_{\mathfrak{p} \in R} H_{\mathfrak{p}} \quad \text{and} \quad J := \prod_{\mathfrak{p} \notin R} J_{\mathfrak{p}}.$$

Then, $\mathbf{G}(F) \cdot G_R \cdot K = \mathbf{G}(\mathbb{A})$, and similarly for \mathbf{H} in place of \mathbf{G} .

25.2. Measures

We equip the quotients $\mathbf{G} := \mathbf{G}(F) \backslash \mathbf{G}(\mathbb{A})$ and $[\mathbf{H}] := \mathbf{H}(F) \backslash \mathbf{H}(\mathbb{A})$ with Tamagawa measures and denote by $\tau(\mathbf{G})$ and $\tau(\mathbf{H})$ their volumes.

We fix a factorization of the associated measures on $\mathbf{G}(\mathbb{A}) = \prod' G_{\mathfrak{p}}$ and $\mathbf{H}(\mathbb{A}) = \prod' H_{\mathfrak{p}}$ in such a way that K and J have volume 1. We always equip products, such as G_R and H_R , with the product of the Haar measures on the corresponding components $G_{\mathfrak{p}}$ and $H_{\mathfrak{p}}$.

25.3. Automorphic forms

For the rest of this paper, the letters Π and Σ denote irreducible square-integrable automorphic representations $\Pi \subseteq L^2([\mathbf{G}])$ and $\Sigma \subseteq L^2([\mathbf{H}])$ that are unramified outside R , i.e., that admit vectors invariant by K and J , respectively. More precisely, we write Π and Σ for the *subspaces spanned by the smooth factorizable vectors* in the corresponding Hilbert space representations, so that we may identify

$$\Pi \cong \bigotimes_{\mathfrak{p}} \Pi_{\mathfrak{p}} \quad \text{and} \quad \Sigma \cong \bigotimes_{\mathfrak{p}} \Sigma_{\mathfrak{p}},$$

where $\Pi_{\mathfrak{p}}$ and $\Sigma_{\mathfrak{p}}$ are smooth irreducible unitarizable representations of $G_{\mathfrak{p}}$ and $H_{\mathfrak{p}}$.

For $\mathfrak{p} \notin R$, the spaces $\Pi_{\mathfrak{p}}^{K_{\mathfrak{p}}}$ and $\Sigma_{\mathfrak{p}}^{J_{\mathfrak{p}}}$ are 1-dimensional, so the fixed subspaces Π^K and Σ^J define irreducible representations of G_R and H_R , respectively. We may identify these fixed subspaces with the products of local components at R :

$$\Pi^K \cong \Pi_R := \bigotimes_{\mathfrak{p} \in R} \Pi_{\mathfrak{p}} \quad \text{and} \quad \Sigma^J \cong \Sigma_R := \bigotimes_{\mathfrak{p} \in R} \Sigma_{\mathfrak{p}}.$$

We fix unitary structures on Π_R and Σ_R such that the above identifications are isometric.

Here, and later, when we (e.g.) sum over Σ , we always have in mind that we sum over a set of representatives of Σ as above, whose Hilbert space direct sum is $L^2([\mathbf{H}])$. (The choice of representatives is ambiguous only in the event of global multiplicity larger than 1.)

25.4. Branching coefficients

Assume that (Π_R, Σ_R) is *distinguished* in that the space of H_R -invariant linear forms on $\Pi_R \otimes \Sigma_R^{\vee}$ is non-zero. That space is then 1-dimensional. The space \mathcal{I} consisting of all H_R -invariant Hermitian forms on

$$\Pi_R \otimes \Pi_R^{\vee} \otimes \Sigma_R^{\vee} \otimes \Sigma_R \longrightarrow \mathbb{C}$$

is likewise 1-dimensional. We may define $\mathcal{P} \in \mathcal{I}$ by the formula

$$\mathcal{P}(v_1 \otimes v_2 \otimes u_1 \otimes u_2) := \left(\int_{[\mathbf{H}]} v_1 \bar{u}_1 \right) \left(\int_{[\mathbf{H}]} \bar{v}_2 u_2 \right). \tag{25.1}$$

If Π_R and Σ_R are tempered, then we may define $\mathcal{H} \in \mathcal{I}$ by

$$\mathcal{H}(v_1 \otimes v_2 \otimes u_1 \otimes u_2) := \int_{h \in H_R} \langle hv_1, v_2 \rangle \langle u_1, hu_2 \rangle, \tag{25.2}$$

using the temperedness assumption to justify convergence (cf. §18).

If \mathcal{H} is non-zero—as is expected (cf. §20)—then it spans the 1-dimensional space \mathcal{I} , so we may define a branching coefficient $\mathcal{L}(\Pi, \Sigma) \in \mathbb{R}_{\geq 0}$ by requiring that

$$\mathcal{P} = \mathcal{L}(\Pi, \Sigma) \cdot \mathcal{H} \quad \text{on } \Pi_R \otimes \Pi_R^\vee \otimes \Sigma_R^\vee \otimes \Sigma_R. \tag{25.3}$$

We have suppressed from our notation the dependence of $\mathcal{L}(\Pi, \Sigma)$ upon the fixed set R of places at which everything is assuming unramified.

25.5. The conjectures of Ichino–Ikeda and N. Harris

See [II] and [HarN].

Conjecture 25.1. If (Π_R, Σ_R) is distinguished and Π_R and Σ_R are tempered, then \mathcal{H} , as defined in §25.4, is non-zero, so $\mathcal{L}(\Pi, \Sigma)$ is defined; it is given by

$$\mathcal{L}(\Pi, \Sigma) = 2^{-\beta} \frac{L^{(R)}\left(\frac{1}{2}, -, \Pi \times \Sigma^\vee\right)}{L^{(R)}(1, \text{Ad}, \Pi \times \Sigma^\vee)} \Delta_G^{(R)}, \tag{25.4}$$

where 2^β is the order of the component group of the Arthur parameter for $\Pi \boxtimes \Sigma$ and $\Delta_G^{(R)}$ is, as in the introduction, the partial L -factor of the L -function whose local factor at almost every prime p equals

$$\frac{p^{\dim(G)}}{\#\mathbf{G}(\mathbf{F}_p)}.$$

Remark 25.2. We expect (but have not attempted to verify rigorously) that

$$2^\beta = \tau(\mathbf{G})\tau(\mathbf{H})$$

holds generically, that is to say, for “typical” Π and Σ . Let us explain where this comes from.

The left-hand side is, by definition, the number of components of the centralizer (in $G^\vee \times H^\vee$) of the Arthur parameter for $\Pi \boxtimes \Sigma$. Now, for “typical” Π and Σ , one expects

that the image of its Arthur parameter meets $G^\vee \times H^\vee$ in a Zariski dense set. Then, the centralizer in question is simply the set of Galois-invariants in the center $Z(G^\vee \times H^\vee)$. The cardinality of the latter is directly related to Tamagawa numbers: a result of Kottwitz [Kot, equation (5.1.1)], building on work of Sansuc, shows that

$$\tau(\mathbf{G})\tau(\mathbf{H}) = \frac{\#Z(G^\vee \times H^\vee)^{\text{Gal}}}{h},$$

where h is the order of the Tate–Shafarevich group for $Z(G^\vee \times H^\vee)$. In the cases at hand, $Z(G^\vee)$ and $Z(H^\vee)$ are either $\{\pm 1\}$ or the torus \mathbb{G}_m , and in the latter case the Galois action is through a quadratic character; in all cases, $h=1$.

25.6. Some unconventional notation

25.6.1. We fix once and for all an archimedean place \mathfrak{q} of F . This place plays a privileged role in both our results and their proofs, so we introduce the otherwise unconventional notation

$$\begin{aligned} G &:= G_{\mathfrak{q}}, \\ G' &:= \prod_{\mathfrak{p} \in R \setminus \{\mathfrak{q}\}} G_{\mathfrak{p}}, \\ \Gamma &:= \mathbf{G}(\mathbb{Z}_F[1/R]) = \mathbf{G}(F) \cap K \hookrightarrow G \times G', \\ [G] &:= \Gamma \backslash (G \times G'), \end{aligned}$$

so that $[G] \cong [\mathbf{G}]/K$. We analogously define H, H', Γ_H and

$$[H] := \Gamma_H \backslash (H \times H') \cong [\mathbf{H}]/J.$$

By our choice of factorization of Haar measures, the quotients $[G]$ and $[H]$ have volumes $\tau(\mathbf{G})$ and $\tau(\mathbf{H})$.

25.6.2. We denote, as in §2.1, by \mathfrak{g} and \mathfrak{h} the Lie algebras of G and H , respectively, and by $\mathfrak{g}^\wedge \cong i\mathfrak{g}^*$ and $\mathfrak{h}^\wedge \cong i\mathfrak{h}^*$ their the Pontryagin duals. We have normalized Haar measures both on H and H' , giving rise to Plancherel measures on the unitary duals of both. The Haar measure on H normalizes Haar measures on \mathfrak{h} and \mathfrak{h}^\wedge , hence a normalized affine measure on the GIT quotient $[\mathfrak{h}^\wedge]$ (see §9 and §17.2).

25.6.3. Recall that for any Π and Σ , we may isometrically identify and embed

$$\Pi_R \cong \Pi^K \hookrightarrow L^2([G]) \quad \text{and} \quad \Sigma_R \cong \Sigma^J \hookrightarrow L^2([H]).$$

We may then unitarily factor

$$\Pi_R = \pi \otimes \pi' \quad \text{and} \quad \Sigma_R = \sigma \otimes \sigma',$$

where

$$\pi \cong \Pi_{\mathfrak{q}}, \quad \sigma = \Sigma_{\mathfrak{q}}, \quad \pi' \cong \prod_{\mathfrak{p} \in R \setminus \{\mathfrak{q}\}} \Pi_{\mathfrak{p}} \quad \text{and} \quad \sigma' \cong \prod_{\mathfrak{p} \in R \setminus \{\mathfrak{q}\}} \Sigma_{\mathfrak{p}}$$

are smooth irreducible unitary representations of G , H , G' and H' , respectively.

25.7. Assumptions

For the remainder of the paper, we fix an individual Π as above. We now impose several assumptions concerning \mathbf{G} , \mathbf{H} , R , \mathfrak{q} and Π :

- (1) The representations π of G and π' of G' are tempered.
 - (2) \mathbf{G} and \mathbf{H} are anisotropic and non-trivial. This has the following consequences:
 - $[\mathbf{G}]$, $[G]$, $[\mathbf{H}]$ and $[H]$ are compact;
 - the pair (\mathbf{G}, \mathbf{H}) is isomorphic either to
 - (U_{n+1}, U_n) with $n \geq 1$, or to
 - (SO_{n+1}, SO_n) with $n \geq 2$,
- and not to (GL_{n+1}, GL_n) . In particular,

$$\tau(\mathbf{G}) = \tau(\mathbf{H}) = 2$$

(see, e.g., [We] and [On]).

- (3) $G_{\mathfrak{p}}$ is compact for every archimedean place $\mathfrak{p} \neq \mathfrak{q}$.

(4) \mathbf{G} is quasi-split at \mathfrak{q} , so that G is quasi-split, and the representation π is generic, and hence satisfies the equivalent conditions of Lemma 11.6. (The assumption concerning π holds if, for instance, it belongs to the principal series.) In particular, the limit orbit \mathcal{O} of π is a *non-empty* union of open G -orbits on the regular subset \mathcal{N}_{reg} of the nilcone in \mathfrak{g}^{\wedge} :

$$\emptyset \neq \mathcal{O} \subseteq \mathcal{N}_{\text{reg}}.$$

Example 25.3. Take $F = \mathbb{Q}(\alpha)$, $\alpha^2 = 5$. Let $n \in \mathbb{Z}_{\geq 4}$. Take for \mathbf{G} the special orthogonal group of the F -quadratic form

$$x_1^2 + \dots + x_m^2 + (1 - \alpha)(x_{m+1}^2 + \dots + x_n^2),$$

where $n=2m$ or $2m+1$.

Similarly, let \mathbf{H} be the special orthogonal group of the quadratic form

$$x_1^2 + \dots + x_m^2 + (1-\alpha)(x_{m+1}^2 + \dots + x_{n-1}^2),$$

embedded in \mathbf{G} as usual. Let \mathfrak{q} be the archimedean place of F sending α to $\sqrt{5}=2.23\dots$, and \mathfrak{q}' the other archimedean place. Then, $G \cong \mathrm{SO}(m, m)$ or $\mathrm{SO}(m, m+1)$ is split, while $G_{\mathfrak{q}'} \cong \mathrm{SO}(n)$ is compact; similarly, $H \cong \mathrm{SO}(m, m-1)$ or $\mathrm{SO}(m, m)$ and $H_{\mathfrak{q}'} \cong \mathrm{SO}(n-1)$.

Remark 25.4. Assumption (1), or some strong bound in that direction, is required by the formulation of the Ichino–Ikeda and N. Harris conjectures; we also exploit it through our use of the Kirillov formula. Assumption (2) is essential for the measure classification step. Assumption (3) is primarily for convenience. Assumption (4) ensures that regular nilpotent elements exist, or equivalently, that the limit coadjoint orbits considered here be non-empty (cf. §11.4.2).

26. Construction of limit states

26.1. Setting

Recall that we have defined the quotient

$$[G] = \Gamma \backslash (G \times G'),$$

where

- G is a reductive group over an archimedean local field,
- G' is an S -arithmetic group, and
- Γ is a cocompact lattice in $G \times G'$.

The groups G and G' arose from “half” of a GGP pair, but the properties just enunciated are what matter here. We denote by $\mu_{[G]}$ the Haar measure that we have normalized on $[G]$. From an automorphic representation Π of \mathbf{G} , we obtained a unitary $G \times G'$ subrepresentation

$$\pi \otimes \pi' \hookrightarrow L^2([G]).$$

We will assume starting in §26.4 that π varies with an infinitesimal parameter $\mathfrak{h} \rightarrow 0$ and admits a regular limit coadjoint orbit (cf. §11.4)

$$(\mathcal{O}, \omega) = \lim_{\mathfrak{h} \rightarrow 0} (\mathfrak{h} \mathcal{O}_\pi, \omega_{\mathfrak{h} \mathcal{O}_\pi}).$$

In applications, \mathcal{O} will be a subset of $\mathcal{N}_{\mathrm{reg}}$, but this feature plays no role in §26.

26.2. Overview

The main aim of this section is to construct, after passing to a subsequence of $\{h\}$, a natural G -equivariant assignment

$$\mathcal{O} \longrightarrow \{\text{probability measures on } [G]\} \quad (26.1)$$

which captures the average limiting behavior of the L^2 -masses $|v|^2 d\mu_{[G]}$ of vectors $v \in \pi$ microlocalized at a point $\xi \in \mathcal{O}$. The construction shares many common features with standard constructions in the pseudodifferential calculus (referred to variously as the quantum limits, semiclassical limits, microlocal defect measures, etc.).

This construction will be achieved in the following three stages.

(1) *From operators to functions.* Recall, from §3, the space $\Psi^{-\infty} := \Psi^{-\infty}(\pi)$ of “smoothing operators” on π ; it contains $\pi \otimes \bar{\pi}$ as the subspace of finite-rank operators. Recall also, from §12.1, the space

$$\mathcal{T}_1 := \mathcal{T}_1(\pi) = \{\text{“smoothly trace class” operators on } \pi\}.$$

The first stage, achieved in §26.3, is to construct a G -equivariant positivity-preserving map

$$\begin{aligned} \mathcal{T}_1 &\longrightarrow C^\infty([G]), \\ T &\longmapsto [T], \end{aligned}$$

such that the trace of T is the integral of $[T]$. (The map will depend upon the choice of a fixed “family of vectors” in the auxiliary representation π' .)

