Math 210B. Frobenius-Schur indicator

1. Introduction

Let G be a finite group, and V a finite-dimensional representation of G over $\mathbf{C}$. If $F \subset \mathbf{C}$ is a subfield we say V is *defined over* F if $V \simeq V_0 \otimes_F \mathbf{C}$ for some $F[G]$-module V_0. Any such V_0 obviously inherits irreducibility (over F) from that of V (over $\mathbf{C}$) and $\chi_{V_0} = \chi_V$, so necessarily χ_V is F-valued. Moreover, such a V_0 is therefore unique up to isomorphism if it exists since its character is uniquely determined (namely, χ_V with values in F). However, this uniqueness for V_0 is non-canonical; i.e., it is not "unique up to unique isomorphism" and so any attempt to use Galois-descent techniques to try to prove the *sufficiency* of the necessary condition $\chi_V(G) \subset F$ for V to be defined over F is doomed to fail, even when V is irreducible

Now assume V is irreducible. Although we know from HW9 Exercise 2 that V is defined over *some* finite extension F of $\mathbf{Q}$, and necessarily $F \supset \mathbf{Q}(\chi_V)$, it is a *very subtle* problem in number theory in general to understand when the descent to a given $F \supset \mathbf{Q}(\chi_V)$ actually exists. Moreover, when $\mathbf{Q}(\chi_V)$ fails to work (i.e., V is not defined over $\mathbf{Q}(\chi_V)$) there is generally no *unique* minimal extension of $\mathbf{Q}(\chi_V)$ that works; rather, there are infinitely many extensions of $\mathbf{Q}(\chi_V)$ that work and share the same minimal degree. This phenomenon can only be understood via the main theorems of global class field theory for number fields, so we don't say anything more about it here.

In some special cases there is no obstruction. For example, the representation theory of symmetric groups is an extremely well-understood subject (relevant to many topics, ranging from cohomology rings of Grassmannians to the combinatorics of symmetric functions and Young tableaux and beyond), and in that setting it is known that all characters are $\mathbf{Q}$-valued and that all V can be directly constructed over $\mathbf{Q}$ (via "Specht modules"). By deeper methods, for $G = \mathrm{GL}_n(\mathbf{F}_q)$ there are again no obstructions: V always descends to the smallest possible field $\mathbf{Q}(\chi_V)$ (which is often larger than $\mathbf{Q}$) but giving a useful description of all irreducible representations is much harder than for symmetric groups. However, for $G = \mathrm{SL}_2(\mathbf{F}_q)$ obstructions to definability over $\mathbf{Q}(\chi_V)$ arise in a variety of ways.

In this handout we take up the most classical instance of this definiability question: when is an irreducible V defined over $\mathbf{R}$? The obvious necessary condition that $\chi := \chi_V$ be $\mathbf{R}$-valued says exactly that $\chi = \overline{\chi}$, and since $\overline{\chi}(g) = \chi_{\overline{V}}$ this is the same as the condition $V \simeq \overline{V}$ as $\mathbf{C}[G]$-modules.

Example 1.1. Consider the quaternion group $Q = \{\pm 1, \pm i, \pm j, \pm k\} \subset \mathbf{H}^\times$. (Read the handout on quaternions to learn about the 4-dimensional division algebra $\mathbf{H}$ over $\mathbf{R}$.) By viewing $\mathbf{H}$ as a 2-dimensional $\mathbf{C}$-vector space V through *right* multiplication on $\mathbf{H}$ by the commutative $\mathbf{R}$-subalgebra $\mathbf{C} = \mathbf{R} \oplus \mathbf{R}i \subset \mathbf{H}$, the *left* multiplication action by Q on $\mathbf{H}$ is $\mathbf{C}$-linear (left multiplication commutes with right multiplication in any associative ring!) and thereby defines a faithful representation $\rho : Q \hookrightarrow \mathrm{GL}_2(\mathbf{C})$. This has $\mathbf{R}$-valued character χ equal to ± 2 on $\pm 1 \in Q$ and vanishing elsewhere, and it is irreducible: either because $\langle \chi, \chi \rangle_Q = 1$ or because Q acting on $\mathbf{H}$ through left multiplications spans $\mathbf{H}$ over $\mathbf{R}$ (and obviously the *division algebra* $\mathbf{H}$ has no nonzero proper $\mathbf{H}$-submodule).

