

TOPOLOGICAL ASPECTS OF $\mathbb{Z}/2\mathbb{Z}$ EIGENFUNCTIONS FOR THE LAPLACIAN ON S^2

C. H. TAUBES[†] & Y. WU[◇]

Abstract

This paper concerns the behavior of the eigenfunctions and eigenvalues of the round sphere's Laplacian acting on the space of sections of a real line bundle which is defined on the complement of an even numbers of points in S^2 . Of particular interest is how these eigenvalues and eigenvectors change when viewed as functions on the configuration spaces of points.

CONTENTS

1. Introduction	379
2. Basic technology	386
3. The lowest eigenvalue as a function on $\mathcal{C}_{2n}$	398
4. The $\mathbb{RP}^\infty$ bundle	409
5. The extension of $\mathbb{RP}$ and $\mathcal{E}$ to $\overline{\mathcal{C}}_{2n}$	421
6. Min-max for $\mathcal{E}$ on $\overline{\mathbb{RP}}$	434
7. The case of $\mathcal{C}_2$	444
Appendix A. The cohomology of a weak $\mathbb{RP}^\infty$ bundle	447
References	461

1. Introduction

This paper describes an eigenvalue problem for the Laplacian on the round 2-sphere which is associated to a configuration of an even number of distinct points on that sphere. Each such point configuration specifies a Hilbert space domain for the Laplacian which makes it a self-adjoint operator much like the standard Sobolev space domain (which is the domain in the case of the zero point configuration). The eigenvalues and eigenfunctions change when the points in the configuration move

[†]CHT was supported in part by the National Science Foundation grant DMS 2002771.

[◇]YW undertook the research reported here while affiliated with the Center of Mathematical Sciences and Applications at Harvard University.

Received August 29, 2021.

about; and our goal is to understand how and why the eigenvalues and eigenfunctions change the way they do. Our goal in this regard is realized only in part; and it is fair to say that our study produces more questions than answers.

1.1. The background. To set the stage for what is to come fix an even, positive integer to be denoted by $2n$ and let $\mathfrak{p}$ denote a chosen collection of $2n$ distinct points in the round 2-sphere. With $\mathfrak{p}$ in hand, let $\mathcal{I}_{\mathfrak{p}}$ denote the real line bundle over $S^2 - \mathfrak{p}$ with monodromy -1 on any embedded circle in $S^2 - \mathfrak{p}$ linking any given point from $\mathfrak{p}$ and no others. This is viewed as an associated vector bundle to the principle $\mathbb{Z}/2\mathbb{Z}$ bundle that is defined by the 2-1 branched cover of S^2 with $\mathfrak{p}$ being the branch locus. (All real line bundles in this paper have implicit fiber metrics that associate them to principle $\mathbb{Z}/2\mathbb{Z}$ bundles). Viewed in this way, the bundle $\mathcal{I}_{\mathfrak{p}}$ has a canonical fiber metric which can and will be used to define the norm of sections and, with the round metric’s inner product, the norm of derivatives of sections. With regards to derivatives: There is a canonical, metric compatible ‘covariant derivative’ for sections of $\mathcal{I}_{\mathfrak{p}}$ and $\mathcal{I}_{\mathfrak{p}}$ valued tensors on S^2 . This is because these are defined locally up to multiplication by the constant -1 . If f denotes a section, then df is used to denote the corresponding covariant derivative which is a section of $\mathcal{I}_{\mathfrak{p}} \otimes T^*S^2$ over $S^2 - \mathfrak{p}$. (The covariant derivative on $\mathcal{I}_{\mathfrak{p}}$ -valued 1-forms and tensors is defined with the help of the round metric’s Levi-Civita connection.)

Let $\mathbb{H}_{\mathfrak{p}}$ denote the Hilbert space completion of the space of smooth, compactly supported sections of $\mathcal{I}_{\mathfrak{p}}$ over $S^2 - \mathfrak{p}$ using the norm whose square is defined by the rule

(1.1)
$$f \longrightarrow \|f\|_{\mathbb{H}}^2 = \int_{S^2 - \mathfrak{p}} |df|^2 + \int_{S^2} |f|^2$$

As explained in [11], there is a largest positive number to be denoted by $E_{\mathfrak{p}}$ such that the inequality

(1.2)
$$\int_{S^2 - \mathfrak{p}} |df|^2 \geq E_{\mathfrak{p}} \int_{S^2} |f|^2$$

holds for every $f \in \mathbb{H}_{\mathfrak{p}}$. (As a consequence, the right most integral on the right hand side of (1.1) is redundant except in the well understood case when $n = 0$.)

Of interest in this article are those elements in $\mathbb{H}_{\mathfrak{p}}$ that are eigensections of the Laplacian: A section f is an eigensection if there is a number $\lambda \in [E_{\mathfrak{p}}, \infty)$ such that

(1.3)
$$\int_{S^2 - \mathfrak{p}} \langle df, dh \rangle = \lambda \int_{S^2 - \mathfrak{p}} \langle f, h \rangle$$

for all $h \in \mathbb{H}_{\mathfrak{p}}$. The number λ is called the eigenvalue. (The notation here and subsequently uses $\langle, \rangle$ to denote the round metric's inner product on both T^*S^2 and on $T^*S^2 \otimes \mathcal{I}_{\mathfrak{p}}$.) Any given eigensection for the Laplacian is a smooth section of $\mathcal{I}_{\mathfrak{p}}$. This is because the eigensection (denoted by f) can be canonically viewed (up to multiplication by -1) as an ordinary function on any disk that is disjoint from $\mathfrak{p}$ where it obeys the spherical metric's Laplace equation

$$(1.4) \quad d^\dagger df = \lambda f$$

with λ denoting the corresponding eigenvalue. (What is denoted by $d^\dagger$ signifies the formal L^2 adjoint of d .)

An eigensection from $\mathbb{H}_{\mathfrak{p}}$ for the Laplacian is said in what follows to be an $\mathcal{I}_{\mathfrak{p}}$ -eigensection and its corresponding eigenvalue is said to be an $\mathcal{I}_{\mathfrak{p}}$ -eigenvalue.

Proposition 2.1 in [11] asserts that $E_{\mathfrak{p}}$ is an $\mathcal{I}_{\mathfrak{p}}$ -eigenvalue. The technology used to prove that proposition can be employed in a straightforward way to prove the proposition that follows:

Proposition 1.1. *Fix a set $\mathfrak{p}$ of distinct points in S^2 . The corresponding Hilbert space $\mathbb{H}_{\mathfrak{p}}$ has an orthonormal basis of $\mathcal{I}_{\mathfrak{p}}$ -eigensections. Moreover, the corresponding set of $\mathcal{I}_{\mathfrak{p}}$ -eigenvalues have finite multiplicities and they form a discrete set in $[E_{\mathfrak{p}}, \infty)$ with no accumulation points.*

Let $\mathcal{C}_{2n}$ denote the space of unordered $2n$ -tuples of points in S^2 . Each such $2n$ -tuple of points has a corresponding set of $\mathcal{I}_{\mathfrak{p}}$ -eigensections and the associated set of $\mathcal{I}_{\mathfrak{p}}$ -eigenvalues: Our goal in this paper is explore how these $\mathcal{I}_{\mathfrak{p}}$ -eigensections and their $\mathcal{I}_{\mathfrak{p}}$ -eigenvalues change as $\mathfrak{p}$ moves in $\mathcal{C}_{2n}$.

1.2. Eigenvalue functions on $\mathcal{C}_{2n}$. We start by investigating the behavior of the eigenvalues as functions of the configurations. Our principle example is the lowest eigenvalue function, the function $\mathfrak{p} \rightarrow E_{\mathfrak{p}}$. This function on $\mathcal{C}_{2n}$ has the following properties:

(1.5)

- The function $E_{(\cdot)}$ is continuous; it is differentiable only where the corresponding eigenspace is 1-dimensional.
- The infimum of $E_{(\cdot)}$ on $\mathcal{C}_{2n}$ is zero, it is only achieved on $\mathcal{C}_{2n}$ if $n = 0$.
- The supremum of $E_{\mathfrak{p}}$ on $\mathcal{C}_{2n}$ is greater than an n -independent multiple of n and it is always achieved. But the function $E_{(\cdot)}$ is not differentiable there (nor is it at any other putative critical point).

By way of an example for the first bullet: The eigenspace for the minimal eigenvalue $E_{\mathfrak{p}}$ on $\mathcal{C}_2$ has dimension 2 when the two points in

the configuration are antipodal (see Section 7). And, as we explain later, there are configurations in any $n > 1$ version of $\mathcal{C}_{2n}$ where the minimal eigenvalue eigenspace has multiplicity at least four. (This multiplicity business is surprising because mathematical physics/differential equation lore says that the lowest eigenvalue of a Laplacian must have multiplicity one.)

To elaborate on the second bullet: The function $E_{(\cdot)}$ limits to zero along paths in $\mathcal{C}_{2n}$ that bring all of the constituent points together. We study the behavior of the lowest eigenvalue function $E_{(\cdot)}$ (and also higher $\mathcal{I}_{(\cdot)}$ -eigenvalues) along sequences where subsets of configuration points coalesce by introducing a natural compactification of $\mathcal{C}_{2n}$ as a stratified space with strata

$$\{\mathcal{C}_{2n}, \mathcal{C}_{2n-2}, \mathcal{C}_{2n-4}, \dots, \mathcal{C}_0\}.$$

We then show that $E_{(\cdot)}$ (in fact, all $\mathcal{I}_{(\cdot)}$ -eigenvalues) extends continuously to this compactification. The compactification is denoted by $\overline{\mathcal{C}}_{2n}$.

To elaborate on the third bullet: Although $E_{(\cdot)}$ is differentiable where the corresponding eigenspace has multiplicity 1, it is not necessarily differentiable where the multiplicity is greater than 1. In this regard, we have an expression for the directional derivative of an eigenvalue function $\mathfrak{p} \longrightarrow \lambda_{\mathfrak{p}}$ where the eigenspace has multiplicity 1 which implies this:

(1.6)

The directional derivatives of $\lambda_{(\cdot)}$ at a given configuration p , are all zero if and only if the corresponding $\mathcal{I}_{\mathfrak{p}}$ -eigensection is $\mathcal{O}\left(\text{dist}(\mathfrak{p}, \cdot)^{\frac{3}{2}}\right)$ near $\mathfrak{p}$.

By way of comparison, the arguments in Part 3 of Section 4 in [11] prove that the most that can be said about an $\mathcal{I}_{\mathfrak{p}}$ -eigensection with no other a priori information is that its norm is $\mathcal{O}\left(\text{dist}(p, \cdot)^{\frac{1}{2}}\right)$ near any given point $p \in \mathfrak{p}$. As we explain below, the behavior in (1.6) implies that the eigenspace multiplicity of $E_{(\cdot)}$ is greater than 1 at any configuration where all of the directional derivatives are zero. That implies, in turn, that $E_{(\cdot)}$ does not have continuous derivatives where it is maximal.

What is said in (1.5) for the minimal eigenvalue function $E_{(\cdot)}$ has an analog for the function that assigns to a configuration its k 'th lowest eigenvalue (for any $k \geq 1$); see Section 3.4 for a precise definition of this function. The basic conclusion is that these eigenvalue functions are differentiable where they have multiplicity 1, and need only be continuous where their multiplicity is greater than 1. Moreover, any putative critical point must occur where the multiplicity is in fact greater than 1 and thus where the function might not be differentiable. (Ljusternick-Schnirrelman constructions – which is to say min-max – can be used to define the notion of a critical value and critical point

for a function that is only continuous.) The preceding remarks about the lack of differentiability of the k 'th eigenvalue functions are proved in Section 3.4.

Looking ahead: Sections 2 and 3 of this paper investigate the properties of these eigenvalue functions on $\mathcal{C}_{2n}$ and $\overline{\mathcal{C}}_{2n}$.

1.3. The projective space bundle. To understand the phenomena depicted in (1.5), we replace the eigenvalue functions on $\mathcal{C}_{2n}$ with a ‘universal’ energy functional on the space of pairs of the form $(\mathbf{p}, f)$ where $\mathbf{p}$ is from $\mathcal{C}_{2n}$ and f is a section of $\mathcal{I}_{\mathbf{p}}$ with the S^2 -integral of f^2 equal to 1. This function is denoted by $\mathcal{E}$ and it is defined by the rule

$$(1.7) \quad (\mathbf{p}, f) \longrightarrow \int_{S^2 - \mathbf{p}} |df|^2.$$

This function $\mathcal{E}$ is (formally) a smooth function of these sorts of pairs. Moreover, a perturbation theoretic argument says this: If $(\mathbf{p}, f)$ is a critical point of $\mathcal{E}$, then f is an $\mathcal{I}_{\mathbf{p}}$ -eigensection (this is Raleigh-Ritz) and f as it nears any point in $\mathbf{p}$ vanishes as $\mathcal{O}\left(\text{dist}(\mathbf{p}, \cdot)^{\frac{3}{2}}\right)$ which is precisely the condition in (1.6).

Given $\mathbf{p} \in \mathcal{C}_{2n}$, let $\mathbb{S}_{\mathbf{p}}$ denote the subset in the Hilbert space $\mathbb{H}_{\mathbf{p}}$ (the norm is depicted in (1.1)) that consists of the elements whose square has S^2 -integral equal to 1. The function $\mathcal{E}$ is formally a function on a ‘fiber bundle’ over $\mathcal{C}_{2n}$ whose fiber of $\mathbf{p}$ is the subspace $\mathbb{S}_{\mathbf{p}}$. But now something interesting appears: There is no such fiber bundle. This is a very pretty example of what the physicists would call an anomaly, this one being a $\mathbb{Z}/2\mathbb{Z}$ version of a integer characteristic class anomaly that was discussed at some length in an unpublished paper by G. Segal [9] in the context of certain observations by L. Fadeev [1]. (There are many papers on anomalies in physics; see the recent paper of L. Müller for other versions of anomalies and references [5].) In the context at hand, there exists an obstruction class in $H^2(\mathcal{C}_{2n}; \mathbb{Z}/2\mathbb{Z})$ that is zero if and only if the assignments $\{\mathbf{p} \longrightarrow \mathbb{S}_{\mathbf{p}}\}_{\mathbf{p} \in \mathcal{C}_{2n}}$ define a fiber bundle. This obstruction class is non-zero; it is denoted by ω .

We can none-the-less define an $\mathbb{RP}^{\infty}$ bundle over $\mathcal{C}_{2n}$ whose fiber over any given configuration $\mathbf{p}$ is $\mathbb{S}_{\mathbf{p}} / \{\pm 1\}$ which is sufficient for the purposes at hand because $\mathcal{E}$ has the same value on $(\mathbf{p}, f)$ as it has on $(\mathbf{p}, -f)$. This bundle is denoted by $\mathbb{RP}$ and $\mathcal{E}$ descends there as a smooth function. In particular, one can look for critical points of $\mathcal{E}$ on $\mathbb{RP}$. This leads us to study the (co)homology of $\mathbb{RP}$ because that (in principle) dictates critical points of $\mathcal{E}$. (This $\mathbb{RP}$ bundle and its function $\mathcal{E}$ are described in Section 4 of this paper.)

By way of a parenthetical remark: In an alternate universe where the obstruction class is zero, one should still consider $\mathcal{E}$ on the quotient $\mathbb{RP}^{\infty}$ to study its critical points.

1.4. Critical values of $\mathcal{E}$. The observation in the third bullet of (1.5) that the minimal eigenvalue function $E_{(\cdot)}$ is not everywhere differentiable is an automatic consequence of the following observation about the function $\mathcal{E}$ on $\mathbb{RP}$:

(1.8)

Suppose that $n > 0$ and that $(\mathfrak{p}, [f])$ is a critical point of the $\mathcal{C}_{2n}$ version of $\mathcal{E}$. Then the corresponding critical value is not the lowest $\mathcal{I}_{\mathfrak{p}}$ -eigenvalue (which is $E_{\mathfrak{p}}$); in fact, it is greater than the n 'th $\mathcal{I}_{\mathfrak{p}}$ -eigenvalue.

As noted in the first bullet of (1.5), the lack of differentiability of the minimal eigenvalue function $E_{(\cdot)}$ is due to the existence of configurations where the minimal eigenvalue has an eigenspace with dimension greater than 1. The condition in (1.8) tells us slightly more; it tells us that there can't be a differentiable section of $\mathbb{RP}$ (as a bundle over $\mathcal{C}_{2n}$) that assigns any given configuration to some $\mathcal{I}_{(\cdot)}$ -eigensection (modulo ± 1) with minimal eigenvalue. As explained in the next paragraph, the lack of a section of $\mathbb{RP}$ is dictated by the existence of a non-zero anomaly class.

With regards to the anomaly and the existence of sections: The key fact about the anomaly class is that it's pull-back to $\mathbb{RP}$ via the projection map is zero; and this can't be the case if there is a section because the composition of first the section and then the projection is the identity map on $\mathcal{C}_{2n}$. (If the anomaly were zero, then there would be a section of the corresponding infinite dimensional sphere bundle; and thus a section of the $\mathbb{RP}^\infty$ bundle. This is because a vector bundle over a finite dimensional manifold whose fiber dimension is greater than the base manifold dimension has a non-zero section.)

Morse theoretic arguments (min-max, for example) using the (co)homology of $\mathbb{RP}$ would yield critical points of $\mathcal{E}$ were $\mathcal{E}$ a proper function on $\mathbb{RP}$. Although $\mathcal{E}$ is proper along the fibers of the projection to $\mathcal{C}_{2n}$, it is not globally proper because $\mathcal{C}_{2n}$ isn't compact.

1.5. Homology classes of $\overline{\mathbb{RP}}$. As noted previously, the space $\mathcal{C}_{2n}$ has a compactification as a stratified space (denoted by $\overline{\mathcal{C}}_{2n}$) that accounts for the coalescence of subsets of configuration points. As it turns out, there is an extension of $\mathbb{RP}$ over this compactification (denoted by $\overline{\mathbb{RP}}$) that maps to the compactification $\overline{\mathcal{C}}_{2n}$ and a corresponding, continuous extension of the function $\mathcal{E}$ as a function over $\overline{\mathbb{RP}}$. This $\overline{\mathbb{RP}}$ space has non-zero $\mathbb{Z}/2\mathbb{Z}$ homology classes in degrees $4n+2m+1$ with m any non-negative integer. It is possible that these classes yield critical points of $\mathcal{E}$ on large n versions of $\overline{\mathcal{C}}_{2n}$ because each such (n, m) class is the image via a fiber preserving map to $\overline{\mathbb{RP}}$ of the $\mathbb{Z}/2\mathbb{Z}$ fundamental class of an $\mathbb{RP}^{2m+1}$ bundle over $\overline{\mathcal{C}}_{2n}$. (The class sees the whole of $\overline{\mathcal{C}}_{2n}$ because of this; in particular, the class is not carried by any single strata in $\overline{\mathcal{C}}_{2n}$.) We show that these classes do give critical points of $\mathcal{E}$ on $\overline{\mathbb{RP}}$ and we

show in Section 6 of this paper that the corresponding critical values (as functions of either n or m) are unbounded. Even so, it remains to be seen whether these critical points all lie in some fixed $\mathcal{C}_{2j}$ strata with j having an n and m independent upper bound. This is work for the future.

By way of a remark: The classes used here can be viewed as a homological work-around for the non-differentiability of the k 'th eigenvalue function on $\mathcal{C}_{2n}$ in the case when k is odd. (There are also classes in the $\mathbb{Z}/2\mathbb{Z}$ homology of $\bar{\mathcal{C}}_{2n}$ with even degree – any even number – but their corresponding critical points are almost surely in the $\mathcal{C}_0$ or $\mathcal{C}_2$ strata. These strata and the corresponding critical points are described in Section 7.)

1.6. A potential application. Critical points of $\mathcal{E}$ on $\mathcal{C}_{2n}$ (should they exist) can be used to construct homogeneous singularity models for $\mathbb{Z}/2$ harmonic 1-forms and spinors on $\mathbb{R}^3$. To explain (see [11] and/or [10] for more), a homogeneous $\mathbb{Z}/2$ harmonic 1-form on $\mathbb{R}^3$ consists of a data set $(Z, \mathcal{I}, \nu)$ with Z being a finite union of rays from the origin, with $\mathcal{I}$ being real line bundle over the complement of Z and with ν being an $\mathcal{I}$ -valued 1-form defined on $\mathbb{R}^3 - Z$ that has the following properties:

(1.9)

- ν is both closed and coclosed.
- $|\nu|$ extends over Z as a Hölder continuous function on $\mathbb{R}^3$.
- ν is homogeneous with respect to pull-back by the rescaling diffeomorphism that multiplies the Euclidean coordinates by a non-zero real number.

As explained in [11], the data sets of this sort with Z being the union of $2n$ rays are in 1-1 correspondence with pairs $(\mathfrak{p}, f)$ with $\mathfrak{p} \in \mathcal{C}_{2n}$ and with $f \in \mathbb{H}_{\mathfrak{p}}$ being a non-zero eigensection of the Laplacian with $|f| \leq \mathcal{O}\left(\text{dist}(\mathfrak{p}, \cdot)^{\frac{3}{2}}\right)$ near $\mathfrak{p}$ (which is exactly the critical point criteria for the function $\mathcal{E}$). This same [11] explains how such a pair can also be used to define a homogeneous, $\mathbb{Z}/2$ harmonic spinor on $\mathbb{R}^3$. To summarize the correspondence: The set Z is the union of the rays from the origin through the points in the configuration $\mathfrak{p}$, the line bundle $\mathcal{I}$ is the pull-back of $\mathcal{I}_{\mathfrak{p}}$ via the map $\pi : x \longrightarrow \frac{x}{|x|}$ from $\mathbb{R}^3 \setminus \{0\}$ to S^2 , and ν is given by the rule

$$(1.10) \quad \nu = d(|x|^{\mu} \pi^* f),$$

where μ is $\frac{1}{2} \left(1 + (1 + 4\lambda)^{\frac{1}{2}}\right)$ with λ denoting the corresponding $\mathcal{I}_{\mathfrak{p}}$ -eigenvalue.

Before [11], the only example of a $\mathbb{Z}/2$ harmonic 1-form on $\mathbb{R}^3$ had Z being a straight line through the origin (which corresponds to a critical

point of $\mathcal{E}$ on $\mathcal{C}_2$). The paper [11] used the symmetry groups of the Platonic solids to construct examples where Z is the union of (respectively) four rays from the origin (pointing to the vertices of a regular tetrahedron), eight rays from the origin (pointing to the vertices of a cube) and 12 rays from the origin (pointing to the vertices of an icosahedron.) Going in the reverse direction, the constructions in [11] give critical points of $\mathcal{E}$ on $\mathcal{C}_4$ and $\mathcal{C}_8$ and $\mathcal{C}_{12}$.

1.7. Two remarks for the reader and a table of contents. The first remark concerns the case of $n = 1$: This case is illuminating because much can be seen explicitly in this case. The $n = 1$ case is discussed in some detail in Section 7 and a reader is advised to refer to that section when new notions are introduced.

The second remark concerns two notational conventions that are used throughout this paper: The first convention has c_0 denoting a number greater than 1 whose value is independent of relevant background such as a configuration in $\mathcal{C}_{2n}$ or the integer n . What precisely it is independent of will be clear from the context. The value of c_0 increases between consecutive incarnations. The second convention has χ denoting a chosen, smooth and non-increasing function on $\mathbb{R}$ that is equal to 1 on $(-\infty, \frac{1}{4}]$ and equal to 0 on $[\frac{3}{4}, \infty)$. This fixed version of χ is implicitly used to construct ‘bump functions’ and ‘cut-off functions’.

Acknowledgments. Before starting, Y-Y. Wu wishes to acknowledge and thank the generous support of the Center of Mathematical Sciences and Applications (CMSA) at Harvard. Meanwhile, C. H. Taubes wishes to acknowledge and thank the generous support of the National Science Foundation (DMS grant number 2002771).

2. Basic technology

This section introduces the basic technology that is used in later sections. To set the stage for what is to come, suppose henceforth that $\mathfrak{p}$ is a chosen $2n$ -tuple from $\mathcal{C}_{2n}$. The convention in what follows has c_0 denoting a number that is greater than 1 and that is independent of any chosen sections of $\mathcal{I}_{\mathfrak{p}}$ in any given appearance. It can depend on $\mathfrak{p}$ unless stated otherwise. As noted in Section 1, the value of c_0 can be assumed to increase between successive appearances.

2.1. $\mathfrak{p}$ -independent a priori bounds. The lemma that follows summarizes some basic a priori bounds on eigensections of the Laplacian for a given $\mathfrak{p} \in \mathcal{C}_{2n}$. These bounds are independent of $\mathfrak{p}$.

Lemma 2.1. *There exists $\kappa > 1$ with the following significance: Fix $\mathfrak{p} \in \mathcal{C}_{2n}$ and let f denote an $\mathcal{I}_{\mathfrak{p}}$ -eigensection. Denote the corresponding $\mathcal{I}_{\mathfrak{p}}$ -eigenvalue by λ . Then the function $|f|$ obeys:*

•
$$\int_{S^2} |d|f||^2 = \lambda \int_{S^2} |f|^2.$$

- $|f| \leq \kappa(\lambda + 1) \left(\int_{S^2} |f|^2 \right)^{\frac{1}{2}}$ on the complement of $\mathfrak{p}$.

Proof of Lemma 2.1. The first bullet inequality follows because any $f \in \mathbb{H}_{\mathfrak{p}}$ obeys the identity $|d|f| = |df|$ almost everywhere, and because the integral of $|df|^2$ for the eigensection f is equal to λ times that of $|f|^2$. To prove the second bullet, note first the function $|f|$ obeys the distributional inequality

$$(2.1) \quad d^\dagger d|f| \leq \lambda|f|.$$

With this understood, fix a point $q \in S^2 - \mathfrak{p}$ and let $G_q(\cdot)$ denote the Green's function for the operator $d^\dagger d + 1$ with pole at q . This Green's function is smooth on $S^2 - q$ and it is bounded near q by an $\mathcal{O}(1)$ multiple of $-\frac{1}{2\pi} \ln(\text{dist}(\cdot, q))$. The plan for what follows is to multiply both sides of (2.1) by G_q , integrate the resulting inequality over S^2 and then integrate by parts twice to bound $|f|(q)$ by λ times the integral of $G_q|f|$. The latter bound will lead to the second bullet's bound via the Cauchy-Schwarz inequality because the S^2 integral of G_q^2 is finite. To implement this strategy, care must be taken near the points of $\mathfrak{p}$. To that end, fix $\epsilon > 0$ but much less than 1, and then let $\chi_\epsilon(\cdot)$ denote the function on S^2 that is given by the rule

$$(2.2) \quad \chi_\epsilon(\cdot) = \chi \left(2 \frac{\ln(\text{dist}(\cdot, \mathfrak{p}))}{\ln(100\epsilon)} - 1 \right).$$

This function is equal to 1 where the distance to $\mathfrak{p}$ is greater than $(100\epsilon)^{\frac{1}{2}}$ and equal to zero where the distance to $\mathfrak{p}$ is less than 100ϵ . If ϵ is much less than the distance from q to $\mathfrak{p}$, then q will sit where χ_ϵ is equal to 1. Note also that the derivative of χ_ϵ has support only in the region where the distance to $\mathfrak{p}$ is between 100ϵ and $(100\epsilon)^{\frac{1}{2}}$ where it is bounded by $c_0 \frac{1}{|\ln \epsilon|} \frac{1}{\text{dist}(\cdot, \mathfrak{p})}$ with c_0 being independent of $\mathfrak{p}$. This implies in particular that the S^2 integral of $|d\chi_\epsilon|^2$ is bounded by $c_0 \frac{1}{|\ln \epsilon|}$ with c_0 again independent of $\mathfrak{p}$. With these last points understood, multiply both sides of (2.1) by $\chi_\epsilon G_q$ with ϵ very small and then integrate by parts. The result of doing that leads to the inequality

$$(2.3) \quad |f|(q) \leq c_0 \lambda \left(\int_{S^2} |f|^2 \right)^{\frac{1}{2}} + c_{\mathfrak{p},q} \frac{1}{\sqrt{|\ln \epsilon|}} \left(\int_{S^2} |df|^2 + \int_{S^2} |f|^2 \right)^{\frac{1}{2}}$$

where c_0 is independent of $\mathfrak{p}$ and q and f and ϵ whereas $c_{\mathfrak{p},q}$ is only independent of f and ϵ (it does depend on the distance from q to $\mathfrak{p}$). Take $\epsilon \rightarrow 0$ in (2.3) to obtain the second bullet's bound for $|f|$. q.e.d.

2.2. Asymptotics near points in $\mathfrak{p}$. As noted in the proof of Lemma 2.1, if f is from $\mathbb{H}_{\mathfrak{p}}$, then its norm is in the standard L^2_1 Sobolev space for the whole of S^2 ; and this implies via Sobolev inequalities that

$|f|^q$ is integrable on S^2 for any finite, non-negative value of q . The following lemma concerns specifically f 's behavior near the points in $\mathfrak{p}$; it says in effect that a function in $\mathbb{H}_{\mathfrak{p}}$ must vanish in a weak sense as a point in $\mathfrak{p}$ is approached.

Lemma 2.2. *Given $\mathfrak{p} \in \mathcal{C}_{2n}$, there exists $\kappa > 1$ with the following significance: If $f \in \mathbb{H}_{\mathfrak{p}}$, then $\int_{S^2-\mathfrak{p}} \frac{1}{\text{dist}(\cdot, \mathfrak{p})^2} |f|^2 \leq \kappa \int_{S^2-\mathfrak{p}} |df|^2$.*

Proof of Lemma 2.2. The asserted bound follows directly from Lemma 3.2 in [11] by invoking it for the annuli from the set $\{\mathcal{A}_n(p) : p \in \mathfrak{p} \text{ and } n \in \mathbb{N}\}$ where any given $\mathcal{A}_n(p)$ is defined by the rule

(2.4) $\mathcal{A}_n(p) \equiv \{q \in S^2 - \mathfrak{p} : 2^{-n-1}\delta < \text{dist}(q, p) < 2^{-n}\delta\}_{n \in \mathbb{N}}$

with δ being $\frac{1}{100}$ times the minimum distance between the points in $\mathfrak{p}$.
q.e.d.