(2) *From symbols to functions.* The second stage, achieved in §26.4, is to compose the map $T \mapsto [T]$ with the operator calculus

$$\text{Op}_h : S^m := S^m(\mathfrak{g}^\wedge) \longrightarrow \Psi^m := \Psi^m(\pi),$$

specialized here to the Schwartz space $S^{-\infty} = \mathcal{S}(\mathfrak{g}^\wedge)$. This gives a family of maps

$$\mathcal{S}(\mathfrak{g}^\wedge) \longrightarrow C^\infty([G]),$$

depending upon h . We will show that the leading order asymptotics of this family of maps are described, after passing to a subsequence, by a limit map

$$[\cdot] : \mathcal{S}(\mathfrak{g}^\wedge) \longrightarrow C^\infty([G]),$$

with several natural properties.

Informally (cf. §1.7), fix a small non-empty open set $E \subseteq \mathcal{O}$, and suppose that $a|_{\mathcal{O}}$ defines a smooth approximation to the normalized characteristic function $\text{vol}(E, d\omega)^{-1}1_E$ of E . Then, $[a]$ roughly describes the limiting *average* of $|v|^2$, taken over all vectors $v \in \pi$ microlocalized in E .

(3) *From points to measures.* The third stage, achieved in §26.5, is to describe the map $[\cdot]$ in terms of measures. The description may be understood as an effective implementation of measure disintegration. It will allow us to analyze our limit states using measure-theoretic techniques, notably Ratner's theorem.

26.3. Stage 1: from operators to functions

We fix a (smooth, finite-rank) non-zero element $T' \in \pi' \otimes \overline{\pi'}$ with the following properties:

- T' is positive definite;
- *unless explicitly stated otherwise*, $\text{tr}(T')=1$.

Equivalently, T' is a finite sum

$$T' = \sum_i c_i u_i \otimes \bar{u}_i, \quad (26.2)$$

where $u_i \in \pi'$ are smooth unit vectors and c_i are non-negative reals summing to 1. The normalization $\text{tr}(T')=1$ serves to simplify notation; in practice, we may reduce to the case in which it is satisfied by multiplying T' by a suitable positive scalar.

For each $T \in \pi \otimes \bar{\pi}$, we obtain an element $T \otimes T' \in L^2([G] \times [G])$. We denote by

$$[T \otimes T'] \in L^1([G])$$

its diagonal restriction. Since the interesting variation will happen in the T variable, we will often drop the T' from the notation by abbreviating

$$[T] := [T \otimes T'] \in L^1([G]).$$

Example 26.1. Suppose that $T = v_1 \otimes \bar{v}_2$ and $T' = w_1 \otimes \bar{w}_2$, with $v_i \in \pi$ and $w_i \in \pi'$, so that $u_i := v_i \otimes w_i$ defines an element of $\pi \otimes \pi' \hookrightarrow L^2([G])$. Then,

$$(T \otimes T')(x, y) = u_1(x) \overline{u_2(y)},$$

while

$$[T](x) = u_1(x) \overline{u_2(x)}.$$

LEMMA 26.2. Fix T' as above. Abbreviate “continuous, uniformly in π ” to “continuous”.

- (i) $\|[T]\|_{L^1([G])} \leq \|T\|_1$, where $\|\cdot\|_1$ denotes the trace norm.
- (ii) The map $T \mapsto T \otimes T'$ extends uniquely to a continuous $G \times G$ -equivariant map

$$\Psi^{-\infty} \longrightarrow C^\infty([G] \times [G]).$$

- (iii) The map $T \mapsto [T]$ extends uniquely to a continuous G -equivariant map

$$\mathcal{T}_1 \longrightarrow C^\infty([G]).$$

- (iv) For $T \in \mathcal{T}_1$, we have

$$\int_{[G]} [T] = \text{tr}(T).$$

- (v) If T is a non-negative self-adjoint operator, then $[T]$ is a non-negative function:

$$T \geq 0 \implies [T] \geq 0.$$

- (vi) $[T]$ is invariant by the action of the center of $\mathbf{G}(\mathbb{A})$ on $[G]$.

Proof. For (i), let $\mathcal{B}(\pi)$ be an orthonormal basis for π consisting of eigenvectors v for the non-negative self-adjoint finite-rank operator $\sqrt{T^*T}$, with eigenvalues $c_v \geq 0$. Writing T' as in (26.2), we then have

$$[T] = \sum_{v \in \mathcal{B}(\pi)} \sum_i c_i (Tv \otimes u_i) \cdot \overline{v \otimes u_i}, \tag{26.3}$$

where each sum is really a finite sum. By the Cauchy–Schwarz inequality,

$$\begin{aligned} \int_{[G]} |(Tv \otimes u_i) \cdot \overline{v \otimes u_i}| &\leq \left(\int_{[G]} |Tv \otimes u_i|^2 \right)^{1/2} = \langle Tv \otimes u_i, Tv \otimes u_i \rangle^{1/2} \\ &= \langle T^*Tv \otimes u_i, v \otimes u_i \rangle^{1/2} = c_v. \end{aligned}$$

Thus by the triangle inequality,

$$\|[T]\|_{L^1([G])} \leq \left(\sum_{v \in \mathcal{B}(\pi)} c_v \right) \left(\sum_i c_i \right) = \sum_{v \in \mathcal{B}(\pi)} c_v = \|T\|_1.$$

For (ii), we note that the map in question is isometric up to a constant factor (given by the Hilbert–Schmidt norm of T'); the conclusion follows from the Sobolev lemma, applied on the homogeneous space $[G] \times [G]$. We analogously deduce (iii) from (i) and the Sobolev lemma on $[G]$.

For (iv) and (v), the proof reduces by continuity to the finite-rank case, and then by linearity to the rank-1 case as in the above example. In that case,

$$\int_{[G]} [T] = \langle u_1, u_2 \rangle = \langle v_1, v_2 \rangle \langle w_1, w_2 \rangle = \text{tr}(T) \text{tr}(T') = \text{tr}(T).$$

If $T \geq 0$, then we may assume moreover that $v_1 = v_2$ and $w_1 = w_2$, hence that $u_1 = u_2 =: u$, and so

$$[T](x) = |u(x)|^2 \geq 0.$$

For (vi), let ω denote the central character of Π . Then, the function $T \otimes T'$ on $[G] \times [G]$ has central character (ω, ω^{-1}) , and so its restriction $[T]$ to $[G]$ has trivial central character. \square

26.4. Stage 2: from symbols to functions

We now allow π to vary with a positive parameter $\hbar \rightarrow 0$ in such a way that we obtain a regular limit coadjoint orbit (\mathcal{O}, π) . We assume that \hbar traverses some sequence $\{\hbar\}$ having zero as its unique accumulation point.

For $k \in \mathbb{Z}_{\geq 1}$ and $(a_1, \dots, a_k) \in \mathcal{S}(\mathfrak{g}^\wedge)^k$, the operator $\text{Op}_\hbar(a_1) \dots \text{Op}_\hbar(a_k)$ belongs to \mathcal{T}_1 , and thus yields a smooth function

$$\begin{aligned} [a_1, \dots, a_k]_\hbar : [G] &\longrightarrow \mathbb{C} \\ [a_1, \dots, a_k]_\hbar &:= \hbar^d [\text{Op}_\hbar(a_1) \dots \text{Op}_\hbar(a_k)], \end{aligned} \tag{26.4}$$

where the exponent d in the normalizing factor \hbar^d denotes half the real dimension of \mathcal{O} , as usual.

LEMMA 26.3. *The linear map*

$$[\cdot]_\hbar : \mathcal{S}(\mathfrak{g}^\wedge)^k \longrightarrow C^\infty([G])$$

defined by (26.4) is \hbar -uniformly continuous. That is to say, for each continuous seminorm ν on $C^\infty([G])$ and each $k \geq 1$, there is a continuous seminorm μ on $\mathcal{S}(\mathfrak{g}^\wedge)$ so that for all $\hbar \in (0, 1]$ and all $a_1, \dots, a_k \in \mathcal{S}(\mathfrak{g}^\wedge)$, we have

$$\nu([a_1, \dots, a_k]_\hbar) \leq \mu(a_1) \dots \mu(a_k).$$

Proof. Indeed, by (12.7) and Lemma 26.2 (iii), $[\cdot]_\hbar$ is a composition

$$[\cdot]_\hbar : \mathcal{S}(\mathfrak{g}^\wedge)^k \xrightarrow{\hbar^d \text{Op}_\hbar^{\otimes k}} \mathcal{T}_1 \xrightarrow{[\cdot]} C^\infty([G])$$

of \hbar -uniformly continuous maps. \square

We aim now to show that this family of maps has a limit satisfying several natural properties.

THEOREM 26.4. *Let notation be as in §26.1 and (26.4). After passing to a subsequence of $\{h\}$, there exist continuous maps (depending upon π and the choice of subsequence)*

$$\begin{aligned} [\cdot]: \mathcal{S}(\mathfrak{g}^\wedge)^k &\longrightarrow C^\infty([G]), \\ (a_1, \dots, a_k) &\longmapsto [a_1, \dots, a_k], \end{aligned}$$

indexed by $k \in \mathbb{Z}_{\geq 1}$ with the following properties:

- (i) for fixed $k \in \mathbb{Z}_{\geq 1}$, the function $[a_1, \dots, a_k]_h$ converges to $[a_1, \dots, a_k]$ in the C^∞ -topology, with continuous dependence upon (a_1, \dots, a_k) ;
- (ii) $[a_1, \dots, a_k] = [a_1 \dots a_k]$; in particular,

$$[a_1 \dots a_k] = \lim_{h \rightarrow 0} h^d [\text{Op}_h(a_1) \dots \text{Op}_h(a_k)] \quad \text{in } C^\infty([G]); \tag{26.5}$$

- (iii) $[a, \bar{a}] \geq 0$;
- (iv) if $a \geq 0$, then $[a] \geq 0$;
- (v) $[\cdot]$ is G -equivariant (for any k);
- (vi)

$$\int_{[G]} [a] = \int_{\mathcal{O}} a \, d\omega. \tag{26.6}$$

Proof. (i) The existence and first property of $[\cdot]$ follows by an Arzela–Ascoli-type argument together with the noted uniform continuity of the maps $[\cdot]_h$. We record the details for convenience.

We observe first that, for any $(a_1, \dots, a_k) \in \mathcal{S}(\mathfrak{g}^\wedge)^k$, the subset

$$\{[a_1, \dots, a_k]_h : h \in (0, 1]\}$$

of $C^\infty([G])$ is *bounded*, i.e., each seminorm has bounded image. This feature is a consequence of the h -uniform continuity of $[\cdot]_h$ noted in the above lemma.

We observe next that each bounded subset of $C^\infty([G])$ is precompact. Indeed, boundedness forces equicontinuity of the subset, and therefore the existence of a subsequence that converges in the space $C([G])$ of continuous functions equipped with the supremum norm. One similarly extracts a further subsequence that converges along with its first derivatives in $C([G])$, and—proceeding in this way and using a diagonal argument—a subsequence that converges in $C^\infty([G])$.

We observe finally that $\mathcal{S}(\mathfrak{g}^\wedge)^k$ is a separable topological vector space. A countable dense subset may be given, for instance, by finite rational linear combinations of Hermite functions.

We now choose a countable dense \mathbb{Q} -subspace S_0 of $\mathcal{S}([\mathfrak{g}])^k$. By the second observation above, we may find a subsequence of $\{h\}$ along which $[a_1, \dots, a_k]_h$ converges in $C^\infty([G])$ for each (a_1, \dots, a_k) in S_0 . Call this limit $L(a_1, \dots, a_k)$. The noted h -uniform continuity of the maps $[\cdot]_h$ then implies that the map $L: S_0 \rightarrow C^\infty([G])$ is continuous, hence extends uniquely to the desired continuous map $\mathcal{S}([\mathfrak{g}])^k \rightarrow C^\infty([G])$. We verify readily that this extension is \mathbb{R} -linear.

(ii) We apply $[\cdot]: \mathcal{T}_1 \rightarrow C^\infty([G])$ to (12.8).

(iii) For every h , the operator

$$\mathrm{Op}_h(a)\mathrm{Op}_h(\bar{a}) = \mathrm{Op}_h(a)\mathrm{Op}_h(a)^*$$

is non-negative-definite, and so $[a, \bar{a}]_h \geq 0$ by the assumed properties of $T \mapsto [T]$. We conclude by taking limits.

(iv) By continuity, it suffices to show that $[a] \geq 0$ for every a of the form $a = b\bar{b}$, which is the content of the previous assertion.

(v) We may assume $k=1$. We then apply $[\cdot]: \mathcal{T}_1 \rightarrow C^\infty([G])$ to (12.9).

(vi) We have

$$\int_{[G]} [a] = \lim_{h \rightarrow 0} \int_{[G]} [a]_h \quad \text{and} \quad \int_{[G]} [a]_h = \mathrm{tr}(h^d \mathrm{Op}_h(a))$$

and (by (12.3))

$$\lim_{h \rightarrow 0} \mathrm{tr}(h^d \mathrm{Op}_h(a)) = \int_{\mathcal{O}} a \, d\omega. \quad \square$$

26.5. Stage 3: from points to measures

THEOREM 26.5. *Let notation and assumptions be as in Theorem 26.4; assume also that we have passed to a subsequence of $\{h\}$ for which that theorem hold. There is then a G -equivariant linear map*

$$\begin{aligned} \mathcal{O} &\longrightarrow \{\text{probability measures on } [G]\}, \\ \xi &\longmapsto [\delta_\xi], \end{aligned} \tag{26.7}$$

so that, for all $a \in C_c^\infty(\mathfrak{g}^\wedge)$ and $\Psi \in C_c([G])$,

$$\int_{[G]} [a]\Psi = \int_{\xi \in \mathcal{O}} a(\xi)([\delta_\xi]\Psi) \, d\omega(\xi). \tag{26.8}$$

An importance consequence of this theorem is that

$$[\delta_\xi] \text{ is invariant by the } G\text{-centralizer of } \xi. \tag{26.9}$$

The measures $[\delta_\xi]$ and the invariance property (26.9) generalize the “microlocal lift” of semiclassical analysis and its invariance under geodesic flow. Informally speaking, $[\delta_\xi]$ describes the *average* behavior of the measures $|v|^2\mu_{[G]}$, where $v \in \pi \subset L^2([G])$ is localized near ξ .

Proof of Theorem 26.5. Consider the distribution η on $\mathfrak{g}^\wedge \times [G]$ given for $a \in C_c^\infty(\mathfrak{g}^\wedge)$ and $\Psi \in C^\infty([G])$ by

$$\eta(a \otimes \Psi) := \int_{[G]} [a] \cdot \Psi. \tag{26.10}$$

(The Schwartz kernel theorem provides the unique extension of this definition from the algebraic tensor product $C_c^\infty(\mathfrak{g}^\wedge) \otimes C^\infty([G])$ of test function spaces to the space $C_c^\infty(\mathfrak{g}^\wedge \times [G])$ of test functions on the product space.) Part (iv) of Theorem 26.4 implies that η is positive on positive functions. Those results and the Riesz representation theorem then imply that η defines a G -invariant positive measure on $\mathfrak{g}^\wedge \times [G]$, pushing forward in the first coordinate to the measure ω on \mathcal{O} .

Since η is positive, we have for $a \in C_c(\mathfrak{g}^\wedge)$ and $\Psi \in C_c([G])$ that

$$|\eta(a \otimes \Psi)| \leq \eta(|a| \otimes |\Psi|) \leq \|\Psi\|_\infty \int_{\mathcal{O}} |a| d\omega. \tag{26.11}$$

Suppose, for a moment, that $\Psi \geq 0$. The rule $a \mapsto \eta(a \otimes \Psi)$ defines then a positive functional on $C_c(\mathfrak{g}^\wedge)$, thus a measure. This measure is absolutely continuous with respect to ω , and thus (by Radon–Nikodym) may be written as $f_\Psi \omega$ for some measurable non-negative function f_Ψ on \mathfrak{g}^\wedge . The bound (26.11) shows that in fact $f_\Psi \leq \|\Psi\|_\infty$ almost everywhere with respect to $d\omega$. The rule $\Psi \mapsto f_\Psi$ defines then a function

$$C_c([G])_{\geq 0} \longrightarrow L^\infty(\mathcal{O})$$

and this extends to a bounded linear map on all of $C_c([G])$ by splitting into positive and negative parts.

