1

Introduction

The purpose of this book is to introduce equivariant stable homotopy theory in a way that will make the methods of Hill, Hopkins and Ravenel (2016) accessible to a well-informed graduate student and facilitate further research in this area.

The research leading to Hill, Hopkins and Ravenel (2016) was an example of the aphorism "computation precedes theory." In 2005, Hopkins and Ravenel set out to study the homotopy fixed point sets of finite subgroups of the Morava stabilizer groups $\mathbf{S}_n$ under their action on the Morava spectra E_n; Hill joined us a short time later. We knew this would be an interesting project, but we did not anticipate that it would lead to a solution to the Kervaire invariant problem, named after Michel Kervaire (1927–2007). We like to say we went hiking in the Alps and found a shortcut up Mount Everest.

After making various assumptions about how things work in equivariant stable homotopy theory, we did the computation that led to our main theorem. Upon further reflection, we realized that the existing literature on the subject did not provide an adequate framework for our calculations. This led to the lengthy appendices in Hill, Hopkins and Ravenel (2016) providing the necessary theoretical infrastructure. Despite their length, they were written as tersely as possible so as to economize journal space.

A similar account will be given here at a more leisurely pace, with more than 150 examples illustrating various concepts. In particular, we do our best to motivate the definition of the model structure we need on the category of equivariant orthogonal spectra, the subject of Chapter 9.

Other works called *Equivariant Stable Homotopy Theory* are Greenlees and May (1995), Lewis et al. (1986) and Segal (1971), and the phrase occurs in numerous other titles.

Nearly every item in the References can be found in the third author's online archive: https://people.math.rochester.edu/faculty/doug/papers.html.

There is a Table of Notations at the end of the book in Part THREE for the reader's convenience.

1.1 The Kervaire Invariant Theorem and the Ingredients of Its Proof

Very briefly, the Kervaire invariant problem concerns the fate of the elements h_j^2 in the classical Adams spectral sequence at the prime 2, originally introduced by J. Frank Adams FRS (1930–1989) in Adams (1958). We refer the reader to Ravenel (2004) for a description of it. A theorem of William Browder (1969) says that h_j^2 is a permanent cycle if and only if there exists a framed manifold of dimension $2^{j+1}-2$ with nontrivial Kervaire invariant. The hypothetical element in $\pi_{2^{j+1}-2}^S$ represented by such a framed manifold is denoted by θ_j.

Here π_k^S denotes the stable k-stem, the value of $\pi_{n+k}S^n$ for large n. It is also the kth homotopy group of the sphere spectrum, which was often denoted by S^0 in early works on the subject. *In this book, we will denote the sphere spectrum by S^{-0} to avoid confusion with the space S^0*; see Remark 1.4.13.

After the publication of Browder's theorem in 1969, there were numerous unsuccessful attempts to prove the existence of θ_j for all $j > 0$. Mark Mahowald (1931–2013) named his sailboat "Thetajay." His colleague and coauthor Michael Barratt (1927–2015) referred to the possibility that they did not all exist as the "Doomsday Hypothesis." More precisely, he gave this name to conjecture, originally due to Joel Cohen (1970), that in the Adams spectral sequence, only a finite number of elements in each filtration were permanent cycles. The first five θ_j were known to exist, the construction of θ_5 being the subject of Barratt, Jones and Mahowald (1984) and recently simplified in Xu (2016).

After 1980, interest in the problem faded as the failed attempts of the 1970s convinced the homotopy theory community that it was beyond its reach. In 2009, just before we announced our theorem, Victor Snaith published Snaith (2009), a witty account of the state of the art at that moment. Three of his statements are worth repeating here.

About the decline of interest in the problem, he said,

> As ideas for progress on a particular mathematics problem atrophy it can disappear. Accordingly I wrote this book to stem the tide of oblivion.

About his own involvement in it, he wrote,

> For a brief period overnight we were convinced that we had the method to make all the sought after framed manifolds – a feeling which must have been shared by many topologists working on this problem. All in all, the temporary high of believing that one had the construction was sufficient to maintain in me at least an enthusiastic spectator's interest in the problem.

Best of all,

> In the light of the above conjecture and the failure over fifty years to construct framed manifolds of Arf–Kervaire invariant one this might turn out to be a book about things which do not exist. This [is] why the quotations which preface each chapter contain a preponderance of utterances from the pen of Lewis Carroll.