More can be said about behavior near the points in $\mathfrak{p}$ when f is an $\mathcal{I}_{\mathfrak{p}}$ -eigensection. In this case, the necessary observations about the behavior of an $\mathcal{I}_{\mathfrak{p}}$ -eigensection near a point $p \in \mathfrak{p}$ invoke observations from Parts 2 and 3 of Section 4 in [11], specifically Equations (4.2)–(4.5) in [11] which hold for a Laplace eigensection whether or not $\mathfrak{p}$ is invariant under a symmetry group action. These observations are summarized by the upcoming Lemma 2.3. This summary of what is said in [11] uses stereographic projection from p 's antipodal twin to define a complex coordinate to be denoted by z near $\mathfrak{p}$ with both $z = 0$ and $|dz| = 1$ at p . By way of notation: The lemma use $\Re(\cdot)$ to denote the real part of the indicated $\mathbb{C}$ -valued function or differential form.

Lemma 2.3. *Let f denote an $\mathcal{I}_{\mathfrak{p}}$ -eigensection. Given any point $p \in \mathfrak{p}$, there exists the following data:*

- A non-zero complex number to be denoted by $\mathfrak{a}$,
- A non-negative integer to be denoted by n .
- A section of $\mathcal{I}_{\mathfrak{p}}$ defined near p to be denoted $\mathfrak{e}$ obeying

$$\lim_{|z| \rightarrow 0} (|\mathfrak{e}| + |z||d\mathfrak{e}|)|z|^{-n-\frac{1}{2}} = 0.$$

These are such when f is depicted using the complex coordinate z and its complex conjugate, it has the form $f|_z = \Re\left(\mathfrak{a}z^{n+\frac{1}{2}}\right) + \mathfrak{e}$.

When the point p is germane to the discussion, then the corresponding versions of $\mathfrak{a}$ and n are written as $\mathfrak{a}_p$ and n_p . When f and $\mathfrak{p}$ are germane, they are written as $\mathfrak{a}_p(f)$ and $n_p(f)$.

By way of a definition: The subset $\mathfrak{p}_f \in \mathfrak{p}$ is defined to be the set of points with the integer n being zero.

2.3. The divergence identity. The identity that follows momentarily uses $\mathfrak{g}$ to denote the round metric on S^2 when viewed as a section of $\text{Sym}^2(T^*S^2)$. Supposing that f denotes an $\mathcal{I}_{\mathfrak{p}}$ -eigensection and λ its eigenvalue, define the section $\mathbb{T}$ of $\text{Sym}^2(T^*S^2)$ on $S^2 - \mathfrak{p}$ by the rule

$$(2.5) \quad \mathbb{T} = df \otimes df - \frac{1}{2}(|df|^2 - \lambda f^2) \mathfrak{g}.$$

By virtue of (1.4), this section is divergence free, thus

$$(2.6) \quad \nabla^\dagger \mathbb{T} = 0.$$

An integrated version of this identity will be presented momentarily in the upcoming Lemma 2.4. An observation is needed first which is this: Let s denote a continuous section of $\text{Sym}^2(T^*S^2)$ that is defined on the whole of S^2 . Then the function $\langle s, \mathbb{T} \rangle$ has finite integral over S^2 and the assignment

$$(2.7) \quad s \longrightarrow \int_{S^2} \langle s, \mathbb{T} \rangle$$

defines a bounded linear functional on the Banach space of continuous sections of $\text{Sym}^2(T^*S^2)$. This follows because $\mathbb{T}$'s norm is bounded by $c_0(|df|^2 + \lambda f^2)$ whose integral over S^2 is bounded by c_0 times the $\mathbb{H}$ -norm of f . (It also follows from Lemma 2.3 that $\mathbb{T}$'s singularities at the points in $\mathfrak{p}$ are integrable.)

The upcoming lemma also uses the convention whereby vector fields and 1-forms are identified using the metric $\mathfrak{g}$.

Lemma 2.4. *Let f denote a given $\mathcal{I}_{\mathfrak{p}}$ -eigensection and let λ denote the corresponding eigenvalue. Supposing that ν is a smooth vector field on S^2 , then*

$$\int_{S^2} \langle \nabla \nu, \mathbb{T} \rangle = -\frac{\pi}{4} \sum_{p \in \mathfrak{p}_f} \langle \nu|_p, \Re \epsilon(\mathfrak{a}_p^2 dz) \rangle.$$

Proof of Lemma 2.4. Fix $\epsilon > 0$ but much less than the distance between any two points in $\mathfrak{p}$. Supposing that $p \in \mathfrak{p}$, use $D_\epsilon(p)$ to denote the disk centered at p with radius ϵ . Integrate the function $\langle \nu, \nabla^\dagger \mathbb{T} \rangle$ over the complement of the union of these radius ϵ disks. This integral is zero because the integrand is zero. Even so, integrate by parts to write that integral (hence 0) as

$$(2.8) \quad \int_{S^2 - \cup_{p \in \mathfrak{p}} D_\epsilon(p)} \langle \nabla \nu, \mathbb{T} \rangle + \sum_{p \in \mathfrak{p}_f} \int_{\partial D_\epsilon(p)} \langle \hat{r} \otimes \nu, \mathbb{T} \rangle$$

with $\hat{r}$ denoting here the outward pointing, length 1 normal vector to the relevant $\partial D_\epsilon(p)$. The $\epsilon \rightarrow 0$ limit of the left hand integral in (2.8) is the integral of $\langle \nabla \nu, \mathbb{T} \rangle$ over S^2 . Use Lemma 2.3 to evaluate the $\epsilon \rightarrow 0$ limit of the right integral in (2.8). The assertion that the sum of these two $\epsilon \rightarrow 0$ limits is zero is Lemma 2.4. q.e.d.

The divergence identity in (2.6) for $\mathbb{T}$ as defined in (2.5) has a generalization that is also useful: To state the more general identity, suppose now that f and f' are two $\mathcal{I}_{\mathbf{p}}$ -eigensections and let λ and λ' denote their respective eigenvalues. Now define the symmetric tensor $\mathbb{S}$ by the rule

(2.9)
$$\mathbb{S} = df \otimes df' + df' \otimes df - \left(\langle df, df' \rangle - \frac{1}{2}(\lambda + \lambda') f f' \right) \mathbf{g}.$$

The tensor $\mathbb{S}$ is divergence free if $\lambda = \lambda'$, but otherwise

(2.10)
$$\nabla^\dagger \mathbb{S} = \frac{1}{2}(\lambda' - \lambda)(f df' - f' df).$$

Much the same argument that led to Lemma 2.4 leads in this case to the following analog of (2.8):

(2.11)
$$\int_{S^2} \langle \nabla \nu, \mathbb{S} \rangle - \frac{1}{2}(\lambda' - \lambda) \int_{S^2} \langle \nu, f df' - f' df \rangle + \sum_{p \in \mathfrak{p}_f} \lim_{\epsilon \rightarrow 0} \int_{\partial D_\epsilon(p)} \langle \hat{r} \otimes \nu, \mathbb{S} \rangle = 0.$$

Now we can invoke the f and f' versions of Lemma 2.3 to evaluate the limit in (2.11) and doing so leads to the following identity:

(2.12)
$$\begin{aligned} \int_{S^2} \langle \nabla \nu, \mathbb{S} \rangle - \frac{1}{2}(\lambda' - \lambda) \int_{S^2} \langle \nu, f df' - f' df \rangle \\ = -\frac{\pi}{2} \sum_{p \in \mathfrak{p}_f \cap \mathfrak{p}_{f'}} \langle \nu|_p, \Re \mathfrak{e}(\mathfrak{a}_p(f) \mathfrak{a}_p(f') dz) \rangle. \end{aligned}$$

This is the promised generalization of Lemma 2.4’s identity. With regards to the derivation of (2.12): The norm of $\mathbb{S}$ is $\mathcal{O}\left(\frac{1}{\epsilon}\right) \|\mathfrak{a}_p(f)\| \|\mathfrak{a}_p(f')\|$ on $\partial D_\epsilon(p)$, which is why the limit in (2.11) is not automatically zero when both $\mathfrak{a}_p(f)$ and $\mathfrak{a}_p(f')$ are non-zero. (Even so, the term $\langle df, df' \rangle \mathbf{g}$ on the right hand side of (2.9) makes no contribution to the limit in (2.11) nor does the $(\lambda + \lambda') f f' \mathbf{g}$ term.)

2.4. Changing the metric. If $\mathbf{p}$ is a given $2n$ -tuple of points in S^2 and $\mathbf{p}'$ is some other $2n$ -tuple, then there are diffeomorphism of S^2 that map the set $\mathbf{p}$ to $\mathbf{p}'$. If ϕ is a diffeomorphism that does this, then $\phi^* \mathcal{I}_{\mathbf{p}'}$ is isomorphic to $\mathcal{I}_{\mathbf{p}}$. As a consequence, the eigenvalues for the round metric’s Laplacian on sections of $\mathcal{I}_{\mathbf{p}'}$ are precisely those of the Laplacian on sections of $\mathcal{I}_{\mathbf{p}}$ that is defined by the ϕ -pull-back of the round metric (which won’t be the round metric unless ϕ is a rotation and/or reflection). By virtue of this pull-back correspondence, questions about the behavior of $\mathcal{I}_{\mathbf{p}}$ -eigenvalues and $\mathcal{I}_{\mathbf{p}}$ -eigenfunctions of the round metric’s Laplacian on the $\mathcal{C}_{2n}$ -dependent family of domains are questions about the metric dependence of the eigenvalues and eigenfunctions for a $\mathcal{C}_{2n}$ -dependent family of Laplace operators on a fixed domain. The latter, alternate point of view is often taken in subsequent sections and it motivates the digression that follows directly about metric Laplacians acting on the sections of the fixed bundle $\mathcal{I}_{\mathbf{p}}$.

Any given smooth Riemannian metric on S^2 has a corresponding version of the Hilbert space $\mathbb{H}_{\mathfrak{p}}$ and a corresponding set of eigensections for its Laplacian. To say more about this, let $\mathfrak{m}$ denote a given Riemannian metric. For convenience (this is not strictly necessary), we assume that the associated area 2-form is the same as that of the standard, round metric. This assumption is implicit in what follows. The metric $\mathfrak{m}$'s inner product on T^*S^2 is denoted by $\langle \cdot, \cdot \rangle_{\mathfrak{m}}$ and the associated norm by $|\cdot|_{\mathfrak{m}}$. The corresponding Hilbert space is the completion of the space of smooth, compactly supported sections of $\mathcal{I}_{\mathfrak{p}}$ using the version of (1.1) with $|df|^2$ replaced by $|df|_{\mathfrak{m}}^2$. An element in this Hilbert space (denoted by f) is an eigensection of the metric $\mathfrak{m}$'s Laplacian with eigenvalue λ when

$$(2.13) \quad \int_{S^2-\mathfrak{p}} \langle df, dh \rangle_{\mathfrak{m}} = \lambda \int_{S^2-\mathfrak{p}} fh$$

for all elements h from the Hilbert space. Eigensections are the elements in the Hilbert space that obey the Laplace equation for $\mathfrak{m}$'s version of the Laplace operator.

With regards to this metric $\mathfrak{m}$ Hilbert space: It is the original round metric's version of $\mathbb{H}_{\mathfrak{p}}$ with a different but equivalent inner product. This is to say that any section in one is in the other and vice versa. This understood, the metric $\mathfrak{m}$'s version of the eigenvalue/eigensection question will be viewed henceforth as a question concerning a bilinear form on the original $\mathbb{H}_{\mathfrak{p}}$ Hilbert space. As was the case with the round metric $\mathfrak{g}$, there is a complete, orthonormal basis for $\mathbb{H}_{\mathfrak{p}}$ of eigensections for the bilinear form in (2.13) with the corresponding eigenvalue set being discrete with no accumulation points and with each eigenvalue having finite multiplicity.

Now suppose that $\mathcal{M}$ is a real analytic, finite dimensional manifold that parameterizes in a real analytic fashion a family of metrics on S^2 subject to the constraint that their area forms are the same as that of $\mathfrak{g}$. There is a corresponding family of bilinear forms on $\mathbb{H}_{\mathfrak{p}}$ and one can ask how the sets of eigenvalues and eigenfunctions change as the domain $\mathcal{M}$ is changed. Exactly this sort of question is addressed in Chapters 6 and 7 of Tosio Kato's classic book *Perturbation theory for linear operators* [4]. The rest of this section summarizes the implications of Kato's discussion.

To start the story, fix a metric $\mathfrak{m}_0 \in \mathcal{M}$ and then fix a smooth coordinate chart for a neighborhood of $\mathfrak{m}_0$ that identifies the chosen neighborhood with a ball about the origin in $\mathbb{R}^N$. (The integer N is the dimension of $\mathcal{M}$.) The identification is chosen so that $\mathfrak{m}_0$ corresponds to the origin in $\mathbb{R}^N$. A point in this ball is denoted by y and the corresponding metric by $\mathfrak{m}_y$. If y is close to the origin in $\mathbb{R}^N$, then $\mathfrak{m}_y^{-1}$ can

be written as

$$(2.14) \quad \mathbf{m}_y^{-1} = \mathbf{m}_0^{-1}(1 + \mathfrak{k}y + \mathfrak{e}_y)$$

with $\mathfrak{k}$ denoting a linear map from $\mathbb{R}^N$ to $C^\infty(S^2; \text{Sym}^2(TS^2))$ and with $\mathfrak{e}_y$ being a map from the coordinate ball in $\mathbb{R}^N$ to $C^\infty(S^2; \text{Sym}^2(TS^2))$ with C^k norm (for any given non-negative integer k) bounded by a k -dependent multiple of $|y|^2$. Now suppose that λ_0 is an eigenvalue for the $\mathbf{m}_0$ Laplacian on $\mathbb{H}_p$. As noted in [4], there will be an eigenvalue of the $\mathbf{m}_y$ Laplacian near to λ_0 if y is sufficiently close to the origin. Moreover, if λ_0 has multiplicity n , then there will be precisely n linearly independent eigensections for the $\mathbf{m}_y$ -Laplacian with eigenvalue very close to λ_0 if y is small (much closer to λ_0 than to any other $\mathbf{m}_0$ eigenvalue.) In particular, the following is true: Given $\epsilon > 0$, there exists $\delta > 0$ such that the eigenvalues of these n eigenvectors of the $\mathbf{m}_y$ Laplacian will differ from λ_0 by at most ϵ if $|y| < \delta$. In addition, if $|y| < \delta$, then the n -dimensional vector space spanned by these n eigensections for the $\mathbf{m}_y$ -Laplacian will have $\mathbb{H}_p$ -distance at most ϵ from the λ_0 eigenspace of the $\mathbf{m}_0$ Laplacian. (This distance is measured by the $\mathbb{H}_p$ -orthogonal projection.) The subsequent paragraphs and the upcoming Proposition 2.5 say more about the relationship between respective $\mathbf{m}_y$ and $\mathbf{m}_0$ Laplacian eigenvalues and eigenspaces.

The simplest case to consider is that where λ_0 has multiplicity 1 in which case there is one nearby eigenvalue when y is near 0 and that eigenvalue varies smoothly with y on a neighborhood of 0 with value λ_0 at $y = 0$. This y -dependent eigenvalue can be written as $\lambda_0 + y \cdot \lambda' + \mathcal{O}(|y|^2)$ with λ' given by the rule

$$(2.15) \quad \lambda' = \int_{S^2} \langle \mathfrak{k}, df_0 \otimes df_0 \rangle$$

In this case, there is one corresponding eigensection of the $\mathbf{m}_y$ Laplacian that is very near to a normalized λ_0 eigensection when y has small norm and it also varies smoothly with y near the origin. (This λ_0 eigensection is denoted by f_0 ; it is normalized so that the S^2 integral of its square is equal to 1.) In particular, this y -dependent element in $\mathbb{H}_p$ can be written as a Taylor's expansion as $f_0 + y \cdot \mathfrak{f} + r_y$ with $\|r_y\|_{\mathbb{H}}$ bounded by $c_0|y|^2$ and with the linear map $\mathfrak{f}$ from $\mathbb{R}^N$ to $\mathbb{H}_p$ determined via the following two rules:

$$(2.16) \quad \begin{aligned} &1) \int_{S^2} \mathfrak{f} f_0 = 0, \\ &2) \int_{S^2} \langle d\mathfrak{f}, dh \rangle_{\mathbf{m}_0} - \lambda_0 \int_{S^2} \mathfrak{f} h + \int_{S^2} \langle \mathfrak{k}, df_0 \otimes dh \rangle = 0 \text{ for all } h \in \mathbb{H} \text{ that} \\ &\quad \text{obey } \int_{S^2} h f_0 = 0. \end{aligned}$$

With this simple case understood, consider now the case when λ_0 has multiplicity $n \geq 1$. To set the stage for this case, introduce by way of

notation $\mathbb{V}$ to denote the corresponding n -dimensional vector space of eigensections of the Laplacian in $\mathbb{H}_p$ with eigenvalue λ_0 . Given a pair f and f' from V , define a vector in $\mathbb{R}^N$ which is denoted by $\mathbf{k}(f, f')$ by the rule

$$(2.17) \quad \mathbf{k}(f, f') \equiv \int_{S^2} \langle \mathfrak{k}, df \otimes df' \rangle.$$

Supposing that $y \in \mathbb{R}^N$, then the assignment of any given pair of elements in $\mathbb{V}$ to the value of $y \cdot \mathbf{k}$ on that pair defines a symmetric bilinear form on $\mathbb{V}$. An element $f \in \mathbb{V}$ is said to be an eigensection for this bilinear form when there exists a number $e(y)$ such that

$$(2.18) \quad y \cdot \mathbf{k}(f, f') = e(y) \int_{S^2} ff' \text{ for all } f' \in \mathbb{V}.$$

With regards to the y -dependence: If f is an eigensection for $y \cdot \mathbf{k}$, then it is likewise for $ry \cdot \mathbf{k}$ for any given real number r , the corresponding eigenvalue $e(ry)$ being $re(y)$.

Proposition 2.5. *Let λ_0 denote an eigenvalue of the $\mathfrak{m}_0$ -metric Laplacian on $\mathbb{H}_p$ and let $\mathbb{V}$ denote the corresponding eigenspace. Given $\epsilon > 0$, there exists a positive number δ with the following significance: If $y \in \mathbb{R}^N$ has norm less than δ , then the span of the eigensections of the $\mathfrak{m}_y$ metric Laplacian on $\mathbb{H}_p$ with eigenvalue within ϵ of λ_0 has dimension $\dim(\mathbb{V})$. Moreover, there is a 1-1 correspondence between these $\mathfrak{m}_y$ -metric Laplacian eigensections and the eigensections of $y \cdot \mathbf{k}$ on $\mathbb{V}$ in the following sense: If f is any one of these $\dim(\mathbb{V})$ eigensections of the metric $\mathfrak{m}$ Laplacian, then this eigensection f can be written as*

$$f = f_0 + y \cdot \mathfrak{f} + \mathcal{O}(|y|^2)$$

where f_0 signifies an eigensection for $y \cdot \mathbf{k}$ on $\mathbb{V}$; and where $\mathfrak{f}$ signifies a homomorphism from $\mathbb{R}^N$ to $\mathbb{H}_p$ whose image is L^2 -orthogonal to $\mathbb{V}$ and which obeys

$$\int_{S^2} \langle d\mathfrak{f}, dh \rangle_{\mathfrak{m}_0} - \lambda_0 \int_{S^2} \mathfrak{f} h + \int_{S^2} \langle \mathfrak{k}, df_0 \otimes dh \rangle = 0$$

whenever $h \in \mathbb{H}_p$ is L^2 -orthogonal to $\mathbb{V}$. Furthermore, the corresponding $\mathfrak{m}$ -metric Laplace eigenvalue of the eigensection f has the form $\lambda_0 + e(y) + \mathcal{O}(|y|^2)$ with $e(y)$ denoting the $y \cdot \mathbf{k}$ eigenvalue of f_0 .

As noted, this proposition summarizes the consequences of various observations from Chapters 6 and 7 of [4]. An important point to keep in mind with regards to this proposition: The correspondence between the eigensections of the $\mathfrak{m}_y$ Laplacian with eigenvalue near λ_0 and those of the bilinear form $y \cdot \mathbf{k}$ is one-to-one when $|y|$ is small if the latter has distinct eigenvalues. If it doesn't, then there might only be a one-to-one correspondence with a particular basis of its eigensections.

2.5. Pulling back by a diffeomorphism. Fix a positive integer to be denoted by N and then fix a linear map from $\mathbb{R}^N$ into the subspace of divergence zero elements in $C^\infty(S^2; TS^2)$; the map is denoted by ν in what follows. Thus, if $y \in \mathbb{R}^N$, then $y \cdot \nu$ is a divergence zero vector field on S^2 . Taking the time 1 point on the integral curves of these vector fields defines a smooth map from a ball about the origin in $\mathbb{R}^N$ to the subspace of area preserving diffeomorphisms of S^2 which is equal to the identity at $y = 0$ and whose differential at $y = 0$ is given by ν . (The differential at $y = 0$ is a linear map from $\mathbb{R}^N$ into the tangent space to the space of area preserving diffeomorphisms of S^2 at the identity element which is the space of divergence zero vector fields on S^2 .) Let $\mathbf{B} \subset \mathbb{R}^N$ denote this ball and, supposing that $y \in \mathbf{B}$, let ϕ_y denote the diffeomorphism that is defined by y . Also: Let $\mathbf{m}_y$ denote the pull-back of the round metric on S^2 by ϕ_y . Because ϕ_y is area preserving, the area form for the metric $\mathbf{m}_y$ is the round metric's area form.

By taking the radius of $\mathbf{B}$ to be smaller if necessary, each $y \in \mathbf{B}$ version of the inverse metric to $\mathbf{m}_y$ can be written as in (2.14) with $\mathfrak{k}$ as follows:

(2.19)
$$\mathfrak{k} = -(\nabla v + (\nabla v)^T)$$

with it understood that the round metric has been used to identify TS^2 with T^*S^2 and to define the transpose that is used in (2.19) to make a symmetric tensor from $\nabla \nu$.

Now suppose that λ_0 is an eigenvalue of the round metric's Laplacian on $\mathbb{H}_{\mathfrak{p}}$ and let $\mathbb{V}$ denote the corresponding eigenspace. Then the $\mathbb{R}^N$ -dependent bilinear form on $\mathbb{V}$ that is depicted in (2.17) can be written as

(2.20)
$$\mathbf{k}(f, f') = - \int_{S^2} \langle \nabla \nu, \mathbb{S} \rangle$$

with $\mathbb{S}$ given by (2.9). (The contributions to $\langle \nabla \nu, \mathbb{S} \rangle$ from the terms proportional to $\mathfrak{g}$ are zero because ν is divergence free.) This being the case, $\mathbf{k}$ can be written using (2.12):

(2.21)
$$\mathbf{k}(f, f') = \frac{\pi}{2} \sum_{p \in \mathfrak{p}_f \cap \mathfrak{p}_{f'}} \langle \nu|_p, \Re \mathfrak{e}(\mathfrak{a}_p(f) \mathfrak{a}_p(f') dz) \rangle.$$

This implies in particular that the eigenvalues of any $y \in \mathbf{B}$ version of $y \cdot \mathbf{k}$ are determined by the asymptotics of the eigensections in $\mathbb{V}$ at the points in $\mathfrak{p}$.

2.6. A set of useful vector fields. The applications of the preceding formulas will use divergence zero vector fields of the sort that are described in this subsection. Each vector field of interest is specified by a data set consisting of a point in S^2 to be denoted by q , a number from $(-1, 1)$ to be denoted by r , a positive number which is denoted by a , and a real number to be denoted by s . To describe the corresponding vector field, consider first the case where q is the north pole. Introduce

spherical coordinates (θ, ϕ) on S^2 by viewing S^2 as the set of length 1 vectors in $\mathbb{R}^3$ and then writing the Cartesian coordinates (x_1, x_2, x_3) of any given point on S^2 as $(x_1 = \sin \theta \cos \varphi, x_2 = \sin \theta \sin \varphi, x_3 = \cos \theta)$. Granted this definition, set

$$(2.22) \quad v_{(q,r,a,s)} \equiv s\chi_{r,a}(\theta) \frac{\partial}{\partial \varphi} \text{ where } \chi_{r,a} \text{ is } \chi \left(\frac{1}{a}(\cos \theta - r) \right) \chi \left(\frac{1}{a}(r - \cos \theta) \right)$$

Note in particular that $v_{(q,r,a,s)}$ is equal to $s \frac{\partial}{\partial \varphi}$ where $\cos \theta = r$ and that it has compact support where $|\cos \theta - r| < a$. The definition of $v_{(q,r,a,s)}$ for any other choice of q is the push-forward of the north polar version of $v_{(q,r,a,s)}$ depicted above by the rotation of S^2 that moves the north pole to q along the short great circle between them. (In the case of the south pole, the formula is given by the analog of (2.22) with θ replaced by $\pi - \theta$ and with the s changed to $-s$.) Let $\mathcal{V}$ denote the set of all such $v_{(q,r,a,s)}$. The next paragraph and the subsequent ones point out some of the salient properties of the vector fields in the set $\mathcal{V}$.

The first point of note is this: Given a point $p \in \mathfrak{p}$, and a non-zero vector $\nu \in TS^2|_p$, there are vectors fields in $\mathcal{V}$ that are zero on a neighborhood in S^2 of $\mathfrak{p} - \{p\}$ and a positive multiple of ν at p . To find such a vector field, fix a Gaussian coordinate chart centered at p . Use this chart to identify a small radius disk around p with the corresponding radius disk about the origin in $\mathbb{R}^2$. Take the radius of this disk to be much less than the distance to any point in $\mathfrak{p} - p$. The Gaussian coordinate chart is uniquely specified by requiring that it identify v with a vector at the origin pointing in the positive direction along the x axis. Now fix a point (which will be q) in the disk in $\mathbb{R}^2$ on the positive y axis and take r to be the distance from the origin to this point. Take the number a to be any positive number that is much less than r . The corresponding $v_{(q,r,a,+)}$ is a non-zero, positive multiple of ν at the point p and it is zero at all other points in $\mathfrak{p}$.

The second point of note is an immediate consequence of the first: If the points of $\mathfrak{p}$ are labeled as $(p_1, \dots, p_{2n})$ and if $(v_1, \dots, v_{2n})$ are chosen vectors in the tangent space to TS^2 at the correspondingly labeled point in $\mathfrak{p}$, then there are choices for $2n$ vector fields from $\mathcal{V}$ (these are labeled as $\nu_1, \dots, \nu_{2n}$) that have disjoint supports, with each ν_k being zero on a neighborhood of each point in $\mathfrak{p} - p_k$ and being a positive multiple of v_k at p_k . (This follows from the construction in the preceding paragraph). This disjoint support property implies in particular that the corresponding 1-parameter family of diffeomorphisms are pairwise commuting.

A third point is that this construction can be smoothly parametrized by certain auxiliary spaces. The first relevant instance is this: Fix a small radius disk about a point $p \in \mathfrak{p}$ such that each point in the disk

is much closer to p than to any other point in $\mathfrak{p}$. Let D_1 denote this disk, and let D_0 denote the concentric disk whose radius is $\frac{1}{100}$ times that of D_1 . The construction described above can be used to produce a smooth map from D_0 to the space of area preserving diffeomorphisms of S^2 such that if $q \in D_0$, then the corresponding diffeomorphism maps p to q and fixes $S^2 - D_1$.

The second relevant instance is a generalization of the first along the following lines: For each $k \in \{1, \dots, 2n\}$, let $D_{k_1} \subset S^2$ denote a very small radius disk centered on $p_k \in \mathfrak{p}$ whose points are much closer to p_k than to any other point in $\mathfrak{p}$. Let $D_{k_0} \subset D_{k_1}$ denote the concentric disk whose radius is $\frac{1}{100}$ times that of D_{k_1} . Then there is a smooth map from $\times_{k=1, \dots, 2n} D_{k_0}$ to the space of area preserving diffeomorphisms of S^2 with the property that if $\mathfrak{q} \equiv (q_1, \dots, q_k) \in \times_{k=1, \dots, 2n} D_{k_0}$ is any given point, then the corresponding diffeomorphism (denoted later by $\phi_{q,p}$) maps p_1 to q_1 and p_2 to q_2 , and so on. The diffeomorphism will also fix any point not in $\times_{k=1, \dots, 2n} D_{k_1}$.

2.7. The $\mathcal{I}_{(\cdot)}$ eigenvalues as functions on $\mathcal{C}_{2n}$. The proposition that follows directly is the gateway to the rest of this paper as it describes the variation of the eigenvalues of the Laplacian on the $\mathfrak{p} \in \mathcal{C}_{2n}$ versions of $\mathbb{H}_{\mathfrak{p}}$. To set the stage for the proposition: Keep in mind that $\mathcal{C}_{2n}$ is a smooth manifold. In particular, its tangent space at any given configuration $\mathfrak{p}$ can be identified with the vector space $\oplus_{p \in \mathfrak{p}} TS^2|_p$. By way of notation: The proposition refers to notation introduced in Lemma 2.3 and in the subsequent paragraph for the asymptotics near points in a given $\mathcal{C}_{2n}$ configuration of an associated $\mathcal{I}_{(\cdot)}$ -eigsection.