The image of the map $\Psi \mapsto f_\Psi$ in fact lies inside the space $C(\mathcal{O})$ of continuous functions, since there is a dense subspace of the source (all convolutions with elements in $C_c(G)$) for which this is true. This uses the following consequence of the regularity of the limit multiorbit \mathcal{O} :

$$\text{every } G\text{-orbit in } \mathcal{O} \text{ is open in } \mathcal{O}. \tag{26.12}$$

Finally, the composition

$$C_c([G]) \xrightarrow{\Psi \mapsto f_\Psi} C(\mathcal{O}) \xrightarrow{\text{eval. at } \xi} \mathbb{R}$$

gives the desired measure δ_ξ . □

27. Equidistribution of limit states

27.1. Overview and statement of result

We aim now to apply Ratner’s classification of invariant measures on a homogeneous space to obtain strong control over the limit states constructed in the prior section.

Recall, thus, that we are considering a real reductive group G , an S -arithmetic group G' , a cocompact lattice $\Gamma < G \times G'$, the corresponding compact quotient

$$[G] := \Gamma \backslash (G \times G'),$$

and a unitary $(G \times G')$ -subrepresentation $\pi \otimes \pi'$ of $L^2([G])$; this data arose from the automorphic representation Π of the reductive group \mathbf{G} over the number field F .

We have constructed, in the prior section, some G -equivariant “limit state” assignments

$$[\cdot]: \mathcal{S}(\mathfrak{g}^\wedge) \longrightarrow C^\infty([G])$$

and

$$\mathcal{O} \xrightarrow{\xi \mapsto [\delta_\xi]} \{\text{measures on } [G]\},$$

capturing the limiting behavior of the functions on $[G]$ obtained from high-frequency vectors in π . Due to the disintegration (26.8), which we abbreviate here to

$$[a] d\mu_{[G]} = \int_{\xi \in \mathcal{O}} a(\xi) [\delta_\xi] d\omega(\xi),$$

the assignments $[\cdot]$ are determined completely by the probability measures $[\delta_\xi]$ on $[G]$. Each measure $[\delta_\xi]$ is invariant by the centralizer $G_\xi \leq G$ of ξ .

We employ now crucially our assumption that $\pi \otimes \pi'$ is *fixed*, independent of the scaling parameter $h \rightarrow 0$. The limit coadjoint orbit \mathcal{O} for π is then contained in the regular nilpotent set $\mathcal{N}_{\text{reg}} \subseteq \mathfrak{g}^\wedge$, so the centralizers G_ξ contain unipotent elements. We might thus hope to apply Ratner’s theorem to deduce that $[\delta_\xi]$ is simply the Haar probability measure on $[G]$, and hence for any symbol $a \in \mathcal{S}(\mathfrak{g}^\wedge)$ that the limit state $[a]: [G] \rightarrow \mathbb{C}$ is simply the constant function satisfying the normalization condition $\int_{[G]} [a] = \int_{\mathcal{O}} a d\omega$. Unwinding the definition of the limit states, this may be understood informally as a form of quantum unique ergodicity for high-frequency vectors in the fixed representation π .

The actual picture is mildly more complicated due to “multiple component” issues arising from the presence of the auxiliary group G' ; the G -action on $[G]/J$, for the compact open $J < G'$ under which our measures have known invariance, is not in general transitive, and so we cannot hope to control the behavior of limit states across all of $[G]$. However, it turns out that the control obtained is enough to understand *integrals of limit states over* $[H]$, in the following sense sufficient for our purposes.

THEOREM 27.1. For any $a \in \mathcal{S}(\mathfrak{g}^\wedge)$,

$$\frac{1}{\tau(\mathbf{H})} \int_{[H]} [a] = \frac{1}{\tau(\mathbf{G})} \int_{[G]} [a]. \tag{27.1}$$

The proof occupies the remainder of this section.

Remark 27.2. The proof makes use of our assumption (§25.7) that G is quasi-split. On the other hand, without this assumption, we have $\mathcal{O} = \emptyset$ by §11.4.2, and hence $[a] = 0$ by Theorem 26.5, so the conclusion (27.1) is uninteresting.

27.2. Reduction to G_∞^+ -invariance

Set $F_\infty := F \otimes \mathbb{R}$, and let G_∞^+ denote the (topological) connected component of the group $\mathbf{G}(F_\infty)$ of archimedean points of \mathbf{G} . Recall that we have extended the rule $a \mapsto [a]$ from $\mathcal{S}(\mathfrak{g}^\wedge)$ to measures on \mathcal{O} in §26.5. We have the following reduction.

LEMMA 27.3. Assume that, for each $\xi \in \mathcal{O}_{\text{reg}}$, the measure $[\delta_\xi]$ is G_∞^+ -invariant. Then, the conclusion of Theorem 27.1 holds.

Before we give the proof, we give some algebraic preliminaries. Let $\tilde{\mathbf{G}}$ denote the simply connected covering group of the derived group $[\mathbf{G}, \mathbf{G}] \subset \mathbf{G}$. Because $\mathbf{G}(F_q)$ is quasi-split, each F -simple factor of $\tilde{\mathbf{G}}$ has non-compact F_q -points. By strong approximation [PR, Theorem 7.12], it follows that the group $\tilde{\mathbf{G}}(F_\infty)$ has dense image in $[\tilde{\mathbf{G}}]$. The image $\mathbf{G}(\mathbb{A})^+$ of $\tilde{\mathbf{G}}(\mathbb{A}) \rightarrow \mathbf{G}(\mathbb{A})$ is readily seen to contain the commutator subgroup of $\mathbf{G}(\mathbb{A})$, and thereby $\mathbf{G}(\mathbb{A})^+$ is a normal subgroup with abelian quotient.

Let U be an open compact subgroup of $\mathbf{G}(\mathbb{A})$ under which $[a]$ is invariant. By the noted density of $\tilde{\mathbf{G}}(F_\infty)$ in $[\tilde{\mathbf{G}}]$, we see that

$$\mathbf{G}(F)G_\infty^+U \supseteq \mathbf{G}(\mathbb{A})^+. \tag{27.2}$$

LEMMA 27.4. Each G_∞^+ -orbit on $[\mathbf{G}]/U$ has the same measure.

Proof. Each such orbit is of the form $\mathbf{G}(F) \backslash \mathbf{G}(F)G_\infty^+gU$, for some $g \in \mathbf{G}(\mathbb{A})$. We claim that the various open sets $\mathbf{G}(F)G_\infty^+gU$ are thus right translates of one another, and their quotients by $\mathbf{G}(F)$ have the same measure. To see this, we note that $\mathbf{G}(F)G_\infty^+gU$ is the right translate by g of $\mathbf{G}(F)G_\infty^+(gUg^{-1})$. This latter set contains $\mathbf{G}(\mathbb{A})^+$, and we deduce that

$$\mathbf{G}(F)G_\infty^+(gUg^{-1}) = \mathbf{G}(F)\mathbf{G}(\mathbb{A})^+G_\infty^+(gUg^{-1})$$

which is independent of g because the quotient $\mathbf{G}(\mathbb{A})/\mathbf{G}(\mathbb{A})^+$ is abelian. □

We may now give the postponed proof of Lemma 27.3.

Proof. The results in §26.5 imply that, if $a \in \mathcal{S}(\mathfrak{g}^\wedge)$, then $[a]$ is G_∞^+ -invariant.

The quotient $[\mathbf{G}]/U$ is a finite union of G_∞^+ -orbits, so it is enough to verify (27.1) when the smooth function $[a]$ is replaced by the characteristic function of one of these G_∞^+ orbits. To that end, it suffices to check that the intersection of $[H]$ with each one of these G_∞^+ orbits has equal measure; once that is shown, both sides of the equality are equal to $1/N$, with N the number of G_∞^+ -orbits.

The quotient $\mathbf{G}(\mathbb{A})/\mathbf{G}(\mathbb{A})^+$ is an abelian group. Thus, the further quotient

$$\varepsilon := \mathbf{G}(F)G_\infty^+U \backslash \mathbf{G}(\mathbb{A}),$$

which indexes the G_∞^+ -orbits on $[\mathbf{G}]/U$, has the structure of finite abelian group. The induced map $\mathbf{H}(\mathbb{A}) \rightarrow \varepsilon$ is a group homomorphism. The subset of $[H]$ meeting a specified component of $[\mathbf{G}]/U$ corresponds to a fiber of the map $\mathbf{H}(F) \backslash \mathbf{H}(\mathbb{A}) \rightarrow \varepsilon$. To see that each such fiber has the same measure, it is enough to show that this map $\mathbf{H}(\mathbb{A}) \rightarrow \varepsilon$ is surjective.

This surjectivity follows if we can verify the surjectivity of

$$\mathbf{G}(F)\mathbf{H}(\mathbb{A}) \twoheadrightarrow \mathbf{G}(\mathbb{A})/\mathbf{G}(\mathbb{A})^+.$$

Now, defining $\mathbf{G}(F_v)^+$ analogously to its adelic counterpart, split into cases as follows:

- $\mathbf{G} = \mathrm{SO}_n$: in this case, the spinor norm injects $\mathbf{G}(F_v)/\mathbf{G}(F_v)^+$ into F_v^\times modulo squares, and our assertion follows from the surjectivity of the spinor norm on SO_{n-1} (for $n-1 \geq 2$).
- $\mathbf{G} = \mathrm{U}_n$: in this case the determinant injects $\mathbf{G}(F_v)/\mathbf{G}(F_v)^+$ into the norm-1 elements of $(E \otimes F_v)^\times$, where E is the quadratic extension defining the unitary group. Again this determinant is surjective on U_{n-1} , whence the conclusion. \square

27.3. Application of Ratner's theorem

We now invoke the full force of our assumptions (see §25.7). Recall, thus, that

- (i) \mathbf{G} is a reductive group over a number field F with adèle ring \mathbb{A} ,
- (ii) \mathbf{G} is anisotropic, so that the quotient $[\mathbf{G}] := \mathbf{G}(F) \backslash \mathbf{G}(\mathbb{A})$ is compact, and
- (iii) \mathfrak{q} is an archimedean place for which the local component $\mathbf{G}_\mathfrak{q} := \mathbf{G} \times_F F_\mathfrak{q}$ is a quasi-split reductive group over the archimedean local field $F_\mathfrak{q}$. We write

$$G := \mathbf{G}(F_\mathfrak{q}) = \mathbf{G}_\mathfrak{q}(F_\mathfrak{q}),$$

as usual.

To simplify terminology in what follows, we regard G as a *real* reductive group; thus, if F_q is complex rather than real, we regard G as the real points of the real algebraic group $\text{Res}_{F_q/\mathbb{R}}(\mathbf{G} \times_{\mathbb{Q}} \mathbb{R})$.

Let $\xi \in \mathcal{O} \subseteq \mathcal{N}_{\text{reg}} \subseteq \mathfrak{g}^\wedge$. As noted above, our goal reduces to showing (for each such ξ) that the measure $[\delta_\xi]$ on $[G]$ is G_∞^+ -invariant. The element ξ is *regular nilpotent*: its G -orbit is an open subset of the nilcone. Such elements may be characterized (with respect to any G -equivariant isomorphism $\mathfrak{g} \cong \mathfrak{g}^\wedge$) by the following result of Kostant [Kos1, Theorem 5.3].

LEMMA 27.5. *Let $x \in \mathfrak{g}$ be nilpotent. Let $\mathfrak{b} = \mathfrak{t} \oplus \mathfrak{n}$ be a Borel subalgebra whose unipotent radical \mathfrak{n} contains x . Then x is regular nilpotent if and only if its component with respect to each simple root is non-zero.*

For instance, if G is a general linear group, then the regular nilpotent elements are the conjugates of the nilpotent Jordan blocks.

Set

- (iv) \mathfrak{u} := centralizer in $\mathfrak{g} := \text{Lie}(G)$ of the regular nilpotent element ξ ,
- (v) $U \leq G$ the connected Lie subgroup generated by $\exp(\mathfrak{u})$, and
- (vi) \mathbf{Z} := center of \mathbf{G} .

Example 27.6. Suppose $G = \text{GL}_3(\mathbb{R})$. Then, we may take ξ to correspond under the trace pairing to an upper-triangular Jordan block, in which case

$$U = \left\{ \begin{pmatrix} a & b & c \\ & a & b \\ & & a \end{pmatrix} : (a, b, c) \in \mathbb{R}_+^\times \times \mathbb{R} \times \mathbb{R} \right\}. \tag{27.3}$$

Under the bijection $[G] \cong [\mathbf{G}]/K$ following from strong approximation and our assumptions on the set of places R , we may identify $[\delta_\xi]$ with μ , where

- (vii) μ is a $(\mathbf{Z}(\mathbb{A}) \cdot U)$ -invariant probability measure on $[\mathbf{G}]$.

In what follows, we write G^+ and G_∞^+ for the topologically connected components of G and G_∞ . Our task reduces to establishing the following.

THEOREM 27.7. *Let notation and assumptions be as in (i)–(vii). Then μ is G_∞^+ -invariant.*

The proof occupies the remainder of this section.

Using the map $[\mathbf{G}]/\mathbf{Z}(\mathbb{A}) \rightarrow [\mathbf{G}/\mathbf{Z}]$, we may replace \mathbf{G} by \mathbf{G}/\mathbf{Z} to reduce to the case that \mathbf{G} is semisimple.

It suffices to establish the G_∞^+ -invariance of μ after push-forward to the quotient of $[\mathbf{G}]$ by an open compact subgroup of $\mathbf{G}(\mathbb{A}_f)$. Any such quotient is a finite union of

spaces of the form $\Gamma \backslash G_\infty$, where

- G_∞ is the Lie group given by

$$G_\infty := \prod_{\mathfrak{p}|\infty} G_{\mathfrak{p}} = G \times G_T, \quad (27.4)$$

where \mathfrak{p} runs over the archimedean places of F and T denotes the set of archimedean places other than \mathfrak{q} . In other words, $G_\infty = \mathbf{G}'(\mathbb{R}) = \mathbf{G}(F_\infty)$, with $\mathbf{G}' := \text{Res}_{F/\mathbb{Q}} \mathbf{G}$. Note that we may regard U as a subgroup of G_∞ .

- $\Gamma \leq G_\infty$ is an arithmetic lattice. (The notation Γ has been used differently in other sections, but this notational overload should introduce no confusion.)

It will thus suffice to show that any U -invariant probability measure ν on $\Gamma \backslash G_\infty$ is in fact G_∞^+ -invariant. In verifying this, we may apply ergodic decomposition (see [EW, §8.7]) to reduce to the case that ν is ergodic. Ratner has proven [Rat] that any ergodic U -invariant ν is the S -invariant measure on a closed S -orbit xS , for a closed connected subgroup $S \leq G_\infty$, with $S \supseteq U$. We will show eventually that $S = G_\infty^+$.

The basic idea of the proof is as follows: the compactness assumption on $[\mathbf{G}]$ will be seen to imply that S contains no non-trivial normal unipotent subgroups, but any connected subgroup S with this property that contains U must also contain G^+ . It may be instructive to consider the example (27.3).

Translating ν by x^{-1} , and replacing U by xUx^{-1} , we may suppose that $x=e$. Then, $\Gamma_S := S \cap \Gamma$ is a lattice inside S . Define the \mathbb{Q} -algebraic group

$$\mathbf{S} := \text{Zariski-closure of } \Gamma_S \text{ inside } \mathbf{G}'.$$

Recall our convention (§1.14) that “reductive” is short for “connected reductive algebraic”.

LEMMA 27.8. \mathbf{S} is reductive, and $\mathbf{S}(\mathbb{R}) \supseteq U$.