The representation ρ cannot be defined over $\mathbf{R}$. Indeed, the evident inclusion $\mathbf{H}^{\mathrm{opp}} \subset \mathrm{End}_{\mathbf{R}[Q]}(V)$ via right multiplication of $\mathbf{H}$ on itself is an equality (as $\mathbf{R}[Q]$ generates $\mathbf{H}$ inside $\mathrm{End}_{\mathbf{R}}(V)$ yet the centralizer of the left action of the division algebra $\mathbf{H}$ on itself is clearly the right action that is the 4-dimensional $\mathbf{H}^{\mathrm{opp}}$) whereas for any finite group G and $\mathbf{C}[G]$-module W of the form $W_0 \otimes_{\mathbf{R}} \mathbf{C}$ for an $\mathbf{R}[G]$-module W_0 we see that W has underlying $\mathbf{R}[G]$-module $W_0 \oplus W_0$ that is reducible and hence $\mathrm{End}_{\mathbf{R}[G]}(W) = \mathrm{Mat}_2(\mathrm{End}_{\mathbf{R}[G]}(W_0))$ is *not* a division algebra.

There are three cases that can arise: (i) χ is not $\mathbf{R}$-valued, (ii) χ is $\mathbf{R}$-valued with V is defined over $\mathbf{R}$, and (iii) χ is $\mathbf{R}$-valued with V not defined over $\mathbf{R}$. We shall give several ways to characterize which of the three cases we are in. The most concrete way is to calculate a certain quantity in terms of χ:

Definition 1.2. The *Frobenius-Schur indicator* of V is the number $\varepsilon_V = (1/\#G)\sum_{g\in G} \chi_V(g^2)$.

We will show that ε_V is equal to either 0, 1, or -1, and that these respectively correspond to cases (i), (ii), and (iii). In Example 1.1 one easily checks $\varepsilon_V = -1$.

2. Main result and preliminary calculations

Here is the result we aim to prove, providing ways to characterize whether we are in case (i), (ii), or (iii).

Theorem 2.1. *Each of the following lists of conditions below consist of equivalent properties:*

(i) χ *is not* $\mathbf{R}$*-valued,* $\mathrm{End}_{\mathbf{R}[G]}(V) = \mathbf{C}$, $\varepsilon_V = 0$, *there does not exist a nonzero* G*-equivariant bilinear form* $B : V \times V \to \mathbf{C}$*;*

(ii) χ *is* $\mathbf{R}$*-valued with* V *defined over* $\mathbf{R}$, $\mathrm{End}_{\mathbf{R}[G]}(V) = \mathrm{Mat}_2(\mathbf{R})$, $\varepsilon_V = 1$, *there exists a symmetric non-degenerate* G*-equivariant bilinear form* $B : V \times V \to \mathbf{C}$*;*

(iii) χ *is* $\mathbf{R}$*-valued with* V *not defined over* $\mathbf{R}$, $\mathrm{End}_{\mathbf{R}[G]}(V) = \mathbf{H}$, $\varepsilon_V = -1$, *there exists a skew-symmetric non-degenerate* G*-equivariant bilinear form* $B : V \times V \to \mathbf{C}$.

In cases (ii) *and* (iii)*, up to* $\mathbf{C}^\times$*-scaling* B *is the unique nonzero* G*-equivariant bilinear form* $B : V \times V \to \mathbf{C}$.

Since $\mathbf{C} \subset \mathrm{End}_{\mathbf{R}[G]}(V)$ through the given $\mathbf{C}$-linear structure on V, in view of the description of $\mathrm{End}_{\mathbf{R}[G]}(V)$ in all three cases we see that (iii) is exactly when the G-equivariant $\mathbf{C}$-linear structure on V can be enhanced to a G-equivariant left $\mathbf{H}$-module structure (via some inclusion of $\mathbf{C}$ as an $\mathbf{R}$-subalgebra of $\mathbf{H}$). For this reason, case (iii) is usually referred to as *quaternionic* (of which Example 1.1 is an example).

As a first step towards proving Theorem 2.1, we relate the conditions on ε_V to the conditions involving equivariant bilinear forms on V. By linear algebra, a bilinear form $B : V \times V \to \mathbf{C}$ is *functorially* the same as a linear map $T : V \to V^*$, so B is G-equivariant if and only if the corresponding T is G-equivariant (this is also part of HW8, Exercise 3(ii)). But V and V^* are *irreducible*, so by Schur's Lemma a nonzero equivariant T must be an isomorphism, and hence a nonzero equivariant B must be a non-degenerate pairing. Thus, a nonzero B exists if and only if $V \simeq V^*$, which in turn is equivalent to $\chi_V = \chi_{V^*}$. But $\chi_{V^*}(g) = \chi(g^{-1}) = \overline{\chi(g)}$, so we conclude that χ is $\mathbf{R}$-valued if and only if a nonzero G-equivariant B exists; that is, we have proved the equivalence of the first and last conditions in part (i) of the Theorem.