Proposition 2.6. *Having fixed $\mathfrak{p} \in \mathcal{C}_{2n}$, let $\lambda_{\mathfrak{p}}$ denote a corresponding $\mathcal{I}_{\mathfrak{p}}$ -eigenvalue.*

- *If $\lambda_{\mathfrak{p}}$ has multiplicity 1: There is an open neighborhood of $\mathfrak{p}$ in $\mathcal{C}_{2n}$ and a smooth function $\lambda_{(\cdot)}$ on this neighborhood with the following properties:*
 - 1) *The value of $\lambda_{(\cdot)}$ at $\mathfrak{p}$ is $\lambda_{\mathfrak{p}}$.*
 - 2) *If $\mathfrak{q}$ is in this neighborhood, then $\lambda_{(\mathfrak{q})}$ is an $\mathcal{I}_{\mathfrak{q}}$ -eigenvalue with multiplicity one.*
 - 3) *The pairing of the differential of the function $\lambda_{(\cdot)}$ at $\mathfrak{p}$ with any given vector $\nu \in \oplus_{p \in \mathfrak{p}} TS^2|_p$ is the number*

$$\frac{\pi}{2} \sum_{p \in \mathfrak{p}_f} \langle \nu|_p, \Re(\mathfrak{a}_p(f)^2 dz) \rangle$$

with f denoting here an $\mathcal{I}_{\mathfrak{p}}$ -eigsection whose eigenvalue is $\lambda_{\mathfrak{p}}$ and whose square has S^2 integral equal to one.

- *If $\lambda_{\mathfrak{p}}$ has multiplicity $N > 1$: There exist $\epsilon > 0$ and a set of N continuous functions on an open neighborhood of $\mathfrak{p}$, these having the following properties:*

- 1) All N of the functions are equal to $\lambda_{\mathbf{p}}$ at the point $\mathbf{p}$.
- 2) If $\mathbf{q}$ is in the aforementioned open neighborhood, then the values of these N functions at $\mathbf{q}$ are the $\mathcal{I}_{\mathbf{q}}$ -eigenvalues that are in the interval $(\lambda_{\mathbf{p}} - \epsilon, \lambda_{\mathbf{p}} + \epsilon)$.
- 3) Let $\mathbb{V}$ denote the span of the $\mathcal{I}_{\mathbf{p}}$ -eigenvectors with eigenvalue $\lambda_{\mathbf{p}}$. Having fixed a vector $\nu \in \oplus_{p \in \mathbf{p}} TS^2|_p$, let $(\eta_1, \dots, \eta_N)$ denote the eigenvalues of the bilinear form on $\mathbb{V}$ whose value on any two $\mathcal{I}_{\mathbf{p}}$ -eigensections f and f' from $\mathbb{V}$ is

$$\frac{\pi}{2} \sum_{p \in \mathbf{p}_f \cap \mathbf{p}_{f'}} \langle \nu|_p, \Re(\mathbf{a}_p(f) \mathbf{a}_p(f') dz) \rangle$$

Let $\gamma : (-\epsilon, \epsilon) \rightarrow \mathcal{C}_{2n}$ parametrize a smooth path in the open set with $\gamma(0) = \mathbf{p}$ and with tangent vector ν at $\mathbf{p}$. The values of the N functions at the configuration $\gamma(t)$ differ from the N numbers $\{\lambda_{\mathbf{p}} + t\eta_1, \dots, \lambda_{\mathbf{p}} + t\eta_N\}$ by $o(t)$.

Proof of Proposition 2.6. This is an instance of Proposition 2.5 given what is said in Sections 2.7 and 2.8. q.e.d.

2.8. Critical points of $\mathcal{I}_{(\cdot)}$ eigenvalues as functions on $\mathcal{C}_{2n}$. Suppose here that $\mathbf{p} \in \mathcal{C}_{2n}$ and that λ is an $\mathcal{I}_{\mathbf{p}}$ -eigenvalue with multiplicity 1. The top bullet of Proposition 2.6 says λ is the restriction to $\mathbf{p}$ of a smooth function on a neighborhood of $\mathbf{p}$ whose restriction to any given point (call it $\mathbf{q}$) in this neighborhood is an $\mathcal{I}_{\mathbf{q}}$ -eigenvalue with multiplicity 1. Denote this function by $\lambda_{(\cdot)}$. The lemma that follows makes a formal assertion to the effect that this function $\lambda_{(\cdot)}$ has no critical points.

Lemma 2.7. *A smooth function on an open set in $\mathcal{C}_{2n}$ whose value at any given configuration is an $\mathcal{I}_{(\cdot)}$ -eigenvalue with multiplicity 1 can not have critical points on that open set.*

Proof of Lemma 2.7. The tangent space to $\mathcal{C}_{2n}$ at a given point $\mathbf{p} \in \mathcal{C}_{2n}$ is $\oplus_{p \in \mathbf{p}} TS^2|_p$. With this understood, suppose that the given eigenvalue function $\mathbf{p} \rightarrow \lambda_{(\mathbf{p})}$ has multiplicity one at each point where it is defined and smoothly varying. Item (3) of the first bullet of Proposition 2.6 depicts the gradient of $\lambda_{(\cdot)}$ at a given configuration $\mathbf{p}$. The important point here is that the gradient is zero if and only if the set $\mathbf{p}_f$ is empty, which is to say that the corresponding $\mathcal{I}_{\mathbf{p}}$ -eigensection f near any given $p \in \mathbf{p}$ is $\mathcal{O}(\text{dist}(\cdot, p)^{\frac{3}{2}})$.

Keeping the preceding in mind, introduce now the three vector fields $\{\partial_1, \partial_2, \partial_3\}$ that generate the action of $SO(3)$ on S^2 . These can be written using the Cartesian coordinates (x_1, x_2, x_3) on $\mathbb{R}^3$ as

$$(2.23) \quad \partial_1 = x_2 \frac{\partial}{\partial x_3} - x_3 \frac{\partial}{\partial x_2}, \quad \partial_2 = x_3 \frac{\partial}{\partial x_1} - x_1 \frac{\partial}{\partial x_3}, \quad \partial_3 = x_1 \frac{\partial}{\partial x_2} - x_2 \frac{\partial}{\partial x_1}.$$

(Although these are depicted as vector fields on $\mathbb{R}^3$, each is tangent to the spheres of constant radius centered at the origin.) Because these vector fields commute with the Laplacian on S^2 , the sections $\{\partial_a f\}_{a=1,2,3}$ of $\mathcal{I}_{\mathfrak{p}}$ obey the equation $d^\dagger d(\cdot) = \lambda_{\mathfrak{p}}(\cdot)$ if f does. Moreover, each of these sections of $\mathcal{I}_{\mathfrak{p}}$ are in $\mathbb{H}_{\mathfrak{p}}$ precisely when $\mathfrak{p}_f = \emptyset$ because this last condition says that f is $\mathcal{O}\left(\text{dist}(\mathfrak{p}, \cdot)^{\frac{3}{2}}\right)$ near $\mathfrak{p}$. (If f were only $\mathcal{O}\left(\text{dist}(p, \cdot)^{\frac{1}{2}}\right)$ near some $p \in \mathfrak{p}$, then this conclusion would be false for at least one section from the set $\{\partial_a f\}_{a=1,2,3}$.) It follows as a consequence that $\lambda_{\mathfrak{p}}$ has multiplicity greater than one if $\mathfrak{p}_f$ is empty. In fact, with a little more work, one sees that $\lambda_{\mathfrak{p}}$ has multiplicity at least four.

Because $\lambda_{\mathfrak{p}}$ is assumed to have multiplicity one, it follows from the preceding observation that $\mathfrak{p}_f$ can't be empty and thus $\mathfrak{p}$ can't be a critical point of $\lambda_{(\cdot)}$. q.e.d.

3. The lowest eigenvalue as a function on $\mathcal{C}_{2n}$

The lemma that follows directly is the gateway for this section. It concerns the function on $\mathcal{C}_{2n}$ that assigns to any given configuration the smallest $\mathcal{I}_{(\cdot)}$ -eigenvalue. (The value of this function at a given configuration $\mathfrak{p}$ is the number $E_{\mathfrak{p}}$ from (1.2).)

Lemma 3.1. *The function $E_{(\cdot)}$ on any given positive integer n version of $\mathcal{C}_{2n}$ is a continuous function which is smooth on a neighborhood of any configuration where $E_{(\cdot)}$ has multiplicity 1 as an $\mathcal{I}_{(\cdot)}$ -eigenvalue. Conversely, it is not differentiable where it has multiplicity greater than 1.*

Proof of Lemma 3.1. The claim that $E_{(\cdot)}$ is smooth on a neighborhood of any configuration where $E_{(\cdot)}$ has multiplicity 1 is an instance of Proposition 2.6. The claim that $E_{(\cdot)}$ is not differentiable where it has multiplicity greater than 1 is a consequence of Proposition 2.6. To see why is this, fix a configuration $\mathfrak{p} \in \mathcal{C}_{2n}$ where $E_{\mathfrak{p}}$ has multiplicity at least 2. Let V denote the linear span of the $\mathcal{I}_{\mathfrak{p}}$ -eigenfunction with the eigenvalue $E_{\mathfrak{p}}$. The first observation is that there is at least one $\mathcal{I}_{\mathfrak{p}}$ -eigensection from V with at least one of the numbers from the set $\{n_p(\cdot)\}_{p \in \mathfrak{p}}$ equal to zero (see Lemma 2.3 for the definition). The reason is this: If f is an $\mathcal{I}_{\mathfrak{p}}$ -eigensection with all of the $\{n_p(f)\}_{p \in \mathfrak{p}}$ being positive, then the action on f of rotation vector fields depicted in (2.23) done some finite number of times will produce an $\mathcal{I}_{\mathfrak{p}}$ -eigensection with the same eigenvalue as f and with at least one of the $n_p(\cdot)$'s equal to zero. (Looking ahead, Lemma 3.9 asserts that at least n of the values of $n_p(\cdot)$ are zero for any non-trivial element from V .) With the preceding understood, it follows then that there exists at least one point $p \in \mathfrak{p}$ and a non-zero $\nu \in TS^2|_p$ for which the corresponding quadratic form from Item (3) of the second bullet in Proposition 2.6 is not identically zero. This quadratic form has either rank 1 or rank 2. In any case, it

follows from that same Item (3) that the function $E_{(\cdot)}$ near $\mathfrak{p}$ is at best a Lipschitz function of local coordinates for $\mathcal{C}_{2n}$ defined near $\mathfrak{p}$.

To elaborate: The vector ν is such that only the component in $TS|_{\mathfrak{p}}$ is non-zero. Supposing that this component has unit length, then it can be written with respect to a chosen stereographic coordinate system based at $\mathfrak{p}$ as $\nu|_{\mathfrak{p}} = \operatorname{Re}\left(e^{i\theta} \frac{\partial}{\partial z}\right)$ with $\theta \in \mathbb{R}/2\pi\mathbb{Z}$. Now let $V_{\mathfrak{p}} \subset V$ denote the kernel of the map sending f to $\mathfrak{a}_{\mathfrak{p}}(f)^2$. By assumption, this kernel is not all of V , so there exists $f \in V$ such that $\mathfrak{a}_{\mathfrak{p}}(f) \neq 0$ and such that f is orthogonal to $V_{\mathfrak{p}}$. Fix such an $\mathcal{I}_{\mathfrak{p}}$ -eigensection f and write $\mathfrak{a}_{\mathfrak{p}}(f)$ as $\mu e^{i\alpha}$ with $\mu > 0$ and $\alpha \in [0, 2\pi]$. The function f very near to p has the form $f = \mu \operatorname{Re}(e^{i\alpha} \sqrt{z}) + \mathcal{O}(|z|^{3/2})$. (The specified ray for the $\sqrt{z}$ branch cut is not relevant in what follows.) Granted this notation, the reader is encouraged to verify that the $\epsilon \rightarrow 0$ limit of the right most integral that appears on the left hand side of the identity in (2.11) when $f' = f$ is equal to $\frac{\pi}{2} \mu^2 \cos(\theta - 2\alpha)$. It follows as a consequence that at least one of the diagonal elements of the quadratic form that is depicted in Item (3) from the second bullet of Proposition 2.6 is non-zero except possibly in the special case where $\theta = 2\alpha \pm \frac{\pi}{2}$. Now if $V/V_{\mathfrak{p}}$ has dimension 1, then that diagonal element is an eigenvalue of the quadratic form which is non-zero unless $\theta = 2\alpha \pm \frac{\pi}{2}$; and all other eigenvalues are zero. If θ is not this special, then it follows from what is said at the end of this same Item (3) that the smallest eigenvalue function $E_{(\cdot)}$ along the path $t \rightarrow \gamma(t)$ (with tangent vector $\nu = \nu|_{\mathfrak{p}}$) is at best a Lipschitz function of t . This is because any given $t\eta_k$ changes sign as t crosses zero when η_k is not zero.

To continue, suppose that $\dim(V/V_{\mathfrak{p}}) = 2$. Then there is a second vector $f' \in V$ which is orthogonal to $V_{\mathfrak{p}}$ and to f such that $\mathfrak{a}_{\mathfrak{p}}(f')$ is not a real multiple of $\mathfrak{a}_{\mathfrak{p}}(f)$. As a consequence, the diagonal element in Item (3)'s quadratic form corresponding to f' has the form $\frac{\pi}{2} \mu'^2 \cos(\theta - 2\alpha')$ with $\mu > 0$ and with $\alpha' \neq \alpha \pm \pi$. This implies that for generic θ , the non-zero diagonal elements in Item(3)'s quadratic form are not equal. As a consequence of that, the two non-zero eigenvalues of that quadratic form are also not equal. As in the previous case, it follows from what is said at the very end of Item (3) in the second bullet of Proposition 2.6 that the smallest eigenvalue function $E_{(\cdot)}$ along that same path $t \rightarrow \gamma(t)$ is at best a Lipschitz function of t . q.e.d.

The remaining parts of this section study the behavior of this function $E_{(\cdot)}$.

3.1. A compactification of $\mathcal{C}_{2n}$. A smooth function on a compact manifold has critical points which reflect the topology of the manifold via Morse theory. By way of comparison, a function on a non-compact manifold need not have critical points. For example: Take the compliment in any compact manifold of the critical points of any given

function. With the preceding pathology in mind, it proves useful to introduce a compactification of $\mathcal{C}_{2n}$ that is compatible with the function $E_{(\cdot)}$. The description of this compactification and its topology has two parts.

Part 1. Introduce some terminology: A $\mathbb{Z}/2\mathbb{Z}$ divisor on S^2 is an equivalence class of finite, formal sums of the form

$$(3.1) \quad \sum_{p \in \mathfrak{p}} m_p p$$

with $\mathfrak{p}$ denoting an unordered, finite set of distinct points and with m_p denoting either 0 or 1. (The set $\mathfrak{p}$ is called a *configuration set*.) The *degree* of the divisor is the number of terms in the sum with $m_p = 1$. Two such formal sums are deemed to be equivalent if they differ by terms of the form $m_p p$ with $m_p = 0$. This understood, each equivalence class is represented in a unique way by a formal sum with $m_p = 1$ for all p . This representative is said to be the *minimal representative* because its configuration set has the minimal possible number of elements. By way of an example: An element in any positive integer n version of $\mathcal{C}_{2n}$ defines a minimal representative of a $\mathbb{Z}/2\mathbb{Z}$ divisor with degree $2n$. Conversely, the configuration set of the minimal representative of any $\mathbb{Z}/2\mathbb{Z}$ divisor defines an element in $\mathcal{C}_{2n'}$ with n' being half the degree of the divisor.

Part 2. Let $\overline{\mathcal{C}}_{2n}$ denote the union of the set of $\mathbb{Z}/2\mathbb{Z}$ divisors with even degree at most $2n$ (including degree 0). Thus, as a point set, this is the disjoint union

$$(3.2) \quad \overline{\mathcal{C}}_{2n} = \mathcal{C}_{2n} \cup \mathcal{C}_{2n-2} \cup \cdots \cup \mathcal{C}_0$$

with $\mathcal{C}_0$ denoting here the unique minimal representative of the equivalence class of $\mathbb{Z}/2\mathbb{Z}$ divisors with degree 0. This part of the subsection defines a topology on $\overline{\mathcal{C}}_{2n}$ that renders it compact.

A basis for the topology of $\overline{\mathcal{C}}_{2n}$ consist of sets labeled by a positive number to be denoted by ϵ , an integer from the set $\{0, \dots, n\}$ to be denoted by k , and a degree $2k$ divisor to be denoted by $[\mathfrak{q}]$. The corresponding set is denoted by $\mathcal{U}(\epsilon, k, [\mathfrak{q}])$; the criteria for membership is given next. A $\mathbb{Z}/2\mathbb{Z}$ divisor $[\mathfrak{p}]$ is in the set $\mathcal{U}(\epsilon, k, [\mathfrak{q}])$ when $[\mathfrak{p}]$ and $[\mathfrak{q}]$ have respective representative sums $\sum_{p \in \mathfrak{p}} m_p p$ and $\sum_{q \in \mathfrak{q}} m_q q$ with $\mathfrak{p}$ having a partition as a disjoint union of non-empty subsets that are indexed by the points in the configuration set $\mathfrak{q}$ representing the divisor $[\mathfrak{q}]$ such that the following is true: The partition subset $\mathfrak{p}(q) \subset \mathfrak{p}$ labeled by any given point $q \in \mathfrak{q}$ obeys

(3.3)

- $\text{dist}(p, q) < \epsilon$ for all $p \in \mathfrak{p}(q)$.
- $\sum_{p \in \mathfrak{p}(q)} m_p = m_q \pmod{2}$.

(The definition of membership in $\mathcal{U}(\epsilon, k, [\mathbf{q}])$ does not depend on the chosen representative $\sum_{p \in \mathfrak{p}} m_p p$ of the equivalence class $[\mathfrak{p}]$ because adding a term $m_{p'} p'$ to the formal sum $\sum_{p \in \mathfrak{p}} m_p p$ with $m_{p'} = 0$ can be accommodated by adding the same $m_{p'} p'$ to the formal sum that represents the equivalence class of the $\mathbb{Z}/2\mathbb{Z}$ divisor $[\mathbf{q}]$.)

A point to note about these $\mathcal{U}(\dots)$ basis sets is that they are all contractible. Indeed, any given $\mathcal{U}(\epsilon, k, [\mathbf{q}])$ deformation retracts on to corresponding $[\mathbf{q}]$; the retraction moves the set of points from each partition subset $\mathfrak{p}(q)$ that appears in (3.3) onto q by moving each point from $\mathfrak{p}(q)$ along the short great circle arc between the point and q .

The proposition that follows summarizes three relevant properties of this space $\overline{\mathcal{C}}_{2n}$.

Proposition 3.2. *Fix an integer n so as to define the spaces $\mathcal{C}_{2n}$ and $\overline{\mathcal{C}}_{2n}$. Then*

- *The space $\overline{\mathcal{C}}_{2n}$ is compact*
- *The subset $\mathcal{C}_{2n}$ is open and dense in $\overline{\mathcal{C}}_{2n}$.*
- *The lowest eigenvalue as a function on $\mathcal{C}_{2n} \cup \mathcal{C}_{2n-2} \cup \dots \cup \mathcal{C}_0$ is continuous with respect to the topology on $\overline{\mathcal{C}}_{2n}$.*

The rest of this subsection is occupied with the proposition's proof.

Proof of Proposition 3.2. Some background is needed in order to prove the first two bullets: The permutation group on k letters (k is a positive integer) acts on $\times_k S^2$ by permuting the factors. The quotient is denoted by $\text{Sym}_k(S^2)$ which is a compact space. It is homeomorphic to the space of non-trivial, homogeneous, degree $2n$ polynomials in the coordinates of $\mathbb{C}^2$ modulo the action of $\mathbb{C}^*$, thus, $\mathbb{CP}^{2n}$. The homeomorphism sends a polynomial to its $2n$ zeros on $\mathbb{CP}^1$. Said differently, any polynomial of the sort under consideration can be factored as

$$(3.4) \quad \prod_{1 \leq j \leq 2n} (\alpha_j z_1 + \beta_j z_2)$$

where the notation has (z_1, z_2) denoting the coordinates for $\mathbb{C}^2$; and where each pair (α_j, β_j) are not both zero and are defined modulo $\mathbb{C}^*$. The $k = 2n$ version is relevant because the manifold $\mathcal{C}_{2n}$ appears as the dense open set in $\mathbb{CP}^{2n}$ that consists of the equivalence classes of polynomials with distinct factors (the $\mathbb{C}^*$ equivalence classes of the pairs $\{(\alpha_j, \beta_j)\}$ are the constituent points of a configuration in $\mathcal{C}_{2n}$).

With regards to the first bullet: Let $\{\mathfrak{p}_j\}_{j \in \mathbb{N}}$ denote a sequence in $\mathcal{C}_{2n}$ (degree $2n$ polynomials with distinct factors modulo $\mathbb{C}^*$) with no convergent subsequence in $\mathcal{C}_{2n}$. It will have a subsequence that converges in $\mathbb{CP}^{2n}$, thus to a homogeneous, degree $2n$ polynomial. That polynomial will have some factors appearing multiple times. Discard all factors that appear an even number of times and replace all factors that appear an odd number of times with that factor appearing just once. This new polynomial corresponds to a point in a $\mathcal{C}_{2k}$ term in (3.1) for some $k < n$. (But for notation, the preceding proves that any sequence in any $\mathcal{C}_{2k}$ term in (3.1) has a limit point in some $\mathcal{C}_{2j}$ term for $j \leq k$.)

With regards to the second bullet and $\mathcal{C}_{2n}$ being dense: Any point in any $k < 2n$ version of $\mathcal{C}_{2k}$ (a homogeneous, degree $2k$ polynomial with distinct roots) can be obtained as a limit of polynomials with $2n$ distinct roots in the manner just described by appending $2(n - k)$ linear factors with nearby but distinct roots in $\mathbb{CP}^1$ and then letting those added roots come together along the sequence. In this way, the original polynomial when viewed as a point in any $\mathcal{C}_{2k}$ term of (3.2), appears as a limit point of some non-convergent sequence in $\mathcal{C}_{2n}$.

With regards to the second bullet and $\mathcal{C}_{2n}$ being open: Let $\mathfrak{p}$ denote a point in $\mathcal{C}_{2n}$. Take $\epsilon > 0$ but less by a factor of $\frac{1}{100}$ than the minimum of the distances between $\mathfrak{p}$'s constituent points. Take $k = n$ and $[\mathfrak{q}] = [\mathfrak{p}]$. The corresponding set $\mathcal{U}(\epsilon, n, [\mathfrak{p}])$ contains only points from $\mathcal{C}_{2n}$.

The proof of the proposition's third bullet has four parts. To set the stage for this, take $\{\mathfrak{p}_i\}_{i \in \mathbb{N}}$ to be a sequence in $\mathcal{C}_{2n}$ that converges in $\bar{\mathcal{C}}_{2n}$ to a point in some $\mathcal{C}_{2k}$ term in (3.2). This configuration is denoted by $\mathfrak{q}$. The third part of the proof of the third bullet explains why $\lim_{i \rightarrow \infty} E_{\mathfrak{p}_i} \geq E_{\mathfrak{q}}$; and the fourth part explains why $\lim_{i \rightarrow \infty} E_{\mathfrak{p}_i} \leq E_{\mathfrak{q}}$. The first two parts of the proof set the stage for the latter two parts.

Part 1. The definition of convergence implies first that there exists a set of distinct points $\mathfrak{k} \subset S^2 - \mathfrak{q}$ with at most $2(n - k)$ members with the following significance: Given $\epsilon > 0$, then all sufficiently large i versions of $\mathfrak{p}_i$ have the following two properties: To state the first, suppose for the moment that $q \in \mathfrak{q}$ and let $D_\epsilon(q)$ denote the radius ϵ disk centered at q . This disk contains some odd number of constituents of $\mathfrak{p}_i$. No generality is lost by assuming that this number is independent of ϵ and the index i when i is sufficiently large with ‘large’ determined by the choice of ϵ . To state the second, let y denote a point from $\mathfrak{k}$. Then the corresponding radius ϵ disk centered on y contains some non-zero, even number of points from $\mathfrak{p}_i$. The number of points in each such disk can be taken to be independent ϵ and i when i is large with ‘large’ again determined by ϵ .

A sequence of choices for ϵ will be made in what follows with the sequence decreasing and converging to zero. An upper bound for the

elements in this sequence is chosen so that each version of $\sqrt{\epsilon}$ is much less than the distance between any two distinct points from the set $\mathfrak{q} \cup \mathfrak{k}$.

Part 2. Now comes a key observation: With ϵ small (with upper bound as just described), then each sufficiently large i version of the real line bundle $\mathcal{I}_{\mathfrak{p}_i}$ is isomorphic to the line bundle $\mathcal{I}_{\mathfrak{q}}$ on the complement of the union of the radius ϵ disks centered at the points in $\mathfrak{q} \cup \mathfrak{k}$. This is because all of the points in any large i version of $\mathfrak{p}_i$ are contained in these disks as are the points of $\mathfrak{q}$; and because a line bundle's isomorphism class on $S^2 - \cup_{z \in \mathfrak{q} \cup \mathfrak{k}} D_\epsilon(z)$ is determined by the isomorphism classes on the boundaries of the disks from the set $\{D_\epsilon(z) : z \in \mathfrak{q} \cup \mathfrak{k}\}$; and because the respective isomorphism classes in the case of $\mathfrak{p}_i$ and $\mathfrak{q}$ are the same when i is large (the Möbius bundle if $z \in \mathfrak{q}$ and the product bundle if $z \in \mathfrak{k}$). By the same token, the line bundles $\mathcal{I}_{\mathfrak{p}_i}$ and $\mathcal{I}_{\mathfrak{q}}$ are isomorphic on the complement of the union of the radius 2ϵ disks centered at the points in $\mathfrak{p}_i$.

Part 3. For each index i , let f_i denote a chosen $\mathcal{I}_{(\cdot)}$ -eigensection for the $\mathfrak{p}_i$ version of the Laplacian with eigenvalue λ . (In the instance of the proposition, $\lambda = E_{\mathfrak{p}_i}$.) It is assumed in what follows that the S^2 integral of each $|f_i|^2$ is equal to 1. Define the function χ_ϵ as in (2.2) with $\mathfrak{p}$ denoting $\mathfrak{p}_i$. This version will be denoted by $\chi_{\mathfrak{p}_i, \epsilon}$. Use the isomorphism from Part 2 to view $\chi_{\mathfrak{p}_i, \epsilon} f_i$ as an element in $\mathbb{H}_{\mathfrak{q}}$. Viewed in this light, the S^2 integrals of the square of the norms of $d(\chi_{\mathfrak{p}_i, \epsilon} f_i)$ and $\chi_{\mathfrak{p}_i, \epsilon} f_i$ obey

(3.5)

$$\begin{aligned} \bullet \int_{S^2} |d(\chi_{\mathfrak{p}_i, \epsilon} f_i)|^2 &= \int_{S^2} |df_i|^2 + \mathfrak{e}_{1,i}, \\ \bullet \int_{S^2} |\chi_{\mathfrak{p}_i, \epsilon} f_i|^2 &= \int_{S^2} |f_i|^2 + \mathfrak{e}_{2,i}, \end{aligned}$$

with $\mathfrak{e}_{1,i}$ and $\mathfrak{e}_{2,i}$ obeying

$$(3.6) \quad |\mathfrak{e}_{1,i}| \leq c_0 \frac{1}{|\ln \epsilon|} n \lambda \quad \text{and} \quad |\mathfrak{e}_{2,i}| \leq c_0 \epsilon n \lambda.$$

To explain: Both bounds are obtained by invoking the bound on $|f_i|$ from the second bullet of Lemma 2.1. The $\mathfrak{e}_{1,i}$ bound follows from that sup-norm bound and the fact that the S^2 integral of $|d\chi_{\mathfrak{p}_i, \epsilon}|^2$ is bounded by $c_0 n \frac{1}{|\ln \epsilon|}$. The $\mathfrak{e}_{2,i}$ bound follows from the sup-norm bound and the fact that the area of the support of $\chi_{\mathfrak{p}_i, \epsilon}$ is bounded by $c_0 n \epsilon$. The $\lambda = E_{\mathfrak{p}_i}$ versions of (3.5) and (3.6) imply directly that $\lim_{i \rightarrow \infty} E_{\mathfrak{p}_i} \geq E_{\mathfrak{q}}$.

Part 4. Let f denote a chosen, minimal eigenvalue eigensection for the $\mathfrak{q}$ version of the Laplacian. (Its eigenvalue is $E_{\mathfrak{q}}$.) Having fixed positive ϵ to be very small, then $\chi_\epsilon f$ can be viewed as a section of each

sufficiently large i version of the line bundle $\mathcal{I}_{\mathbf{p}_i}$ by choosing an isomorphism between $\mathcal{I}_{\mathbf{q}}$ and the large i versions of $\mathcal{I}_{\mathbf{p}_i}$ on the complement of the support of χ_ϵ . But for terminology, the argument that lead to (3.5) and (3.6) can be repeated to see that these same inequalities hold with f_i replaced by f and with respective error terms ϵ_1 and ϵ_2 that obey the bounds in (3.6) with λ replaced by $E_{\mathbf{q}}$. Granted that these inequalities hold for each sufficiently small ϵ , it then follows directly that $\lim_{i \rightarrow \infty} E_{\mathbf{p}_i} \leq E_{\mathbf{q}}$ also. q.e.d.

3.2. How large can $E_{(\cdot)}$ get? Proposition 3.2 implies that the supremum of $E_{(\cdot)}$ on $\mathcal{C}_{2n}$ is a non-decreasing function of the integer n . The following two propositions say more about this supremum.

Proposition 3.3. *There are configurations in $\mathcal{C}_{2n}$ with smallest Laplace eigenvalue greater than $\kappa^{-1}n$ with $\kappa > 1$ and independent of n .*

This proposition is proved momentarily. A consequence of Propositions 3.2 and 3.3 follows directly.

Corollary 3.4. *There exists an unbounded set of positive integers with the following property: Supposing that n is from this set, there is a configuration in $\mathcal{C}_{2n}$ where $E_{(\cdot)}$ is equal to its supremum of $E_{(\cdot)}$ over the whole of $\mathcal{C}_{2n}$.*

The rest of this subsection is dedicated to proving Proposition 3.3.