Proof. We show first that $\mathbf{S}(\mathbb{R})$ contains S . To see this, recall that the ergodicity of the U -action on $\Gamma_S \backslash S$, with respect to ν , implies that almost every orbit is equidistributed, and in particular dense. We may thus find $s \in S$ such that $\Gamma_S sU$ is dense in S . Then,

$$\Gamma_S (sUs^{-1}) \text{ is dense in } S. \quad (27.5)$$

On the other hand, Borel’s density theorem [Mo, Proposition 4.7.1] implies that $\mathbf{S}(\mathbb{R})$ contains all unipotent elements of S ; in particular, $\mathbf{S}(\mathbb{R}) \supseteq sUs^{-1}$. By (27.5), we deduce that $\mathbf{S}(\mathbb{R})$ contains a dense set of elements of S and hence S itself, as required.

Since \mathbf{S} contains the topologically connected set S as a Zariski-dense subset, it follows in particular that \mathbf{S} is connected.

The unipotent radical of \mathbf{S} is trivial. Otherwise $\mathbf{S}(\mathbb{Q})$ would contain non-trivial unipotent elements, as would $\mathbf{G}'(\mathbb{Q})=\mathbf{G}(F)$, contradicting the standard compactness criterion for the compact quotient $[\mathbf{G}]$. \square

LEMMA 27.9. $\mathbf{S}(\mathbb{R}) \geq G$.

Proof. The splitting (27.4) comes from a splitting of real algebraic groups

$$\mathbf{G}' \times_{\mathbb{Q}} \mathbb{R} = \mathbf{G}'_{\mathfrak{q}} \times \mathbf{G}'_T, \tag{27.6}$$

where $\mathbf{G}'_{\mathfrak{q}} = \text{Res}_{F_{\mathfrak{q}}/\mathbb{R}}(\mathbf{G}_{\mathfrak{q}})$; in particular, $\mathbf{G}'_{\mathfrak{q}}$ is quasi-split. Let $\mathbf{S}_{\mathfrak{q}}$ denote the kernel of $\mathbf{S} \rightarrow \mathbf{G}'_T$; it defines a reductive subgroup of $\mathbf{G}'_{\mathfrak{q}}$. It will suffice to verify that $\mathbf{S}_{\mathfrak{q}} = \mathbf{G}'_{\mathfrak{q}}$. We prove this below, in Lemma 27.12, in its natural generality. \square

LEMMA 27.10. $S \geq G^+$.

Proof. Let $\mathfrak{s} \leq \text{Lie}(G_{\infty})$ denote the Lie algebra of the Lie group $S \leq G_{\infty}$. We may regard \mathfrak{g} as a subalgebra of $\text{Lie}(G_{\infty})$. We must show that \mathfrak{s} contains \mathfrak{g} .

Observe first that \mathfrak{s} is invariant under the adjoint action of S , hence that of Γ_S . Since \mathbf{S} is the Zariski closure of Γ_S , it follows that \mathfrak{s} is invariant under $\mathbf{S}(\mathbb{R})$. By Lemma 27.9, we deduce that \mathfrak{s} is invariant by $\mathbf{G}'_{\mathfrak{q}}(\mathbb{R})=G$.

Since \mathfrak{s} is G -invariant and contains \mathfrak{u} , we deduce that $\mathfrak{h} := \mathfrak{s} \cap \mathfrak{g}$ is a normal Lie subalgebra of \mathfrak{g} , containing \mathfrak{u} . Splitting the semisimple real Lie algebra \mathfrak{g} as a sum $\bigoplus_{i \in I} \mathfrak{g}_i$ of quasi-split simple real Lie algebras \mathfrak{g}_i , we must have by normality $\mathfrak{h} = \bigoplus_{j \in J \subset I} \mathfrak{g}_j$; but the projection of \mathfrak{u} to each factor \mathfrak{g}_i is non-trivial, so in order that \mathfrak{h} contains \mathfrak{u} , we must have $J=I$, i.e., $\mathfrak{h}=\mathfrak{g}$. This implies $\mathfrak{s} \leq \mathfrak{g}$ as desired. \square

LEMMA 27.11. *Any orbit of G^+ on any connected component of $\Gamma \backslash G_{\infty}$ is dense.*

Proof. We apply strong approximation. Recall that we have reduced to the case that \mathbf{G} is semisimple. Since G^+ is normal in G_{∞} , we reduce further to verifying that the G^+ orbit of the identity coset in $\Gamma \backslash G_{\infty}$ is dense in its connected component.

Let \mathbf{G}^{sc} be the simply connected covering group of \mathbf{G} . By definition, \mathbf{G}^{sc} is an F -algebraic group. Every F -simple factor of \mathbf{G}^{sc} must have non-compact $F_{\mathfrak{q}}$ -points, because $\mathbf{G}(F_{\mathfrak{q}})$ is quasi-split. By strong approximation [PR, Theorem 7.12], the orbits of $\mathbf{G}^{\text{sc}}(F_{\mathfrak{q}})$ on $[\mathbf{G}^{\text{sc}}]$ are dense; in particular, the closure of each $\mathbf{G}^{\text{sc}}(F_{\mathfrak{q}})$ -orbit is actually stable under $\mathbf{G}^{\text{sc}}(F_{\infty})$. By a theorem of Cartan, the groups $\mathbf{G}^{\text{sc}}(F_{\mathfrak{q}})$ and $\mathbf{G}^{\text{sc}}(F_{\infty})$ are topologically connected. Their images in $\mathbf{G}(\mathbb{A})$ are thus the connected components G^+ and G_{∞}^+ , and the desired conclusion readily follows. \square

By Lemmas 27.10 and 27.11, we see that the closed orbit $\Gamma_S \backslash S$ coincides with a component of $\Gamma \backslash G_{\infty}$, and hence that $S = G_{\infty}^+$. This completes the proof of Theorem 27.7

modulo the following lemma, postponed above, which we have found convenient to formulate using notation independent from that in the rest of this section.

LEMMA 27.12. *Let \mathbf{G} be a quasi-split real reductive group, and let $\mathbf{S} \subset \mathbf{G}$ be a real reductive algebraic subgroup. (By our usual conventions, \mathbf{G} and \mathbf{S} are connected.) Let $\mathfrak{g} = \mathrm{Lie}(\mathbf{G}(F))$ and $\mathfrak{s} := \mathrm{Lie}(\mathbf{S}(F))$ denote the corresponding Lie algebras. Let ξ be a regular nilpotent element of \mathfrak{g}^* , with centralizer $\mathfrak{g}_\xi \leq \mathfrak{g}$. Suppose that $\mathfrak{s} \supseteq \mathfrak{g}_\xi$. Then $\mathbf{S} = \mathbf{G}$.*

Proof. We may without loss replace \mathbf{G} , \mathbf{S} , \mathfrak{g} and \mathfrak{s} by their complexifications. Since \mathbf{S} is connected, it is enough to show that $\mathfrak{s} = \mathfrak{g}$.

We may identify the coadjoint action on \mathfrak{g}^* with the adjoint action on \mathfrak{g} , via the trace pairing with respect to a faithful linear representation. Thus $\xi \in \mathfrak{g}$. Our hypothesis implies also that $\xi \in \mathfrak{s}$, so it makes sense to speak of the regularity of ξ both with respect to \mathbf{S} and \mathbf{G} .

The regularity of ξ is equivalent (by Springer, Kurtzke, see [Kur]) to its Lie algebra centralizer being abelian. Since ξ is regular for \mathbf{G} , it follows that it is likewise regular for \mathbf{S} . Thus,

$$\mathrm{rank}(\mathbf{S}) = \dim(\mathfrak{g}_\xi) = \mathrm{rank}(\mathbf{G}),$$

that is to say, \mathbf{S} is an equal-rank subgroup of \mathbf{G} .

Since $\xi \in \mathfrak{s}$ is nilpotent, we may find

- a maximal torus $\mathbf{T} \leq \mathbf{S}$, with Lie algebra $\mathfrak{t} \subseteq \mathfrak{s}$, and
- a Borel subalgebra $\mathfrak{t} \oplus \mathfrak{n}$ of \mathfrak{s}

such that $\xi \in \mathfrak{n}$. Since \mathbf{S} and \mathbf{G} have equal rank, \mathbf{T} is likewise a maximal torus for \mathbf{G} . Let Φ denote the set of roots for \mathbf{T} acting on \mathfrak{g} , so that

$$\mathfrak{g} = \mathfrak{t} \oplus \bigoplus_{\alpha \in \Phi} \mathfrak{g}_\alpha,$$

with each \mathfrak{g}_α 1-dimensional. Since \mathfrak{s} is a \mathfrak{t} -stable subspace of \mathfrak{g} that contains \mathfrak{t} , we have

$$\mathfrak{s} = \mathfrak{t} \oplus \bigoplus_{\alpha \in \Phi'} \mathfrak{g}_\alpha$$

for some $\Phi' \subseteq \Phi$. We must show that $\Phi' = \Phi$.

We may find a positive system $\Phi^+ \subseteq \Phi$ for which $\mathfrak{n} \subseteq \sum_{\alpha \in \Phi^+} \mathfrak{g}_\alpha$ (e.g., by considering a generic 1-parameter subgroup of \mathbf{T} having positive eigenvalues on \mathfrak{n}).

In particular, $\xi = \sum_{\alpha \in \Phi^+} \xi_\alpha$, with $\xi_\alpha \in \mathfrak{g}_\alpha$. Let $\Delta \subseteq \Phi^+$ denote the corresponding simple system. Since ξ is regular nilpotent for \mathfrak{g} , we have by Lemma 27.5 that $\xi_\alpha \neq 0$ for all $\alpha \in \Delta$. Since $\xi \in \mathfrak{s}$, we deduce that Φ' contains Δ . Since \mathbf{S} is reductive, we conclude that $\Phi' = \Phi$. \square

28. Recap and overview of the proof

We pause to recall (from §25) some aspects of the basic setup, as well as what we have shown thus far. We will then discuss how the remainder of the proof proceeds from this point.

28.1. Recap

We are considering an inclusion of compact quotients

$$[H] = \Gamma_H \backslash (H \times H') \hookrightarrow [G] = \Gamma \backslash (G \times G')$$

arising from a Gross–Prasad pair; $H < G$ and $H' < G'$ are inclusions of real reductive and S -arithmetic groups, respectively. From the fixed automorphic representation Π of \mathbf{G} we obtained a $G \times G'$ -subrepresentation $\pi \otimes \pi' \hookrightarrow L^2([G])$. From a fixed positive-definite tensor $T' \in \pi' \otimes \bar{\pi}'$ of trace 1, we constructed in §26 an assignment from Schwartz functions $a \in \mathcal{S}(\mathfrak{g}^\wedge)$ to “limit states” $[a] \in C^\infty([G])$, describing the limiting behavior of L^2 -masses of certain families of vectors in $\pi \otimes \pi'$.

Each limit state $[a]$, for (say) real-valued a , is (informally) a limiting weighted average of L^2 -masses of vectors $v \otimes v' \in \pi \otimes \pi'$ for which v is microlocalized within the support of $a|_{\mathcal{O}}$. Thus, $\int_{[H]} [a]$ is an average of integrals $\int_{[H]} |v \otimes v'|^2$, each of which may be decomposed spectrally as a sum of contributions $\mathcal{P}_\Sigma(v \otimes v')$ from automorphic representations Σ as above; more generally and precisely, we denote by

$$\mathcal{P}_\Sigma: \Psi^{-\infty} \longrightarrow \mathbb{C} \tag{28.1}$$

the (uniformly in Σ) continuous maps obtained by the following composition:

$$\Psi^{-\infty} \xrightarrow{T \mapsto T \otimes T'} C^\infty([G]^2) \xrightarrow{\text{restrict}} C^\infty([H]^2) \xrightarrow{\text{project}} (\Sigma_R \hat{\otimes} \Sigma_R^\vee)^\infty \xrightarrow{\int_{[H]}} \mathbb{C}. \tag{28.2}$$

We retain the convention from §26.3 of dropping T' from the notation; when we wish to indicate it explicitly, we write $\mathcal{P}_\Sigma(T \otimes T')$. For example, if $T = v \otimes \bar{v}$ and $T' = w \otimes \bar{w}$, so that $u := v \otimes w \in \pi \otimes \pi' \hookrightarrow C^\infty([G])$, then

$$\mathcal{P}_\Sigma(T) = \|\text{projection to } \Sigma \text{ of the restriction } u|_{[H]}\|^2.$$

We will verify below that, for any $T \in \Psi^{-\infty}$,

$$\int_{[H]} [T] = \sum_{\Sigma} \mathcal{P}_\Sigma(T), \tag{28.3}$$

where we sum over Σ as in §25.3. In particular, exchanging the limit $[a] = \lim_{\hbar \rightarrow 0} [a]_{\hbar}$ taken in $C^{\infty}([G])$ with integration over the compact set $[H]$ gives

$$\int_{[H]} [a] = \lim_{\hbar \rightarrow 0} \hbar^d \sum_{\Sigma} \mathcal{P}_{\Sigma}(\text{Op}_{\hbar}(a)). \quad (28.4)$$

The period formula (28.5) says that, under assumptions on Π and Σ to be recalled below,

$$\mathcal{P}_{\Sigma}(T) = \mathcal{L}(\Pi, \Sigma) \mathcal{H}_{\sigma}(T) \mathcal{H}_{\sigma'}(T'), \quad (28.5)$$

where the maps $\mathcal{H}_{\sigma}: \Psi^{-\infty} \rightarrow \mathbb{C}$ and $\mathcal{H}_{\sigma'}: \pi' \otimes \bar{\pi}' \rightarrow \mathbb{C}$ are as in §18. To summarize, then

$$\lim_{\hbar \rightarrow 0} \hbar^d \sum_{\Sigma: (28.5) \text{ holds}} \mathcal{L}(\Pi, \Sigma) \mathcal{H}_{\sigma}(\text{Op}_{\hbar}(a)) \mathcal{H}_{\sigma'}(T') + \sum_{\text{remaining } \Sigma} = \int_{[H]} [a]. \quad (28.6)$$

Moreover, we showed already in §27—using Ratner’s theorem—that the limit state construction $a \mapsto [a]$ is *essentially trivial* in the examples of interest: modulo “connectedness issues”, it produces constant functions on $[G]$, and the integral $\int_{[H]} a$ is proportional to the trace of $T \otimes T'$. Because of this, (28.6) gives an asymptotic of the desired nature, but it still requires some cleanup.

28.2. Cleanup

We now outline what is required to massage (28.6) into the required shape (cf. §1.6, §1.10). The most important issue is, of course, to choose T and T' so that $\sigma \mapsto \mathcal{H}_{\sigma}(T)$ and $T' \mapsto \mathcal{H}_{\sigma'}(T')$ approximate desired test functions on the unitary duals of H and H' . This problem has already been solved, at least enough for our purposes, in Part IV of this paper addressing inverse branching. We focus on the other issues that arise (although we will recall a part of this analysis below):

The second summand of (28.6) arises from situations where $\mathcal{L}(\Pi, \Sigma)$ is undefined, in particular:

- (i) either σ or σ' is non-tempered;
- (ii) $\mathcal{H}_{\sigma} = 0$ (even though (π, σ) is distinguished).

The first of these represents an actual possibility that does occur in practice (e.g., when Σ is the trivial representation!); the second is not expected to occur, but has not yet been ruled out in the literature. Since we aim in this paper to prove an unconditional theorem, we must show that these bad cases yield negligible contributions to (28.4).

The machinery to handle (i) was already set up in the course of our analysis of inverse branching. Informally speaking, we may construct \tilde{T} and \tilde{T}' with the property

that $\mathcal{P}_\Sigma(\tilde{T} \otimes \tilde{T}')$ majorizes $\mathcal{P}_\Sigma(T \otimes T')$ on the non-tempered spectrum, but such that the trace of $\tilde{T} \otimes \tilde{T}'$ is small. In the archimedean case, for example, this was proved in Theorem 22.2 by cutting off the symbol a to a small neighborhood of the locus of irregular infinitesimal characters. In any case, once \tilde{T} and \tilde{T}' have been constructed, we just apply (28.4) to control the total contribution of the non-tempered spectrum.