1.1A The Main Theorem

Indeed the sought after framed manifolds (with a small number of exceptions) *do not exist*. The following was first announced by the second author in April 2009, in a lecture at a conference in Edinburgh honoring the 80th birthday of Sir Michael Atiyah (1929–2019).

Main Theorem *The Arf–Kervaire elements $\theta_j \in \pi_{2^{j+1}-2}^S$ do not exist for $j \geqslant 7$.*

The status of θ_6 in the 126-stem remains open.

Figure 1.1 Fenway's dream.

In Ravenel (1978) (see also Ravenel (2004, §6.4)) the third author showed long ago that the cohomology of the subgroup of order p in $\mathbf{S}_{p-1}$ could be used to show that odd primary analogs of the Kervaire invariant elements do not exist for $p \geqslant 5$.

Here $\mathbf{S}_n$ denotes the nth *Morava stabilizer group*, which plays a critical role in chromatic homotopy theory. We refer the reader to Ravenel (2004, Chapter 6) for its definition and properties. It is a pro-p-group that is the strict automorphism group of a height n formal group law over a sufficiently large finite field of characteristic p. Its cohomology in some sense controls the nth chromatic layer of the Adams–Novikov E_2-term, as explained first in Miller, Ravenel and Wilson (1977) and later in Ravenel (2004, Chapter 5). It is known to have elements of order p^{i+1} precisely when $(p-1)p^i$ divides n. In particular, $\mathbf{S}_{p-1}$ has a cyclic subgroup of order p, and for $p = 2$, $\mathbf{S}_4$ has one of order 8.

This odd primary Kervaire invariant problem was easier (and hence solved 30 years earlier) than the 2-primary case because Hirosi Toda (1967, 1968) had shown a decade earlier that $\theta_2 \in \pi^S_{2p^2(p-1)-2}$ does not exist. This could be reinterpreted as a proof that the corresponding element in the Adams–Novikov spectral sequence, $\beta_{p/p}$, supports a nontrivial differential hitting $\alpha_1\theta_1^p = \alpha_1\beta_1^p$. The cohomology of $C_p \subseteq \mathbf{S}_{p-1}$ then provided a way to leverage this into a proof that $\theta_j = \beta_{p^{j-1}/p^{j-1}}$ supports a differential hitting $\alpha_1\theta_{j-1}^p$ for all $j \geqslant 2$.

At the prime 2, there was no analog of Toda's theorem; there was no θ_j that was known not to exist. We also know that while the θ_j themselves can be detected in the cohomology of $C_8 \subseteq \mathbf{S}_4$, their products cannot be. This means that the leverage of Ravenel (1978) is not available. The methods of Hill, Hopkins and Ravenel (2016), which include the use of equivariant stable homotopy theory, are quite different.

We have a much simpler way of defining the action of the group C_8. In chromatic homotopy theory (for background on this topic, see Lurie's 2010 Harvard course, Lurie

(2010), Barthel and Beaudry (2020) with its numerous references, Ravenel (2004) and Ravenel (1992)) we learn that $\mathbf{S}_n$, the strict automorphism group of a height n formal group law F_n over the field $\mathbf{F}_{p^n}$, acts on the ring over which its universal deformation (lifting to characteristic zero) is defined. The same goes for $\mathbf{G}_n$, the extension of $\mathbf{S}_n$ by the Galois group of $\mathbf{F}_{p^n}$ over $\mathbf{F}_p$. This ring turns out to be $\pi_0 E_n$, where E_n is the nth Morava E-theory, a variant of the Johnson–Wilson spectrum $E(n)$. These considerations leads to an "action" of $\mathbf{S}_n$ on the spectrum E_n, but it is only defined *up to homotopy*.