Proof of Proposition 3.3. The proof has three parts.

Part 1. Fix a small positive number to be denoted by R and then fix a maximal (in number) configuration of points in S^2 such that the distance between any two distinct points is no less than R . Let $\mathbf{p}_0$ denote this configuration. Each point in S^2 has distance at most R from some point from the set $\mathbf{p}_0$ because $\mathbf{p}_0$ is maximal. The number of elements in $\mathbf{p}_0$ (denoted by $|\mathbf{p}_0|$) is such that

(3.7)
$$c_0^{-1} \frac{1}{R^2} \leq |\mathbf{p}_0| \leq c_0 \frac{1}{R^2}.$$

Note that the disks of radius R centered at the points in $\mathbf{p}_0$ cover S^2 . Also, the disks of radius $\frac{1}{2}R$ centered at these points are pairwise disjoint.

Assign to each point from $\mathbf{p}_0$ a pair of points that lie at distance $\frac{1}{8}R$ from the given point on opposite sides of a chosen great circle through that point. When q denotes the chosen point from $\mathbf{p}_0$, then these other two points are denoted by q_+ and q_- . Let $\mathbf{p}$ denote the configuration consisting of all such (q_-, q_+) pairs coming from the points in $\mathbf{p}_0$. Since the number of points in the configuration $\mathbf{p}$ is twice the number from $\mathbf{p}_0$, this number also obeys the inequality in (3.7) with a possibly different c_0 . An important point for what follows: The disks of radius $\frac{1}{8}R$ centered at the points from the set $\mathbf{p}$ are pairwise disjoint.

Part 2. What follows is the fundamental lemma about sections of $\mathcal{I}_{\mathfrak{p}}$:

Lemma 3.5. *There exists $\kappa > 1$ such that the following is true when $R < \frac{1}{\kappa}$. Fix $q \in \mathfrak{p}_0$ and let D denote the radius R disk centered at q . If f is any section of $\mathcal{I}_{\mathfrak{p}}$ over D with $|df|^2$ having finite integral over D , then*

$$\int_D |df|^2 \geq \frac{1}{\kappa R^2} \int_D f^2.$$

Proof of Lemma 3.5. Let D' denote the radius $\frac{1}{8}R$ disk centered at either q_+ or q_- and let f denote a section of $\mathcal{I}$ over D' . The integral of $|df|^2$ over D' is no less than $c_0^{-1} \frac{1}{R^2}$ times that of f^2 because f must vanish at one or more points on any concentric circle about the center point of D' . This is because $\mathcal{I}_{\mathfrak{p}}$ restricts to any of these circles as the Möbius line bundle.

With the preceding understood, there are now two cases to consider: The first is when the integral of f^2 over the union of the radius $\frac{1}{8}R$ disks centered at q_+ and q_- is greater than $\frac{1}{1000}$ times its integral over D . In this case, the assertion in the lemma follows from what was said in the previous paragraph about the integral of $|df|^2$ over the q_+ and q_- versions of D' . In the second case, the integral of f^2 over D' is no greater than $\frac{1}{1000}$ times its integral over the union of the q_+ and q_- versions of D' . In this case, the assertion of the lemma follows from the fact that $|f|$ isn't uniformly close to the constant function on D in a suitable topology. Indeed, by rescaling D to have radius 1, the bound amounts to a straightforward observation about the Dirichlet energy of honest functions on the radius 1 disk in $\mathbb{R}^2$. q.e.d.

Part 3. Proposition 3.3 follows directly from the next lemma and (3.7).

Lemma 3.6. *There exists $\kappa > 1$ such that the following is true when $R < \frac{1}{\kappa}$. Construct $\mathfrak{p}$ as described above and let f denote any given section of $\mathcal{I}_{\mathfrak{p}}$ from the Hilbert space $\mathbb{H}_{\mathfrak{p}}$. Then*

$$\int_{S^2} |df|^2 \geq \frac{1}{\kappa R^2} \int_{S^2} f^2.$$

Proof of Lemma 3.6. The integral of $|df|^2$ over S^2 is no less than c_0^{-1} times the sum of its integrals over the various $q \in \mathfrak{p}$ versions of D . This is because at most c_0 of these disks intersect at any one point. Thus, the sum of the integrals of $|df|^2$ over these disks overcounts the integral of $|df|^2$ by a factor of at most c_0^{-1} .

$$(3.8) \quad \int_{S^2} |df|^2 \geq c_0^{-1} \sum_{q \in \mathfrak{p}_0} \int_{D(q)} |df|^2,$$

with $D(q)$ denoting here the radius R disk centered at q . With (3.8) in hand, appeal to Lemma 3.5 to see that

(3.9)
$$\int_{S^2} |df|^2 \geq c_0^{-1} \frac{1}{R^2} \sum_{q \in \mathfrak{p}_0} \int_{D(q)} f^2,$$

Meanwhile, the sum of the various $D(q)$ integrals of f^2 for $q \in \mathfrak{p}_0$ is no less than the integral of f^2 over S^2 because the union of these $D(\cdot)$'s covers S^2 . This last observation with (3.9) implies directly what is asserted by the lemma. q.e.d.

3.3. No critical values of $E_{(\cdot)}$. Corollary 3.4 asserts in effect that there are infinitely many values of n where $E_{(\cdot)}$ on $\mathcal{C}_{2n}$ has a global maximum. However, according to Lemma 2.7, these are not configurations where the lowest eigenvalue has multiplicity 1; and so there is no guarantee that $E_{(\cdot)}$ has a differential at these points (let alone, one that vanishes). Even so, with the second bullet of Proposition 2.6 in mind, one can ask whether $E_{(\cdot)}$ has a critical value in the sense of the next definition. This definition refers to the collection of numbers $\{n_p(\cdot)\}_{p \in \mathfrak{p}}$ from Lemma 2.3. To set the notation: Supposing that k is a positive integer, the definition refers to the k 'th *eigenvalue function*, λ_k , whose value at any given configuration $\mathfrak{p}$ is the k 'th lowest $\mathcal{I}_{\mathfrak{p}}$ -eigenvalue of the Laplacian. Thus $\lambda_1(\mathfrak{p}) = E_{\mathfrak{p}}$.

Definition 3.7. A configuration $\mathfrak{p} \in \mathcal{C}_{2n}$ is k -critical if there exists an $\mathcal{I}_{\mathfrak{p}}$ -eigenfunction f with eigenvalue $\lambda_k(\mathfrak{p})$ with all numbers from the set $\{n_p(f)\}_{p \in \mathfrak{p}}$ being positive.

The following lemma makes a formal statement to the effect that there are no 1-critical points in $\mathcal{C}_{2n}$.

Lemma 3.8. *If $\mathfrak{p} \in \mathcal{C}_{2n}$, and if f is an $\mathcal{I}_{\mathfrak{p}}$ -eigensection with eigenvalue $E_{\mathfrak{p}}$, then at least $n + 1$ integers from the set $\{n_p(f)\}_{p \in \mathfrak{p}}$ are equal to zero.*

With a look towards finding homogeneous $\mathbb{Z}/2\mathbb{Z}$ harmonic 1-forms on $\mathbb{R}^3$: They can't correspond (as described in Section 1.6) to $\mathcal{I}_{\mathfrak{p}}$ -eigensections with the lowest eigenvalue $E_{\mathfrak{p}}$ because f is an $\mathcal{I}_{\mathfrak{p}}$ -eigensection and if $n_p(f) = 0$ for a given $p \in \mathfrak{p}$, then $|f|$ near p is greater than a non-zero, constant multiple of $\text{dist}(\cdot, p)^{1/2}$.

Proof of Lemma 3.8. Let f denote a given $\mathcal{I}_{\mathfrak{p}}$ -eigensection with eigenvalue $E_{\mathfrak{p}}$, and given f , let $\mathfrak{c}$ denote the set of points in $S^2 - \mathfrak{p}$ where both f and df are zero. Also, let $\mathfrak{Z}$ denote the zero locus of $|f|$ (including the points in $\mathfrak{p}$). As explained directly, the set $\mathfrak{Z}$ is an embedded graph in S^2 with $\mathfrak{p} \cup \mathfrak{c}$ being its vertex set. To see this, note first that the complement of $\mathfrak{p} \cup \mathfrak{c}$ in $\mathfrak{Z}$ consists of a disjoint union of embedded open arcs because it is a level set of f where $df \neq 0$. With regards to $\mathfrak{Z}$ near

points in $\mathfrak{p}$: Lemma 2.3 depicts f near any given point in $\mathfrak{p}$, and the depiction implies that $\mathfrak{Z}$ intersects a very small, radius disk centered at p as the union of $2n_p + 1$ embedded, closed arcs with one endpoint at p and the other on the boundary of the disk. (The angle at p between consecutive arcs impinging on p is $2\pi/(2n_p + 1)$.) Thus, $\mathfrak{Z}$ near any given $p \in \mathfrak{p}$ is homeomorphic to a neighborhood of a vertex in a graph with that vertex having $2n_p + 1$ incident half-edges. With regards to $\mathfrak{Z}$ near points in $\mathfrak{c}$: If $p \in \mathfrak{c}$, then f near p (which is an $f^{-1}(0)$ critical point) has a Taylor's expansion with respect to a complex coordinate centered at that point (it is denoted by z below) which looks like

$$(3.10) \quad f = \Re(\alpha z^{m_p}) + \mathcal{O}(|z|^{m_p+1})$$

with m_p being a positive integer greater than 1 and with α being a non-zero complex number. In this case, $\mathfrak{Z}$ near p is homeomorphic to a vertex in a graph with that vertex having $2m_p$ incident half-edges.

Given this graph structure for $\mathfrak{Z}$, it follows that its Euler characteristic is

$$(3.11) \quad \chi_3 = 2n + |\mathfrak{c}| - \frac{1}{2} \sum_{p \in \mathfrak{p}} (2n_p + 1) - \frac{1}{2} \sum_{p \in \mathfrak{c}} 2m_p.$$

If χ_3 is non-positive, then $\mathfrak{Z}$ must have a closed cycle because a tree has Euler characteristic 1 and the Euler characteristic of a union of trees is the number of trees in the union. As explained momentarily, $\mathfrak{Z}$ can't have any closed cycles. Assuming this, then positivity of χ_3 requires that n be greater than $\sum_{p \in \mathfrak{p}} n_p$ which requires in turn that at least $n+1$ points in $\mathfrak{p}$ have $n_p = 0$.

To see why $\mathfrak{Z}$ can't have any closed cycles, assume to the contrary that it does to generate nonsense. If it does, then its complement in S^2 is disconnected. Assuming that, let Ω denote one of the components if $S^2 - \mathfrak{Z}$. Now let f_Ω denote the section of $\mathcal{I}_{\mathfrak{p}}$ that is equal to f on $S^2 - \Omega$ and equal to $-f$ on Ω . Supposing that the S^2 integral of f^2 is 1, then this is also the case for the S^2 integral of f_Ω^2 . And, it is also the case that the S^2 integral of $|df_\Omega|^2$ is equal to $E_{\mathfrak{p}}$. Even so, f_Ω can't be an eigensection for the Laplacian because $f + f_\Omega$ would then be an eigensection and $f + f_\Omega$ is zero on an open set. But if f_Ω isn't an eigensection, then the integral of $|df_\Omega|^2$ must be greater than $E_{\mathfrak{p}}$ since $E_{\mathfrak{p}}$ is the minimal eigenvalue of the $\mathcal{I}_{\mathfrak{p}}$ -Laplacian and since f_Ω can, in any event, be written as a linear combination of its eigensections. This last conclusion is required. q.e.d.

3.4. The k 'th lowest eigenvalue. Supposing that k denotes a positive integer, introduce the function λ_k on $\mathcal{C}_{2n}$ whose value at any given configuration $\mathfrak{p}$ is the k 'th lowest $\mathcal{I}_{\mathfrak{p}}$ -eigenvalue. To be precise about the definition: The value of $\lambda_k(\cdot)$ at $\mathfrak{p}$ is the supremum of the set of numbers $s \in [0, \infty)$ such that there are at most $k-1$ linearly independent $\mathcal{I}_{\mathfrak{p}}$ -eigensections with eigenvalue less than s . It follows from this

definition that $\lambda_k(\mathbf{p})$ is an $\mathcal{I}_{\mathbf{p}}$ -eigenvalue. It also follows from this definition that if the number of linearly independent $\mathcal{I}_{\mathbf{p}}$ -eigensections with eigenvalue less than $\lambda_k(\mathbf{p})$ is less than $k - 1$, then that number plus the multiplicity of $\lambda_k(\mathbf{p})$ is greater than or equal to k . The function $\lambda_k(\cdot)$ is continuous on $\mathcal{C}_{2n}$ and it is smooth on some neighborhood of any configuration where it has multiplicity 1. (These assertions follow from Proposition 2.6.) Lemma 3.8 says that there are no configurations where λ_k is 1-critical (in the sense of Definition 3.7). The next lemma says that $\lambda_k(\cdot)$ can't have a $\mathbf{p}$ -critical configuration for small values of k .

Lemma 3.9. *Fix $k \leq n$; and suppose that $\mathbf{p} \in \mathcal{C}_{2n}$, and that f is an $\mathcal{I}_{\mathbf{p}}$ -eigensection with eigenvalue λ_k . Then at least $n+1-k$ integers from the set $\{n_p(f)\}_{p \in \mathbf{p}}$ are equal to zero.*

Proof of Lemma 3.9. Assume that f has eigenvalue λ and that $n_p(f) \geq 1$ for a subset of n or more points $p \in \mathbf{p}$. Write the size of this set as $n + \ell$. It then follows from (3.11) that the graph $\mathfrak{Z}$ (which is the $f = 0$ locus in S^2) has Euler characteristic $-\ell$ or less. This implies in turn that the complement of $\mathfrak{Z}$ in S^2 has at least $\ell + 2$ components. Letting K denote the number of components of $S^2 - \mathfrak{Z}$, label these components as $\{\Omega_\alpha\}_{\alpha=1,2,\dots,K}$. Now define f_α for $\alpha \in \{1, \dots, K\}$ to be the section of $\mathcal{I}_{\mathbf{p}}$ that is given by the rule $f_\alpha \equiv f$ on Ω_α and $f_\alpha \equiv 0$ on $S^2 - \Omega_\alpha$. These K -sections span a subspace of $\mathbb{H}_{\mathbf{p}}$ where the inequality

$$(3.12) \quad \int_{S^2} |dh^2| \leq \lambda \int_{S^2} |h|^2$$

holds. Since f is the only actual eigensection in this span (otherwise, there would be an eigensection vanishing on an open set), the complement of f in this subspace must project injectively to the span of the $\mathcal{I}_{\mathbf{p}}$ -eigensections with eigenvalue less than λ . Thus, there must be at least $\ell + 1$ of the latter since $K \geq \ell + 2$. Thus λ can't be the k 'th eigenvalue for $k \leq \ell + 1$. Turning this around, if λ is the k 'th eigenvalue, then $\ell \leq k - 1$ and thus at least $n - (k - 1)$ of the integers from the set $\{n_p(f)\}_{p \in \mathbf{p}}$ are zero. q.e.d.

The last lemma in this section concerns the differentiability of the k^{th} -eigenvalue functions; it is a $k > 1$ analog of Lemma 3.1 which is about the $k = 1$ case.

Lemma 3.10. *For any positive integer n and positive integer k : The k 'th eigenvalue function λ_k is differentiable where λ_k has multiplicity 1. On the other hand, it is not differentiable at any configuration in $\mathcal{C}_{2n}$ where the eigenvalue λ_k has multiplicity greater than 1 and where there are $k-1$ linearly independent $\mathcal{I}_{\mathbf{p}}$ -eigensections with eigenvalue strictly less than λ_k .*

Proof of Lemma 3.10. The proof is almost the same as the proof of Lemma 3.1. q.e.d.

4. The $\mathbb{RP}^\infty$ bundle

This section introduces an $\mathbb{RP}^\infty$ fiber bundle over $\mathcal{C}_{2n}$ whose topological properties explain (in part) why $E_{(\cdot)}$ and the other k 'th eigenvalue functions behave the way they do.

4.1. No universal bundle over $\mathcal{C}_{2n}$. To set the stage for what is to come: Let $\mathcal{F} \rightarrow \mathcal{C}_{2n}$ denote the fiber bundle whose fiber over any given configuration $\mathbf{p} \in \mathcal{C}_{2n}$ is $S^2 - \mathbf{p}$. To elaborate on the fiber bundle structure: Given a configuration $\mathbf{p}$ from $\mathcal{C}_{2n}$, Section 2.6 describes a map (denoted there and here by $\phi_{\mathbf{p}}$) from the product of small radius disks about the points comprising $\mathbf{p}$ to the space of area preserving diffeomorphisms of S^2 . By way of a reminder: If the points that comprise $\mathbf{p}$ are labeled as $(p_1, \dots, p_{2n})$, then the disk around a given p_k from $\mathbf{p}$ was denoted by D_{0k} . (Note that the radii of these disks can depend on k .) Set $U_{\mathbf{p}} \equiv \times_{k=1, \dots, 2n} D_{0k}$, this projecting (after symmetrization) diffeomorphically onto an open neighborhood of $\mathbf{p}$ in $\mathcal{C}_{2n}$. (Borrowing terminology from complex analysis: Sets of this form will be called *poly-disk neighborhoods of $\mathbf{p}$*). The map $\phi_{\mathbf{p}}$ from $U_{\mathbf{p}}$ into the space of area preserving diffeomorphisms of S^2 has the following property: Supposing that $\mathbf{q} = (q_1, \dots, q_{2n})$ denotes a given configuration in $U_{\mathbf{p}}$, then the corresponding diffeomorphism (which was denoted in Section 2.6 by $\phi_{\mathbf{q}, \mathbf{p}}$) sends each p_k to the corresponding q_k . This map $\phi_{\mathbf{p}}$ gives $\mathcal{F}$ its local product structure as a fiber bundle because it gives a identification between $U_{\mathbf{p}} \times (S^2 - \mathbf{p})$ and $\mathcal{F}|_{U_{\mathbf{p}}}$ that is compatible with the projection map to $U_{\mathbf{p}}$.

Each $\mathbf{p} \in \mathcal{C}_{2n}$ has its corresponding line bundle $\mathcal{I}_{\mathbf{p}}$ sitting over the fiber of $\mathcal{F}$ at $\mathbf{p}$ (which is $S^2 - \mathbf{p}$); and one might hope that these fit together to define a global line bundle over $\mathcal{F}$. As explained in what follows, this is not the case.

To see that there is no such universal line bundle, note first that $\mathcal{F}$ has a second fibration with base S^2 : The fibering map sends $(\mathbf{p}, x)$ with $\mathbf{p}$ from $\mathcal{C}_{2n}$ and x from $S^2 - \mathbf{p}$ to x . The fiber is diffeomorphic to the space of $2n$ unordered, distinct points in $\mathbb{R}^2$. (This space is denoted by C_{2n}^0 .) Let D_+ and D_- denote the (closed) northern and southern hemispheres of S^2 . Their intersection is the equatorial circle which will be denoted by C . Restricting the fibration $\mathcal{F} \rightarrow S^2$ to D_+ and D_- and C leads to a Mayer-Vietoris sequence for the $\mathbb{Z}/2\mathbb{Z}$ homology of $\mathcal{F}$ that starts with

$$(4.1) \quad 0 \longrightarrow H^1(\mathcal{F}) \longrightarrow H^1(\mathcal{F}|_{D_+}) \oplus H^1(\mathcal{F}|_{D_-}) \longrightarrow H^1(\mathcal{F}|_C) \longrightarrow \dots$$

(Note that all cohomology henceforth is defined also using $\mathbb{Z}/2\mathbb{Z}$ coefficients.)

To exploit this exact sequence, note first that D_+ retracts onto any given point in D_+ (and similarly for D_-), and that these retractions are covered by corresponding deformation retractions of $\mathcal{F}$ onto the fiber at that point. (A deformation retraction of $\mathcal{F}|_{D_+}$ can be constructed with parameter $t \in [0, 1]$ that acts on any given $(\mathbf{p}, x)$ by first pushing the points of $\mathbf{p}$ radially away from x to lie outside a disk concentric to D_+ with slightly larger radius; and having done that, the deformation retraction then moves x via a linear deformation retract of D_+ onto any give point in D_+ .)

There are two immediate consequences: First, $H^1(\mathcal{F}|_{D_+}) = H^1(C_{2n}^0)$ which is $\mathbb{Z}/2\mathbb{Z}$ (see [2]), and likewise $H^1(\mathcal{F}|_{D_-}) = \mathbb{Z}/2\mathbb{Z}$. The second is that their respective homomorphisms to $H^1(\mathcal{F}|_C)$ are injective. (To prove injectivity: The fiber preserving deformation retraction of $\mathcal{F}|_{D_+}$ onto $\mathcal{F}|_{\text{point}}$ can be made when the point in question is in the circle C . Because of this, the isomorphism $H^1(\mathcal{F}|_{D_+}) \rightarrow H^1(\mathcal{F}|_{\text{point} \in C})$ factors as the chain of pull-back homomorphisms $H^1(\mathcal{F}|_{D_+}) \rightarrow H^1(\mathcal{F}|_C) \rightarrow H^1(\mathcal{F}|_{\text{point} \in C})$ induced by the inclusions of first the point in C and then C in D_+ .) These two consequences imply in turn that the rank of $H^1(\mathcal{F})$ is at most 1. (This argument is more general: Suppose that $n \geq 2$ and that $\mathcal{F}$ is a fiber bundle over S^n with path connected fiber whose restriction to the respective northern and southern hemispheres of S^n is diffeomorphic as a fiber bundle to the product bundle. There is a corresponding version of (4.1) in this case which implies that the rank of the first homology of X is no greater than the rank of the first homology of the fiber.)

Keeping in mind that the rank of $H^1(\mathcal{F})$ is at most 1, now consider the fibration $\pi : \mathcal{F} \rightarrow \mathcal{C}_{2n}$. The group $H^1(\mathcal{C}_{2n})$ is also isomorphic to $\mathbb{Z}/2\mathbb{Z}$ (see [8], Corollary 3.1). Denote its generator by α . The element $\pi^*\alpha$ is necessarily non-zero in $H^1(\mathcal{F})$ because the fiber of π is path connected. Thus, $\pi^*\alpha$ generates $H^1(\mathcal{F})$ and, because it is a pull-back, it restricts as 0 to each fiber. Because of this, its restriction to any $\mathbf{p} \in \mathcal{C}_{2n}$ version of $\mathcal{F}|_{\mathbf{p}} = S^2 - \mathbf{p}$ can not be the first Stieffel-Whitney class of the line bundle $\mathcal{I}_{\mathbf{p}}$.

By way of a parenthetical remark: The obstruction to the existence of this universal $\mathcal{I}$ bundle is carried by a class in $H^2(\mathcal{C}_{2n})$ which can be constructed in the Čech version as follows: Fix a locally finite open cover of $\mathcal{C}_{2n}$ by sets of the form $U_{\mathbf{p}}$. Having done that, define $\hat{I}_{\mathbf{p}}$ over $\pi^{-1}U_{\mathbf{p}}$ to be the pull-back of $\mathcal{I}_{\mathbf{p}}$ by the diffeomorphism from $\pi^{-1}U_{\mathbf{p}}$ to $U_{\mathbf{p}} \times (S^2 - \mathbf{p})$ whose inverse sends any given pair $(\mathbf{q}, x)$ to $\phi_{\mathbf{q}, \mathbf{p}}(x)$. Now suppose that $\mathbf{p}'$ is a second configuration from $\mathcal{C}_{2n}$ and suppose that $U_{\mathbf{p}}$ and $U_{\mathbf{p}'}$ intersect. The bundles $\hat{I}_{\mathbf{p}}$ and $\hat{I}_{\mathbf{p}'}$ are isomorphic over $U_{\mathbf{p}} \cup U_{\mathbf{p}'}$ because they have the same first Stieffel-Whitney class. Let $\eta_{\mathbf{p}' \mathbf{p}}$ denote an isometry between them. (There are precisely two such isometries.) Now if $\mathbf{p}''$ is a third configuration from $\mathcal{C}_{2n}$ with $U_{\mathbf{p}''}$ sharing points with

$U_p \cap U_{p'}$, then there are corresponding isomorphisms $\eta_{p''p'}$ and $\eta_{pp''}$. Composing these three leads to a self-isometry of $\hat{I}_p$ which is denoted by $z_{pp'p''} \in \{\pm 1\}$. The collection of these indexed by 3-tuples of points whose sets define the locally finite open cover form a 2-dimensional Čech cocycle on $\mathcal{C}_{2n}$ whose cohomology class is independent of the choices for various versions of $\eta_{pp'}$. This cohomology class is the obstruction class in $H^2(\mathcal{C}_{2n})$.

By way of another remark: Fix a point $x \in \mathcal{C}_{2n}$ and let $\mathcal{C}_{2n}^x$ denote the codimension 2 submanifold consisting of configurations that contain x . The obstruction class for the universal line bundle vanishes over $\mathcal{C}_{2n} - \mathcal{C}_{2n}^x$ because it is the Alexander dual to $\mathcal{C}_{2n}^x$. Note in this regard that $H^2(\mathcal{C}_{2n})$ is isomorphic to $\mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z}$ with $\mathcal{C}_{2n}^x$ being dual to a generator (see Corollary 3.1 in [8]). As explained momentarily, the vanishing of this obstruction class on $\mathcal{C}_{2n} - \mathcal{C}_{2n}^x$ is an instance of the upcoming Lemma 4.1.

To set the stage for the lemma, suppose that X denotes a path connected manifold and that $\pi : \mathcal{Z} \rightarrow X$ denotes a smooth fiber bundle with path connected fiber. There is an associated locally constant sheaf over X whose stalk at any given point is the $\mathbb{Z}/2\mathbb{Z}$ first cohomology of its inverse image via π ; thus $H^1(\pi^{-1}(x))$. This sheaf is denoted by $\mathcal{H}^1$. A section of this sheaf is a continuous assignment over X of a class in the first cohomology of the fibers. For example, the restriction of a class in $H^1(\mathcal{Z})$ to each fiber defines a section of $\mathcal{H}^1$. Lemma 4.1 gives a partial converse, a condition to guarantee that a section of $\mathcal{H}^1$ comes from an H^1 -class (With regards to notation: Keep in mind that all cohomology uses $\mathbb{Z}/2\mathbb{Z}$ coefficients.)

Lemma 4.1. *Let $\pi : \mathcal{Z} \rightarrow X$ denote a fiber bundle over X . A global section of the sheaf $\mathcal{H}^1$ comes from a class in $H^1(\mathcal{Z})$ when the fiber bundle $\mathcal{Z}$ admits a section.*

In the instance of the fiber bundle $\pi : \mathcal{F} \rightarrow \mathcal{C}_{2n}$, the section of the sheaf $\mathcal{H}^1$ at any given $p \in \mathcal{C}_{2n}$ is the class in $H^1(\mathcal{F}|_p)$ that has non-zero pairing with all sufficiently small radius linking circles around the points from p . Because the restriction $\mathcal{F}$ to $\mathcal{C}_{2n} - \mathcal{C}_{2n}^x$ admits the section $p \rightarrow (p, x)$, the relevant section of the sheaf $\mathcal{H}^1$ comes from a class in $H^1(\mathcal{Z})$. Since classes in H^1 correspond to real line bundles, there is a real line bundle over $\mathcal{C}_{2n} - \mathcal{C}_{2n}^x$ whose restriction to any given fiber of $\mathcal{F}$ is isomorphic to the corresponding real line bundle $I_{(\cdot)}$ on that fiber.

Proof of Lemma 4.1. The argument has three parts. By way of a heads up for what is to come: Because the classes in H^1 with $\mathbb{Z}/2\mathbb{Z}$ coefficients are in 1-1 correspondence with real line bundles, the arguments can be given in terms of real line bundles. This is what is done in the four parts of the proof. (Also, the existence of a real line bundle is of ultimate interest in the special case where $X = \mathcal{C}_{2n}$ and $\mathcal{Z} = \mathcal{C}_{2n}$.)

Part 1. Let Z denote a given fiber of $\mathcal{Z}$. Now fix a locally finite, good open cover of X (to be denoted by $\mathfrak{U}$) such that each set $U \in \mathfrak{U}$ comes with a diffeomorphism $\Psi_U : \mathcal{Z}|_U \rightarrow U \times Z$ that intertwines the projection map π with the product projection to U . (A good cover is one where all of its sets and all non-empty intersections of its sets is contractible.) The product space $U \times Z$ has the class from $H^1(Z)$ that restricts to each fiber as the pull-back via the Ψ_U^{-1} of the section from $\mathcal{H}^1$. Fix a real line bundle on $U \times Z$ that represents this cohomology class (for example, a product cocycle). The pull-back of this real line bundle by Ψ_U defines a real line bundle on $\mathcal{Z}|_U$ whose first Stiefel-Whitney class restricts to each fiber as the given section of $\mathcal{H}^1$. This real line bundle is denoted by I_U in what follows.

Part 2. Suppose that U and U' are sets from $\mathfrak{U}$ that share points. The respective real line bundles I_U and $I_{U'}$ on $\mathcal{Z}|_{U \cap U'}$ have the same Stiefel-Whitney classes; and so they are isomorphic, and hence isometric. Fix an isometry $\eta_{UU'} : I_{U'} \rightarrow I_U$. Where three sets from $\mathfrak{U}$ overlap, say U , U' , and U'' , the composition of their corresponding versions of η determine a self-isometry of I_U . This is

$$(4.2) \quad \eta_{UU'U''} \equiv \eta_{UU'} \eta_{U'U''} \eta_{UU''}.$$

Since there are just two self-isometries (multiplication by $+1$ and by -1), the collection $\{\eta_{UU'U''}\}_{U,U',U'' \in \mathfrak{U}}$ defines a degree 2 Čech cocycle on X (it is a priori closed by virtue of its definition via (4.2)). The corresponding class in the second Čech cohomology of X is the obstruction class in the context of the fiber bundle $\mathcal{F}$. (Changing the sign of any number of $\eta_{UU'}$'s changes this cocycle by a coboundary which has no effect on the cohomology class).