Part (ii) is dealt with by means of symbol calculus: such (π, σ) cannot be orbit-distinguished (in the sense of §22.2). Now orbit-distinction is the semiclassical analogue of distinction, and if (π, σ) fails to be orbit-distinguished, then microlocal analysis shows that any H -invariant functional on π must at least be negligibly small (in a suitable sense). To formalize this intuition with symbol calculus, we construct a suitable symbol on \mathfrak{h}^\wedge with the property that it misses entirely the orbit for σ , but on the other hand is identically 1 on the support of a (or rather its projection under $\mathfrak{g}^\wedge \rightarrow \mathfrak{h}^\wedge$); see Theorem 29.1 for further details.

There remain other contributions to clean up, e.g., from when

(iii) $h\lambda_\sigma$ is very large, or

(iv) the pair $(\lambda_\pi, \lambda_\sigma)$ is not \mathbf{H} -stable (cf. §14), in which case the asymptotic formulas in §19 for \mathcal{H}_σ do not apply.

We do not discuss these in detail here, but just observe that much of the difficulty of (iv) is avoided by fiat: we consider only those test functions on the unitary dual of H which are supported above the stable locus.

As should be clear from this outline, all of this analysis makes heavy use of the microlocal calculus from Parts I and II.

29. Spectral expansion and truncation of the H -period

29.1. Spectral decomposition

Here we verify the Parseval-type identity (28.3). More generally, let f be a smooth function on $[H] \times [H]$. Then,

$$f = \sum_{\Sigma_1, \Sigma_2} f_{\Sigma_1 \otimes \Sigma_2^\vee},$$

where $f_{\Sigma_1 \otimes \Sigma_2^\vee}$ belongs to the $(J \times J)$ -fixed subspace of the smooth completion of $\Sigma_1 \otimes \Sigma_2^\vee$, and the sum converges in the $C^\infty([H] \times [H])$ -topology, hence commutes with integration over the diagonal copy of $[H]$ inside $[H] \times [H]$:

$$\int_{[H]} f = \sum_{\Sigma} \int_{[H]} f_{\Sigma_1 \otimes \Sigma_2^\vee}. \tag{29.1}$$

Taking for f the restriction of $T \otimes T'$ to $[H] \times [H]$, one obtains $f|_{[H]} = [T]|_{[H]}$ and

$$\int_{[H]} f_{\Sigma \otimes \Sigma'} = \mathcal{P}_{\Sigma}(T),$$

so (29.1) specializes to (28.3).

29.2. Weyl law upper bound

The map \mathcal{P}_{Σ} is identically zero unless σ' is T' -distinguished in the sense that there is a non-zero equivariant map $\pi \otimes \bar{\pi}' \rightarrow \sigma \otimes \bar{\sigma}'$ and it does not vanish on T' . By our assumption that $G_{\mathfrak{p}}$ is compact for all archimedean $\mathfrak{p} \neq \mathfrak{q}$, the latter condition forces σ' to belong to some compact subset of the unitary dual of H' depending only upon T' . A weak form of the Weyl law then reads: for $x \geq 1$,

$$\#\{\Sigma : \mathcal{P}_{\Sigma} \neq 0 \text{ and } |\lambda_{\sigma}| \leq x\} \ll x^{O(1)}. \quad (29.2)$$

This follows from the *usual* Weyl law on $[H]$, using §9.8.

29.3. Truncated spectral decomposition

The main result of this section is the following. As explained above, it allows us to discard all terms in the spectral expansion (28.4) with large eigenvalue, and also those microlocally separated from the symbol a .

Let U be an open subset of $[\mathfrak{h}^{\wedge}]$ and $W \subseteq \mathfrak{h}^{\wedge}$ be an open subset of the preimage of U . We say that σ is *bad* (relative to the scaling parameter h and the choice of U and W) if either

- $h\lambda_{\sigma} \notin U$, or
- σ is tempered and $h\mathcal{O}_{\sigma} \cap W = \emptyset$.

We say otherwise that σ is *good*.

THEOREM 29.1. *Let $a \in C_c^{\infty}(\mathfrak{g}^{\wedge})$ be supported in the preimage of W . Then,*

$$\int_{[H]} [a] = \lim_{h \rightarrow 0} h^d \sum_{\Sigma : \sigma \text{ is good}} \mathcal{P}_{\Sigma}(\text{Op}_h(a)). \quad (29.3)$$

More generally, for a_1, \dots, a_k satisfying the same assumptions as a ,

$$\int_{[H]} [a_1 \dots a_k] = \lim_{h \rightarrow 0} h^d \sum_{\Sigma : \sigma \text{ is good}} \mathcal{P}_{\Sigma}(\text{Op}_h(a_1) \dots \text{Op}_h(a_k)). \quad (29.4)$$

Proof. Note that \mathcal{P}_Σ factors as an $(H \times H)$ -equivariant composition of \mathfrak{h} -uniformly continuous maps

$$\Psi^{-\infty} = \Psi^{-\infty}(\pi) \longrightarrow \Psi^{-\infty}(\sigma) \xrightarrow{\text{tr}_\sigma} \mathbf{C}. \quad (29.5)$$

We may choose $\varepsilon > 0$, depending only upon the support of a , such that

$$\xi \in \text{supp}(a) \text{ and } \mathfrak{h}\lambda_\sigma \notin U \implies \text{dist}([\xi|_{\mathfrak{h}}], \mathfrak{h}\lambda_\sigma) \geq \varepsilon.$$

By (10.9), it follows that, for each fixed $N \in \mathbb{Z}_{\geq 0}$,

$$\mathfrak{h}\lambda_\sigma \notin U \implies \mathcal{P}_\Sigma(\text{Op}_{\mathfrak{h}}(a)) \ll \mathfrak{h}^N \langle \mathfrak{h}\lambda_\sigma \rangle^{-N}. \quad (29.6)$$

From this and (29.2), we see that the contribution to (28.4) from those Σ with $\mathfrak{h}\lambda_\sigma \notin U$ is negligible.

It remains to estimate the contribution from when σ is tempered and $\mathfrak{h}\lambda_\sigma \in U$ but $\mathfrak{h}\mathcal{O}_\sigma \cap W = \emptyset$. By (29.2), the number of such Σ is $\mathfrak{h}^{-O(1)}$, so it will suffice to show for each such Σ that $\mathcal{P}_\Sigma(\text{Op}_{\mathfrak{h}}(a)) = O(\mathfrak{h}^\infty)$.

The image in \mathfrak{h}^\wedge of the support of a is a compact subset of W . By the smooth version of Urysohn's lemma, we may thus choose a real-valued $b \in C_c^\infty(W)$ such that $b \equiv 1$ on the image of the support of a . By the composition formula (8.36), we then have

$$\text{Op}_{\mathfrak{h}}(a) \equiv \text{Op}_{\mathfrak{h}}(b : \pi) \text{Op}_{\mathfrak{h}}(a) \pmod{\mathfrak{h}^\infty \Psi^{-\infty}}.$$

By the continuity of (29.5), it follows that

$$\mathcal{P}_\Sigma(\text{Op}_{\mathfrak{h}}(a)) = \mathcal{P}_\Sigma(\text{Op}_{\mathfrak{h}}(b : \pi) \text{Op}_{\mathfrak{h}}(a)) + O(\mathfrak{h}^\infty).$$

By the equivariance of (29.5), we have

$$\mathcal{P}_\Sigma(\text{Op}_{\mathfrak{h}}(b : \pi) \text{Op}_{\mathfrak{h}}(a)) = \text{tr}(T_1 T_2),$$

where $T_1 := \text{Op}_{\mathfrak{h}}(b : \sigma)$ and $T_2 \in \Psi^{-\infty}(\sigma)$ denotes the image of $\text{Op}_{\mathfrak{h}}(a)$. By Cauchy–Schwarz for the Hilbert–Schmidt inner product, we have

$$|\text{tr}(T_1 T_2)|^2 \leq \text{tr}(T_1^* T_1) \text{tr}(T_2^* T_2).$$

Using the composition formula, the Kirillov formula for σ and the assumption on the support of b , we see that $\text{tr}(T_1^* T_1) \ll \mathfrak{h}^\infty$. On the other hand, we obtain using Theorem 5.8 and (12.6) and the continuity of (29.5) that $\text{tr}(T_2^* T_2) \ll \mathfrak{h}^{-O(1)}$.

This completes the proof of (29.3).

For (29.4), note first by (26.5) and (28.4) that

$$\int_{[H]} [a_1 \dots a_k] = \lim_{h \rightarrow 0} h^d \sum_{\Sigma} \mathcal{P}_{\Sigma}(\text{Op}_h(a_1) \dots \text{Op}_h(a_k)).$$

We fix $N \geq 0$ large enough and apply the composition formula (5.5) to write

$$\text{Op}_h(a_1) \dots \text{Op}_h(a_k) \equiv \sum_{0 \leq j < J} h^j \text{Op}_h(b_j) \pmod{h^N \Psi^{-N}},$$

where J is large but fixed and the b_j satisfy the same assumptions as the a_i . We apply (10.9) as before to the contribution from the $\text{Op}_h(b_j)$. We clean up the remainder contribution using (10.7). □

30. The smoothly weighted asymptotic formula

30.1. Overview

We retain the notation and setup of §25. This section contains the main automorphic result in this paper: an asymptotic formula for the averaged GGP branching coefficient $\mathcal{L}(\Pi, \Sigma)$, with the automorphic representation Π of \mathbf{G} fixed and the automorphic representation Σ of \mathbf{H} traversing a smoothly-weighted family. (We will refine this in minor ways in §31, by extending to unweighted families and then dividing through by their cardinalities.) The proof involves three main inputs developed hitherto:

- The “inverse branching” results of Part IV (especially §22.3 and §24.2), which allow us to pick off any reasonable family of representations Σ , say defined by weights $w(\Sigma)$, using a suitable family of vectors $v \in \Pi$:

$$\int_{[H]} |v|^2 \approx \sum_{\Sigma} \mathcal{L}(\Pi, \Sigma) w(\Sigma) \quad \text{on average over } v. \tag{30.1}$$

- The “truncated spectral formula” in §29, which allows us to discard some of the contributions of undesirable automorphic forms Σ implicit in (30.1).
- The “equidistribution” result, proved in §27 using Ratner’s theorem, by which we deduce that (on average)

$$\frac{1}{\text{vol}([H])} \int_{[H]} |v|^2 \approx \frac{1}{\text{vol}([G])} \int_{[G]} |v|^2,$$

leading to the required asymptotic formula for

$$\sum_{\Sigma} \mathcal{L}(\Pi, \Sigma) w(\Sigma).$$

30.2. Function spaces

We now define the precise sets of weights $w(\Sigma)$ to be summed against.

30.2.1. Spaces of representations of H

We may apply the notation of §20 to each place \mathfrak{p} of F , thus the sets

$$\widehat{H}_{\mathfrak{p}} \supseteq (\widehat{H}_{\mathfrak{p}})_{\text{temp}} \supseteq (\widehat{H}_{\mathfrak{p}})_{\text{temp}}^{\pi_{\mathfrak{p}}}$$

are respectively the unitary dual of $\widehat{H}_{\mathfrak{p}}$, the tempered dual and its $\pi_{\mathfrak{p}}$ -distinguished subset. We omit the index when $\mathfrak{p}=\mathfrak{q}$, and use a superscripted prime to denote a product over all $\mathfrak{p} \in R \setminus \{\mathfrak{q}\}$.

30.2.2. The distinguished archimedean place

Recall (§9.2) that the geometric quotient $[\mathfrak{h}^{\wedge}]$ of $\mathfrak{h}^{\wedge} = i\mathfrak{h}^*$ is isomorphic to an affine space. We denote as in §22 by $\mathcal{O}_{\text{stab}} \subseteq \mathcal{O} \subseteq \mathfrak{g}^{\wedge}$ the subset of \mathbf{H} -stable elements of the limit coadjoint orbit \mathcal{O} of π , and set

$$\mathcal{K} := C_c^{\infty}([\mathfrak{h}^{\wedge}] \cap \text{image}(\mathcal{O}_{\text{stab}})).$$

In view of the stability characterization given in §14.2, an element of \mathcal{K} is just a smooth compactly-supported function k on the space

$$\{\mu \in [\mathfrak{h}^{\wedge}] : \text{no eigenvalue of } \mu \text{ equals zero}\}.$$

(We used the fact that \mathcal{O} is contained in the nilcone.)

As in §22.3, we assign to each fixed $k \in \mathcal{K}$ and all sufficiently small $h > 0$ a function

$$k_h: \widehat{H} \longrightarrow \mathbb{C},$$

as follows:

- if σ is tempered and $\mathcal{O}_{\pi, \sigma} \neq \emptyset$ (i.e., σ is “orbit-distinguished” by π ; see §22.2), then we evaluate on the rescaled infinitesimal character:

$$k_h(\sigma) := k(h\lambda_{\sigma});$$

- otherwise, we set $k_h(\sigma) := 0$.

30.2.3. The auxiliary archimedean places

Let \mathfrak{p} be an archimedean place other than \mathfrak{q} . Since the groups $(G_{\mathfrak{p}}, H_{\mathfrak{p}})$ are assumed compact, the set $(\widehat{H}_{\mathfrak{p}})_{\text{temp}}^{\pi_{\mathfrak{p}}}$ is finite. We denote by $\mathcal{K}_{\mathfrak{p}}$ the set of complex-valued functions $k_{\mathfrak{p}}: (\widehat{H}_{\mathfrak{p}})_{\text{temp}}^{\pi_{\mathfrak{p}}} \rightarrow \mathbb{C}$. (Compare with §21.)

30.2.4. The auxiliary p -adic places

We now let \mathfrak{p} be a finite place in R . We denote by $\mathcal{K}_{\mathfrak{p}}$ the space of allowable functions $k_{\mathfrak{p}}: (\widehat{H}_{\mathfrak{p}})_{\text{temp}}^{\pi_{\mathfrak{p}}} \rightarrow \mathbb{C}$ (see §24.2 for the definition of “allowable”, and recall that Theorem 24.5, stated in §24.3, provides a large supply of allowable functions).

30.2.5. The auxiliary places, grouped together

We denote by \mathcal{K}' the space of functions

$$k': (\widehat{H}')_{\text{temp}}^{\pi'} \rightarrow \mathbb{C}$$

spanned by the pure tensors

$$k'(x) := \prod_{\mathfrak{p} \in R \setminus \{\mathfrak{q}\}} k_{\mathfrak{p}}(x_{\mathfrak{p}})$$

for $k_{\mathfrak{p}} \in \mathcal{K}_{\mathfrak{p}}$. We extend each such k' by zero to a function on the unitary dual \widehat{H}' of H' .

30.3. Main result

We denote by $d \in \mathbb{Z}_{\geq 0}$ half the real dimension of \mathcal{O} , as usual (see §31.4 for numerics). We fix $k \in \mathcal{K}$, $k' \in \mathcal{K}'$, and set

$$\ell := \mathfrak{h}^d \sum_{\substack{\sigma \text{ and } \sigma' \text{ tempered} \\ \mathcal{O}_{\pi, \sigma} \neq \emptyset}} \mathcal{L}(\Pi, \Sigma)_{k_{\mathfrak{h}}(\sigma)} k'(\sigma').$$

Here and henceforth we sum over automorphic representation Σ of \mathbf{H} , as in §25, whose components σ and σ' satisfy the displayed conditions, so that $\mathcal{L}(\Pi, \Sigma)$ is defined.

We write $\int k$ and $\int k'$ for integrals taken with respect to the normalized affine measure on $[\mathfrak{h}^{\wedge}]$ and the Plancherel measure on $\widehat{H}'_{\text{temp}}$, respectively. We write $A \simeq B$ for $A = B + o_{\mathfrak{h} \rightarrow 0}(1)$.

THEOREM 30.1.

$$\ell \simeq \frac{\tau(\mathbf{H})}{\tau(\mathbf{G})} \int k \int k'.$$

The proof occupies the remainder of this section. We may and shall assume that k and k' are *non-negative*.