Let us probe more deeply for properties of a nonzero equivariant B when it exists. The space of bilinear forms on V is $(V \otimes V)^* \simeq V^* \otimes V^*$, and by *functoriality* of this isomorphism it must be G-equivariant (you can also check that by hand, but it is better to argue by functoriality considerations alone). We also have the functorial decomposition

$$V^* \otimes V^* \simeq \mathrm{Sym}^2(V^*) \oplus \wedge^2(V^*),$$

and by functoriality (or by hand, but you really should argue by functoriality) this is a G-equivariant isomorphism. The two summands on the right are exactly the spaces of symmetric bilinear forms and skew-symmetric bilinear forms respectively: although Sym^n and $\wedge^n$ are defined as *quotients* of the nth tensor power, in characteristic 0 these are *canonically* identified with analogous subspaces too. Thus,

$$((V \otimes V)^*)^G = (\mathrm{Sym}^2(V^*))^G \oplus (\wedge^2(V^*))^G.$$

But we have seen above that when $((V \otimes V)^*)^G$ is nonzero it is 1-dimensional (i.e., a nonzero equivariant B is unique up to $\mathbf{C}^\times$-scaling if it exists), so when the left side is nonzero we conclude that *exactly one* of the two summands on the right side is 1-dimensional and the other vanishes. That is, if there exists a nonzero equivariant B then it is either symmetric or alternating and there is no nonzero B' of the other type. To relate these conditions to χ_V, we need to compute the characters of Sym^2 and $\wedge^2$:

Lemma 2.2. *For any* G*-representation* (W, ρ), $\chi_{\mathrm{Sym}^2(W)}(g) = (1/2)(\chi_W(g)^2 + \chi_W(g^2))$ *and* $\chi_{\wedge^2(W)}(g) = (1/2)(\chi_W(g)^2 - \chi_W(g^2))$.

Proof. Consider an upper-triangular form for $\rho(g)$, say with eigenvalues $\lambda_1, \dots, \lambda_n$ with multiplicity. Thus, $\chi := \chi_W$ satisfies $\chi(g) = \sum_j \lambda_j$ and $\chi(g^2) = \sum_j \lambda_j^2$. Computing the symmetric and exterior squares relative to the same bases, it is easy to see that we again get an upper triangular form for the action of g, and more specifically that (listing with multiplicity) its eigenvalues on $\mathrm{Sym}^2(W)$ are $\{\lambda_i\lambda_j\}_{i \le j}$ and its eigenvalues on $\wedge^2(W)$ are $\{\lambda_i\lambda_j\}_{i<j}$. Thus, the asserted identities say exactly

$$\sum_{i \le j} \lambda_i\lambda_j \overset{?}{=} (1/2)((\sum_j \lambda_j)^2 + \sum_j \lambda_j^2), \quad \sum_{i<j} \lambda_i\lambda_j \overset{?}{=} (1/2)((\sum_j \lambda_j)^2 - \sum_j \lambda_j^2),$$

both of which are trivial to verify. ∎

The dimension of the space of G-invariants in a representation space with character ψ is $\langle \psi, 1\rangle_G = (1/\#G)\sum_{g\in G}\psi(g)$. Thus, by using the Lemma for $W = V^*$, averaging over all $g \in G$, and applying complex conjugation to the two resulting averaged formulas, we conclude that the case when no nonzero equivariant B exists is exactly when

$$a_+ := \frac{1}{2\cdot \#G}\sum_{g\in G}(\chi(g)^2 + \chi(g^2)) = 0,\ \ a_- := \frac{1}{2\cdot \#G}\sum_{g\in G}(\chi(g)^2 - \chi(g^2)) = 0,$$

whereas if a nonzero B exists then it is non-degenerate symmetric precisely when $a_+ = 1$ and $a_- = 0$, and it is non-degenrate skew-symmetric precisely when $a_+ = 0$ and $a_- = 1$. Since $\varepsilon_V = a_+ - a_-$, we see by checking each case that the final asserted equivalences in (i), (ii), and (iii) are established (and that no other values for ε_V are possible).