If this degree 2 obstruction class is zero, then the collection $\{\eta_{UU'}\}_{U,U' \in \mathfrak{U}}$ are a collection of transition functions that glue the various $U \in \mathfrak{U}$ versions of I_U together so as to define a line bundle over the whole $\mathcal{Z}$.

Part 3. The task now is to show that the class defined by the cocycle $\{\eta_{UU'U''}\}_{U,U',U'' \in \mathfrak{U}}$ is the zero class if $\mathcal{Z}$ admits a section. To this end, it is important to keep in mind that an isometry between two real line bundles is determined globally by its identification between the corresponding lines over just a single point. (The reason is this: If I and I' are isomorphic line bundles, then an isometry between them is a unit length section of $I \otimes I'$ which is isomorphic to the product bundle. And, there are only two unit length sections of the product bundle, respective multiplication by $+1$ and by -1 .)

With the preceding understood, fix a point x_U in each set U from the cover, and then fix a unit length point $\hat{u}_U$ in the fiber of I_U over that point. Use the section $s : X \rightarrow \mathcal{Z}$ and the fact that U is contractible to

define a product structure for I_U over $s(U)$ with $\hat{i}_U$ defining a preferred component. Granted these product structures, choose the isomorphism $\eta_{UU'}$ to identify the $\hat{i}_{U'}$ component of $I_{U'}$ with the $\hat{i}_U$ component of I_U .

Now consider what happens in a triple intersections, $U \cap U' \cap U''$: The right most $\eta_{U''U}$ isomorphism in (4.2) identifies the $\hat{i}_U$ component of I_U with the $\hat{i}_{U''}$ component of $I_{U''}$; and the middle isomorphism in (4.2), which is $\eta_{U'U''}$, identifies that component of $I_{U''}$ with the $\hat{i}_{U'}$ component of $I_{U'}$; and then the left most isomorphism in (4.2) identifies the $\hat{i}_{U'}$ component of $I_{U'}$ with the $\hat{i}_U$ component of I_U . Hence, the full composition in (4.2) identifies the $\hat{i}_U$ component in I_U with itself; it is multiplication by $+1$. Since this is the case for all of the various $\eta_{UU'U''}$, the degree 2 obstruction class vanishes. q.e.d.

4.2. No universal Hilbert space bundle. The observation in Section 4.1 regarding the non-existence of a universal line bundle over $\mathcal{C}_{2n}$ implies in turn that there is no Hilbert space bundle over $\mathcal{C}_{2n}$ whose fiber over $\mathfrak{p}$ is the corresponding $\mathbb{H}_{\mathfrak{p}}$. To see this, suppose for the sake of argument that such a bundle exists (call it $\mathbb{H}$). Given $\mathbb{H}$, one could construct (using a partition of unity) a finite set of sections of $\mathbb{H}$ to be denoted by $\{s_1, \dots, s_N\}$ with the following properties:

(4.3)

- For any given $\mathfrak{p} \in \mathcal{C}_{2n}$ each section from $\{s_1, \dots, s_N\}$ at $\mathfrak{p}$ is a smooth element in the corresponding $\mathbb{H}_{\mathfrak{p}}$.
- Given $\mathfrak{p} \in \mathcal{C}_{2n}$ and $x \in S^2 - \mathfrak{p}$, there exists $s \in \{s_1, \dots, s_N\}$ with $s|_{\mathfrak{p}}$ being non-zero at x .

These conditions imply that the assignment of $(\mathfrak{p}, x)$ to the N -tuple

$$(s_1|_{\mathfrak{p}}(x), \dots, s_N|_{\mathfrak{p}}(x))$$

defines a smooth map from $\mathcal{F}$ to $\mathbb{RP}^N$. Let $\mathcal{I}$ denote the pull-back by this map of the tautological line bundle over $\mathbb{RP}^N$. The restriction of $\mathcal{I}$ to each $\mathfrak{p} \in \mathcal{C}_{2n}$ version of $\mathcal{F}|_{\mathfrak{p}}$ is necessarily isomorphic to $\mathcal{I}_{\mathfrak{p}}$. The existence of $\mathcal{I}$ and, by implication, the existence of $\mathbb{H}$, runs afoul of the non-existence result from Section 4.1.

By way of a summary: The obstruction to the existence of a universal $\mathbb{H}$ bundle is the obstruction class in $H^2(\mathcal{C}_{2n})$ that was described in the preceding subsection.

4.3. A universal $\mathbb{RP}^\infty$ bundle. Although there is no universal Hilbert space bundle, there is a universal $\mathbb{RP}^\infty$ bundle over $\mathcal{C}_{2n}$ whose fiber over any given $\mathfrak{p} \in \mathcal{C}_{2n}$ is the quotient by ± 1 of the subspace $\mathbb{S}_{\mathfrak{p}} \subset \mathbb{H}_{\mathfrak{p}}$ that is defined by the rule whereby $f \in \mathbb{S}_{\mathfrak{p}}$ if and only if

$$(4.4) \quad \int_{S^2} f^2 = 1.$$

The $\mathbb{RP}^\infty$ bundle will be denoted by $\mathbb{RP}$.

The topology on this bundle is generated by a certain basis of open neighborhoods which is described momentarily. To set notation for this: Suppose for the moment that $\mathfrak{p}$ is a given point in $\mathcal{C}_{2n}$. Section 4.1 introduced the notion of a polydisk neighborhood of $\mathfrak{p}$, denoted there by $U_{\mathfrak{p}}$. By way of a reminder, the polydisk neighborhood $U_{\mathfrak{p}}$ is the projection to $\mathcal{C}_{2n}$ of a product of disks in $\times_{2n} S^2$ with these disks centered at the constituents of $\mathfrak{p}$ and chosen so that the set of concentric disks with 100 times the radius have pairwise disjoint closure. (The disks in this product can have different radii.) A convention for what follows: The radii of these disks are constrained by an upper bound that guarantees the existence of a disk of radius $\frac{1}{100n}$ that is disjoint from each of them.

For the purpose of what is to come, the collection of disks that define $U_{\mathfrak{p}}$ is denoted by $\{D_{0k}\}_{k=1,\dots,2n}$ and the set of respective 100 times larger radius concentric disks is denoted by $\{D_{1k}\}_{k=1,\dots,2n}$. An important point to keep in mind: If $\mathfrak{q}$ is a configuration from $U_{\mathfrak{p}}$, then the bundles $\mathcal{I}_{\mathfrak{q}}$ and $\mathcal{I}_{\mathfrak{p}}$ are isomorphic on $S^2 - \cup_{k=1,\dots,2n} D_{1k}$. (This is because they have the same first Stieffel-Whitney classes on this set.) Also keep in mind that there are precisely two isometries between these bundles on $S^2 - \cup_{k=1,\dots,2n} D_{1k}$, one is obtained from the other by fiber-wise multiplication by -1 .

By way of more notation: Supposing that $\mathfrak{p}$ again denotes a configuration in $\mathcal{C}_{2n}$, and supposing that f is from $\mathbb{S}_{\mathfrak{p}}$, then the $\{\pm 1\}$ equivalence class of f is denoted in what follows by $[f]$.

With preceding notation understood, suppose that $(\mathfrak{p}, [f])$ is as just described. A basis for the open neighborhoods in $\mathbb{RP}$ of $(\mathfrak{p}, [f])$ is given by sets labeled by data consisting a polydisk neighborhood in $\mathcal{C}_{2n}$ of $\mathfrak{p}$ which is denoted by $U_{\mathfrak{p}}$, an open set in $S^2 - \cup_{k=1,\dots,2n} D_{1k}$ which is denoted by Ω containing a disk of radius $\frac{1}{100n}$, and positive numbers δ and E with E greater than the integral that appears on the right hand side of (1.7), the S^2 integral of $|df|^2$. The corresponding neighborhood in $\mathbb{RP}$ consists of the set of pairs $(\mathfrak{q}, [h])$ with $\mathfrak{q} \in U_{\mathfrak{p}}$, and with $h \in \mathbb{S}_{\mathfrak{q}}$ satisfying two constraints:

- $\int_{S^2 - \mathfrak{p}} |dh|^2 < E$.
- *There is an isometry between $\mathcal{I}_{\mathfrak{p}}$ and $\mathcal{I}_{\mathfrak{q}}$ on $S^2 - \cup_{k=1,\dots,2n} D_{1k}$ that identifies h as a section of $\mathcal{I}_{\mathfrak{p}}$ obeying*

(4.5)
$$\int_{S^2 - \cup_{k=1,\dots,2n} D_{1k}} (|d(f - h)|^2 + |f - h|^2) < \delta.$$

The next lemma makes the formal assertion to the effect that $\mathbb{RP}$ fibers over $\mathcal{C}_{2n}$.

Lemma 4.2. *The space $\mathbb{RP}$ has the structure of a (topological) fiber bundle over $\mathcal{C}_{2n}$ with the projection map sending any given pair $(\mathbf{p}, [f])$ to $\mathbf{p}$.*

A notion of a differentiable structure is introduced in the next subsection.

Proof of Lemma 4.2. The proof of the lemma has three parts. The proof requires the specification of a local product structure and then the verification that the transition functions between overlapping open sets obey the required cocycle condition.

Part 1. A product structure for $\mathbb{RP}$ over a given set $U_{\mathbf{p}}$ in $\mathcal{C}_{2n}$ can be defined using the map $\phi_{\mathbf{p}}$ from Sections 2.6 and 4.1 as follows: If $\mathbf{q} \in U_{\mathbf{p}}$, then pull-back by $\phi_{\mathbf{q},\mathbf{p}}$ defines an invertible, linear map from the Hilbert space $\mathbb{H}_{\mathbf{q}}$ to the Hilbert space of sections of $\phi_{\mathbf{q},\mathbf{p}}^* \mathcal{I}_{\mathbf{q}}$ which preserves the integral in (4.4). Composing this linear map with any given isometry from $\phi_{\mathbf{q},\mathbf{p}}^* \mathcal{I}_{\mathbf{q}}$ to $\mathcal{I}_{\mathbf{p}}$ defines an invertible, linear map from $\mathbb{S}_{\mathbf{q}}$ to $\mathbb{S}_{\mathbf{p}}$. Choosing a different isometry from $\phi_{\mathbf{q},\mathbf{p}}^* \mathcal{I}_{\mathbf{q}}$ to $\mathcal{I}_{\mathbf{p}}$ multiplies this linear map by ± 1 . As a consequence, there is no ambiguity if pull-back by $\phi_{\mathbf{q},\mathbf{p}}$ is viewed as a map from $\mathbb{S}_{\mathbf{q}}/\{\pm 1\}$ to $\mathbb{S}_{\mathbf{p}}/\{\pm 1\}$. As $\mathbf{q}$ varies in $U_{\mathbf{p}}$, the map $\phi_{\mathbf{p}}$ gives a fiberwise identification between $\mathbb{RP}|_{U_{\mathbf{p}}}$ and $U_{\mathbf{p}} \times \mathbb{RP}|_{\mathbf{p}}$. As explained in Parts 2 and 3 of the proof, this identification is a homeomorphism of topological spaces when $U_{\mathbf{p}} \times \mathbb{RP}|_{\mathbf{p}}$ is given the product topology using the quotient topology for $\mathbb{RP}|_{\mathbf{p}}$ from the latter's identification with $\mathbb{S}_{\mathbf{p}}/\{\pm 1\}$ and with $\mathbb{S}_{\mathbf{p}}$ viewed as a submanifold in $\mathbb{H}_{\mathbf{p}}$. Accept that this is the case for what is said in the next paragraph.

The local product structure defines a fiber bundle topology on $\mathbb{RP}$ if and only if the corresponding transition functions obey the required cocycle condition on triple intersections (these being $U_{\mathbf{p}} \cap U_{\mathbf{p}'} \cap U_{\mathbf{p}''}$). To see this, suppose for the moment that $\mathbf{q} \in U_{\mathbf{p}} \cap U_{\mathbf{p}'}$. On the one hand, $\mathbb{RP}|_{\mathbf{q}}$ has been identified with $\mathbb{RP}|_{\mathbf{p}}$ by the rule

$$(4.6) \quad [f] \rightarrow [\phi_{\mathbf{p},\mathbf{q}}^* f].$$

(Keep in mind that pull-back commutes with multiplication by ± 1 .) There is a similar formula for the identification between $\mathbb{RP}|_{\mathbf{q}}$ and $\mathbb{RP}|_{\mathbf{p}'}$. Composing the inverse of the one depicted in (4.6) with its primed version leads to the $\mathbf{q}$ -dependent linear, invertible map from $\mathbb{RP}|_{\mathbf{p}}$ to $\mathbb{RP}|_{\mathbf{p}'}$ that is given by the rule

$$(4.7) \quad [f] \rightarrow [(\phi_{\mathbf{q},\mathbf{p}'} \circ \phi_{\mathbf{p},\mathbf{q}}^{-1})^* f].$$

Supposing that $\mathbf{q}$ is in a triple intersection, $U_{\mathbf{p}} \cap U_{\mathbf{p}'} \cap U_{\mathbf{p}''}$, then it follows from the three versions of (4.7) that the relevant cocycle for this triple intersection when evaluated at $\mathbf{q}$ is the identity map (this is because $(\phi_{\mathbf{q},\mathbf{p}} \circ \phi_{\mathbf{q},\mathbf{p}''}^{-1}) \circ ((\phi_{\mathbf{q},\mathbf{p}'} \circ \phi_{\mathbf{q},\mathbf{p}}^{-1}) \circ (\phi_{\mathbf{q},\mathbf{p}'} \circ \phi_{\mathbf{q},\mathbf{p}}^{-1}))$ is the identity map.)

Part 2. It remains now to prove that the identification between $\mathbb{R}P|_{U_{\mathfrak{p}}}$ and $U_{\mathfrak{p}} \times \mathbb{R}P|_{\mathfrak{p}}$ given by (4.6) defines a homeomorphism between the two topological spaces. The proof that such is the case is facilitated by first reinterpreting the Hilbert space topology on any given $\mathfrak{p} \in \mathcal{C}_{2n}$ version of $\mathbb{H}_{\mathfrak{p}}$. This part of the proof does that.

Consider the topology on $\mathbb{H}_{\mathfrak{p}}$ whereby a basis for the open neighborhoods of a given element f is labeled by an open neighborhood with compact closure in $S^2 - \mathfrak{p}$ that contains a disk of radius $\frac{1}{100n}$ (to be denoted by Ω); and positive numbers δ and E with E greater than the S^2 integral of $|df|^2$. The corresponding open set consists of the subset of elements $h \in \mathbb{H}_{\mathfrak{p}}$ that obey the top bullet in (4.5), and also obey the $\mathfrak{q} = \mathfrak{p}$ version of the second bullet in (4.5) with the isometry between $\mathcal{I}_{\mathfrak{p}}$ and itself being multiplication by $+1$. This basis set is denoted by $\mathcal{V}_f(\Omega, \delta, E)$.

As explained directly, this new topology is equivalent to the Hilbert space topology that is defined by the $\|\cdot\|_{\mathbb{H}}$ -metric balls (the norm $\|\cdot\|_{\mathbb{H}}$ is depicted in (1.1)). Indeed, the set $\mathcal{V}_f(\Omega, \delta, E)$ is open with respect to the Hilbert Space topology because if h is in $\mathbb{H}_{\mathfrak{p}}$ and if $\|f - h\|_{\mathbb{H}}$ is sufficiently small (given δ and E), then both the conditions in (4.5) will hold for h .

To see that the Hilbert space balls are open in this new topology, fix $\epsilon > 0$ and consider in particular the ball in $\mathbb{H}_{\mathfrak{p}}$ where $\|f - (\cdot)\|_{\mathbb{H}} < \epsilon$. This ball will be open in the new topology if each of its elements is contained in a suitable $\mathcal{V}_{(\cdot)}(\cdot)$ set from the new topology which is entirely inside the $\|f - (\cdot)\|_{\mathbb{H}} < \epsilon$ ball. In this regard: No generality is lost by proving that there is a $\mathcal{V}_f(\cdot)$ set inside the $\|f - (\cdot)\|_{\mathbb{H}} < \epsilon$ ball. To do this, fix $\delta > 0$ for the moment and then an open set with compact closure in $S^2 - \mathfrak{p}$ (denoted by Ω) such that

(4.8)
$$\int_{S^2 - \Omega} |df|^2 < \delta.$$

Let $\mathcal{E}$ denote the S^2 integral of $|df|^2$. With these definitions in hand, suppose that h is from the set $\mathcal{V}_f(\Omega, \delta, E = \mathcal{E} + \delta)$. (These basis sets for the topology are defined at the start of Section 2.) If this is the case, then

(4.9)

$$\begin{aligned} &\bullet \int_{\Omega} (|d(f - h)|^2 + |f - h|^2) < \delta. \\ &\bullet \int_{S^2 - \Omega} |d(f - h)|^2 < 2 \int_{S^2 - \Omega} (|df|^2 + |dh|^2) < 2\delta + 2 \int_{S^2 - \Omega} |dh|^2. \end{aligned}$$

To exploit this: By virtue of the top bullet in (4.9), the Ω integral of $|dh|^2$ is greater than $\mathcal{E} - 2\sqrt{\delta}\sqrt{\mathcal{E}}$. Therefore, because the integral of $|dh|^2$

over $S^2 - \mathfrak{p}$ is at most $\mathcal{E} + \delta$ (from the definition of $\mathcal{V}_f(\Omega, \delta, \mathcal{E} = E + \delta)$), its integral over $S^2 - \Omega$ must obey

$$(4.10) \quad \int_{S^2 - \Omega} |dh|^2 \leq \delta + 2\sqrt{\delta}\sqrt{\mathcal{E}}.$$

And thus, the left hand side of the inequality in the second bullet of (4.9) is no greater than $4\delta + 4\sqrt{\delta}\sqrt{\mathcal{E}}$. This last observation plus the top bullet in (4.9) implies the following: If δ is sufficiently small given $\mathcal{E}$ (and an appropriate Ω is chosen to obey (4.8)), then the S^2 integral of $|d(f - h)|^2$ will be less than $\epsilon\delta^{\frac{1}{3}}$. Meanwhile, the latter integral plus the Ω integral of $|f - h|^2$ bounds a positive multiple of the S^2 integral of $|f - h|^2$ (the multiple is independent of Ω because of the constraint that Ω must contain a disk of radius $\frac{1}{100n}$). With the preceding understood, it follows that the $\|\cdot\|_{\mathbb{H}}$ norm of $f - h$ will be less than ϵ if δ is sufficiently small and if Ω is chosen so that (4.8) holds. This says in effect that the set $\mathcal{V}_f(\Omega, \delta, E = \mathcal{E} + \delta)$ is entirely inside the $\|\cdot\|_{\mathbb{H}}$ radius ϵ ball in $\mathbb{H}_{\mathfrak{p}}$ centered at f .

Part 3. Return now to the assertion that pull-back using the map $\phi_{\mathfrak{p}}$ defines a homeomorphism between $\mathbb{RP}|_{U_{\mathfrak{p}}}$ and $U_{\mathfrak{p}} \times \mathbb{RP}|_{\mathfrak{p}}$. To do this, it is sufficient to prove that a basis set in $\mathbb{RP}|_{U_{\mathfrak{p}}}$ corresponds to an open set in $U_{\mathfrak{p}} \times \mathbb{RP}|_{\mathfrak{p}}$ and vice-versa. That this is so follows by unwinding definitions from the analysis in Part 2 and from the three observations that follow.

The first observation is that a pair consisting of a configuration $\mathfrak{q} \in U_{\mathfrak{p}}$ and an element $[f] \in \mathbb{RP}|_{\mathfrak{p}}$ is identified by the local product map with the element $(\mathfrak{q}, [(\phi_{\mathfrak{q}, \mathfrak{p}}^{-1})^* f])$ in $\mathbb{RP}$ (keep in mind that $[(\phi_{\mathfrak{q}, \mathfrak{p}}^{-1})^* f]$ is in $\mathbb{RP}|_{\mathfrak{q}}$). More to the point, the transformation given by the rule $f \rightarrow (\phi_{\mathfrak{q}, \mathfrak{p}}^{-1})^* f$ defines a norm-bounded linear map from $\mathbb{H}_{\mathfrak{p}}$ to $\mathbb{H}_{\mathfrak{q}}$ given an isometry between $(\phi_{\mathfrak{q}, \mathfrak{p}}^{-1})^* \mathcal{I}_{\mathfrak{p}}$ and $\mathcal{I}_{\mathfrak{q}}$. Conversely, if $[h] \in \mathbb{RP}|_{\mathfrak{q}}$, then $\phi_{\mathfrak{q}, \mathfrak{p}}^* h$ is in $\mathbb{H}_{\mathfrak{p}}$ given an identification between $\phi_{\mathfrak{q}, \mathfrak{p}}^* \mathcal{I}_{\mathfrak{q}}$ and $\mathcal{I}_{\mathfrak{p}}$.

The second observation is this: If Ω has compact closure in $S^2 - \mathfrak{p}$, then $\phi_{\mathfrak{q}, \mathfrak{p}}(\Omega)$ has compact closure in $S^2 - \mathfrak{p}$.

To state the third observation, suppose that $f \in \mathbb{H}_{\mathfrak{p}}$. Then (as explained in Section 2.6), a change of coordinates identifies the S^2 integral of $|d(\phi_{\mathfrak{q}, \mathfrak{p}}^{-1})^* f|^2$ with an S^2 integral that has the form

$$(4.11) \quad \int_{S^2 - \mathfrak{p}} \langle df, df \rangle_{\mathfrak{m}_{\mathfrak{q}}}$$

with $\langle \cdot, \cdot \rangle_{\mathfrak{m}_{\mathfrak{q}}}$ denoting the inner product of df with itself as defined by the $\phi_{\mathfrak{q}, \mathfrak{p}}$ -pull back of the round metric. The point is that this integral defines a bounded quadratic form on $\mathbb{H}_{\mathfrak{p}}$ that depends smoothly on $\mathfrak{q}$. Moreover, the value of the integral in (4.11) differs from the corresponding integral of $|df|^2$ by at most $c_0 \text{dist}(\mathfrak{q}, \mathfrak{p})$ times the latter integral. Thus, if the

S^2 integral of $|df|^2$ is less than a given number (call it E), then the S^2 integral of $|d(\phi_{q,p}^{-1})^* f|^2$ will be bounded by $(1 + c_0 \operatorname{dist}(q, p))E$. q.e.d.

4.4. On the differentiability of maps to and from $\mathbb{R}P$. The notion of a differentiable map to or from $\mathbb{R}P$ is problematic because the fiber-bundle structure on $\mathbb{R}P$ was defined in the topological category but not in the C^1 (let alone smooth) category. This was done because the transition function between two product neighborhood charts, versions of $U_p \times \mathbb{R}P_p$ and $U_{p'} \times \mathbb{R}P_{p'}$, are not differentiable with respect to variations of the base configuration $q \in U_p \cap U_{p'}$. To elaborate: Differentiability fails because the transition functions (depicted in (4.7)) involve the composition of fiber elements with diffeomorphisms that depend on the configuration q . And, the problem with this dependence is that the derivative with respect to variations in q of the fiber-wise linear transition map depicted in (4.7) when acting on a given element $f \in \mathbb{R}P|_p$ costs one derivative of f (this is the chain rule). As a consequence, the derivative of the linear transition map when viewed from the context of $\mathbb{H}_p$ does not define a bounded linear operator from $\mathbb{H}_p$ to $\mathbb{H}_q$. It is only bounded as a linear map from $\mathbb{H}_p$ to the L^2 completion of $\mathbb{H}_{p'}$. By way of an example: If $h \in \mathbb{H}_p$, then the linear function

(4.12)
$$f \rightarrow \int_{S^2} h f$$

pulls back over $U_p \cap U_{p'}$ via the product structure transition functions to a q -dependent linear functional on $\mathbb{H}_{p'}$ which is differentiable with respect to q (but not twice differentiable).

One can also define the notion of a *smooth* function on $\mathbb{R}P$: This is a function that is smooth when written using any of the product charts $\{U_p \times \mathbb{R}P|_p\}_{p \in \mathcal{C}_{2n}}$. The next subsection gives the most relevant example of a smooth function. An example in the meantime is defined by the rule

(4.13)
$$f \longrightarrow \int_{S^2} G f^2$$

with G being a smooth function on $\mathcal{C}_{2n}$.

With regards to maps into $\mathbb{R}P$: A map into a product chart $U_p \times \mathbb{R}P|_p$ is given by a pair $(q(\cdot), [f(\cdot)])$ of maps into U_p and into $\mathbb{R}P|_p$ respectively. Although both might be smooth, if $q(\cdot)$ maps to $U_p \cap U_{p'}$, then the image of this map via the bundle transition function need not be differentiable as a map into $U_p \cap U_{p'} \times \mathbb{R}P|_{p'}$. (This can happen even when the $[f(\cdot)]$ component is constant.) However, such a map will be differentiable (at least of class C^1) if the $[f(\cdot)]$ component over U_p lifts to a map from the domain space into a certain dense subvector space in $\mathbb{H}_p$. This vector subspace is the closure of the space of smooth sections of $\mathcal{I}_p$ with compact support in $S^2 - p$ using the norm whose square is defined by

the rule

$$(4.14) \quad f \longrightarrow \int_{S^2 - \mathfrak{p}} (|\nabla df|^2 + |df|^2 + |f|^2).$$

In this regard: The fiber-wise linear bundle transition map at each $\mathfrak{q} \in U_{\mathfrak{p}} \cap U_{\mathfrak{p}'}$ sends this dense domain in $\mathbb{H}_{\mathfrak{p}}$ in a 1-1 and onto fashion to the $\mathfrak{p}'$ version of this domain in $\mathbb{H}_{\mathfrak{p}'}$. And in so doing, the $\mathfrak{q}$ -derivatives of the bundle transition maps give a bounded linear operator from the dense domain to $\mathbb{H}_{\mathfrak{p}'}$. Said differently: The bundle transition functions for $\mathbb{RP}$ are differentiable when restricted to the vector subspace in $\mathbb{H}_{\mathfrak{p}}$ where the integral depicted in (4.14) is finite.

4.5. The function $\mathcal{E}$. Let $\mathcal{E}$ denote the function on $\mathbb{RP}$ whose value at any given $(\mathfrak{p}, [f])$ is this:

$$(4.15) \quad \mathcal{E}(\mathfrak{p}, [f]) \equiv \int_{S^2 - \mathfrak{p}} |df|^2.$$

The upcoming Lemma 4.3 uses $\mathcal{E}$ to reinterpret part of what is said by Lemma 2.3. It says in effect that k -critical points in $\mathcal{C}_{2n}$ in the sense of Definition 3.7 are the π -images of the critical points of $\mathcal{E}$ on $\mathbb{RP}$. (In this regard: With $\mathcal{E}$ viewed as a function on $\mathbb{RP}$ as is done here and below, the critical points of $\mathcal{E}$ are critical points subject to the constraint that the integral of f^2 on $S^2 - \mathfrak{p}$ is equal to 1.)

Lemma 4.3. *The function $\mathcal{E}$ from (4.15) is a smooth function on $\mathbb{RP}$. As such, its critical points are precisely the set of $\{\pm 1\}$ equivalence classes of pairs $(\mathfrak{p}, f)$ such that $\mathfrak{p} \in \mathcal{C}_{2n}$ and $f \in \mathbb{S}_{\mathfrak{p}}$ is an $\mathcal{I}_{\mathfrak{p}}$ -eigensection with all numbers from Lemma 2.3's set $\{n_p(f)\}_{p \in \mathfrak{p}}$ being positive.*

Proof of Lemma 4.3. By definition, the function $\mathcal{E}$ is smooth on $\mathbb{RP}$ if and only if it is smooth when viewed on the restriction of $\mathbb{RP}$ to any given $\mathfrak{p} \in \mathcal{C}_{2n}$ version of $U_{\mathfrak{p}}$ with $\mathbb{RP}$ identified there as $U_{\mathfrak{p}} \times (\mathbb{S}_{\mathfrak{p}}/\{\pm 1\})$. To see about that, suppose that $\mathfrak{q} \in U_{\mathfrak{p}}$ and $[f] \in \mathbb{S}_{\mathfrak{p}}/\{\pm 1\}$. The function $\mathcal{E}$ on $(\mathfrak{q}, [f])$ has the form depicted in (4.11) which is observedly jointly smooth with respect to variations in $\mathfrak{q}$ and f (see the discussion in Section 2.4).

With regards to the critical points: The function is critical at $(\mathfrak{p}, [f])$ with respect to variations in the fiber direction iff and only if f is an $\mathcal{I}_{\mathfrak{p}}$ -eigensection (this follows by definition). The claim that it is also critical with respect to variations in $\mathfrak{p}$ if and only if all numbers from the set $\{n_p(f)\}_{p \in \mathfrak{p}}$ are positive follows from Proposition 2.5 and the formula in (2.21). q.e.d.

If $\mathcal{E}$ were a Palais-Smale function on $\mathbb{RP}$, which is to say proper, then its critical points would account for topological invariants of $\mathbb{RP}$ in a Morse theoretic sense (see e.g [7] and the review [6]). Although the function $\mathcal{E}$ is proper along the fibers of $\mathbb{RP}$, it is not globally proper

because $\mathcal{C}_{2n}$ is not compact. As a consequence, it could be that $\mathcal{E}$ does not have any critical points notwithstanding the fact that $\mathbb{R}\mathbb{P}$ has plenty of homology/cohomology. (The cohomology of $\mathbb{R}\mathbb{P}$ is described in the next subsection.) What follows are some comments about this.