It will be convenient to introduce the following otherwise unusual notation.

Definition 30.2. Let $p: X \rightarrow Y$ be a continuous map between topological spaces. Let $U \subseteq X$ and $V \subseteq Y$ be subsets. We write $U \prec V$ if $\overline{p(U)} \subseteq V^0$, and similarly $V \prec U$ if $\overline{V} \subseteq p(U)^0$. (Here \overline{V} and V^0 denote closure and interior.)

The precise choice of p (typically a “projection”) to which one should apply this notation should be clear by context in what follows. We caution that \prec is not in any sense “transitive”.

LEMMA 30.3. *There are precompact open subsets $U \subset [\mathfrak{h}^\wedge]$, $W \subseteq \mathfrak{h}^\wedge$ and $V \subset \mathfrak{g}^\wedge$ such that, for small enough $\mathfrak{h} > 0$,*

- (i) \overline{V} consists of \mathbf{H} -stable elements,
- (ii) $\text{supp}(k) \prec V \prec W \prec U \prec \mathcal{O}_{\text{stab}}$, and
- (iii) if $\sigma \in \widehat{H}_{\text{temp}}$ satisfies $\mathcal{O}_{\pi, \sigma} = \emptyset$, then $\mathfrak{h} \mathcal{O}_\sigma \cap W = \emptyset$.

Informally, this says that

- V is large enough to support symbols a suitable for approximating k via

$$k(\mu) \approx \int_{\mathcal{O}(\mu)} a,$$

but not much larger;

- W is large enough to majorize the image of the support of any such symbol a , but small enough to avoid any multiorbit \mathcal{O}_σ for which $\mathcal{O}_{\pi, \sigma} = \emptyset$;
- U is large enough to majorize all of the above, but not too large.

Proof. We start by choosing U as indicated with $\text{supp}(k) \prec U \prec \mathcal{O}_{\text{stab}}$. We then find a compact subset K of $\mathcal{O}_{\text{stab}}$ so that $\text{supp}(k) \subseteq \text{image}(K) \subset U$. Fix $\xi \in K$. We may find $\xi_{\mathfrak{h}} \in \mathfrak{h} \mathcal{O}_\pi$ tending to ξ . Since ξ is \mathbf{H} -stable, its \mathfrak{h} -stabilizer is trivial. It follows readily that the map $G \rightarrow G \cdot \xi \rightarrow \mathfrak{h}^\wedge$ has surjective differential $\mathfrak{g} \rightarrow T_\xi(G \cdot \xi) \rightarrow \mathfrak{h}^\wedge$. The same holds true for all ξ' in a small neighborhood of ξ ; in particular, for \mathfrak{h} small enough, it holds for $\xi_{\mathfrak{h}}$. We may thus find a small precompact open neighborhood $\widetilde{W}_\xi \subseteq \mathfrak{g}^\wedge$ of ξ whose image $W_\xi \subseteq \mathfrak{h}^\wedge$ satisfies $W_\xi \prec \mathfrak{h} \mathcal{O}_\pi$ for all small \mathfrak{h} . In particular, $W_\xi \cap \mathfrak{h} \mathcal{O}_\sigma = \emptyset$ whenever $\mathcal{O}_{\pi, \sigma} = \emptyset$. We may assume, moreover, having chosen \widetilde{W}_ξ small enough, that $W_\xi \prec U$.

Choose a small precompact open neighborhood V_ξ of ξ , with \overline{V}_ξ consisting of \mathbf{H} -stable elements, such that $V_\xi \prec W_\xi$. Since K is compact, we may find $\xi_1, \dots, \xi_n \in K$ such that $K \subseteq V := V_{\xi_1} \cup \dots \cup V_{\xi_n}$. We then take $W := W_{\xi_1} \cup \dots \cup W_{\xi_n}$. \square

We henceforth fix such U , W and V . We note that U and V satisfy the conditions enunciated in §22.3.

LEMMA 30.4. *Let $a_1, \dots, a_k \in C_c^\infty(\mathfrak{g}^\wedge)$. Let $T' \in \pi' \otimes \overline{\pi'}$ be positive definite, but not necessarily of trace 1. Then,*

$$h^d \sum_{\Sigma} \mathcal{P}_{\Sigma}(\mathrm{Op}_h(a_1) \dots \mathrm{Op}_h(a_k) \otimes T') \simeq \frac{\tau(\mathbf{H})}{\tau(\mathbf{G})} \cdot \left(\int_{\mathcal{O}} a_1 \dots a_k d\omega \right) \cdot \mathrm{tr}(T'). \quad (30.2)$$

If the a_i are supported in the preimage of W , then the same holds after restricting to Σ for which σ is good in the sense of §29.3.

Proof. We may normalize T' to have trace 1. We then construct “limit states” $[a_1 \dots a_k]$ as in §26. (This involves passing to subsequences of $\{h\}$, which we may do after having assumed for the sake of contradiction that the estimate fails for some infinite sequence of h tending to zero.) The left-hand side tends to $\int_{[H]} [a_1 \dots a_k]$, by (26.5) and (28.3). We then use the equidistribution statement (27.1), together with (26.6), to get to the right-hand side.

For the last statement we use Theorem 29.1. □

LEMMA 30.5. *For each $\varepsilon > 0$ and $N \in \mathbb{Z}_{\geq 0}$ there exist non-negative*

$$a, a_1, a_2, a_{\mathrm{nt}} \in C_c^\infty(V)$$

and (smooth, finite-rank) positive-definite tensors $T', T'_1, T'_2 \in \pi' \otimes \overline{\pi'}$ such that

$$\left| \int k \int k' - \left(\int_{\mathcal{O}} a^2 d\omega \right) \mathrm{tr}(T') \right| \leq \varepsilon, \quad (30.3)$$

$$\sum_{j=1}^2 \left(\int_{\mathcal{O}} a_j^2 d\omega \right) \mathrm{tr}(T'_j) \leq \varepsilon, \quad (30.4)$$

$$\left(\int_{\mathcal{O}} a_{\mathrm{nt}}^2 d\omega \right) \mathrm{tr}(T') \leq \varepsilon, \quad (30.5)$$

and, for σ tempered with $h\lambda_\sigma \in U$,

$$|k_h(\sigma)k'(\sigma') - \mathcal{H}_\sigma(\mathrm{Op}_h(a)^2)\mathcal{H}_{\sigma'}(T')| \leq \sum_{j=1}^2 |\mathcal{H}_\sigma(\mathrm{Op}_h(a_j)^2)|\mathcal{H}_{\sigma'}(T'_j) \quad (30.6)$$

and, for σ non-tempered with $h\lambda_\sigma \in U$,

$$\mathcal{P}_{\Sigma}(\mathrm{Op}_h(a)^2 \otimes T') = \mathcal{P}_{\Sigma}(\mathrm{Op}_h(a_{\mathrm{nt}})^2 \otimes T') + O(h^N). \quad (30.7)$$

Proof. We combine together the analogous approximation results obtained in §21, §22.3 and §24.2. We record details below.

In what follows, we write (e.g.) $C(x, y, z)$ for some constant ≥ 1 depending only upon the quantities x, y and z . We allow the precise definitions of such constants to vary from one invocation to the next. For instance, we have

$$\int k \leq C(k), \quad \int k' \leq C(k') \quad \text{and} \quad \int k_{\mathfrak{p}} \leq C(k_{\mathfrak{p}}).$$

We fix $\eta \in (0, 1)$; at the end of the proof, η will be chosen small enough in terms of ε and the intervening constants.

We may assume that k' is a pure tensor

$$\prod_{\mathfrak{p} \neq \mathfrak{q}} k_{\mathfrak{p}},$$

with each factor $k_{\mathfrak{p}}$ non-negative; here and henceforth \mathfrak{p} is restricted to the set R of relevant places. By §21 and §24.2 (see (24.3) and (24.5)), we may find for each auxiliary place $\mathfrak{p} \in R \setminus \{\mathfrak{q}\}$ some positive-definite tensors $T'_{\mathfrak{p}}, T_{\mathfrak{p}}^{(1)}$ and $T_{\mathfrak{p}}^{(2)}$, with

$$\mathrm{tr}(T_{\mathfrak{p}}^{(1)}) \leq C(k_{\mathfrak{p}}), \quad |k_{\mathfrak{p}}(\Sigma_{\mathfrak{p}})| \leq \mathcal{H}_{\Sigma_{\mathfrak{p}}}(T_{\mathfrak{p}}^{(1)})$$

and

$$\mathrm{tr}(T_{\mathfrak{p}}^{(2)}) \leq \eta, \quad |k_{\mathfrak{p}}(\Sigma_{\mathfrak{p}}) - \mathcal{H}_{\Sigma_{\mathfrak{p}}}(T'_{\mathfrak{p}})| \leq \mathcal{H}_{\Sigma_{\mathfrak{p}}}(T_{\mathfrak{p}}^{(2)}).$$

By (24.7), we may assume also that

$$\left| \int k_{\mathfrak{p}} - \mathrm{tr}(T'_{\mathfrak{p}}) \right| \leq \eta,$$

and hence, in particular, that $\mathrm{tr}(T'_{\mathfrak{p}}) \leq C(k_{\mathfrak{p}})$. We set

$$T' := \bigotimes_{\mathfrak{p} \neq \mathfrak{q}} T'_{\mathfrak{p}}.$$

By Theorem 22.2, we may find a, a_1, a_2 and a_{nt} of the required form for which

- $\int_{\mathcal{O}} a_1^2 \leq C(k, V)$,
- $\int_{\mathcal{O}} a_2^2, \int_{\mathcal{O}} a_{\mathrm{nt}}^2$ and $|\int k - \int_{\mathcal{O}} a^2|$ are bounded by η (thus $\int_{\mathcal{O}} a^2 \leq C(k)$),
- if σ is tempered and $\mathcal{O}^{\pi, \sigma}$ is non-empty, then

$$\begin{aligned} |k_{\mathrm{h}}(\sigma)| &\leq |\mathcal{H}_{\sigma}(\mathrm{OP}_{\mathrm{h}}(a_1)^2)|, \\ |k_{\mathrm{h}}(\sigma) - \mathcal{H}_{\sigma}(\mathrm{OP}_{\mathrm{h}}(a)^2)| &\leq |\mathcal{H}_{\sigma}(\mathrm{OP}_{\mathrm{h}}(a_2)^2)|, \end{aligned}$$

- if σ is non-tempered, then (30.7) holds (note that the Hermitian form $\mathcal{P}_{\Sigma}(- \otimes T')$ satisfies the hypotheses indicated around (22.8)).

For tempered σ with $h\lambda_\sigma \in U$, the left-hand side of (30.6) is bounded by

$$\sum_{j=(j_p)_{p \in R}}^{\#} \mathcal{H}_\sigma(\text{Op}_h(a_{j_q})^2) \prod_{p \neq q} \mathcal{H}_{\Sigma_p}(T_p^{(j_p)}), \quad (30.8)$$

where the sum is taken over all tuples j as indicated, for which

- $j_p \in \{1, 2\}$ for all $p \in R$, and
- $j_p = 2$ for at least one $p \in R$.

For $r=1, 2$, define

$$T'_r := \sum_{j=(j_p)_{p \in R}: j_q=r}^{\#} \bigotimes_{p \neq q} T_p^{(j_p)}.$$

Then, (30.8) equals the right-hand side of (30.6), so that (30.6) holds.

It remains to verify the estimates (30.3)–(30.5). Combining together the estimates noted above, we obtain the inequalities

$$\begin{aligned} \left| \int k \int k' - \left(\int_{\mathcal{O}} a^2 d\omega \right) \text{tr}(T') \right| &\leq \eta 2^{\#R} C(k) \prod_{p \neq q} C(k_p), \\ \left(\int_{\mathcal{O}} a_1^2 d\omega \right) \text{tr}(T'_1) &\leq \eta 2^{\#R-1} C(k, V) \prod_{p \neq q} C(k_p), \\ \left(\int_{\mathcal{O}} a_2^2 d\omega \right) \text{tr}(T'_2) &\leq \eta 2^{\#R-1} \prod_{p \neq q} C(k_p), \\ \left(\int_{\mathcal{O}} a_{\text{nt}}^2 d\omega \right) \text{tr}(T') &\leq \eta \prod_{p \neq q} C(k_p). \end{aligned} \quad (30.9)$$

We conclude by choosing η small enough that the right-hand side of each of the above inequalities is bounded by ε . \square

We turn now to the proof of the theorem. We retain the definition of “good” (relative to h , U and W) from §29.3; recall that this excises all σ for which $h\lambda_\sigma \notin U$, as well as those which are not orbit-distinguished. Note also that every σ in the support of k_h is good. Let $\varepsilon > 0$ be small; we eventually let it tend to zero sufficiently slowly with respect to h .

LEMMA 30.6. *We have*

$$|\ell - (\mathcal{M} - \mathcal{M}_{\text{nt}})| \leq \varepsilon, \quad (30.10)$$

where ℓ is as defined in §30.3 and

$$\begin{aligned}\mathcal{M} &:= \mathfrak{h}^d \sum_{\Sigma: \sigma \text{ is good}} \mathcal{P}_{\Sigma}(\mathrm{Op}_{\mathfrak{h}}(a)^2 \otimes T'), \\ \mathcal{M}_{\mathrm{nt}} &:= \mathfrak{h}^d \sum_{\substack{\Sigma: \sigma \text{ is good} \\ \text{and non-tempered}}} \mathcal{P}_{\Sigma}(\mathrm{Op}_{\mathfrak{h}}(a)^2 \otimes T'), \\ \mathcal{E} &:= \mathfrak{h}^d \sum_{\Sigma: \sigma \text{ is good}} \sum_{j=1}^2 \mathcal{P}_{\Sigma}(\mathrm{Op}_{\mathfrak{h}}(a_j)^2 \otimes T'_j).\end{aligned}$$

Proof. By the definitions and the period formula (28.5), we have

$$\ell - (\mathcal{M} - \mathcal{M}_{\mathrm{nt}}) = \mathfrak{h}^d \sum_{\substack{\Sigma: \sigma \text{ is good} \\ \text{and tempered}}} \mathcal{L}(\Pi, \Sigma) (k_{\mathfrak{h}}(\sigma)k'(\sigma') - \mathcal{H}_{\sigma}(\mathrm{Op}_{\mathfrak{h}}(a)^2)\mathcal{H}_{\sigma'}(T')).$$

By (30.6), the above is bounded in magnitude by

$$\sum_{j=1}^2 \mathfrak{h}^d \sum_{\substack{\Sigma: \sigma \text{ is good} \\ \text{and tempered}}} \underbrace{\mathcal{L}(\Pi, \Sigma) |\mathcal{H}_{\sigma}(\mathrm{Op}_{\mathfrak{h}}(a_j)^2)| |\mathcal{H}_{\sigma'}(T'_j)|}_{= \mathcal{P}_{\Sigma}(\mathrm{Op}_{\mathfrak{h}}(a_j)^2 \otimes T'_j)} \leq \mathcal{E}.$$

Note that the absolute values surrounding \mathcal{H}_{σ} on the right-hand side of (30.6), which are expected to be extraneous (see the remark following Theorem 22.2), have disappeared, due to the manifest positivity of \mathcal{P}_{Σ} . \square

By Lemma 30.4 and (30.3), we have (with \simeq as defined in §30.3)

$$\mathcal{M} \simeq \frac{\tau(\mathbf{H})}{\tau(\mathbf{G})} \left(\int_{\mathcal{O}} a^2 d\omega \right) \mathrm{tr}(T') = \frac{\tau(\mathbf{H})}{\tau(\mathbf{G})} \int k \int k' + o_{\varepsilon \rightarrow 0}(1). \quad (30.11)$$

The estimate

$$\mathcal{M}_{\mathrm{nt}} = o_{\varepsilon \rightarrow 0}(1). \quad (30.12)$$

follows from Lemma 30.4, using the weak Weyl law (29.2) to discard error terms:

$$\begin{aligned}\mathcal{M}_{\mathrm{nt}} &\stackrel{(a)}{=} \mathfrak{h}^d \sum_{\substack{\Sigma: \sigma \text{ is good} \\ \text{and non-tempered}}} (\mathcal{P}_{\Sigma}(\mathrm{Op}_{\mathfrak{h}}(a_{\mathrm{nt}})^2 \otimes T') + O(\mathfrak{h}^N)) \\ &\stackrel{(b)}{\leq} \mathfrak{h}^d \sum_{\Sigma} \mathcal{P}_{\Sigma}(\mathrm{Op}_{\mathfrak{h}}(a_{\mathrm{nt}})^2 \otimes T') + O(\mathfrak{h}^{N-O(1)}) \\ &\stackrel{(c)}{=} \frac{\tau(\mathbf{H})}{\tau(\mathbf{G})} \underbrace{\left(\int_{\mathcal{O}} a_{\mathrm{nt}}^2 d\omega \right) \mathrm{tr}(T')}_{\leq \varepsilon, \text{ by (30.5)}} + o_{\mathfrak{h} \rightarrow 0}(1) + O(\mathfrak{h}^{N-O(1)}) \\ &= o_{\varepsilon \rightarrow 0}(1),\end{aligned}$$

where (a) is due to (30.7), (b) to (29.2) and (c) to Lemma 30.4. We note that in the second step, we used the positivity of \mathcal{P}_Σ to drop the condition “ σ non-tempered”, and that in the final step, the implied constant $O(1)$ in the exponent is independent of N , and so by taking N large enough we may arrange that the difference $N - O(1)$ is positive.