This awkward state of affairs was the motivating issue for the Goerss–Hopkins–Miller theorem in the early 1990s; see Rezk (1998) and Goerss and Hopkins (2004). Morava's E_n was known to be an E_∞-ring spectrum, meaning that it has a multiplication that is homotopy commutative in the strongest possible sense. They showed that for an E_∞-ring spectrum R, there is a *space* of E_∞-ring automorphisms $Aut(R)$. This required a deeper understanding of the stable homotopy category than was prevalent at the time. In the case of $R = E_n$, we knew that the set of path components of this space had to be $\mathbf{G}_n$. *They showed that each path component is contractible.*

This means that $Aut(E_n)$ is homotopy equivalent to $\mathbf{G}_n$ and that, for any closed subgroup $G \subseteq \mathbf{G}_n$, one can define the homotopy fixed point spectrum E_n^{hG}. In particular, $E_n^{h\mathbf{G}_n} = L_{K(n)} S^0$, the Bousfield localization of the sphere spectrum with respect to the nth Morava K-theory. The calculation of Ravenel (1978) could be reinterpreted as a calculation with $E_{p-1}^{hC_p}$.

The proof of this gratifying result is quite technical. *Fortunately, we do not have to deal with it here.* We have a much more direct way of mapping $\pi_* S^0$ to the cohomology of a cyclic 2-group using equivariant stable homotopy theory.

1.1B The Equivariant Approach

The starting point is the action of C_2 on the complex cobordism spectrum MU via complex conjugation. The resulting C_2-spectrum is denoted by $MU_{\mathbf{R}}$, and known as "real cobordism."

This terminology derives from Atiyah's definition of real K-theory in Atiyah (1966). (The reader hoping for a definition of "reality" as a technical term will be disppointed to find that the word only appears in the title of the paper.) For him, a "real" space is a topological space X equipped with an involution τ. For $x \in X$, he denotes $\tau(x)$ by $\overline{x}$. A "real" vector bundle E over a real space X was not a bundle of real vector spaces but a complex vector bundle equipped with an involution compatible with that on X such that the induced map from the fiber over x to that over $\overline{x}$ is conjugate linear.

A key example of a real space is the set of complex points of an algebraic variety X defined over the real numbers, which comes equipped with an involution related to complex conjugation. Its fixed point set is the space of real points of X. In particular, X could be the Grassmannian variety $G_{n,k}$, whose real and complex points are respectively the spaces of linear k-dimensional subspaces of an n-dimensional vector space over the real and complex numbers. Taking the colimit as n and k go to infinity, we get the classifying space BU

equipped with an involution induced by complex conjugation. We denote this object by $BU_{\mathbf{R}}$. We can Thomify this and get a C_2-equivariant spectrum $MU_{\mathbf{R}}$, the *real cobordism spectrum*. Its precise construction is the subject of Chapter 12. It was first studied by Peter Landweber in Landweber (1968) and subsequently by Michikazu Fujii (1975/76), Shôrô Araki (1930–2005) (1979) and Po Hu and Igor Kriz (2001).

The next step is to elevate the C_2-spectrum $MU_{\mathbf{R}}$ to a C_{2^n}-spectrum. More generally, when H is a subgroup of G, we define a norm functor N_H^G from the category of H-spectra to that of G-spectra; see Definition 9.7.3. Roughly speaking, for an H-spectrum E, the G-spectrum $N_H^G E$ is $E^{\wedge |G/H|}$ with G permuting the H-invariant factors. A recent theorem of Jeremy Hahn and XiaoLin Danny Shi (2017) implies that there is a map $N_{C_2}^{C_{2^n}} MU_{\mathbf{R}} \to E_{2^n-1}$ that is equivariant with respect to the action of C_{2^n} as a subgroup of $\mathbf{S}_{2^n-1}$.

Classically, there is a way to derive Atiyah's real K-theory spectrum $K_{\mathbf{R}}$ from $MU_{\mathbf{R}}$, and the former is 8-periodic, meaning that $\pi_i K_{\mathbf{R}}$ and its equivariant variants only depend on the congruence class of i modulo 8. It is a retract of a mapping telescope obtained from $MU_{\mathbf{R}}$ by inverting a certain element in its equivariant homotopy group.

There are similar spectra $K_{\mathbf{H}}$ and $K_{\mathbf{O}}$ that are retracts of telescopes related to $N_{C_2}^{C_4} MU_{\mathbf{R}}$ and $N_{C_2}^{C_8} MU_{\mathbf{R}}$ that are respectively 32- and 256-periodic. The use of the symbols **H** and **O** here is purely a matter of convenience, as these spectra have very little to do with the quaternions or octonions. The spectrum $K_{\mathbf{H}}$ is studied extensively in Hill, Hopkins and Ravenel (2017c), where it and $K_{\mathbf{O}}$ are denoted by $K_{[2]}$ and $K_{[3]}$.