With regards to topology and critical points: Let C denote a non-zero class in $H_*(\mathbb{R}\mathbb{P})$. Set $\mathcal{E}_C$ to denote the infimum over all singular cycles representing the class C of the maximal value of $\mathcal{E}$ on the simplices that appear with non-zero coefficient in the cycle. This number would be a critical value of $\mathcal{E}$ were $\mathcal{E}$ a proper function. Since homology and cohomology in any given degree are dual to each other using $\mathbb{Z}/2\mathbb{Z}$ coefficients (all homology/cohomology here and in what follows uses $\mathbb{Z}/2\mathbb{Z}$ coefficients), there is also a cohomological point of view: Letting C denote a non-zero cohomology class, define $\mathcal{E}_C$ to be the infimum of values of $E \in (0, \infty)$ such that the class C is non-zero on the $\mathcal{E} < E$ part of $\mathbb{R}\mathbb{P}$. This version of $\mathcal{E}_C$ would be a critical value of $\mathcal{E}$ were $\mathcal{E}$ proper on $\mathbb{R}\mathbb{P}$. Homotopy invariants of $\mathbb{R}\mathbb{P}$ also lead to putative critical values for $\mathcal{E}$: Take the infimum over a homotopy class of maps from a compact space to $\mathbb{R}\mathbb{P}$ of the maximum value of $\mathcal{E}$ on the image of the map. (This last version of a min-max value is used in Section 6.)

Unfortunately $\mathcal{E}$ is not proper and therefore none of these hypothetical critical values of $\mathcal{E}$ must correspond to an actual critical points of $\mathcal{E}$. Examples of this pathology are provided in Section 5.2.

The homology/cohomology of $\mathbb{R}\mathbb{P}$ is the subject of the next subsection.

4.6. The homology/cohomology of $\mathbb{R}\mathbb{P}$. The upcoming Lemma 4.4 describes some non-zero classes in the $H^*(\mathbb{R}\mathbb{P})$. Since the lemma refers to classes in $H^*(\mathcal{C}_{2n})$, a digression comes first to describe the latter.

The computations in [8] (see Corollary 3.1) imply that $H^*(\mathcal{C}_{2n})$ can be depicted as follows: Let ω denote the obstruction class in $H^2(\mathcal{C}_{2n})$ from Sections 4.1 and 4.2. Then $H^*(\mathcal{C}_{2n})$ can be written as

(4.16)
$$H^*(\mathcal{C}_{2n}) \cong H^*(\mathcal{C}_{2n}) \oplus (\omega \wedge H^{*-2}(\mathbb{C}_{2n-1}))$$

where $\mathbb{C}_k$ denotes the space of k unordered points in $\mathbb{C}$. The identifications in (4.16) come from the splitting of the long exact sequence for the pair $(\mathcal{C}_{2n}, \mathcal{C}_{2n} - \mathcal{C}_{2n}^x)$ with x being any chosen point in S^2 (and recall that $\mathcal{C}_{2n}^x$ denotes the set of configurations in $\mathcal{C}_{2n}$ that do not contain x). In terms of homology, (4.16) asserts in effect that each class in $H_*(\mathcal{C}_{2n})$ can be written uniquely as a sum of a class from $\mathcal{C}_{2n} - \mathcal{C}_{2n}^x$ and a class with non-zero cap-product pairing with ω . With regards to $H^*(\mathbb{C}_k)$: These modules are computed by [2], they are zero above a degree that is on the order of n and, in general, non-zero in most degrees below the maximal one.

The cohomology of $\mathbb{R}\mathbb{P}$ can be computed using the fact that this space is a fiber bundle over $\mathcal{C}_{2n}$ with the infinite dimensional real projective

space as fiber. The next lemma summarizes. The lemma refers to the class ω that appears in (4.16). Keep in mind that $\omega \wedge \omega$ is zero in cohomology.

Lemma 4.4. *The cohomology of $\mathbb{RP}$ is (non-canonically) isomorphic to the vector space of polynomials in a degree 2-class τ with coefficient ring $\mathcal{K}$ as described in the subsequent bullets. This is to say that*

$$H^*(\mathbb{RP}) = \mathcal{K}^* \oplus (\mathcal{K}^{*-2} \wedge \tau) \oplus (\mathcal{K}^{*-4} \wedge \tau \wedge \tau) \oplus \cdots$$

- *The definition of τ : Let $\hat{\omega}$ denote a given representative cocycle for ω . Then $\pi^*\hat{\omega}$ can be written as $d\nu$ with ν being a 1-cochain whose restriction to each fiber algebraically generates the cohomology of the fiber. The cochain $\nu \wedge \nu$ is closed and non-zero in $H^*(\mathbb{RP})$. The class of $\nu \wedge \nu$ is τ .*
- *The definition of $\mathcal{K}$: This module can be written as*

$$\mathcal{K}^* = \pi^*H^*(\mathbb{C}_{2n}) \oplus (\Theta \wedge \pi^*H^{*-3}(\mathbb{C}_{2n}))$$

where π^* and Θ are as follows:

- a) *What is denoted by π^* signifies the pull-back by the projection map π to $\mathbb{C}_{2n}$. This pull-back homomorphism is injective on the $H^*(\mathbb{C}_{2n})$ summand in (4.16) and zero on the other summand.*
- b) *To define Θ , let $\hat{o}$ denote a 3-cocycle on $\mathbb{C}_{2n}$ obeying $d\hat{o} = \hat{\omega} \wedge \hat{\omega}$. The 3-cocycle $\nu \wedge \pi^*\hat{\omega} + \pi^*\hat{o}$ on $\mathbb{RP}$ is closed and non-zero in $H^*(\mathbb{RP})$. Its class is independent of the choice for the cocycle $\hat{o}$. This cohomology class is Θ .*

Some implications of Lemma 4.4 concerning the critical point set of $\mathcal{E}$ are derived in Section 5.2.

Proof of Lemma 4.4. This lemma is an instance of Proposition A.1 in the appendix. q.e.d.

5. The extension of $\mathbb{RP}$ and $\mathcal{E}$ to $\bar{\mathbb{C}}_{2n}$

This section describes the sense in which the space $\mathbb{RP}$ and its function $\mathcal{E}$ extend to the whole of $\mathbb{C}_{2n}$'s compactification $\bar{\mathbb{C}}_{2n}$ from (3.2). To say more about this extension, write the space $\bar{\mathbb{C}}_{2n}$ as done in (3.1). Then each $k \in \{0, \dots, n\}$ version of $\mathbb{C}_{2k}$ has its corresponding version of $\mathbb{RP}$ which will be denoted here as $\mathbb{RP}_{2k}$. (The $k = 0$ version, $\mathbb{RP}_0$, is the quotient by the multiplicative action of $\{\pm 1\}$ on the subspace of functions in the standard L_1^2 Sobolev space on S^2 that obey (4.4).) The extension of $\mathbb{RP}$ over $\bar{\mathbb{C}}_{2n}$ is denoted by $\overline{\mathbb{RP}}$ and it is (as a point set)

$$(5.1) \quad \overline{\mathbb{RP}} = \mathbb{RP}_{2n} \cup \mathbb{RP}_{2n-2} \cup \cdots \cup \mathbb{RP}_0.$$

This extension maps to $\bar{\mathcal{C}}$ in the obvious way (the map is denoted by π). The topology on $\overline{\mathbb{RP}}$ is defined from a basis of neighborhood open sets that are analogous to those used in Section 4.3 for the topology on

$\mathbb{RP}$. The details appear momentarily. But note in advance the space $\overline{\mathbb{RP}}$ does not meet the technical definition of a fiber bundle because there aren't local product neighborhoods for the points in $\overline{\mathcal{C}}_{2n} - \mathcal{C}_{2n}$. (Even so, the map π is continuous and surjective.) There is, however, a weak notion of a local product structure for $\overline{\mathcal{C}}_{2n}$ which is described in the upcoming Section 5.3.

5.1. A topology for $\overline{\mathbb{RP}}$. A convention concerning elements in $\overline{\mathcal{C}}_{2n}$ is used for the specification of topology. The statement of the convention refers to the elements in $\overline{\mathcal{C}}_{2n}$ as $\mathbb{Z}/2\mathbb{Z}$ divisors as done in Section 3.1: Supposing that $[q]$ is such a divisor, then its minimal representative is implicitly viewed as a configuration in the relevant $k \in \{0, \dots, n\}$ version of $\mathcal{C}_{2k}$. This number k is said to be the *degree* of the divisor.

To define the topology on $\overline{\mathbb{RP}}$, first recall from Part 2 of Section 3.1 how the topology on $\overline{\mathcal{C}}_{2n}$ was defined by a basis of open neighborhoods of any given $\mathbb{Z}/2\mathbb{Z}$ divisor: Let k denote an integer from $\{0, \dots, n-1\}$ and let $[q]$ denote a given degree k divisor from $\overline{\mathcal{C}}_{2n}$. Fix a positive number to be denoted by ϵ which is chosen to be very small (much less than $\frac{1}{1000}$ times the distance between the points in $[q]$'s minimal representative). The corresponding open neighborhood was denoted there by $\mathcal{U}(\epsilon, k, [q])$ and it consists of the configurations that obey the conditions in (3.3). (Remember that the minimal representative of a divisor as depicted in (3.1) is the configuration $\mathfrak{p}$ with all coefficient numbers $\{m_p\}_{p \in \mathfrak{p}}$ equal to 1.)

With the preceding in mind, let $\mathfrak{q} \in \mathcal{C}_{2k}$ denote the minimal representative of the divisor $[q]$, and let $[f]$ denote an equivalence class in $\mathbb{RP}_{2k} \mid_{\mathfrak{q}}$. A basis set for the open neighborhoods of the pair $([q], [f])$ in $\overline{\mathbb{RP}}$ are labeled by positive numbers ϵ, δ and $\mathcal{E}$ with ϵ less than $\frac{1}{1000n}$ and with E greater than the S^2 integral of $|df|^2$. This basis set consists of the pairs $([c], [h]) \in \overline{\mathbb{RP}}$ that obey the following constraints:

(5.2)

- $[c] \in \mathcal{U}(\epsilon, k, [q])$.
- $\int_{S^2 - c} |dh|^2 < E$.
- Let $\mathfrak{c}$ denote the minimal divisor of $[c]$ and let Ω denote the $\text{dist}(\cdot, \mathfrak{c}) > 200\epsilon$ part of S^2 . There is an isometry between $\mathcal{I}_{\mathfrak{c}}$ and $\mathcal{I}_{\mathfrak{q}}$ on Ω that identifies h with a section of $\mathcal{I}_{\mathfrak{q}}$ obeying $\int_{\Omega} (|d(f-h)|^2 + |f-h|^2) < \delta$.

It is left to the reader to verify that this topology restricts to each stratum in (5.1) to give the topology that is described in Section 4.3 for the relevant $k \in \{0, \dots, 2n\}$ version of $\mathbb{RP}_{2k}$. It is also left to the

reader to verify that the map π from $\overline{\mathbb{RP}}$ to $\overline{\mathcal{C}}_{2n}$ is continuous with this topology on $\overline{\mathbb{RP}}$.

5.2. The extension of $\mathcal{E}$. The lemma that follows makes a formal assertion to the effect that the function $\mathcal{E}$ extends as a continuous function to the whole of $\overline{\mathbb{RP}}$.

Lemma 5.1. *The function on $\overline{\mathbb{RP}}$ that is defined by the rule*

$$([q], [f]) \rightarrow \mathcal{E}([q], [f]) \equiv \int_{S^2-q} |df|^2$$

is continuous in general and smooth along any stratum in (5.1).

Proof of Lemma 5.1. The assertion that $\mathcal{E}$ is smooth along the strata of $\overline{\mathbb{RP}}$ are instances of Lemma 4.3. The proof that $\mathcal{E}$ is continuous at a given pair $([q], [f])$ invokes the following observation:

(5.3)

Given a positive number μ at most $\frac{1}{1000n}$, there exists $\epsilon > 0$ such that the integral of $|df|^2$ on the union of any n disks of radius ϵ is at most μ .

By way of a proof: There is a smooth section of $\mathcal{I}_q$ with compact support in S^2-q to be denoted by f_μ with the S^2 integral of $|d(f-f_\mu)|^2 + |f-f_\mu|^2$ being less than $\frac{1}{100}\mu$. (This follows from the definition of $\mathbb{H}_q$.) The assertion in (5.3) holds for f_μ because f_μ is smooth and supported away from q . Thus, it holds also for f .

With (5.3) in hand, then the proof of Lemma 5.1 can be completed by copying with only notational changes what is said by (4.8)–(4.10) and what is said by the surrounding discussion in the last paragraph of Part 2 of Lemma 4.2’s proof. q.e.d.

To motivate the subsequent constructions and by way of an application of Lemmas 4.4 and 5.1, return to the context of the min-max values for $\mathcal{E}$ on classes in $H_*(\mathbb{RP})$ as defined in the last paragraph in Section 4.5. It follows from Lemma 5.1 that these min-max value for $\mathcal{E}$ are zero on the π -pull back of classes from the $H^*(\mathbb{C}_{2n})$ summand in (4.16). This is because there are representative cocycles for the duals of these classes in $H_*(\mathbb{C}_{2n})$ that are entirely in $\mathbb{C}_{2n}$ ’s intersection with any given open neighborhood of the $\mathcal{C}_0$ stratum in $\overline{\mathcal{C}}_{2n}$. To explain: These are classes from the base space $\mathbb{C}_{2n}$ that can be represented by continuously parametrized families of configurations that all lie in any given disk in the sphere. In particular, the disk can have arbitrarily small radius. If the radius is small enough, then the whole family can be lifted to $\mathbb{RP}$ via a section that assigns to each configuration in the family the (unique) eigenfunction with the smallest eigenvalue (which is nearly zero). The eigenvalue is nearly zero and the eigenfunction is unique because all of the points in all of the configurations that comprise the family are very

close together (near the $\mathcal{C}_0$ stratum in $\overline{\mathcal{C}}_{2n}$). Said differently, the duals of these classes are sent to zero by the inclusion induced homomorphism from $H_*(\mathcal{C}_{2n})$ into $H_*(\overline{\mathcal{C}}_{2n})$. By the same token, with τ and Θ as defined in Lemma 4.4, and supposing that m is a non-negative integer, then the min-max values for $\mathcal{E}$ on the cup product of either $\wedge^m \tau$ or $\wedge^m \tau \wedge \Theta$ with the π -pull back of any class in the $H^*(\mathcal{C}_{2n})$ summand in (4.16) is the same as the min-max value of $\mathcal{E}$ on just the class $\wedge^m \tau$ or $\wedge^m \tau \wedge \Theta$ the case may be. These min-max values aren't known, but they are, in any event, no greater than their min-max values on $\overline{\mathbb{R}\mathbb{P}}^2$. In any event, a corollary of Proposition 3.3 is this: Having fixed an integer m , there exists an integer n_m with the following significance:

(5.4) *If $n \geq n_m$, then the min-max value of $\mathcal{E}$ for either $\wedge^m \tau$ or $\wedge^m \tau \wedge \Theta$ on the $\mathcal{C}_{2n}$ version of $\mathbb{R}\mathbb{P}$ is the min-max value of $\mathcal{E}$ for these classes on some stratum of $\overline{\mathbb{R}\mathbb{P}} - \mathbb{R}\mathbb{P}$.*

The point of the preceding is that $\overline{\mathbb{R}\mathbb{P}}$'s cohomology is the relevant cohomology with regards to the relation between $\mathcal{E}$'s min-max values on $\mathbb{R}\mathbb{P}$ and $\mathcal{E}$'s critical points on $\mathbb{R}\mathbb{P}$.

5.3. $\overline{\mathbb{R}\mathbb{P}}$ as a weak projective space bundle over $\overline{\mathcal{C}}_{2n}$. The upcoming Lemma 5.3 describes a weak $\mathbb{R}\mathbb{P}^\infty$ bundle structure for the space $\overline{\mathbb{R}\mathbb{P}}$. This notion of a weak $\mathbb{R}\mathbb{P}^\infty$ bundle is defined in the context where there are topological spaces $\mathcal{R}$ and X with a surjective map $\pi : \mathcal{R} \rightarrow X$ whose fibers are homeomorphic to $\mathbb{R}\mathbb{P}^\infty$. A weak $\mathbb{R}\mathbb{P}^\infty$ bundle structure for $(\mathcal{R}, X, \pi)$ consists of the following data.

- (5.5)
- For every non-negative integer k : A locally finite open cover of X (to be denoted by $\mathfrak{U}_k$) such that each $U \in \mathfrak{U}_k$ has a vector bundle $V_U \rightarrow U$ of rank greater than k with a fiberwise embedding Ψ_U from $(V_U - 0)/\mathbb{R}^*$ to $\mathcal{R}|_U$ that covers the projections to U

$$\begin{array}{ccc} (V_U - 0)/\mathbb{R}^* & \xrightarrow{\Psi_U} & \mathcal{R}|_U \\ & \searrow \pi & \swarrow \pi \\ & U & \end{array}$$

- and induces an isomorphism on the first homology of each fiber.
- For each pair of positive integers k and k' : If $U \in \mathfrak{U}_k$ and $U' \in \mathfrak{U}_{k'}$ intersect, and supposing that $\dim(V_{U'}) \geq \dim(V_U)$, then Ψ_U factors through $\Psi_{U'}$ over $U \cap U'$ via an injective bundle map from $V_U|_{U \cap U'}$ to $V_{U'}|_{U \cap U'}$.
 - Each $\mathbb{Z}/2\mathbb{Z}$ cohomology class of $\mathcal{R}$ is non-zero when restricted to some sufficiently large k version of $\mathcal{R}$'s subspace $\bigcup_{U \in \mathfrak{U}_k} \Psi_U((V_U - 0)/\mathbb{R}^*)$.

A preliminary lemma is needed in order to define the weak $\mathbb{RP}^\infty$ structure for $\overline{\mathbb{RP}}$. To set the stage: The preliminary lemma (Lemma 5.2) refers to a $\mathbb{Z}/2\mathbb{Z}$ divisor from a stratum of $\overline{\mathcal{C}}_{2n}$ (it denotes the latter by $[q]$) and its minimal representative (denoted by $\mathfrak{q}$) which is a configuration from some version of $\mathcal{C}_{2k}$ (for $k \in \{0, \dots, n\}$). With $\mathfrak{q}$ in hand, the lemma refers to a positive number that is not an $\mathcal{I}_{\mathfrak{q}}$ -eigenvalue. Letting E denote that number, Lemma 5.2 uses $\mathbb{H}_{\mathfrak{q},E}$ to denote the span in the Hilbert space $\mathbb{H}_{\mathfrak{q}}$ of the $\mathcal{I}_{\mathfrak{q}}$ -eigensections with eigenvalue less than E . The lemma also refers to L^2 -orthogonal projection on $\mathbb{H}_{\mathfrak{q}}$. This is the orthogonal projection defined by the L^2 inner product:

$$(5.6) \quad (f, f') \rightarrow \int_{S^2} f f'.$$

The instances of L^2 -orthogonal projection in the lemma are to the subspaces $\mathbb{H}_{(\cdot),E}$.

Lemma 5.2. *Fix a degree $2k$ divisor in $\overline{\mathcal{C}}_{2n} - \mathcal{C}_{2n}$ to be denoted by $[q]$ and let $\mathfrak{q}$ denote the minimal representative. Then fix a positive number to be denoted by E which is not an $\mathcal{I}_{\mathfrak{q}}$ -eigenvalue, and fix a positive number to be denoted by δ . Granted this data, there exists an open neighborhood of $[q]$ in $\overline{\mathcal{C}}_{2n}$ with the following properties: Denote the neighborhood by $\overline{U}$. Supposing that $\mathfrak{p}$ is the minimal representative of a $\mathbb{Z}/2\mathbb{Z}$ divisor in $\overline{U}$, then E is not an $\mathcal{I}_{\mathfrak{p}}$ -eigenvalue; and there exists a bounded linear map from $\mathbb{H}_{\mathfrak{q}}$ to $\mathbb{H}_{\mathfrak{p}}$ and a bounded linear map from $\mathbb{H}_{\mathfrak{p}}$ to $\mathbb{H}_{\mathfrak{q}}$ with the properties listed next (both linear maps are denoted by L).*

- Both versions of L are injective on the domain $\mathbb{H}_{(\cdot),E}$ subspace, and their composition with the L^2 -orthogonal projection to the range $\mathbb{H}_{(\cdot),E}$ subspace is an isomorphism.
- If f is from the domain $\mathbb{H}_{(\cdot),E}$, then

$$\begin{aligned} \text{a) } & (1 - \delta) \int_{S^2} |f|^2 \leq \int_{S^2} |Lf|^2 < (1 + \delta) \int_{S^2} |f|^2, \\ \text{b) } & \int_{S^2} |df|^2 - \delta \int_{S^2} |f|^2 \leq \int_{S^2} |d(Lf)|^2 < \int_{S^2} |df|^2 + \delta \int_{S^2} |f|^2. \end{aligned}$$

A note about $\overline{U}$: The subspace $\overline{U} \cap \mathcal{C}_{2m}$ is empty when m is less than k (which is half of the number of points that comprise $\mathfrak{q}$) and it is non-empty otherwise.

This lemma is proved momentarily. A weak $\mathbb{RP}^\infty$ structure for $\overline{\mathbb{RP}}$ can be constructed using Lemma 5.2 according to the rules laid out below in Lemma 5.3.

Lemma 5.3. *What follows describes a weak $\mathbb{RP}^\infty$ structure for $\overline{\mathbb{RP}}$. Fix a positive integer k . For any given $\mathbb{Z}/2\mathbb{Z}$ divisor $[q] \in \overline{\mathcal{C}}_{2n}$ first fix a positive number E which is not an $\mathcal{I}_{\mathfrak{q}}$ -eigenvalue and with $\dim(\mathbb{H}_{\mathfrak{q},E}) > k$. Then fix some very small, positive number δ which is much less than*

the distance from E to the nearest $\mathcal{I}_q$ -eigenvalue. With the preceding data as input, use Lemma 5.2 to define a version of the open set $\overline{U}$ (which is denoted here by $\overline{U}_q$). Having done this for each element in $\overline{\mathcal{C}}_{2n}$, take a locally finite cover of $\overline{\mathcal{C}}_{2n}$ from the collection $\{\overline{U}_q : q \in \overline{\mathcal{C}}_{2n}\}$ and denote this cover by $\mathfrak{U}_k$. The required vector bundle $V_{\overline{U}_q}$ over a set $\overline{U}_q$ from the cover $\mathfrak{U}_k$ is the bundle whose fiber over any $[p] \in \overline{U}_q$ is $\mathbb{H}_{p,E}$. The embedding $\Psi_{\overline{U}_q}$ for the first bullet of (5.5) is the tautological inclusion map.

This lemma is proved using the next lemma.

By way of terminology for the upcoming lemma: A parametrized family of maps from one subset of $\overline{\mathbb{R}P}$ to another is said to be *fiber preserving* when the projection via π to $\overline{\mathcal{C}}_{2n}$ is not changed by any map from the family.

Lemma 5.4. *Fix a positive number to be denoted by $\mathfrak{C}$ and let $\overline{\mathbb{R}P}^{\mathfrak{C}}$ denote the $\mathcal{E} < \mathfrak{C}$ part of $\overline{\mathbb{R}P}$. If $k > 2\mathfrak{C}$, then there exists a fiber preserving deformation retract of $\overline{\mathbb{R}P}^{\mathfrak{C}}$ onto an $\overline{\mathbb{R}P}$ version of the subspace $\bigcup_{U \in \mathfrak{U}_k} \Psi_U((V_U - 0)/\mathbb{R}^*)$.*

Lemma 5.4 is proved after Lemma 5.2.

Proof of Lemma 5.3. The conditions in the first two bullets of (5.5) follow directly from the construction. The third bullet of (5.5) holds for $\overline{\mathbb{R}P}$ if any given (singular) homology class is represented by a closed chain whose image lies entirely in some $\overline{\mathbb{R}P}$ version of $\bigcup_{U \in \mathfrak{U}_k} \Psi_U((V_U - 0)/\mathbb{R}^*)$ for k sufficiently large. To establish this, note first that any given closed singular chain has compact support in $\overline{\mathbb{R}P}$ which implies that $\mathcal{E}$ has an upper bound on that chain. Let $\mathfrak{C}$ denote this upper bound. Granted this upper bound, use the deformation retract from Lemma 5.4 to obtain a homologous chain that lies in a $\bigcup_{U \in \mathfrak{U}_k} \Psi_U((V_U - 0)/\mathbb{R}^*)$ subspace of $\overline{\mathbb{R}P}$.
q.e.d.

Proof of Lemma 5.2. The construction of L borrows much by way of strategy and notation from Parts 1–4 of the proof of Proposition 3.2. There are three parts to the construction of L and the verification of its properties.

Part 1. Fix a positive number to be denoted by ϵ which is chosen to be very small (much less than $\frac{1}{1000n}$), and then reintroduce from Part 2 of Section 3.1 the set $\mathcal{U}(\epsilon, k, [q])$. Let p denote a configuration of $2j$ points in this set for some $j \in \{k, \dots, n\}$. Define the function χ_ϵ from (2.2) using p , but denote it here by $\chi_{p,\epsilon}$. Note in particular that the function is zero at all points from q and that the line bundles $\mathcal{I}_p$ and $\mathcal{I}_q$ are isomorphic on the support of $\chi_{p,\epsilon}$.

By virtue of ϵ being small, there is a point to be denoted by x where $\chi_{p,\epsilon}$ is equal to 1. Use this point to define an isomorphism between the

line bundles $\mathcal{I}_{\mathbf{p}}$ and $\mathcal{I}_{\mathbf{q}}$ on the support of $\chi_{\mathbf{p},\epsilon}$. With this isomorphism in hand, the map $L : \mathbb{H}_{\mathbf{q}} \rightarrow \mathbb{H}_{\mathbf{p}}$ is given by the rule

$$(5.7) \quad f \longrightarrow \chi_{\mathbf{p},\epsilon} f.$$

This same formula likewise defines the corresponding version of L sending $\mathbb{H}_{\mathbf{p}}$ to $\mathbb{H}_{\mathbf{q}}$. In either direction, the arguments for (3.5) and (3.6) can be repeated with only notational changes to see that

$$(5.8)$$

$$\begin{aligned} \bullet \int_{S^2} |d(\chi_{\mathbf{p},\epsilon} f)|^2 &= \int_{S^2} |df|^2 + \mathfrak{e}_1 \int_{S^2} |f|^2, \\ \bullet \int_{S^2} |\chi_{\mathbf{p},\epsilon} f|^2 &= (1 + \mathfrak{e}_2) \int_{S^2} |f|^2, \end{aligned}$$

where $|\mathfrak{e}_1| \leq c_0 \frac{1}{|\ln \epsilon|} n(1 + \lambda)$ and $|\mathfrak{e}_2| \leq c_0 \epsilon n(1 + \lambda)$ when f is in the span of the respective $\mathcal{I}_{\mathbf{q}}$ -eigensections or $\mathcal{I}_{\mathbf{p}}$ -eigensections with eigenvalue less than λ .

Part 2. Choose a positive number to be denoted by μ so that there are no $\mathcal{I}_{\mathbf{q}}$ -eigenvalues in the interval between $E - \mu$ and $E + \mu$. Supposing that ϵ is small, assume that there exists a configuration $\mathbf{p}$ from $\mathcal{U}(\epsilon, k, [\mathbf{q}])$ with an $\mathcal{I}_{\mathbf{p}}$ -eigenvalue in the interval between $E - \frac{1}{2}\mu$ and $E + \frac{1}{2}\mu$. The plan is to derive nonsense from this assumption. To this end, let f denote such an eigensection with the S^2 integral of f^2 equal to 1. Write $\chi_{\mathbf{p},\epsilon} f$ as a sum of $\mathcal{I}_{\mathbf{q}}$ -eigensections. This depiction has the form $f_{<} + f_{>}$ where $f_{<}$ denotes the contribution to this sum from the eigensections with eigenvalue less than $E - \mu$ (and thus $f_{>}$ contains the contribution from the eigensections with eigenvalue greater than $E + \mu$). Apply (5.8) to $\chi_{\mathbf{p},\epsilon} f_{<}$ (viewed as a section of $\mathcal{I}_{\mathbf{p}}$) to see that it is, for the most part, a sum of $\mathcal{I}_{\mathbf{p}}$ -eigensections with eigenvalue less than $E - \frac{1}{2}\mu$ when ϵ is very small. Therefore, the L^2 -inner product of $\chi_{\mathbf{p},\epsilon} f_{<}$ with f is very small when ϵ is very small. Thus, $f_{<}$ itself must be very small when ϵ is very small because the L^2 inner product between $\chi_{\mathbf{p},\epsilon} f_{<}$ and f is the same as that between $f_{<}$ and $\chi_{\mathbf{p},\epsilon} f$ and the latter inner product is the S^2 integral of $f_{<}^2$ (by definition). As explained directly, the detailed analysis leads to this:

$$(5.9) \quad \int_{S^2} |f_{<}|^2 \leq c_0 \frac{1}{\mu |\ln \epsilon|} n(1 + E)$$

To derive the preceding inequality, start with the observation that

$$(5.10) \quad \int_{S^2} |df_{<}|^2 \leq (E - \mu) \int_{S^2} |f_{<}|^2.$$

With this in hand, write its left hand side as the S^2 integral of $\langle df_{<}, d(\chi_{\mathbf{p},\epsilon} f) \rangle$ (the integral of the inner product between $df_{<}$ and $d(\chi_{\mathbf{p},\epsilon} f)$).