We have already shown that the first and last conditions in (i) are equivalent, so to complete the proof of the equivalences in (i) we have to relate the condition $\mathrm{End}_{\mathbf{R}[G]}(V) = \mathbf{C}$ to the rest. The $\mathbf{C}$-structure on V defines an inclusion $\mathbf{C} \subset \mathrm{End}_{\mathbf{R}[G]}(V)$, so we are exploring when this is an equality, and that is just a dimenson condition over $\mathbf{R}$. Such a dimension can be calculated after scalar extension to $\mathbf{C}$, and (as in the hint to Exercise 5(ii) in HW8)

$$\mathbf{C}\otimes_{\mathbf{R}} \mathrm{End}_{\mathbf{R}[G]}(V) = \mathrm{End}_{\mathbf{C}[G]}(\mathbf{C}\otimes_{\mathbf{R}} V).$$

As $\mathbf{C}[G]$-modules we have

$$\mathbf{C}\otimes_{\mathbf{R}} V = \mathbf{C}\otimes_{\mathbf{R}} (\mathbf{C}\otimes_{\mathbf{C}} V) = (\mathbf{C}\otimes_{\mathbf{R}} \mathbf{C})\otimes_{\mathbf{C}} V,$$

and the map $\mathbf{C}\otimes_{\mathbf{R}}\mathbf{C} \simeq \mathbf{C}\otimes\mathbf{C}$ via $a\otimes b \mapsto (ab, a\overline{b})$ is an isomorphism of $\mathbf{C}$-algebras when using the $\mathbf{C}$-structure on the *left* tensor factor whereas it transports the $\mathbf{C}$-structure from the right tensor factor over to the one on $\mathbf{C}\times\mathbf{C}$ that is the usual one on the first factor and the action through conjugation on the second factor. Thus,

$$\mathbf{C}\otimes_{\mathbf{R}} V \simeq V \oplus \overline{V}$$

as $\mathbf{C}[G]$-modules, so

$$\mathrm{End}_{\mathbf{C}[G]}(\mathbf{C}\otimes_{\mathbf{R}} V) = \mathrm{End}_{\mathbf{C}[G]}(V\oplus\overline{V}).$$

In case $V \not\simeq \overline{V}$ (i.e., χ is not $\mathbf{R}$-valued), this collapses to $\mathrm{End}_{\mathbf{C}[G]}(V)\times\mathrm{End}_{\mathbf{C}[G]}(\overline{V}) = \mathbf{C}\times\mathbf{C}$ by Schur's Lemma. This is 2-dimensional over $\mathbf{C}$. In case $V \simeq \overline{V}$ (i.e., χ is $\mathbf{R}$-valued), this becomes

$$\mathrm{End}_{\mathbf{C}[G]}(V\oplus V) = \mathrm{Mat}_2(\mathrm{End}_{\mathbf{C}}[G](V)) = \mathrm{Mat}_2(\mathbf{C})$$

by Schur's Lemma and so we get a 4-dimensional endomorphism algebra. In particular, when $\mathrm{End}_{\mathbf{R}[G]}(V)$ is not 4-dimensional we *must* be in case (i). This completes the proof of the equivalences in (i), and shows that $\mathrm{End}_{\mathbf{R}[G]}(V)$ is 4-dimensional over $\mathbf{R}$ if and only if χ is $\mathbf{R}$-valued.

3. The R-valued case

In view of what has already been done, it remains to focus on the cases with $\mathbf{R}$-valued χ (so $\mathrm{End}_{\mathbf{R}[G]}(V)$ is 4-dimensional over $\mathbf{R}$) and to prove that each of the first two conditions in (ii) and (iii) is equivalent to the third and fourth conditions (whose equivalence has already been shown).

First assume V is defined over $\mathbf{R}$, say $V = V_0\otimes_{\mathbf{R}}\mathbf{C}$ for some $\mathbf{R}[G]$-module V_0. Thus, the inclusion $\mathbf{R}\subset\mathrm{End}_{\mathbf{R}[G]}(V_0)$ is an equality because it becomes an equality after scalar extension to $\mathbf{C}$ (thanks to Schur's Lemma), and as an $\mathbf{R}[G]$-module we have $V = V_0\oplus V_0$, so

$$\mathrm{End}_{\mathbf{R}[G]}(V) = \mathrm{End}_{\mathbf{R}[G]}(V_0\oplus V_0) = \mathrm{Mat}_2(\mathrm{End}_{\mathbf{R}[G]}(V_0)) = \mathrm{Mat}_2(\mathbf{R}).$$

It follows that the first condition in (ii) implies the second.