There is a similar telescope associated with $N_{C_2}^{C_{2^n}} MU_{\mathbf{R}}$ for each $n \geqslant 1$. It is obtained by inverting an element D specified for the case $n = 3$ in Corollary 13.3.25. Theorem 13.3.23 shows that it has periodicity $2^{n+1+2^{n-1}}$. Passing from the telescope to its retract $K_{[n]}$ simplifies explicit calculations of homotopy groups but is not needed for our current purposes.

1.1C The Spectrum Ξ

The fixed point spectrum Ξ (denoted by Ω in Hill, Hopkins and Ravenel, 2016) of the telescope for $N_{C_2}^{C_8} MU_{\mathbf{R}}$, which we denote by $\Xi_{\mathbf{O}}$, is the central object in the solution to the Kervaire invariant problem. It is a nonconnective ring spectrum with a unit map $S^0 \to \Xi$. It has the following properties:

Key Properties of the C_8 Fixed Point Spectrum Ξ

(i) **Detection Theorem.** *It has an Adams–Novikov spectral sequence (which is a device for calculating homotopy groups) in which the image of each θ_j is nontrivial.* This means that if θ_j exists, we will see its image in $\pi_*(\Xi)$.

(ii) **Periodicity Theorem.** *It is 256-periodic, meaning that $\pi_k(\Xi)$ depends only on the reduction of k modulo 256. As in the case of Bott periodicity, we have a stable equivalence $\Omega^{256}\Xi \simeq \Xi$.*

(iii) **Gap Theorem.** $\pi_k(\Xi) = 0$ *for* $-4 < k < 0$.

These will be proved in Chapter 13, after developing the necessary machinery in the intervening 11 chapters. We will identify Ξ in Definition 13.3.27. Property (iii) is our zinger. Its proof involves a new tool we call the *slice spectral sequence*.

If $\theta_7 \in \pi_{254}(S^0)$ exists, (i) implies it has a nontrivial image in $\pi_{254}(\Xi)$. On the other hand, (ii) and (iii) imply that $\pi_{254}(\Xi) = 0$, so θ_7 cannot exist. The argument for θ_j for larger j is similar, since $|\theta_j| = 2^{j+1} - 2 \equiv -2 \bmod 256$ for $j \geqslant 7$. (Historical note: the third author spent part of his undergraduate career living in a rented room at 254 Elm Street near Oberlin College. It was there that he first became acquainted with homotopy theory, but at that time, he did not appreciate the significance of his street number. In 2002, he lived in a rented house at 62 Eden Street in Cambridge, UK.)

At the present time, the three theorems listed above are just about *all* we know about Ξ, which is just enough to prove the main theorem. If we could show that $\pi_{126}\Xi = 0$, we would know that θ_6 does not exist. This appears to be a daunting calculation. We computed $\pi_* K_{\mathbf{H}}^{C_4}$ in Hill, Hopkins and Ravenel (2017c) as a warm-up exercise for it.

The reader may wonder *why we chose the group* C_8. Briefly, the argument for the Detection Theorem, §1.1C (i), would break down were we to use C_2 or C_4. We will say more about this in §13.4, specifically in Remark 13.4.18. It would go through for any larger cyclic 2-group, but the period would be greater, which would lead to a weaker theorem. For C_{16}, the period is 8192, so the resulting theorem would say that θ_j does not exist for $j \geqslant 12$ rather than for $j \geqslant 7$. The Gap Theorem holds for any cyclic 2-group.

1.2 Background and History

1.2A Pontryagin's Early Work on Homotopy Groups of Spheres

The Arf–Kervaire invariant problem has its origins in the early work of Lev Pontryagin (1908–1988) on a geometric approach to the homotopy groups of spheres (Pontryagin, 1938; 1950 and 1955).