Then move the $\chi_{\mathfrak{p},\epsilon}$ from multiplying f to multiplying $f_{<}$ so as to write that integral as

$$(5.11) \qquad \int_{S^2} \langle d(\chi_{\mathfrak{p},\epsilon} f_{<}), df \rangle + \tau_1.$$

where τ_1 is a term whose norm is bounded by c_0 times the S^2 -integral of the function $|d\chi_{\mathfrak{p},\epsilon}|(|f_{<}||df| + |f||df_{<}|)$; and thus by $c_0 \frac{1}{|\ln \epsilon|} n$ times the S^2 integral of $|df|^2$. (Use Lemma 2.1 and the fact that $|d\chi_{\mathfrak{p},\epsilon}|$ is bounded by $c_0 \frac{1}{|\ln \epsilon|} \frac{1}{\text{dist}(\mathfrak{p}, \cdot)}$. The factor of n appears because $\mathfrak{p}$ has $2n$ points. Let λ denote f 's eigenvalue and write (5.11) as

$$(5.12) \qquad \lambda \int_{S^2} |f_{<}|^2 + \tau_1.$$

Because $\lambda > E - \frac{\mu}{2}$, the inequality in (5.10) with (5.12) imply that

$$(5.13) \qquad \frac{\mu}{2} \int_{S^2} |f_{<}|^2 \leq c_0 \frac{1}{|\ln \epsilon|} n(1 + E)$$

which is what is claimed by (5.9).

The inequality in (5.9) implies (because $\chi_{\mathfrak{p},\epsilon} f$ is nearly $f_{>}$) that

$$(5.14) \qquad \int_{S^2} |d(\chi_{\mathfrak{p},\epsilon} f)|^2 \geq (E + \mu) \left(1 - c_0 \mu^{-1} \frac{1}{|\ln \epsilon|} n(1 + E) \right)$$

The latter inequality implies in turn (via 5.8) that

$$(5.15) \qquad \int_{S^2} |df|^2 \geq (E + \mu) \left(1 - c_0 \mu^{-1} \frac{1}{|\ln \epsilon|} n(1 + E) \right)$$

with the version of c_0 in this equation being larger than the version of c_0 in (5.14) because of the appeal to (5.8). This last inequality exhibits the desired nonsense if ϵ is sufficiently small (given μ and E) because the left hand side of this inequality is, by assumption, no greater than $E + \frac{1}{2\mu}$.

Part 3. Let $\Pi_{\mathfrak{p},E}$ denote the projection operator on $\mathbb{H}_{\mathfrak{p}}$ onto $\mathbb{H}_{\mathfrak{p},E}$ for the L^2 -inner product on $\mathbb{H}_{\mathfrak{p}}$. It follows from the $\mathbb{H}_{\mathfrak{q}}$ to $\mathbb{H}_{\mathfrak{p}}$ version of (5.8) that the composition of first L and then $\Pi_{\mathfrak{p},E}$ is injective on $\mathbb{H}_{\mathfrak{q},E}$ when ϵ is small (given E and μ); and nearly isometric with respect to the $\mathbb{H}_{(\cdot)}$ and the L^2 inner products as indicated in the statement of the proposition.

Much the same arguments can be used with the $\mathbb{H}_{\mathfrak{p}}$ to $\mathbb{H}_{\mathfrak{q}}$ version of L to see that the composition of first this version of L and then $\Pi_{\mathfrak{q},E}$ maps $\mathbb{H}_{\mathfrak{p},E}$ injectively into $\mathbb{H}_{\mathfrak{q},E}$ when ϵ is small (given E and μ); and that this composition is nearly an isometry when ϵ is small (given E and μ).
q.e.d.

Proof of Lemma 5.4. The three parts of the proof that follow explain how to construct the desired deformation retract.

Part 1. Here is the key observation: Suppose that $[q]$ is a $\mathbb{Z}/2\mathbb{Z}$ -divisor and let q denote its minimal representative. Suppose also that $E_q > 2\mathfrak{C}$ and that E_q is not an $\mathcal{I}_q$ -eigenvalue. Let f denote an element in $\mathbb{H}_q$ with $\mathfrak{C}$ bounding the S^2 -integral of $|df|^2$. Then the S^2 -integral of $|\Pi_{q,E_q} f|^2$ is greater than $\frac{1}{2}$ if the S^2 -integral of $|f|^2$ is equal to 1.

Given a positive integer k , then E_q can be chosen larger if necessary so that the dimension of $\mathbb{H}_{q,E_q}$ is greater than k . This lower bound is assumed in what follows.

Part 2. Keeping Part 1 in mind, there is an open neighborhood of $[q]$ in $\overline{\mathcal{C}}_{2n}$ to be denoted by $\overline{U}_q$ such that if $[q']$ is from this neighborhood, then E_q is not an $\mathcal{I}_{q'}$ -eigenvalue. Now each $[q] \in \overline{\mathcal{C}}_{2n}$ has its version of $\overline{U}_q$ and since the union of all of these sets covers $\overline{\mathcal{C}}_{2n}$, there is a finite subcover. Choose such a finite cover and denote this finite collection of sets by $\mathfrak{U}_k$. Let $\{\chi_q : \overline{U}_q \in \mathfrak{U}_k\}$ denote a chosen partition of unity subordinate to the cover.

Part 3. Supposing that $([p], [f])$ is in $\overline{\mathbb{RP}}^{\mathfrak{C}}$, define the new $\mathcal{I}_p$ section f_1 by the rule

$$(5.16) \quad f_1 \equiv \sum_{U_q \in \mathfrak{U}_k} \chi_q \Pi_{p,E_q} f.$$

What follows are two important facts about f_1 . To state the first, let E denote the minimum of those E_q that appear with non-zero χ_q in (5.16). Then the S^2 -integral of $|\Pi_{p,E} f_1|^2$ is greater than $\frac{1}{2}$ because $\Pi_{p,E} f_1$ is the same as $\Pi_{p,E} f$. This implies in particular that f_1 is not identically zero. With this last point understood, introduce Z_1 to denote the square root of the S^2 -integral of $|f_1|^2$.

To state the second fact, let E' denote the maximum of those E_q that appear with non-zero χ_q in (5.16). Then $\Pi_{p,E'} f_1 = f_1$ which is to say that f_1 is in $\mathbb{H}_{p,E'}$. This second fact has implications with regards to Lemma 5.3's weak $\mathbb{RP}^\infty$ structure for $\overline{\mathbb{RP}}$: The $\overline{\mathbb{RP}}$ element $([p], \frac{1}{Z_1}[f_1])$ is in the image of at least one $\overline{U}_q \in \mathfrak{U}_k$ version of $\Psi_{\overline{U}_q}$. (Remember that $\Psi_{\overline{U}_q}$ is a fiber preserving embedding of $(V_{\overline{U}_q} - 0)/\mathbb{R}^*$ to $\overline{\mathbb{RP}}$), and remember that the fiber of $V_{\overline{U}_q}$ over p is Π_{p,E_q} .

It follows from the preceding fact that the map from $\overline{\mathbb{RP}}^{\mathfrak{C}}$ to $\overline{\mathbb{RP}}$ that sends any given element $([p], f)$ to $([p], f_1)$ maps $\overline{\mathbb{RP}}^{\mathfrak{C}}$ into the subspace $\bigcup_{\overline{U}_q \in \mathfrak{U}} \Psi_{\overline{U}_q}(V_{\overline{U}_q} - 0)/\mathbb{R}$.

Part 4. The desired deformation retract is a map from $[0, 1] \times \overline{\mathbb{RP}}^{\mathfrak{C}}$ to $\overline{\mathbb{RP}}^{\mathfrak{C}}$ whose $t = 1$ end member is the map from Part 3 that send $([p], f)$ to $([p], [f_1])$. The time $t \in [0, 1]$ member of the family sends $([p], f)$ to $([p], [f_1 + (1 - t)(f - f_1)])$. This is a well defined map from $\overline{\mathbb{RP}}^{\mathfrak{C}}$ because f_1 is non-zero and it is fiber preserving because the $[p]$ component of an

given element doesn't change. Note also that any $t \in [0, 1]$ version maps $\overline{\mathbb{R}\mathbb{P}}^{\mathfrak{C}}$ to itself. Indeed, supposing that $([p], [f])$ is from $\overline{\mathbb{R}\mathbb{P}}^{\mathfrak{C}}$, then the projection of f onto the span of the $\mathcal{I}_p$ -eigensections with eigenvalue less than or equal to $\mathfrak{C}$ does not change as t increases whereas the L^2 norm of the projection of f to the span of the $\mathcal{I}_p$ -eigensections with eigenvalues greater than $\mathfrak{C}$ is a non-increasing function of t . And, by construction, the $t = 1$ end member of this family maps into $\bigcup_{\overline{U}_q \in \mathfrak{U}} \Psi_{\overline{U}_q}(V_{\overline{U}_q} - 0)/\mathbb{R}$.
q.e.d.

5.4. The obstruction class. Returning to the context of (5.5), suppose that $(\mathcal{R}, X, \pi)$ has a weak $\mathbb{R}\mathbb{P}^\infty$ bundle structure. This weak $\mathbb{R}\mathbb{P}^\infty$ bundle structure comes from a *weak vector bundle* structure when the collection $\{V_U : U \in \mathfrak{U}_k\}_{k \in \mathbb{N}}$ from (5.5) has the additional property described below in (5.17). To set notation, suppose that k, k' are non-negative integers, and then suppose that U is from $\mathfrak{U}_k$ and U' is from $\mathfrak{U}_{k'}$, that these sets with non-empty intersection, and that their corresponding vector bundles obey $\dim(V_{U'}) \geq \dim(V_U)$. Let $T_{U',U}$ denote the injective bundle map from $V_U|_{U \cap U'}$ to $V_{U'}|_{U \cap U'}$ that is required by the second bullet of (5.5).

(5.17)

For any triple $k, k', k'' \in \mathbb{N}$ and sets U, U' and U'' from $\mathfrak{U}_k, \mathfrak{U}_{k'}$ and $\mathfrak{U}_{k''}$ that share points and are such that $\dim(V_{U''}) \geq \dim(V_{U'}) \geq \dim(V_U)$, the corresponding vector bundle maps obey the identity $T_{U''U} = T_{U''U'}T_{U'U}$.

Now a weak $\mathbb{R}\mathbb{P}^\infty$ bundle structure need not come from weak vector bundle structure (which turns out to be our situation with our bundle over $\overline{\mathcal{C}}_{2n}$). To see the obstruction, note first that whether or not the weak $\mathbb{R}\mathbb{P}^\infty$ bundle structure comes from a weak vector bundle structure, the bundle maps for the weak $\mathbb{R}\mathbb{P}^\infty$ structure do obey

(5.18) $T_{U''U} = x_{U''U'U} T_{U''U'} T_{U'U}$

with $x_{U''U'U}$ being a nowhere zero, continuous map from $U'' \cap U' \cap U$ to $\mathbb{R}$. Let $z_{U''U'U}$ denote the sign (either 1 or -1) of $x_{U''U'U}$. This is a locally constant map from $U'' \cap U' \cap U$ to $\mathbb{Z}/2\mathbb{Z}$. These locally constant maps are such that the collection

(5.19) $\left\{ z_{U''U'U} : U'', U', U \in \bigcup_{k \in \mathbb{N}} \mathfrak{U}_k \right\}$

when restricted to any locally finite subcover defines a (multiplicative) degree two Čech cocycle for the subcover whose image in $H^2(X; \mathbb{Z}/2\mathbb{Z})$ is independent of the chosen subcover. This cohomology class is said to be the *obstruction class*. As explained in the next paragraphs, if the first $\mathbb{Z}/2\mathbb{Z}$ cohomology of all sets from $\bigcup_{k \in \mathbb{N}} \mathfrak{U}_k$ is zero, then this weak $\mathbb{R}\mathbb{P}^\infty$ structure comes from a weak vector bundle structure if and only

if this class is zero. (In general, if the class is zero, then the collection of $T_{U',U}$ maps can be modified by multiplying each by a suitable nowhere zero function so that they obey the identity in (5.17).)

To explain why the data in (5.19) always defines a cochain (a closed cocycle): Consider four sets U_1, U_2, U_3 and U_4 that share points and with corresponding vector bundles whose dimensions don't decrease with increasing label. Write T_{ji} for the $U = U_i$ and $U' = U_j$ version of $T_{U',U}$ and consider $T_{43}T_{32}T_{21}$. This can be written on the one hand as $x_{321}x_{431}T_{41}$ and on the other as $x_{432}x_{421}T_{41}$. Thus, $z_{321}z_{431}z_{432}z_{421} = 1$ which is the statement that the collection in (5.19) defines a closed cocycle.

With regards to changing (5.19) by a (multiplicative) coboundary: Changing (5.19) by the differential of a 1-cocycle $\{z_{U',U} : U, U' \in \bigcup_{k \in \mathbb{N}} \mathcal{U}_k\}$, accounts for the change of each $T_{U',U}$ to $z_{U',U}T_{U',U}$. If this change can be made so as to make all $x_{U''U',U}$ positive, then a partition of unity can be employed to modify each $T_{U',U}$ so that the condition in (5.17) holds. (To do this, write $x_{U''U',U}$ as the exponential of a real valued function.)

If the first cohomology of the sets from the collection $\bigcup_{k \in \mathbb{N}} \mathcal{U}_k$ are not all trivial, then there is an additional equivalence that can be used to obtain the given weak $\mathbb{RP}^\infty$ structure from a weak vector bundle structure: Each $U \in \bigcup_{k \in \mathbb{N}} \mathcal{U}_k$ version of V_U can be modified by tensoring it with a real line bundle over U with the proviso that any two respective U and U' line bundles are isomorphic on $U \cap U'$. Assuming this, let $\tilde{V}_U$ denote the modified V_U . The modified version of $T_{U',U}$ is obtained from the original by tensoring with a chosen isomorphism from the respective U line bundle to the respective U' line bundle. Let $\eta_{U',U}$ denote a chosen isomorphism between these line bundles. When U, U' and U'' are three sets that share points, then $\eta_{U'',U}$ can be written as $\hat{x}_{U''U',U}\eta_{U''U'}\eta_{U',U}$ with $\hat{x}_{U''U',U}$ mapping to $\mathbb{R}^*$. Let $\hat{z}_{U''U',U}$ denote the sign of $\hat{x}_{U''U',U}$. The collection of $\hat{z}$'s also defines a multiplicative, degree 2 Čech cocycle. Modifying each $U \in \bigcup_{k \in \mathbb{N}} \mathcal{U}_k$ versions of V_U in this way changes each $T_{U',U}$ to $\eta_{U',U}T_{U',U}$ and this corresponds to changing each $z_{U''U',U}$ to the product $\hat{z}_{U''U',U}z_{U''U',U}$.

The relevant example is $\overline{\mathbb{RP}}$ where the obstruction class in $H^2(\overline{\mathcal{C}}_{2n})$ is non-zero (its restriction to $\mathbb{RP}$ is already non-zero.) As was the case with $\mathbb{RP}$, the $\overline{\mathbb{RP}}$ obstruction class is denoted by ω .

The next lemma says more about this class (it is needed for the next subsection).

Lemma 5.5. *The self-cup product ω^{2n} is non-zero in $H^{4n}(\overline{\mathcal{C}}_{2n})$.*

Proof of Lemma 5.5. A preliminary step is a proof that $H^{4n}(\overline{\mathcal{C}}_{2n})$ is non-zero. To prove that this is so, consider that the exact sequence for

the pair $(\overline{\mathcal{C}}_{2n}, \overline{\mathcal{C}}_{2n-2})$ appears in part as

$$(5.20) \quad H^{4n-1}(\overline{\mathcal{C}}_{2n-2}) \rightarrow H^{4n}(\overline{\mathcal{C}}_{2n}, \overline{\mathcal{C}}_{2n-2}) \rightarrow H^{4n}(\overline{\mathcal{C}}_{2n}) \rightarrow H^{4n}(\overline{\mathcal{C}}_{2n-2})$$

whose left-most and right-most terms are zero. Thus, $H^{4n}(\overline{\mathcal{C}}_{2n})$ is isomorphic to $H^{4n}(\overline{\mathcal{C}}_{2n}, \overline{\mathcal{C}}_{2n-2})$. Meanwhile, $H^{4n}(\overline{\mathcal{C}}_{2n}, \overline{\mathcal{C}}_{2n-2})$ is isomorphic to the compactly supported cohomology of $\mathcal{C}_{2n}$ which is isomorphic to $\mathbb{Z}/2\mathbb{Z}$ since $\mathcal{C}_{2n}$ is a manifold.

The calculation for the $2n$ -fold self-cup product of ω has three parts.

Part 1. The standard compactification of $\mathcal{C}_{2n}$ is to view it as a subset of $\mathbb{CP}^{2n}$ as follows: First view S^2 as $\mathbb{CP}^1$ which is to say $(\mathbb{C}^2 - 0)/\mathbb{C}^*$. Letting (z, w) denote the coordinates for $\mathbb{C}^2$ and thus homogeneous coordinates for $\mathbb{CP}^1$, a point in $\mathbb{CP}^1$ is determined by homogeneous coordinates $[\alpha, \beta]$ as the zeros of the linear function $\alpha w - \beta z$. Then, $2n$ points in $\mathbb{CP}^1$ (counting multiplicities) are the zeros of the $2n$ 'th order homogeneous polynomial

$$(5.21) \quad (\alpha_1 w - \beta_1 z)(\alpha_2 w - \beta_2 z) \cdots (\alpha_{2n} w - \beta_{2n} z)$$

Multiplying this out, then the coefficients of the various powers of w are the homogeneous coordinates for $\mathbb{CP}^{2n}$:

$$(5.22) \quad a_{2n} w^{2n} + a_{2n-1} w^{2n-1} z + \cdots + a_0 z^{2n}.$$

Doing this identifies $\mathcal{C}_{2n}$ as the open subspace in $\mathbb{CP}^{2n}$ where the polynomial in (5.22) has distinct zeros (the zeros are the homogeneous coordinates $([\alpha_1, \beta_1], \dots, [\alpha_{2n}, \beta_{2n}])$). Viewed in this light, there is a continuous map (to be denoted by $\mathbb{F}$) from $\mathbb{CP}^{2n}$ to $\overline{\mathcal{C}}_{2n}$ that is the identity map on the top strata $\mathcal{C}_{2n}$ and collapses various lower dimensional varieties in $\mathbb{CP}^{2n}$ corresponding to points where (5.22) has roots with multiplicity onto corresponding strata in the $\overline{\mathcal{C}}_{2n}$ compactification depicted in (3.2).

Part 2. Let w denote the generator of $H^2(\mathbb{CP}^{2n})$ with it understood that the coefficient ring is $\mathbb{Z}/2\mathbb{Z}$. The powers of w generate $H^*(\mathbb{CP}^{2n})$, and in particular, w^{2n} is non-zero. If $\mathbb{F}^* \omega = w$, then ω^{2n} can't be zero (because pull-back and cup-product commute). Now the class $\mathbb{F}^* \omega$ is equal to w if and only if its restriction to a complex line in $\mathbb{CP}^{2n}$ is non-zero. To see about that, consider first the degree $2n-1$ homogeneous polynomial $w^{2n-1} - z^{2n-1}$ which has $2n-1$ distinct roots which are the points where $z = \eta w$ with η being a $(2n-1)$ 'st root of unity. Supposing that $a \in \mathbb{C}$, then the polynomial

$$(5.23) \quad (aw - z)(w^{2n-1} - z^{2n-1}) = aw^{2n} - w^{2n-1}z + awz^{2n-1} - z^{2n}$$

has distinct roots as long as $a^{2n-1} \neq 1$. In any event, as a varies in $\mathbb{C} \cup \infty$, the corresponding locus in $\mathbb{CP}^{2n}$ traces out a degree 1 rational curve. Let Σ denote this curve.

Part 3. The $\mathbb{F}$ -image of this curve in $\mathcal{C}_{2n}$ is a set of configurations where $2n-1$ points are fixed to be the points in $\mathbb{CP}^1$ where $z = \eta w$

with $\eta^{2n-1} = 1$, and with the last point (parametrized in homogeneous coordinates by $[a, 1]$) varying at will except that it must avoid the $2n - 1$ fixed points. The $\mathbb{F}$ -image of Σ in $\overline{\mathcal{C}}_{2n}$ adds $2n - 1$ configurations in the stratum $\mathcal{C}_{2n-2}$ from (3.2) which correspond to the collisions between the a -parametrized point and each of the $2n - 1$ fixed points.

Now note that if x denotes the $z = 0$ point in $\mathbb{CP}^1$, then $\mathbb{F}(\Sigma) \cap \mathcal{C}_{2n-2}$ is disjoint from $\mathcal{C}_{2n}^x$. With that understood, recall that the class ω is dual in $\mathcal{C}_{2n}$ to $\mathcal{C}_{2n}^x$. Since the image $\mathbb{F}(\Sigma)$ intersects $\mathcal{C}_{2n}^x$ just once (where the parameter $a = 0$), the image curve $\mathbb{F}(\Sigma)$ has pairing 1 with ω in $\mathcal{C}_{2n}$.

The preceding calculation is not the end of the story because the closure of the divisor $\mathcal{C}_{2n}^x$ in $\overline{\mathcal{C}}_{2n}$ consists of the whole of $\overline{\mathcal{C}}_{2n-2}$ and in particular, the stratum $\mathcal{C}_{2n-2}$. Since $\mathbb{F}(\Sigma)$ intersects $\mathcal{C}_{2n-2}$ an odd number of times, it is necessary to show that each such intersection counts as zero when evaluating the pairing of $\mathbb{F}(\Sigma)$ with ω .

This is done by deforming the stratum $\overline{\mathcal{C}}_{2n-2}$ so that a neighborhood of the points in $\mathbb{F}(\Sigma) \cap \mathcal{C}_{2n-2}$ are pushed into $\mathcal{C}_{2n}$. Here is how to do that: Let β denote a chosen function with compact support on $\mathcal{C}_{2n-2} - \mathcal{C}_{2n-2}^x$ that is equal to 1 at the $2n - 1$ configurations that comprise the intersection of $\mathbb{F}(\Sigma)$ with $\mathcal{C}_{2n-2}$, and equal to zero on any configuration that contains a point with distance less than $\frac{1}{100n}$ from x . Define $\mathcal{F} : [0, 1] \times \overline{\mathcal{C}}_{2n-2} \rightarrow \overline{\mathcal{C}}_{2n}$ by the rule whereby

$$(5.24) \quad \mathcal{F}(t, \mathbf{q}) = \mathbf{q} \cup \left(\left[\frac{1}{1000n} t \beta(\mathbf{q}), 1 \right], \left[-\frac{1}{1000n} t \beta(\mathbf{q}), 1 \right] \right).$$

Of particular note is that $\mathbb{F}(\Sigma)$ is disjoint from $\mathcal{F}(t = 1, \mathbf{q})$. Thus, the $\mathbb{F}$ -image of Σ in $\overline{\mathcal{C}}_{2n}$ also has pairing 1 with the class ω . This is to say that $\mathbb{F}^* \omega$ has pairing 1 with Σ . As noted previously, this implies that $\mathbb{F}^* \omega = w$. q.e.d.

5.5. Some non-zero classes in the cohomology the space $\overline{\mathbb{RP}}$.

A part of the cohomology of $\overline{\mathbb{RP}}$ is described in the next lemma. The lemma uses ω to denote the obstruction class in $H^2(\overline{\mathcal{C}}_{2n})$ as was done in the previous subsection.

Lemma 5.6. *The cohomology of $\overline{\mathbb{RP}}$ contains the vector space of polynomials in a degree two class τ with coefficient ring $\mathcal{Q}$ as described in the subsequent bullets. Thus,*

$$\mathcal{Q} \oplus \mathcal{Q}\tau \oplus \mathcal{Q}\tau^2 \oplus \mathcal{Q}\tau^3 \oplus \cdots$$

- *The definition of τ : Let $\hat{\omega}$ denote a given representative cocycle for ω . Then $\pi^* \hat{\omega}$ can be written as $d\nu$ with ν being a 1-cochain whose restriction to each fiber algebraically generates the cohomology of the fiber. The cochain $\nu \wedge \nu$ is closed and non-zero in $H^*(\overline{\mathbb{RP}})$. The class of $\nu \wedge \nu$ is the class τ .*
- *The definition of $\mathcal{Q}$: What is denoted by $\mathcal{Q}$ is the 2-dimensional vector space spanned by the degree zero generator 1 and a degree*

$4n + 1$ class to be denoted by Ξ . To define the latter, let $\hat{i}$ denote a degree $(4n + 1)$ -cocycle on $\overline{\mathcal{C}}_{2n}$ obeying $d\hat{i} = \hat{\omega}^{2n+1}$. The degree $(4n + 1)$ -cocycle $\nu \wedge \pi^*\hat{\omega}^{2n} + \pi^*\hat{i}$ on $\overline{\mathbb{R}\mathbb{P}}$ is closed and non-zero in $H^{2n+1}(\overline{\mathbb{R}\mathbb{P}})$. Its class is independent of the choice for the cocycle $\hat{i}$ and it is independent of the choice for the cochain ν and the choice for the representative cocycle $\hat{\omega}$ for ω . This cohomology class is Ξ .

Proof of Lemma 5.6. Except for the assertion to the effect that Ξ is independent of the various choices, the assertion of the lemma is an instance of Proposition A.1 in the appendix. To see about invariance: Changing any of $\hat{\omega}$, ν and $\hat{i}$ maintaining the relations between them changes the cocycle representing Ξ by something that is exact (by inspection) plus the π -pull back of a closed cocycle on X of degree $4n + 1$. Meanwhile, all of those are exact because the top dimensional strata of X has dimension $4n$.

Note that if the cohomology of $\overline{\mathbb{R}\mathbb{P}}$ is computed using a CW decomposition for $\overline{\mathcal{C}}_{2n}$, then $\hat{\omega}^{2n+1}$ is zero and then $\hat{i} \equiv 0$ too. If simplicial cohomology is used, then there is no guarantee that $\hat{\omega}^{2n+1}$ is the zero cochain; but it is cohomologous to zero. q.e.d.

6. Min-max for $\mathcal{E}$ on $\overline{\mathbb{R}\mathbb{P}}$

There are some hints that the classes from the set $\{\Xi \wedge \tau^m\}_{m \in \{0,1,\dots\}}$ have min-max values that correspond to critical points of $\mathcal{E}$ on large n version of $\mathcal{C}_{2n}$. There are two indications, the first being this: If it were the case that the $(2m + 1)$ 'st $\mathcal{I}_p$ -eigenvalue for $\mathcal{E}$ had multiplicity 1 at each configuration (which can't actually happen), then its supremum on $\mathcal{C}_{2n}$ would be the min-max value for the cohomology class $\Xi \wedge \tau^m$. Thus, min-max using the classes $\{\Xi \wedge \tau^m\}$ constitute a cohomological work-around for the fact that the functions on $\mathcal{C}_{2n}$ given by ordering of the $\mathcal{I}_{(\cdot)}$ -eigenvalues are not everywhere differentiable. The second reason for hope concerns the appearance of the $\hat{\omega}^{2n}$ factor in the definition of Ξ : This factor sees the whole of $\mathcal{C}_{2n}$, not just some low codimension subset that can be pushed into $\overline{\mathcal{C}}_{2n} - \mathcal{C}_{2n}$.

6.1. The modified definition of min-max value. What follow directly is the definition of ‘min-max value’ for the context at hand. To set the stage, let $\overline{\mathcal{Y}}$ denote a fiber bundle over $\overline{\mathcal{C}}_{2n}$ with compact, finite dimensional fiber (in practice, the fiber will be $\mathbb{R}\mathbb{P}^{2m+1}$) which is smooth along each stratum of $\overline{\mathcal{C}}_{2n}$. A continuous map from $\overline{\mathcal{Y}}$ to $\overline{\mathbb{R}\mathbb{P}}$ is said to be *strata preserving* if, for each $k \in \{0, \dots, n\}$, the map sends $\overline{\mathcal{Y}}|_{\mathcal{C}_{2k}}$ to $\overline{\mathbb{R}\mathbb{P}}|_{\mathcal{C}_{2k}}$. Let $\mathfrak{F}$ denote a chosen homotopy class of strata preserving maps from $\overline{\mathcal{Y}}$ to $\overline{\mathbb{R}\mathbb{P}}$. The associated min-max value of $\mathcal{E}$ for the set $\mathfrak{F}$ is this:

(6.1) $\inf_{\Psi \in \mathfrak{F}} \max_{y \in \overline{\mathcal{Y}}} \mathcal{E}(\Psi(y)).$

The upcoming Lemma 6.1 is used to obtain the relevant version of $\overline{\mathcal{Y}}$ and set $\mathfrak{F}$. With regards to terminology: When the lemma speaks of an $\mathbb{RP}^{2m+1}$ fiber bundle, this refers to a special sort of bundle: Each point in the base has a neighborhood that identifies the fiber bundle over the neighborhood with the $\mathbb{R}^*$ quotient of the complement of the zero section in a $(2m+2)$ -dimensional vector bundle. In addition, the transition functions between two such intersecting open sets are induced by fiberwise linear maps that identify the corresponding vector bundles over the intersection of the two base sets.

Lemma 6.1. *Fix a positive integer for n and a non-negative integer for m . There is an $\mathbb{RP}^{2m+1}$ -fiber bundle $\pi_{\mathcal{Y}} : \overline{\mathcal{Y}}_{n,m} \rightarrow \overline{\mathcal{C}}_{2n}$ with a fiber preserving map to $\overline{\mathbb{RP}}$ that pushes forward the fundamental class of $\overline{\mathcal{Y}}_{n,m}$ to a class that has non-zero pairing with $\nu^m \wedge \Xi$. (The space $\overline{\mathcal{Y}}_{n,m}$ is the $X = \overline{\mathcal{C}}_{2n}$ instance of what is denoted by $\mathcal{Y}_m$ in the appendix.)*

Proof of Lemma 6.1. This is an instance of Lemma A.6 in the appendix (Section A.6). q.e.d.

Granted Lemma, 6.1, take $\mathfrak{F}$ to be the homotopy class of maps Ψ as described by Lemma 6.1. The corresponding version of the min-max value in (6.1) is denoted by $\mathcal{E}_{n,m}$.

6.2. Min-max values $\{\mathcal{E}_{n,m}\}_{n>0,m\geq 0}$. Fix a positive integer for n . The lemma that follows concerns the size of the min-max values from the set $\{\mathcal{E}_{n,m}\}_{m=0,1,\dots}$.