The estimate

$$\mathcal{E} = o_{\varepsilon \rightarrow 0}(1) \tag{30.13}$$

follows similarly from Lemma 30.4 and (30.4):

$$\mathcal{E} = \frac{\tau(\mathbf{H})}{\tau(\mathbf{G})} \underbrace{\sum_{j=1}^2 \left(\int_{\mathcal{O}} a_j^2 d\omega \right) \text{tr}(T_j')}_{\leq \varepsilon} + o_{\mathfrak{h} \rightarrow 0}(1) = o_{\varepsilon \rightarrow 0}(1).$$

By combining these last estimates (30.11)–(30.13) with (30.10), we conclude the proof of Theorem 30.1.

31. The normalized asymptotic formula

The hard work having been completed, we explain here how Theorem 30.1 may be applied to sharply-truncated sums over families. By dividing out the cardinalities of those families, we then obtain the normalized asymptotic formulas promised in §1.

31.1. Approximating nice sets by continuous functions

Let (X, μ) be a normal topological space equipped with a Borel measure.

Definition 31.1. We say that a subset $U \subseteq X$ is *nice* if it is open, precompact, and has measure-zero boundary.

Definition 31.2. Given a class \mathcal{C} of integrable functions on X , we say that an integrable function w on X is *approximable by \mathcal{C}* if

- (i) there exists $k_0 \in \mathcal{C}$ such that $|w| \leq k_0$, and
- (ii) for each $\varepsilon > 0$, there exist $k, k_+ \in \mathcal{C}$ such that

$$|w - k| \leq k_+ \quad \text{and} \quad \int k_+ \leq \varepsilon.$$

By an exercise in applying Urysohn’s lemma, we have the following result.

LEMMA 31.3. *The characteristic function of any nice subset $U \subseteq X$ is approximable by the class $C_c(X)$ of continuous compactly-supported functions.*

The notion of approximability is well behaved with respect to products. Indeed, suppose given spaces X_1, \dots, X_n as above, each equipped with a Borel measure. For $j=1, \dots, n$, let w_j be an integrable function on X_j that is approximable by some class \mathcal{C}_j of integrable functions. Assume also that $k_1+k_2 \in \mathcal{C}_j$, whenever $k_1, k_2 \in \mathcal{C}_j$. Then, the function $X_1 \times \dots \times X_n \ni x \mapsto w_1(x_1) \dots w_n(x_n)$ is approximable by the class \mathcal{C} consisting of all sums of functions of the form $x \mapsto k_1(x_1) \dots k_n(x_n)$, with $k_j \in \mathcal{C}_j$.

31.2. Summing against approximable weights

The spaces $[\mathfrak{h}^\wedge] \cap \text{image}(\mathcal{O}_{\text{stab}})$ and $(\widehat{H}')_{\text{temp}}^{\pi'}$ come with natural topologies, given in the former case by identifying $[\mathfrak{h}^\wedge]$ with an affine space, in the latter by the discussion in §21 and §24.1. They also come with natural measures: normalized affine measure and the restriction of Plancherel measure, respectively. The terminology of §31.1 thus applies, and we readily derive from Theorem 30.1 the following.

COROLLARY 31.4. *Let $k: [\mathfrak{h}^\wedge] \rightarrow \mathbb{C}$ and $k': (\widehat{H}')_{\text{temp}}^{\pi'} \rightarrow \mathbb{C}$ be any functions approximable by the function spaces \mathcal{K} and \mathcal{K}' defined in §30.2. Then, with notation as in §30.3,*

$$h^d \sum_{\substack{\sigma \text{ and } \sigma' \text{ tempered} \\ \mathcal{O}^{\pi, \sigma} \neq \emptyset}} \mathcal{L}(\Pi, \Sigma) k_h(\sigma) k'(\sigma') \simeq \frac{\tau(\mathbf{H})}{\tau(\mathbf{G})} \int k \int k'.$$

31.3. Summing over unweighted families

Definition 31.5. We say that a subset $U \subset [\mathfrak{h}^\wedge] \cap \text{image}(\mathcal{O}_{\text{stab}})$ is *admissible* if it is nice (cf. §31.1), and that a subset $U' \subset (\widehat{H}')_{\text{temp}}^{\pi'}$ is *admissible* if it is nice and if its projection onto $(\widehat{H}_{\mathfrak{p}})_{\text{temp}}^{\pi_{\mathfrak{p}}}$, for a finite place $\mathfrak{p} \in R$, is contained in the union of the allowable components (see §24).

(The smooth version of) Lemma 31.3 implies the following result.

LEMMA 31.6. *If U (resp. U') as above is admissible, then its characteristic function is approximable by \mathcal{K} (resp. \mathcal{K}').*

We henceforth fix some non-empty admissible sets U and U' above, and define the family \mathcal{F}_h as follows.

Definition 31.7. We let \mathcal{F}_h be the set of Σ such that

- σ and σ' are tempered,
- σ is orbit-distinguished, i.e. $\mathcal{O}^{\pi, \sigma} \neq \emptyset$,
- $h\lambda_\sigma \in U$ and $\sigma' \in U'$.

Corollary 31.4 specializes as follows.

COROLLARY 31.8. *We have*

$$h^d \sum_{\Sigma \in \mathcal{F}_h} \mathcal{L}(\Pi, \Sigma) \simeq \frac{\tau(\mathbf{H})}{\tau(\mathbf{G})} \text{vol}(U) \text{vol}(U'). \quad (31.1)$$

31.4. Family size

To interpret the left-hand side of (31.1) as a normalized average of $\mathcal{L}(\Pi, \Sigma)$, we need to know the approximate cardinality of the family \mathcal{F}_h .

LEMMA 31.9. *One has*

$$h^d |\mathcal{F}_h| = \tau(\mathbf{H}) \text{vol}(U) \text{vol}(U') + o_{h \rightarrow 0}(1), \quad (31.2)$$

Proof. This can be deduced from the methods of Duistermaat, Kolk, Varadarajan [DKV] and Sauvageot [Sau], but it is simpler for us to give a more direct argument. We apply the trace formula for the compact quotient $[H]$ in a standard way, using the approximation arguments given by

- Lemma 22.4, and
- Theorem 24.5, via the consequences (24.4) and (24.6).

For small h , the support condition $1 + o_{h \rightarrow 0}(1)$ on the test function at \mathfrak{q} implies that the only non-zero contribution on the geometric side comes from the identity element. We readily obtain the smoothly-weighted variant of (31.2) from which (31.2) itself then follows as in the proof of (31.1). \square

31.5. Analytic conductors

Let $\Sigma \in \mathcal{F}_h$. We denote by $C(\Pi, \Sigma)$ the analytic conductor of $L(s, -, \Pi \times \Sigma^\vee)$ at $s = \frac{1}{2}$.

LEMMA 31.10. $C(\Pi, \Sigma) \asymp |\mathcal{F}_h|^4$.

Proof. Our assumptions imply that the contribution to the analytic conductor from places other than \mathfrak{q} is bounded.

By the discussion in §15 (especially (15.3)), we have

$$C(\Pi, \Sigma) \asymp h^{-\varepsilon m_H m_G},$$

with notation as in Table 1. On the other hand, we have seen that $|\mathcal{F}_h| \asymp h^{-\dim B_H}$. By inspection, $4 \dim(B_H) = \varepsilon n_G n_H$ in all cases. The required estimate follows. \square

label	\mathbf{G}	\mathbf{H}	$\dim(B_H)$	m_G	m_H	ε
(i)	GL_{n+1}	GL_n	$\frac{n(n+1)}{2}$	$n+1$	n	2
(ii)	U_{n+1}	U_n	$\frac{n(n+1)}{2}$	$n+1$	n	2
(iii)	SO_{2n+2}	SO_{2n+1}	$n(n+1)$	$2n+2$	$2n$	1
(iv)	SO_{2n+1}	SO_{2n}	n^2	$2n$	$2n$	1

Table 1. Numerology: m_G is the dimension of the standard representation of \widehat{G} , similarly for H .

31.6. Main result

Dividing (31.1) by (31.2), we conclude the following.

THEOREM 31.11. *Let notation and assumptions be as above. In particular, we assume the following:*

- (\mathbf{G}, \mathbf{H}) is a GGP pair over a number field F . We have fixed a large enough finite set of places R .
- We have fixed an automorphic representation Π on \mathbf{G} , unramified outside of R , and satisfying the conditions enunciated in §25.7.
- We have fixed an archimedean place $\mathfrak{q} \in R$, and set

$$H = \mathbf{H}(F_{\mathfrak{q}}) \quad \text{and} \quad H' = \prod_{\mathfrak{p} \in R \setminus \{\mathfrak{q}\}} \mathbf{H}(F_{\mathfrak{p}}).$$

For an automorphic representation Σ on \mathbf{H} , unramified outside of R , we have denoted by σ and σ' the associated representations of H and H' .

Let U and U' be admissible subsets of the tempered distinguished spectra of H and H' , as above, and let $\mathcal{F}_{\mathfrak{h}}$ denote the family of representations on \mathbf{H} associated with these subsets, as in Definition 31.7.

Let $\mathcal{L}(\Pi, \Sigma)$ denote the branching coefficient as defined in §25.4. Then

$$\frac{1}{|\mathcal{F}_{\mathfrak{h}}|} \sum_{\Sigma \in \mathcal{F}_{\mathfrak{h}}} \mathcal{L}(\Pi, \Sigma) = \frac{1}{\tau(\mathbf{G})} + o_{\mathfrak{h} \rightarrow 0}(1). \tag{31.3}$$

If we assume the conjectures of Ichino–Ikeda and N. Harris (§25.5), or restrict to cases

in which those conjectures are known (e.g., [Z1]), then

$$\frac{1}{|\mathcal{F}_h|} \sum_{\Sigma \in \mathcal{F}_h} \frac{\tau(\mathbf{G})\tau(\mathbf{H})}{2^\beta} \Delta_G^{(R)} \frac{L^{(R)}(\frac{1}{2}, -, \Pi \times \Sigma^\vee)}{L^{(R)}(1, \text{Ad}, \Pi \times \Sigma^\vee)} = \tau(\mathbf{H}) + o_{h \rightarrow 0}(1). \quad (31.4)$$

We may “simplify” the final formulas by recalling (§25.7) that $\tau(\mathbf{G}) = \tau(\mathbf{H}) = 2$.

References

- [AV] ADAMS, J. & VOGAN, D. A., JR., Contragredient representations and characterizing the local Langlands correspondence. *Amer. J. Math.*, 138 (2016), 657–682.
- [ALTV] ADAMS, J. D., VAN LEEUWEN, M. A. A., TRAPA, P. E. & VOGAN, D. A. J, Unitary representations of real reductive groups. *Astérisque*, 417 (2020), 188 pp.
- [BV] BARBASCH, D. & VOGAN, D. A., JR., The local structure of characters. *J. Funct. Anal.*, 37 (1980), 27–55.
- [Ba] BARYSHNIKOV, YU., GUEs and queues. *Probab. Theory Related Fields*, 119 (2001), 256–274.
- [Bea1] BEALS, R., A general calculus of pseudodifferential operators. *Duke Math. J.*, 42 (1975), 1–42.
- [Bea2] — Characterization of pseudodifferential operators and applications. *Duke Math. J.*, 44 (1977), 45–57.
- [BZ] BERNSTEIN, I. N. & ZELEVINSKY, A. V., Induced representations of reductive p -adic groups. I. *Ann. Sci. École Norm. Sup.*, 10 (1977), 441–472.
- [BDK] BERNSTEIN, J., DELIGNE, P. & KAZHDAN, D., Trace Paley–Wiener theorem for reductive p -adic groups. *J. Analyse Math.*, 47 (1986), 180–192.
- [BR1] BERNSTEIN, J. & REZNIKOV, A., Sobolev norms of automorphic functionals. *Int. Math. Res. Not.*, 40 (2002), 2155–2174.
- [BR2] — Subconvexity bounds for triple L -functions and representation theory. *Ann. of Math.*, 172 (2010), 1679–1718.
- [Ber] BERNSTEIN, J. N., Le “centre” de Bernstein, in *Representations of Reductive Groups over a Local Field*, pp. 1–32. Hermann, Paris, 1984.
- [BP1] BEUZART-PLESSIS, R., Endoscopie et conjecture locale raffinée de Gan–Gross–Prasad pour les groupes unitaires. *Compos. Math.*, 151 (2015), 1309–1371.
- [BP2] — La conjecture locale de Gross–Prasad pour les représentations tempérées des groupes unitaires. *Mém. Soc. Math. Fr.*, 149 (2016), 191 pp.
- [BP3] — A local trace formula for the Gan–Gross–Prasad conjecture for unitary groups: the Archimedean case. *Astérisque*, 418 (2020), 299 pp.
- [BP4] — Comparison of local spherical characters and the Ichino–Ikeda conjecture for unitary groups. Preprint, 2016. [arXiv:1602.06538](https://arxiv.org/abs/1602.06538) [math.NT].
- [BGW] BHARGAVA, M., GROSS, B. H. & WANG, X., Arithmetic invariant theory II: Pure inner forms and obstructions to the existence of orbits, in *Representations of Reductive Groups*, Progr. Math., 312, pp. 139–171. Birkhäuser/Springer, Cham, 2015.
- [Bl] BLOMER, V., Period integrals and Rankin–Selberg L -functions on $GL(n)$. *Geom. Funct. Anal.*, 22 (2012), 608–620.
- [BFK+] BLOMER, V., FOUVRY, E., KOWALSKI, E., MICHEL, P. & MILIĆEVIĆ, D., On moments of twisted L -functions. *Amer. J. Math.*, 139 (2017), 707–768.