Pontryagin's approach to maps $f\colon S^{n+k} \to S^n$ is to assume that f is smooth and that the base point y_0 of the target is a regular value. (Any continuous f can be continuously deformed to a map with this property.) This means that $f^{-1}(y_0)$ is a closed smooth k-manifold M in S^{n+k}. Let D^n be the closure of an open ball around y_0. If it is sufficiently small, then $V^{n+k} = f^{-1}(D^n) \subset S^{n+k}$ is an $(n+k)$-manifold homeomorphic to $M \times D^n$ with boundary homeomorphic to $M \times S^{n-1}$. It is also a tubular neighborhood of M^k and comes equipped with a map $p\colon V^{n+k} \to M^k$ sending each point to the nearest point in M. For each $x \in M$, $p^{-1}(x)$ is homeomorphic to a closed n-ball B^n. The pair $(p, f|V^{n+k})$ defines an explicit homeomorphism

$$V^{n+k} \xrightarrow[\approx]{(p,f|V^{n+k})} M^k \times D^n.$$

This structure on M^k is called a *framing*, and M is said to be *framed in* $\mathbf{R}^{n+k}$. A choice of basis of the tangent space at $y_0 \in S^n$ pulls back to a set of linearly independent normal vector fields on $M \subset \mathbf{R}^{n+k}$. These will be indicated in Figures 1.2–1.3.

Conversely, suppose we have a closed sub-k-manifold $M \subset \mathbf{R}^{n+k}$ with a closed tubular neighborhood V and a homeomorphism h to $M \times D^n$ as above. This is called a *framed sub-k-manifold* of $\mathbf{R}^{n+k}$. Some remarks are in order here.

- The existence of a framing puts some restrictions on the topology of M. All of its characteristic classes must vanish. In particular, it must be orientable.
- A framing can be twisted by a map $g\colon M \to SO(n)$, where $SO(n)$ denotes the group of orthogonal $n \times n$ matrices with determinant 1. Such matrices act on D^n in an obvious way. The twisted framing is the composite

$$V \xrightarrow{h} M^k \times D^n \longrightarrow M^k \times D^n$$
$$(m, x) \longmapsto (m, g(m)(x)).$$

When $M^k = S^k$, this leads to the Hopf–Whitehead J-homomorphism of Remark 1.2.2.
- If we drop the assumption that M is framed, then the tubular neighborhood V is a (possibly nontrivial) disk bundle over M. The map $M \to y_0$ needs to be replaced by a map to the classifying space for such bundles, $BO(n)$. This leads to unoriented bordism theory, which was analyzed by René Thom (1923–2002) in Thom (1954). Two helpful references for this material are the books by Milnor and Stasheff (1974) and Robert Stong (1936–2008) (Stong, 1968a).

Pontryagin constructs a map $P(M,h)\colon S^{n+k} \to S^n$ as follows. We regard S^{n+k} as the one-point compactification of $\mathbf{R}^{n+k}$ and S^n as the quotient $D^n/\partial D^n$. This leads to the following diagram.

$$\begin{array}{ccccc}
(V, \partial V) & \xrightarrow{h} & M \times (D^n, \partial D^n) & \xrightarrow{p_2} & (D^n, \partial D^n) \\
\downarrow & & & & \downarrow \\
(\mathbf{R}^{n+k}, \mathbf{R}^{n+k} - \mathrm{int} V) & \longrightarrow & (S^{n+k}, S^{n+k} - \mathrm{int} V) & \xrightarrow{P(M,h)} & (S^n, \{\infty\}).
\end{array}$$

The map $P(M,h)$ is the extension of $p_2 h$ obtained by sending the compliment of V in S^{n+k} to the point at infinity in S^n. For $n > k$, the choice of the embedding (but not the choice of framing) of M into the Euclidean space is irrelevant. Any two embeddings (with suitably chosen framings) lead to the same map $P(M,h)$ up to continuous deformation.

To proceed further, we need to be more precise about what we mean by continuous deformation. Two maps $f_1, f_2\colon X \to Y$ are *homotopic* if there is a continuous map $h : X \times [0,1] \to Y$ (called a *homotopy between f_1 and f_2*) such that

$$h(x,0) = f_1(x) \qquad \text{and} \qquad h(x,1) = f_2(x).$$

Now suppose $X = S^{n+k}$, $Y = S^n$, and the map h (and hence f_1 and f_2) is smooth with y_0 as a regular value. Then $h^{-1}(y_0)$ is a framed $(k+1)$-manifold N whose boundary is the disjoint union of $M_1 = f^{-1}(y_0)$ and $M_2 = g^{-1}(y_0)$. This N is called a *framed cobordism* between M_1 and M_2, and when it exists, the two closed manifolds are said to be *framed cobordant*. An example is shown in Figure 1.2.