Conversely, the second condition in (ii) implies the first. Indeed, if $\mathrm{End}_{\mathbf{R}[G]}(V) = \mathrm{Mat}_2(\mathbf{R})$ then this is not a division algebra and hence V *cannot be irreducible* as an $\mathbf{R}[G]$-module. Thus, we can pick a nonzero proper $\mathbf{R}[G]$-submodule $V_0\subset V$. Using the G-equivariant $\mathbf{C}$-linear structure on V, we thereby get a canonical *nonzero* map of $\mathbf{C}[G]$-modules

$$T : \mathbf{C}\otimes_{\mathbf{R}} V_0 \to V,$$

so this must be surjective since V is irreducible. As an $\mathbf{R}[G]$-module the left side is $V_0 \oplus V_0$, so we conclude that its quotient $\mathbf{R}[G]$-module V must be V_0-isotypic with multiplicity at most 2. But $\dim_{\mathbf{R}} V_0 < \dim_{\mathbf{R}} V$, so the multiplicity must be exactly 2 and so T must be an isomorphism. This shows that V_0 is an $\mathbf{R}[G]$-module descent of V; i.e., V is defined over $\mathbf{R}$.

To summarize, the first two conditions in (ii) are equivalent and the last two conditions in (ii) are equivalent. We shall now show that the first condition in (ii) implies the fourth. Pick any positive-definite inner product $\langle \cdot, \cdot \rangle$ on V_0, and then average it to achieve G-equivariance:

$$\langle \cdot, \cdot \rangle' := (1/\#G) \sum_{g \in G} \langle g(\cdot), g(\cdot) \rangle.$$

This average is still positive-definite and hence also *non-degenerate*. It is also symmetric, so extending scalars to $\mathbf{C}$ then gives B as in the fourth condition of (ii). Thus, to finish the proof of the equivalences in (ii) it suffices to show that the fourth condition implies the first. This involves a more serious argument in linear algebra over $\mathbf{R}$, and we refer to the final paragraph of the proof of Theorem 31 in §13.2 of Serre's book *Linear representations of finite groups* for this argument (due to Frobenius and Schur).

With the proof of the equivalences in (ii) finished, the equivalences in (iii) essentially follow immediately by exhaustion of possibilities as follows. We already know the third and fourth conditions are equivalent, and that the fourth condition in (iii) accounts for exactly those situations when the fourth conditions in (i) and (ii) fail. Consequently, it also accounts for exactly those situations when the first conditions in (i) and (ii) fail, as well as exactly those situations when the second conditions in (i) and (ii) fail. But the complement of the union of the first conditions in (i) and (ii) is exaclty the first condition in (iii), so the equivalence of the fourth (and hence third) condition in (iii) with the first is established.

When $\mathrm{End}_{\mathbf{R}[G]}(V) = \mathbf{H}$ we cannot be in any of the situations covered by the equivalent conditions in (i), nor in (ii), so by tracking the fourth conditions (or the third, or the first) throughout we see that the fourth condition in (iii) must hold (or likewise for the third or first conditions in (iii)). Hence, it remains to show that when the equivalent first, third, and fourth conditions hold in (iii) then necessarily $\mathrm{End}_{\mathbf{R}[G]}(V) \simeq \mathbf{H}$. The analysis of case (ii) actually showed that if V is reducible as an $\mathbf{R}[G]$-module then we must be in case (ii) (and more specifically, that the first condition in (ii) holds). Thus, we know that V must be irreducible as an $\mathbf{R}[G]$-module, so $\mathrm{End}_{\mathbf{R}[G]}(V)$ is a division algebra. We have also seen that this endomorphism algebra must be 4-dimensional over $\mathbf{R}$ (using that χ is $\mathbf{R}$-valued). Now to conclude we have to bring in a theorem of Frobenius: the only finite-dimensional division algebras over $\mathbf{R}$ are: $\mathbf{R}$, $\mathbf{C}$, and $\mathbf{H}$. In particular, the *only* 4-dimensional one is $\mathbf{H}$!

[Frobenius' theorem is proved in many textbooks via a variety of elementary methods; also see the Wikipedia page for "Frobenius theorem (real division algebras)" for such a proof. Those elementary proofs tend to be somewhat gritty, and ultimately not so illuminating. There is a short conceptual proof in terms of general principles in group cohomology applicable to all fields, the special features of $\mathbf{R}$ being that $\mathrm{Gal}(\mathbf{C}/\mathbf{R})$ is cyclic and the norm map $\mathrm{N}_{\mathbf{C}/\mathbf{R}} : \mathbf{C}^\times \to \mathbf{R}^\times$ from the multiplicative group of its algebraic closure $\mathbf{C}$ has image of index 2 in $\mathbf{R}^\times$. This conceptual argument rests on the double-periodicity of Tate cohomology of cyclic groups.]