- [BFS+] BLOMER, V., FOUVRY, É., KOWALSKI, E., MICHEL, P., MILIĆEVIĆ, D. & SAWIN, W., The second moment theory of families of L -functions. To appear in *Mem. Amer. Math. Soc.*
- [BM] BLOMER, V. & MILIĆEVIĆ, D., The second moment of twisted modular L -functions. *Geom. Funct. Anal.*, 25 (2015), 453–516.
- [Bo1] BOREL, A., Automorphic L -functions, in *Automorphic Forms, Representations and L-Functions* (Corvallis, OR, 1977), Part 2, Proc. Sympos. Pure Math., 33, pp. 27–61. Amer. Math. Soc., Providence, RI, 1979.
- [Bo2] — *Linear Algebraic Groups*. Graduate Texts in Mathematics, 126. Springer, New York–Heidelberg, 1991.
- [BK] BUI, H. M. & KEATING, J. P., On the mean values of L -functions in orthogonal and symplectic families. *Proc. Lond. Math. Soc.*, 96 (2008), 335–366.
- [Cah] CAHEN, B., Invariant symbolic calculus for compact Lie groups. *Arch. Math. (Brno)*, 55 (2019), 139–155.
- [Cas1] CASSELMAN, W., Canonical extensions of Harish-Chandra modules to representations of G . *Canad. J. Math.*, 41 (1989), 385–438.
- [Cas2] — Introduction to the theory of admissible representations of reductive p -adic groups. Course notes, 1995.
<https://www.math.ubc.ca/~cass/research/pdf/p-adic-book.pdf>.
- [Ch] CHEVALLEY, C., Invariants of finite groups generated by reflections. *Amer. J. Math.*, 77 (1955), 778–782.
- [CF] CONREY, J. B. & FARMER, D. W., Mean values of L -functions and symmetry. *Internat. Math. Res. Notices*, 17 (2000), 883–908.
- [CFK+] CONREY, J. B., FARMER, D. W., KEATING, J. P., RUBINSTEIN, M. O. & SNAITH, N. C., Integral moments of L -functions. *Proc. London Math. Soc.*, 91 (2005), 33–104.
- [CHH] COWLING, M., HAAGERUP, U. & HOWE, R., Almost L^2 matrix coefficients. *J. Reine Angew. Math.*, 387 (1988), 97–110.
- [DN] DIMITROV, M. & NYSSÉN, L., Test vectors for trilinear forms when at least one representation is not supercuspidal. *Manuscripta Math.*, 133 (2010), 479–504.
- [DKV] DUISTERMAAT, J. J., KOLK, J. A. C. & VARADARAJAN, V. S., Spectra of compact locally symmetric manifolds of negative curvature. *Invent. Math.*, 52 (1979), 27–93. Correction in *Invent. Math.*, 54 (1979), 101.
- [EW] EINSIEDLER, M. & WARD, T., *Ergodic Theory with a View towards Number Theory*. Graduate Texts in Mathematics, 259. Springer, London, 2011.
- [FLO] FEIGON, B., LAPID, E. & OFFEN, O., On representations distinguished by unitary groups. *Publ. Math. Inst. Hautes Études Sci.*, 115 (2012), 185–323.
- [FS1] FRAHM, J. & SU, F., Upper bounds for geodesic periods over rank one locally symmetric spaces. *Forum Math.*, 30 (2018), 1065–1077.
- [FS2] — Rankin–Selberg periods for spherical principal series. Preprint, 2020.
[arXiv:1706.05167](https://arxiv.org/abs/1706.05167) [math.NT].
- [GGP] GAN, W. T., GROSS, B. H. & PRASAD, D., Symplectic local root numbers, central critical L values, and restriction problems in the representation theory of classical groups. *Astérisque*, 346 (2012), 1–109.
- [Gr] GROSS, B. H., On the motive of a reductive group. *Invent. Math.*, 130 (1997), 287–313.
- [GP] GROSS, B. H. & PRASAD, D., On the decomposition of a representation of SO_n when restricted to SO_{n-1} . *Canad. J. Math.*, 44 (1992), 974–1002.
- [Gu] GUTT, S., An explicit $*$ -product on the cotangent bundle of a Lie group. *Lett. Math. Phys.*, 7 (1983), 249–258.

- [Hai] HAINES, T. J., The stable Bernstein center and test functions for Shimura varieties, in *Automorphic Forms and Galois Representations*. Vol. 2, London Math. Soc. Lecture Note Ser., 415, pp. 118–186. Cambridge Univ. Press, Cambridge, 2014.
- [HR] HAINES, T. J. & ROSTAMI, S., The Satake isomorphism for special maximal parahoric Hecke algebras. *Represent. Theory*, 14 (2010), 264–284.
- [HarB] HARRIS, B., Wave front sets of reductive Lie group representations II. *Trans. Amer. Math. Soc.*, 370 (2018), 5931–5962.
- [HarN] HARRIS, R. N., The refined Gross–Prasad conjecture for unitary groups. *Int. Math. Res. Not. IMRN*, 2 (2014), 303–389.
- [HS] HECHT, H. & SCHMID, W., A proof of Blattner’s conjecture. *Invent. Math.*, 31 (1975), 129–154.
- [Hoh] HOHENWARTER, M. ET AL., GeoGebra. Online calculator, April 2018.
<http://www.geogebra.org>.
- [Hör] HÖRMANDER, L., *The Analysis of Linear Partial Differential Operators*. III. Classics in Mathematics. Springer, Berlin–Heidelberg, 2007.
- [HNS] HU, Y., NELSON, P. D. & SAHA, A., Some analytic aspects of automorphic forms on $GL(2)$ of minimal type. *Comment. Math. Helv.*, 94 (2019), 767–801.
- [II] ICHINO, A. & IKEDA, T., On the periods of automorphic forms on special orthogonal groups and the Gross–Prasad conjecture. *Geom. Funct. Anal.*, 19 (2010), 1378–1425.
- [KaSa] KATZ, N. M. & SARNAK, P., Zeroes of zeta functions and symmetry. *Bull. Amer. Math. Soc.*, 36 (1999), 1–26.
- [KeSn1] KEATING, J. P. & SNAITH, N. C., Random matrix theory and $\zeta(1/2+it)$. *Comm. Math. Phys.*, 214 (2000), 57–89.
- [KeSn2] — Random matrix theory and L -functions at $s=1/2$. *Comm. Math. Phys.*, 214 (2000), 91–110.
- [Ki1] KIRILLOV, A. A., Unitary representations of nilpotent Lie groups. *Uspehi Mat. Nauk*, 17 (1962), 57–110 (Russian).
- [Ki2] — Merits and demerits of the orbit method. *Bull. Amer. Math. Soc.*, 36 (1999), 433–488.
- [Kn] KNAPP, A. W., *Representation Theory of Semisimple Groups*. Princeton Mathematical Series, 36. Princeton Univ. Press, Princeton, NJ, 1986.
- [Kos1] KOSTANT, B., The principal three-dimensional subgroup and the Betti numbers of a complex simple Lie group. *Amer. J. Math.*, 81 (1959), 973–1032.
- [Kos2] — Lie group representations on polynomial rings. *Amer. J. Math.*, 85 (1963), 327–404.
- [Kos3] — On Whittaker vectors and representation theory. *Invent. Math.*, 48 (1978), 101–184.
- [Kot] KOTTWITZ, R. E., Stable trace formula: cuspidal tempered terms. *Duke Math. J.*, 51 (1984), 611–650.
- [KMS] KOWALSKI, E., MICHEL, P. & SAWIN, W., Bilinear forms with Kloosterman sums and applications. *Ann. of Math.*, 186 (2017), 413–500.
- [Kur] KURTZKE, J. F., JR., Centralizers of irregular elements in reductive algebraic groups. *Pacific J. Math.*, 104 (1983), 133–154.
- [Küs] KÜSTER, B., On the semiclassical functional calculus for h -dependent functions. *Ann. Global Anal. Geom.*, 52 (2017), 57–97.
- [L] LUO, Z., A local trace formula for the local Gan–Gross–Prasad conjecture for special orthogonal groups. Preprint, 2020. [arXiv:2009.13947\[math.RT\]](https://arxiv.org/abs/2009.13947).
- [MV] MICHEL, P. & VENKATESH, A., The subconvexity problem for GL_2 . *Publ. Math. Inst. Hautes Études Sci.*, 111 (2010), 171–271.

- [Mi] MILNE, J. S., Lectures on étale cohomology (v2.21). Course notes, 2013.
www.jmilne.org/math/CourseNotes/lec.html.
- [MW] MÈGLIN, C. & WALDSPURGER, J.-L., *Sur les conjectures de Gross et Prasad. II*. Astérisque, 347. Société Mathématique de France, Paris, 2012.
- [Mo] MORRIS, D. W., *Ratner's Theorems on Unipotent Flows*. Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 2005.
- [MP] MOY, A. & PRASAD, G., Jacquet functors and unrefined minimal K -types. *Comment. Math. Helv.*, 71 (1996), 98–121.
- [MFK] MUMFORD, D., FOGARTY, J. & KIRWAN, F., *Geometric Invariant Theory*. Ergebnisse der Mathematik und ihrer Grenzgebiete, 34. Springer, Berlin–Heidelberg, 1994.
- [NE] NELSON, E., Analytic vectors. *Ann. of Math.*, 70 (1959), 572–615.
- [NS] NELSON, E. & STINESPRING, W. F., Representation of elliptic operators in an enveloping algebra. *Amer. J. Math.*, 81 (1959), 547–560.
- [NPD1] NELSON, P. D., Microlocal lifts and quantum unique ergodicity on $GL_2(\mathbb{Q}_p)$. *Algebra Number Theory*, 12 (2018), 2033–2064.
- [NPD2] — Subconvex equidistribution of cusp forms: reduction to Eisenstein observables. *Duke Math. J.*, 168 (2019), 1665–1722.
- [NPS] NELSON, P. D., PITALE, A. & SAHA, A., Bounds for Rankin–Selberg integrals and quantum unique ergodicity for powerful levels. *J. Amer. Math. Soc.*, 27 (2014), 147–191.
- [OI] OLSHANSKI, G., Projections of orbital measures, Gel'fand–Tsetlin polytopes, and splines. *J. Lie Theory*, 23 (2013), 1011–1022.
- [On] ONO, T., On the Tamagawa number of algebraic tori. *Ann. of Math.*, 78 (1963), 47–73.
- [PR] PLATONOV, V. & RAPINCHUK, A., *Algebraic Groups and Number Theory*. Pure and Applied Mathematics, 139. Academic Press, Boston, MA, 1994.
- [RS] RALLIS, S. & SCHIFFMANN, G., Multiplicity one conjectures. Preprint, 2007.
[arXiv:0705.2168](https://arxiv.org/abs/0705.2168) [math.RT].
- [Rao] RANGA RAO, R., Orbital integrals in reductive groups. *Ann. of Math.*, 96 (1972), 505–510.
- [Rat] RATNER, M., On Raghunathan's measure conjecture. *Ann. of Math.*, 134 (1991), 545–607.
- [Re] REZNIKOV, A., Rankin–Selberg without unfolding and bounds for spherical Fourier coefficients of Maass forms. *J. Amer. Math. Soc.*, 21 (2008), 439–477.
- [Ri] RIEFFEL, M. A., Lie group convolution algebras as deformation quantizations of linear Poisson structures. *Amer. J. Math.*, 112 (1990), 657–685.
- [Rob] ROBERTS, J., Classical $6j$ -symbols and the tetrahedron. *Geom. Topol.*, 3 (1999), 21–66.
- [Ros1] ROSSMANN, W., Kirillov's character formula for reductive Lie groups. *Invent. Math.*, 48 (1978), 207–220.
- [Ros2] — Limit characters of reductive Lie groups. *Invent. Math.*, 61 (1980), 53–66.
- [Ros3] — Limit orbits in reductive Lie algebras. *Duke Math. J.*, 49 (1982), 215–229.
- [Ros4] — Tempered representations and orbits. *Duke Math. J.*, 49 (1982), 231–247.
- [Ros5] — Picard–Lefschetz theory and characters of a semisimple Lie group. *Invent. Math.*, 121 (1995), 579–611.
- [SVe] SAKELLARIDIS, Y. & VENKATESH, A., Periods and harmonic analysis on spherical varieties. *Astérisque*, 396 (2017), 360 pp.
- [SVO] SALAMANCA-RIBA, S. A. & VOGAN, D. A., JR., On the classification of unitary representations of reductive Lie groups. *Ann. of Math.*, 148 (1998), 1067–1133.
- [Sar] SARNAK, P., Fourth moments of Grössencharakteren zeta functions. *Comm. Pure Appl. Math.*, 38 (1985), 167–178.

- [SST] SARNAK, P., SHIN, S. W. & TEMPLIER, N., Families of L -functions and their symmetry, in *Families of Automorphic Forms and the Trace Formula*, Simons Symp., pp. 531–578. Springer, Cham, 2016.
- [Sau] SAUVAGEOT, F., Principe de densité pour les groupes réductifs. *Compositio Math.*, 108 (1997), 151–184.
- [Se] SERRE, J.-P., *A Course in Arithmetic*. Graduate Texts in Mathematics, 7. Springer, New York–Heidelberg, 1973.
- [ShTs] SHALIKA, J. & TSCHINKEL, Y., Height zeta functions of equivariant compactifications of unipotent groups. *Comm. Pure Appl. Math.*, 69 (2016), 693–733.
- [ShTe] SHIN, S. W. & TEMPLIER, N., Sato–Tate theorem for families and low-lying zeros of automorphic L -functions. *Invent. Math.*, 203 (2016), 1–177.
- [So] SOUNDARARAJAN, K., Weak subconvexity for central values of L -functions. *Ann. of Math.*, 172 (2010), 1469–1498.
- [SoTh] SOUNDARARAJAN, K. & THORNER, J., Weak subconvexity without a Ramanujan hypothesis. *Duke Math. J.*, 168 (2019), 1231–1268.
- [SZ] SUN, B. & ZHU, C. B., Multiplicity one theorems: the Archimedean case. *Ann. of Math.*, 175 (2012), 23–44.
- [T] TAYLOR, M. E., Noncommutative microlocal analysis. I. *Mem. Amer. Math. Soc.*, 52 (1984), 182 pp.
- [Ve] VENKATESH, A., Sparse equidistribution problems, period bounds and subconvexity. *Ann. of Math.*, 172 (2010), 989–1094.
- [Vo1] VOGAN, D. A., JR., Gel’fand–Kirillov dimension for Harish-Chandra modules. *Invent. Math.*, 48 (1978), 75–98.
- [Vo2] — The algebraic structure of the representation of semisimple Lie groups. I. *Ann. of Math.*, 109 (1979), 1–60.
- [Vo3] — The method of coadjoint orbits for real reductive groups, in *Representation Theory of Lie Groups* (Park City, UT, 1998), IAS/Park City Math. Ser., 8, pp. 179–238. Amer. Math. Soc., Providence, RI, 2000.
- [Wld] WALDSPURGER, J.-L., La formule de Plancherel pour les groupes p -adiques (d’après Harish-Chandra). *J. Inst. Math. Jussieu*, 2 (2003), 235–333.
- [Wll] WALLACH, N. R., *Real Reductive Groups*. I. Pure and Applied Mathematics, 132. Academic Press, Boston, MA, 1988.
- [Wa] WATSON, T. C., *Rankin Triple Products and Quantum Chaos*. Ph.D. Thesis, Princeton University, Princeton, NJ, 2002.
- [We] WEIL, A., *Adeles and Algebraic Groups*. Progress in Mathematics, 23. Birkhäuser, Boston, MA, 1982.
- [Wh] WHITNEY, H., Elementary structure of real algebraic varieties. *Ann. of Math.*, 66 (1957), 545–556.
- [Wo] WOODBURY, M. C., *Explicit Trilinear Forms and Subconvexity of the Triple Product L -Function*. Ph.D. Thesis, The University of Wisconsin–Madison, Madison, WI.
- [Z1] ZHANG, W., Fourier transform and the global Gan–Gross–Prasad conjecture for unitary groups. *Ann. of Math.*, 180 (2014), 971–1049.
- [Z2] — Periods, cycles, and L -functions: a relative trace formula approach, in *Proceedings of the International Congress of Mathematicians* (Rio de Janeiro, 2018), Vol. II, pp. 487–521. World Sci. Publ., Hackensack, NJ, 2018.

PAUL D. NELSON
Departement Mathematik
Eidgenössische Technische Hochschule
HG G 32.5 Rämistrasse 101
CH-8092 Zurich
Switzerland
nelson.paul.david@gmail.com

AKSHAY VENKATESH
Institute for Advanced Study
1 Einstein Drive
Princeton, NJ 08540
U.S.A.
akshay.venkatesh@gmail.com

Received February 12, 2019

Received in revised form August 27, 2020