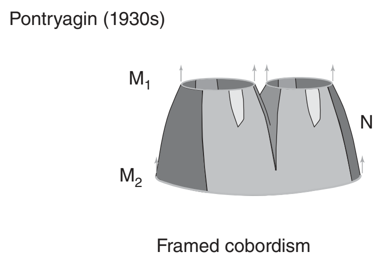

Figure 1.2 A framed cobordism between $M_1 = S^1 \coprod S^1 \subset \mathbf{R}^2$ and $M_2 = S^1 \subset \mathbf{R}^3$ with $N \subset [0,1] \times \mathbf{R}^2$. The normal framings on the circles can be chosen so they extend over N.

Let $\Omega^{\mathrm{fr}}_{k,n}$ denote the cobordism group of framed k-manifolds in $\mathbf{R}^{n+k}$. The above construction leads to Pontryagin's isomorphism

$$\Omega^{\mathrm{fr}}_{k,n} \xrightarrow{\approx} \pi_{n+k}(S^n).$$

First consider the case $k = 0$. Here the 0-dimensional manifold M is a finite set of points in $\mathbf{R}^n$. Each comes with a framing that can be obtained from a standard one by an element in the orthogonal group $O(n)$. We attach a sign to each point corresponding to the sign of the associated determinant. With these signs, we can count the points algebraically and get an integer called the *degree of f*. Two framed 0-manifolds are cobordant if and only if they have the same degree.

Now consider the case $k = 1$. M is a closed 1-manifold – i.e., a disjoint union of circles. Two framings on a single circle differ by a map from S^1 to the group $SO(n)$, and it is known that

$$\pi_1(SO(n)) = \begin{cases} 0 & \text{for } n = 1 \\ \mathbf{Z} & \text{for } n = 2 \\ \mathbf{Z}/2 & \text{for } n > 2. \end{cases}$$

It turns out that any disjoint union of framed circles is cobordant to a single framed circle. This can be used to show that

$$\pi_{n+1}(S^n) = \begin{cases} 0 & \text{for } n = 1 \\ \mathbf{Z} & \text{for } n = 2 \\ \mathbf{Z}/2 & \text{for } n > 2. \end{cases}$$

The case $k = 2$ is more subtle. As in the 1-dimensional case, we have a complete classification of closed 2-manifolds, and it is only necessary to consider path connected ones. The existence of a framing implies that the surface is orientable, so it is characterized by its genus.

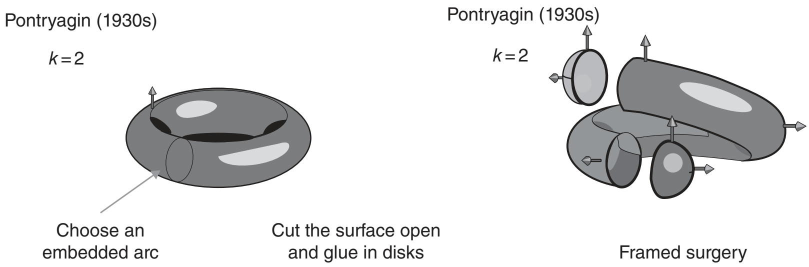

Figure 1.3 The case $k = 2$ and genus 1. If the framing on the embedded arc extends to a disk, then there is a cobordism (called a framed surgery) that converts the torus to a 2-sphere, as shown.

If the genus is zero, namely if $M = S^2$, then there is a framing that extends to a 3-dimensional ball. This makes M cobordant to the empty set, which means that the map is *null homotopic* (or, more briefly, *null*), meaning that it is homotopic to a constant map. Any two framings on S^2 differ by an element in $\pi_2(SO(n))$. This group is known to vanish, so any two framings on S^2 are equivalent, and the map $f: S^{n+2} \to S^n$ is null.

Now suppose the genus is one, as shown in Figure 1.3. Suppose we can find an embedded arc as shown on which the framing extends to a disk. Then there is a cobordism that effectively cuts along the arc and attaches two disks, as shown. This process is called *framed surgery*. If we can do this, then we have converted the torus to a 2-sphere, and we have shown that the map $f: S^{n+2} \to S^n$ is null.

When can we find such a closed curve in M? It must represent a generator of $H_1(M)$ and carry a trivial framing. This leads to a map

(1.2.1) $$\varphi: H_1(M; \mathbf{Z}/2) \to \mathbf{Z}/2,$$

defined as follows. Each class in H_1 can be represented by a closed curve that is framed either trivially or nontrivially. It can be shown that homologous curves have the same framing invariant, so φ is well defined. At this point, Pontryagin made a famous mistake which went undedected for over a decade: *he assumed that φ was a homomorphism*. We now know this is not the case, and we will say more about it in §1.2C.

On that basis he argued that φ must have a nontrivial kernel, since the source group is $(\mathbf{Z}/2)^2$. Therefore, there is a closed curve along which we can do the surgery shown in Figure 1.3. It follows that M can be surgered into a 2-sphere, leading to the erroneous conclusion that $\pi_{n+2}(S^n) = 0$ for all n. Freudenthal (1938) and later George Whitehead (1950) both proved that it is $\mathbf{Z}/2$ for $n \geqslant 2$. Pontryagin corrected his mistake in Pontryagin (1950), and in Pontryagin (1955) he gave a complete account of the relation between framed cobordism and homotopy groups of spheres.

Remark 1.2.2 (The Hopf–Whitehead J-homomorphism) Suppose our framed manifold is S^k with a framing that extends to a D^{k+1}. This will lead to the trivial element in $\pi_{n+k}(S^n)$, but twisting the framing can lead to nontrivial elements. The twist is determined up to homotopy by an element in $\pi_k(SO(n))$. Pontryagin's construction thus leads to the homomorphism

$$\pi_k(SO(n)) \xrightarrow{J} \pi_{n+k}(S^n),$$

introduced by Hopf (1935) and Whitehead (1942). Both source and target are known to be independent of n for $n > k+1$.

In this case, the source group for each k (denoted simply by $\pi_k(SO)$ since n is irrelevant) was determined by Bott (1959) in his remarkable periodicity theorem. He showed

$$\pi_k(SO) = \begin{cases} \mathbf{Z} & \text{for } k \equiv 3 \text{ or } 7 \bmod 8 \\ \mathbf{Z}/2 & \text{for } k \equiv 0 \text{ or } 1 \bmod 8 \\ 0 & \text{otherwise.} \end{cases}$$

Here is a table showing these groups for $k \leqslant 10$.

k	1	2	3	4	5	6	7	8	9	10
$\pi_k(SO)$	$\mathbf{Z}/2$	0	$\mathbf{Z}$	0	0	0	$\mathbf{Z}$	$\mathbf{Z}/2$	$\mathbf{Z}/2$	0

In each case where the group is nontrivial, the image under J of its generator is known to generate a direct summand; see Adams (1966, Theorems 1.1, 1.3, 1.5 and 1.6). In the jth case, we denote this image by β_j and its dimension by $\phi(j)$, which is roughly $2j$. (They will figure in Hypothesis 1.2.4.) The first three of these are the Hopf maps $\eta \in \pi_1^S$, $\nu \in \pi_3^S$ and $\sigma \in \pi_7^S$. After that, we have $\beta_4 \in \pi_8^S$, $\beta_5 \in \pi_9^S$, $\beta_6 \in \pi_{11}^S$, and so on.

For the case $\pi_{4m-1}(SO) = \mathbf{Z}$, the image under J is known to be a cyclic group whose order a_m is the denominator of $B_m/4m$, where B_m is the mth Bernoulli number. Details can be found in Adams (1966, Theorems 1.5 and 1.6) and Milnor and Stasheff (1974, Appendix B). Here is a table showing these values for $m \leqslant 8$.

m	1	2	3	4	5	6	7	8
a_m	24	240	504	480	264	65,520	24	16,320

1.2B Our Main Result

Our main theorem can be stated in three different but equivalent ways:

- **Manifold formulation:** It says that a certain geometrically defined invariant $\Phi(M)$ (the Arf–Kervaire invariant, to be defined later) on certain manifolds M is always zero.
- **Stable homotopy theoretic formulation:** It says that certain long-sought hypothetical maps between high-dimensional spheres do not exist.
- **Unstable homotopy theoretic formulation:** It says something about the EHP sequence (to be defined below), which has to do with unstable homotopy groups of spheres.