Lemma 6.2. *Fix a positive integer for n .*

- *For any non-negative integer m : The min-max value $\mathcal{E}_{n,m}$ is not less than the maximum over $\overline{\mathcal{C}}_{2n}$ of the function $E_{(\cdot)}$ from (1.2). Therefore, it is greater than $\kappa^{-1}n$ with κ positive and independent of n .*
- *The set $\{\mathcal{E}_{n,m}\}_{m\in\{0,1,\dots\}}$ is unbounded.*

Proof of Lemma 6.2. With regards to the lemma's first bullet: Supposing that the min-max value $\mathcal{E}_{n,m}$ is as least as large as the supremum of $E_{(\cdot)}$ on $\mathcal{C}_{2n}$, then it is a priori greater than $c_0^{-1}n$ because the supremum of $E_{(\cdot)}$ on $\mathcal{C}_{2n}$ is greater than $c_0^{-1}n$. (See Proposition 3.3.)

To prove that the min-max value is not less than $E_{(\cdot)}$, suppose that this isn't the case so as to derive some nonsense. If Ψ is from $\mathfrak{F}$, then the push-forward by Ψ of the fundamental class of $\overline{\mathcal{Y}}$ is a closed, $2m+4n+1$ dimensional chain in $\overline{\mathbb{RP}}$ that has non-zero pairing with $\nu^m \wedge \Xi$. Suppose that the maximum of the values of $\mathcal{E}$ on the image of Ψ is less than the maximum of value of $E_{(\cdot)}$ on $\mathcal{C}_{2n}$. If this is the case, then the π image of $\Psi(\overline{\mathcal{Y}})$ misses a configuration in $\mathcal{C}_{2n}$. (Remember that the maximum of $E_{(\cdot)}$ on $\overline{\mathcal{C}}_{2n}$ is at configurations in $\mathcal{C}_{2n}$.) Now comes an important point: The cocycle $\hat{\omega}^{2n}$ is cohomologous to zero in the complement of any

configuration in $\overline{\mathcal{C}}_{2n}$ because the fundamental class of the complement of any configuration is zero. This is to say that $\hat{\omega}^{2n}$ on the complement of a configuration can be written as $d\beta$ with β being a degree $4n - 1$ cochain. Also, $\hat{\omega}^{2n+1}$ can be written there as $d(\hat{\omega} \wedge \beta)$. This implies that the pairing between a cocycle representative of $\nu^m \wedge \Xi$ and the Ψ –push-forward of the fundamental class of $\overline{\mathcal{Y}}$ is the same as the pairing between $\nu^{2m+1} \wedge d\pi^*\beta - \nu^{2m} \wedge d\pi^*(\hat{\omega} \wedge \beta)$ and that class, which is zero because the class is closed and

(6.2)
$$\nu^{2m+1} \wedge d\pi^*\beta - \nu^{2m} \wedge d\pi^*(\hat{\omega} \wedge \beta) = -d(\nu^{2m+1} \wedge \beta).$$

This conclusion is the desired nonsense because it runs afoul of the assumptions about Ψ .

With regards to the second bullet: Let $\mathfrak{c}_{2m}$ denote a homology class on $\overline{\mathbb{R}\mathbb{P}}$ that is dual to ν^{2m} . This class is represented by the fundamental class of any embedding of $\mathbb{R}\mathbb{P}^{2m}$ in any fiber of $\overline{\mathbb{R}\mathbb{P}}$ that induces an isomorphism on the first homology of the fiber. An important point is that the min-max value of $\mathcal{E}$ for $\mathfrak{c}_{2m}$ is an unbounded function of m . This follows from instances of Proposition 1.1 and Lemmas 5.2 and 5.3 which together imply the following: Given any positive number E , there exists an integer (call it k_E) such that the k_E ’th lowest $\mathcal{I}_p$ -eigenvalue is greater than E at each configuration in $\overline{\mathcal{C}}_{2n}$. This implies that the min-max value for $\mathfrak{c}_{2m}$ is greater than E if $2m > k_E$. As explained next, the number E (assuming $2m > k_E$) is a lower bound for $\mathcal{E}_{n,m}$.

To prove that $\mathcal{E}_{n,m} \geq E$, suppose to the contrary that $\mathcal{E}_{n,m}$ is less than E to generate nonsense. In this event, there is a map $\Psi : \overline{\mathcal{Y}} \rightarrow \overline{\mathbb{R}\mathbb{P}}$ from $\mathfrak{F}$ with $\Psi^*\nu^{2m}$ evaluating to zero on the generator of $2m$ ’th homology of the fibers of the projection from $\overline{\mathcal{Y}}$ to $\overline{\mathcal{C}}_{2m}$. Meanwhile, the cohomology of $\overline{\mathcal{Y}}$ is described by Lemma A.2 in the appendix, and it follows from that description, and from what is said in Sections A.2 and A.3, that Lemma A.2’s cohomology class τ can be taken to be that of $\Psi^*\nu^2$. And, according to Lemma A.2, the class τ^m does have non-zero pairing with the $2m$ ’th homology of the fibers of the projection to $\overline{\mathcal{C}}_{2n}$. q.e.d.

6.3. Convergence to a critical point. The proposition that follows directly makes an assertion to the effect that the min-max value of $\mathcal{E}$ for any $m \geq 0$ version of $\nu^m \wedge \Xi$ is a critical value of $\mathcal{E}$ on some strata of $\mathcal{C}_{2n}$. This is to say that there exists $k \in \{0, 1, \dots, n\}$ and a critical point of $\mathcal{E}$ on $\mathcal{C}_{2k}$ where $\mathcal{E}$ ’s value is equal to the min-max value $\mathcal{E}_{n,m}$.

Proposition 6.3. *For each positive integer n and non-negative integer m , there exists $k \in \{0, \dots, n\}$ and a critical point of $\mathcal{E}$ on $\mathbb{R}\mathbb{P}|_{\mathcal{C}_{2k}}$ where $\mathcal{E}$ is equal to $\mathcal{E}_{n,m}$.*

Proof of Proposition 6.3. The proof of this proposition has eight parts. Part 1 has a preliminary observation. Parts 2–4 focus on just the differential of $\mathcal{E}$ along the fibers of the projection to $\overline{\mathcal{C}}_{2n}$. The final parts

consider the behavior of the differential of $\mathcal{E}$ on lifts to $\overline{\mathbb{RP}}$ of vector fields along the strata of $\overline{\mathcal{C}}_{2n}$.

Part 1. Here is an important point to keep in mind: Fix a positive number to be denoted by E and let $\overline{\mathbb{RP}}^E$ denote the subspace of elements in $\overline{\mathbb{RP}}$ that are characterized as follows: An element $(\mathbf{p}, [f])$ is in $\overline{\mathbb{RP}}^E$ if and only if f is in the span of the $\mathcal{I}_{\mathbf{p}}$ -eigenvectors with eigenvalue at most E . The point to keep in mind is that the subspace $\overline{\mathbb{RP}}^E$ is compact.

Part 2. Suppose for the moment that $(\mathbf{p}, [f])$ is a given element in $\overline{\mathbb{RP}}$. Let E denote a positive number that is not an $\mathcal{I}_{\mathbf{p}}$ -eigenvalue. Write f as $f_{<} + f_{>}$ where $f_{<}$ is the L^2 orthogonal projection of f onto the subspace $\mathbb{H}_{\mathbf{p}, E}$ (the subspace in $\mathbb{H}_{\mathbf{p}}$ that is spanned by the $\mathcal{I}_{\mathbf{p}}$ -eigensections with eigenvalue less than E). An observation follows: If $\mu > 0$ and if the S^2 -integral of $|f_{>}|^2$ is greater than μ , then $\mathcal{E}(\mathbf{p}, [f])$ will be greater than μE .

The preceding observation has the following consequence: Suppose that $\delta > 0$ and that Ψ is a map from $\mathfrak{F}$ that maps the whole of $\overline{\mathcal{Y}}$ to the $\mathcal{E}(\cdot) < \mathcal{E}_{n,m} + \delta$ part of $\overline{\mathbb{RP}}$. If $(\mathbf{p}, [f])$ is from the image of Ψ , and if E is greater than $\mathcal{E}_{n,m} + \delta$, then the S^2 integral of the corresponding $|f_{>}|^2$ is strictly less than 1 (it is at most $\frac{\mathcal{E}_{n,m} + \delta}{E}$). This implies, in particular, that $f_{<}$ is non-zero.

As explained momentarily, the observation in the previous paragraph can be used to continuously deform the map Ψ to a new strata preserving map that has the properties listed in the next lemma.

Lemma 6.4. *Fix $\delta > 0$ and suppose that Ψ is a map from $\mathfrak{F}$ when $\mathcal{E}(\Psi(\cdot)) < \mathcal{E}_{n,m} + \delta$. Then there exists a homotopy of Ψ in $\mathfrak{F}$ whose end member (denoted by Ψ') obeys the following:*

- The map Ψ' is from $\mathfrak{F}$.
- Supposing that $y \in \overline{\mathcal{Y}}$, then $\mathcal{E}(\Psi'(y)) \leq \mathcal{E}(\Psi(y))$
- Supposing that $y \in \overline{\mathcal{Y}}$ write $\Psi(y)$ as $(\mathbf{p}, [f])$. Then f is in the span of the $\mathcal{I}_{\mathbf{p}}$ -eigensections with eigenvalue less than $\mathcal{E}_{n,m} + 2\delta$.

Proof of Lemma 6.4. To set up the construction of Ψ' , suppose for the moment that $\mathbf{p}$ is a given element in $\overline{\mathcal{C}}_{2n}$. Fix a number to be denoted by $E_{\mathbf{p}}$ that is greater than $\mathcal{E}_{n,m} + \delta$ but less than $\mathcal{E}_{n,m} + 2\delta$ and which is not an $\mathcal{I}_{\mathbf{p}}$ -eigenvalue. According to Lemma 5.2, there is an open neighborhood of $\mathbf{p}$ in $\overline{\mathcal{C}}_{2n}$ with the property that if $\mathbf{q}$ is in this neighborhood, then $E_{\mathbf{p}}$ is not an $\mathcal{I}_{\mathbf{q}}$ -eigenvalue. Fix such a neighborhood and denote it by $U_{\mathbf{p}}$. The various $\mathbf{p} \in \overline{\mathcal{C}}_{2n}$ versions of $U_{\mathbf{p}}$ form an open cover of $\overline{\mathcal{C}}_{2n}$ and thus there is a finite subcover. This is to say that there is a finite set of points whose corresponding $\overline{U}_{(\cdot)}$ sets cover $\overline{\mathcal{C}}_{2n}$. Choose such a finite set and denote it by Λ . Then choose a partition of

unity subordinate to the associated cover. When $\mathfrak{p}$ is from Λ , then the corresponding function from the partition of unity is denoted below by $\chi_{\mathfrak{p}}$.

Supposing now that $y \in \overline{\mathcal{Y}}$, write $\Psi(y)$ as $(\mathfrak{q}, [f])$. Fix for the moment an element $\mathfrak{p} \in \Lambda$ whose corresponding set $\overline{U}_{\mathfrak{p}}$ contains $\mathfrak{q}$. Then write f as $f_{\mathfrak{p},<} + f_{\mathfrak{p},>}$ with the former being the projection of f to the span of the $\mathcal{I}_{\mathfrak{q}}$ -eigenvalues with eigenvalue less than $E_{\mathfrak{p}}$. For any given $t \in [0, 1]$, define f_t by the rule whereby

(6.3)
$$f_t \equiv \frac{1}{Z_t} \left(\sum_{\mathfrak{p} \in \Lambda} \chi_{\mathfrak{p}} f_{\mathfrak{p},<} + (1-t) \sum_{\mathfrak{p} \in \Lambda} \chi_{\mathfrak{p}} f_{\mathfrak{p},>} \right)$$

where Z_t is positive and chosen so that the S^2 -integral of $|f_t|^2$ is equal to 1. Because $f_{\mathfrak{p},<}$ is non-zero, this deformation is continuous with respect to t .

Lemma 5.2 guarantees that the assignment of $t \in [0, 1]$ and $y \in \overline{\mathcal{Y}}$ to $\Psi_t(y)$ defines a continuous map from $[0, 1] \times \overline{\mathcal{Y}}$ into $\overline{\mathbb{R}\mathbb{P}}$. Since it is strata preserving, this is a homotopy of Ψ in $\mathfrak{F}$. Set Ψ' to be the $t = 1$ member. The construction is such that if $y \in \overline{\mathcal{Y}}$, then the map $t \rightarrow \mathcal{E}(\Psi_t(y))$ (which is $\mathcal{E}(\mathfrak{q}, f_t)$) is non-increasing and it is strictly decreasing if $f_{\mathfrak{p},>}$ is non-zero for at least one $\mathfrak{p} \in \Lambda$. By construction, $\mathcal{E}(\Psi_1(y)) \leq \mathcal{E}_{n,m} + 2\delta$ because f_1 has zero projection to the span of the $\mathcal{I}_{\mathfrak{q}}$ -eigenvectors with eigenvalue greater than $\mathcal{E}_{n,m} + 2\delta$. q.e.d.

Part 3. The preceding lemma has two important implication with the first being this: Having fixed δ for Lemma 6.4, suppose that $(\mathfrak{p}, [f])$ is in the image of Lemma 6.4's map Ψ' . If there are *no* $\mathcal{I}_{\mathfrak{p}}$ -eigenvalues between $\mathcal{E}_{n,m} - \delta'$ and $\mathcal{E}_{n,m} + 2\delta$ with δ' any given positive number, then $\mathcal{E}(\mathfrak{p}, [f])$ is at most $\mathcal{E}_{n,m} - \delta'$.

Here is the second implication:

(6.4) *With Ψ' as described by Lemma 6.4, there exists $(\mathfrak{p}, [f])$ from the image of Ψ' with $\mathcal{E}(\mathfrak{p}, [f]) \geq \mathcal{E}_{n,m}$ such that the projection of f to the span of the $\mathcal{I}_{\mathfrak{p}}$ -eigensections with eigenvalues between $\mathcal{E}_{n,m} - \sqrt{\delta}$ and $\mathcal{E}_{n,m} + 2\delta$ has L^2 norm greater than $(1 - 3\sqrt{\delta})^{\frac{1}{2}}$.*

To prove this: Supposing that (6.4) fails for a given element $(\mathfrak{p}, [f])$ from the image of Ψ' , write $f = f_{<} + f_{>}$ with $f_{<}$ denoting the L^2 orthogonal projection of f to the span of the $\mathcal{I}_{\mathfrak{p}}$ -eigensections with eigenvalue less than $\mathcal{E}_{n,m} - \sqrt{\delta}$. Using this decomposition, one has that

(6.5)
$$\mathcal{E}(\mathfrak{p}, [f]) \leq (\mathcal{E}_{n,m} - \sqrt{\delta}) \|f_{<}\|_{\mathbb{L}}^2 + (\mathcal{E}_{n,m} + 2\delta)(1 - \|f_{<}\|_{\mathbb{L}}^2);$$

and this is strictly less than $\mathcal{E}_{n,m} - \delta$ if $\|f_{<}\|_{\mathbb{L}}^2$ is greater than $3\sqrt{\delta}$ which is when $\|f_{>}\|^2$ is less than $1 - 3\sqrt{\delta}$. Thus, (6.4) can't fail for every f

with $\mathcal{E}(\mathfrak{p}, [f]) \geq \mathcal{E}_{n,m}$ because the maximum value of $\mathcal{E}$ on the image of Ψ' would then be less than $\mathcal{E}_{n,m}$ which violates the assumptions.

Part 4. What is said by Part 1 about compactness, and what is said in Lemma 5.6 and (6.4) have the following implications: Fix any decreasing sequence of positive numbers with limit zero (denoted by $\{\delta_j\}_{j=1,2,\dots}$) and then a corresponding sequence of maps from $\mathfrak{F}$ (denoted by $\{\Psi_j\}_{j=1,2,\dots}$) such that for each index j , the conditions in Lemma 6.4 are obeyed with $\delta = \delta_j$ and with $\Psi' = \Psi_j$. For each j , set A_j to be the $\mathcal{E} \geq \mathcal{E}_{n,m}$ part of $\Psi_j(\overline{\mathcal{Y}})$. The set of limit points of the sequence $\{A_j\}_{j=1,2,\dots}$ contains at least one $(\mathfrak{p}, [f]) \in \overline{\mathbb{RP}}$ with f being an $\mathcal{I}_{\mathfrak{p}}$ -eigenvector with eigenvalue is $\mathcal{E}_{n,m}$, which is to say that $-\Delta f = \mathcal{E}_{n,m}f$. (Remember that the sequence $\{A_j\}_{j=1,2,\dots}$ is a sequence in a compact subset of $\overline{\mathbb{RP}}$. As a consequence, any corresponding sequence of elements (the first from $\mathcal{A}_1$, the second from $\mathcal{A}_2$, and so on) has a convergent subsequence. By way of terminology: The set of limits of these subsequences is the set of limit points of $\{A_j\}_{j=1,2,\dots}$.)

Part 5. This part of the proof is a digression to present some background and then introduce notation for the remaining parts. To start the digression, suppose for the moment that $\{\phi_t\}_{t \in \mathbb{R}}$ is a 1-parameter family of area preserving diffeomorphisms of S^2 with the $t = 0$ member being the identity diffeomorphism. This family acts in the obvious way on $\overline{\mathcal{C}}_{2n}$ so as to preserve the stratification in (3.2). To elaborate: Supposing that $k \in \{0, \dots, n\}$ and $\mathfrak{p}$ is in the stratum $\mathcal{C}_{2k}$, then $\phi_t(\mathfrak{p})$ is the configuration whose points are the ϕ_t -images of the points in $\mathfrak{p}$. This 1-parameter family can be lifted to act on $\overline{\mathbb{RP}}$ as follows: The parameter t member ϕ_t sends any given element $(\mathfrak{p}, [f])$ to the element $(\phi_t(\mathfrak{p}), [f \circ \phi_t^{-1}])$. This lifted map is denoted by $\bar{\phi}_t$.

What follows are some remarks about this definition. The first remark is that an identification between respective line bundles $(\phi_t^{-1})^* \mathcal{I}_{\mathfrak{p}}$ and $\mathcal{I}_{\phi_t(\mathfrak{p})}$ is not required to define $\bar{\phi}_t$ because the two possible isometric identifications differ by a sign. The second remark is that the lifted map $\bar{\phi}_t$ does not map $\mathcal{I}_{\mathfrak{p}}$ -eigensections to $\mathcal{I}_{\phi_t(\mathfrak{p})}$ -eigensections. The third remark is that the lifted map is not differentiable with respect to t ; it is only continuous. This is because the t -derivative of the ϕ_t^{-1} -pull back of a generic element in $\mathbb{H}_{\phi_t(\mathfrak{p})}$ will not have finite $\mathbb{H}_{\mathfrak{p}}$ -norm. (The $\mathbb{H}_{\mathfrak{p}}$ norm of the t -derivative of the pull-back will be finite however if the integral on the right hand side of (4.14) is finite for the element in question. See Section 4.4 for more about this issue.)

Now suppose that $\Psi : \overline{\mathcal{Y}} \rightarrow \overline{\mathbb{RP}}$ is a map from the set $\mathfrak{F}$ and that $\tau : \overline{\mathcal{Y}} \rightarrow [0, 1]$ is a continuous map. The preceding constructions can be used to obtain a 1-parameter family of deformations of Ψ with any given $t \in \mathbb{R}$ version being the map (denoted by $\Psi_{\tau,t}$) that is defined by the rule

(6.6)
$$y \rightarrow \Psi_{\tau,t}(y) \equiv \bar{\phi}_{\tau(y)t}(\Psi(y)).$$

Each member of this family is from $\mathfrak{F}$. (Each member is strata preserving because both $\bar{\phi}_{(\cdot)}$ and Ψ are strata preserving.

Part 6. This part of the proof starts with a digression to serve as a reminder with regards to observations from Section 3. In particular, remember that when $(\mathfrak{p}, [f])$ is a given element in $\overline{\mathbb{RP}}$, then the value of $\mathcal{E}$ on $\bar{\phi}_t(\mathfrak{p}, [f])$ (the latter is $(\phi_t(\mathfrak{p}), [f \circ \phi_t^{-1}])$) can be written as

(6.7)
$$\mathcal{E}(\bar{\phi}_t(\mathfrak{p}, [f])) = \int_{S^2} \langle df, df \rangle_{\mathfrak{m}_t}$$

where $\langle \cdot, \cdot \rangle_{\mathfrak{m}_t}$ denotes the ϕ_t -pull back of the round metric’s inner product on T^*S^2 .

An important conclusion from (6.7) is that the map $t \rightarrow \mathcal{E}(\bar{\phi}_t(\cdot))$ is differentiable with respect to t , this is notwithstanding the fact that $\bar{\phi}_t$ is only continuous. Moreover, the t -derivative of $\mathcal{E}(\bar{\phi}_t(\mathfrak{p}, [f]))$ can be written using (6.7) as

(6.8)
$$\int_{S^2} \langle \mathfrak{k}_t, df \otimes df \rangle_{\mathfrak{m}_t}$$

where $\mathfrak{k}_t$ denotes the Lie derivative of $\mathfrak{m}_t$ along the vector field on S^2 that is obtained by first viewing $\phi_{(\cdot)}$ as a map from $\mathbb{R} \times S^2$ to S^2 and then using this map to push forward the tangent vector to the $\mathbb{R}$ factor.

The next lemma supplies a key observation about (6.7) and (6.8).

Lemma 6.5. *Suppose that δ is a positive number and that the $t = 0$ and $(\mathfrak{p}, [f]) \in \overline{\mathbb{RP}}$ version of (6.8) is more negative than $-\delta$. Then there is an open neighborhood of $(\mathfrak{p}, [f])$ in $\overline{\mathbb{RP}}$ (denoted by $\overline{U}$) and a positive number (denoted by ϵ) with the following significance: If $\xi \in \overline{U}$ and $t \in (-\epsilon, \epsilon)$, then $\mathcal{E}(\bar{\phi}_t(\xi)) \leq \mathcal{E}(\xi) - \frac{1}{4}\delta t$.*

Proof of Lemma 6.5. The first observation for the proof concerns (6.8): If this is more negative than $-\delta$ at $(\mathfrak{p}, [f])$, then there is an open neighborhood of $(\mathfrak{p}, [f])$ in $\overline{\mathbb{RP}}$ (this will be the neighborhood $\overline{U}$) such that the $t = 0$ and $(\mathfrak{p}', [f'])$ version of (6.8) is more negative than $-\frac{1}{2}\delta$ when $(\mathfrak{p}', [f'])$ is in $\overline{U}$. Except for notation, the proof of that this is so is identical to the proof of Lemma 5.1. Then, looking at the t -dependence, it follows from the continuity of $\mathfrak{k}_{(\cdot)}$ that there exists $\epsilon > 0$ such that any $t \in [0, \epsilon)$ and $(\mathfrak{p}', [f']) \in \overline{U}$ version of (6.8) is more negative than $-\frac{1}{4}\delta$. The assertion in the lemma follows directly from this last observation by integrating (6.8) with respect to t . q.e.d.

Part 7. Take $(\mathfrak{p}, [f])$ and δ and then $\overline{U}$ and ϵ from Lemma 6.5. Suppose that Ψ is from the set $\mathfrak{F}$. Let $\tau_*(\cdot)$ denote a continuous function on $\overline{\mathcal{Y}}$ mapping to the interval $[0, 1]$ with compact support on $\Psi^{-1}(\overline{U})$. A path of deformations of Ψ parametrized by $[0, \epsilon]$ is defined by the rule

whereby $(t, y) \in [0, \epsilon] \times \overline{\mathcal{Y}}$ is sent to $\bar{\phi}_{\tau_*(y)t}(\Psi(y))$. The $t \in [0, \epsilon]$ member of this deformation family is denoted by $\Psi_{\tau_*,t}$. The behavior of $\mathcal{E}$ with respect to these deformations is the topic of the next paragraphs.

With regards to $\mathcal{E}$: If y is a point from $\overline{\mathcal{Y}}$ that is not mapped by Ψ to $\overline{U}$, then $\Psi_{\tau_*,t}(y)$ is $\Psi(y)$ because $\tau_*(y)$ is zero; and thus $\mathcal{E}(\Psi_{\tau_*,t}(y)) = \mathcal{E}(\Psi(y))$. With this understood, suppose now that y is mapped by Ψ to $\overline{U}$. In this case, Lemma 6.5 finds

$$(6.9) \quad \mathcal{E}(\Psi_{\tau_*,t}(y)) \leq \mathcal{E}(y) - \frac{1}{4}\tau_*(y)t\delta.$$

This implies (among other things) that the maximum of $\mathcal{E}$ on the image of $\Psi_{\tau_*,t}$ is no greater than its maximum on the image of Ψ .

Part 8. Let $\mathfrak{C}_{n,m}$ denote the subset in $\overline{\mathbb{RP}}$ consisting of elements $(\mathfrak{p}, [f])$ with f being an $\mathcal{I}_{\mathfrak{p}}$ -eigensection with eigenvalue $\mathcal{E}_{n,m}$. A key point to keep in mind for what follows that $\mathfrak{C}_{n,m}$ is a closed, compact set.

Suppose for the sake of argument that there is not one $k \in \{0, 1, \dots, n\}$ such that the intersection of $\mathfrak{C}_{n,m}$ with the stratum $\mathcal{C}_{2k}$ contains a critical point of $\mathcal{E}$ on $\mathcal{C}_{2k}$. The four steps that follow use this assumption to generate nonsense.

Step 1: Suppose now that $k \in \{0, 1, \dots, n\}$ and that $\mathfrak{p}$ is in the stratum $\mathcal{C}_{2k}$ of the compactification $\overline{\mathcal{C}}_{2k}$. Let v denote a tangent vector to $\mathfrak{p}$ along this stratum (this being a vector in $\otimes_{p \in \mathfrak{p}} TS^2|_p$). Section 2.6 describes how to extend v from the points in $\mathfrak{p}$ as a divergence free vector field on S^2 . Denote this vector field by $\mathfrak{v}$. Integrating this vector field gives a 1-parameter family of area preserving diffeomorphisms of S^2 . The family is parameterized by the coordinate t on the real line $\mathbb{R}$ with the $t = 0$ member being the identity diffeomorphism and with the t -derivative at $t = 0$ being $\mathfrak{v}$. The diffeomorphism labeled by any given $t \in \mathbb{R}$ is denoted by ϕ_t .

Step 2: Because of the assumptions, for each $k \in \{0, 1, \dots, n\}$ the following is true: If $(\mathfrak{p}, [f])$ is from $\mathfrak{C}_{n,m} \cap \mathbb{RP}_{2k}$, then there is a corresponding ϕ_t from Step 1 for which (6.8) is negative. Therefore, each $(\mathfrak{p}, [f])$ in $\mathfrak{C}_{n,m}$ has a corresponding ϵ, δ and $\overline{U}$ as described in Lemma 6.5. Since $\mathfrak{C}_{n,m}$ is compact, there exists positive ϵ_* and δ_* such that each $(\mathfrak{p}, [f]) \in \mathfrak{C}_{n,m}$ version of ϵ and δ is greater than the respective ϵ_* and δ_* . In addition, there exists a finite set of elements in $\mathfrak{C}_{n,m}$ whose corresponding versions of $\overline{U}$ supply an open cover of $\mathfrak{C}_{n,m}$. Let Λ denote this finite set of elements and let $\overline{U}$ denote the union of these Λ -versions of $\overline{U}$.

Step 3: By virtue of what is said in Part 4, there exists some positive number to be denoted by $\epsilon_{n,m}$ with the following significance: If Ψ is from $\mathfrak{F}$ and if the maximum of $\mathcal{E}$ on Ψ is less than $\mathcal{E}_{n,m} + \epsilon_{n,m}$, then the part of the image of Ψ where $\mathcal{E} > \mathcal{E}_{n,m} - \epsilon_{n,m}$ is contained in $\overline{U}$.

With this understood, suppose henceforth that the maximum of $\mathcal{E}$ on the given version of Ψ is indeed less than $\mathcal{E}_{n,m} + \epsilon_{n,m}$.

Let τ denote the function on $\overline{\mathcal{Y}}$ given by the rule

(6.10)
$$\tau_0(\cdot) \rightarrow \chi \left(\frac{2}{\epsilon_{n,m}} \left(\mathcal{E}_{n,m} - \frac{1}{2} \epsilon_{n,m} - E(\Psi(\cdot)) \right) \right).$$

This function is equal to 1 on the part of $\overline{\mathcal{Y}}$ where $\mathcal{E}(\Psi(\cdot)) \geq \mathcal{E}_{n,m} - \frac{1}{2} \epsilon_{n,m}$ and it is equal to 0 on the part where $\mathcal{E}(\Psi(\cdot)) < \mathcal{E}_{n,m} - \epsilon_{n,m}$. In particular, this implies that the support of τ is in $\overline{U}$.

Step 4: Label the elements in Λ consecutively starting from 1 and use the corresponding labels for the Λ -versions of $\overline{U}$. Fix for each label i from Λ , fix an open set V_i in U_i having compact closure in U_i and such that the union of these sets also covers $\mathfrak{C}_{n,m}$. Then, for each index i from Λ , fix a continuous function with values between 0 and 1 to be denoted by σ_i which is compactly supported on U_i and is equal to 1 on V_i .

Apply the homotopy from Part 7 to Ψ using $\overline{U}_1$ and with τ_* being $\sigma_1 \tau_0$ with τ_0 defined via (6.10). Let Ψ_1 denote the $t = 1$ end member of this homotopy. Now apply the homotopy from Part 7 again using $\overline{U}_2$ starting from Ψ_1 and with τ_* being $\sigma_2 \tau_0$. Continue in this vein until all of the Λ -versions of $\overline{U}$ are accounted for. Doing that results in a map from $\mathfrak{F}$ (denoted by Ψ_*) that has non-zero pairing with $\tau^{2m} \wedge \Xi$.

By virtue of (6.9), this map Ψ_* obeys

(6.11)
$$\max_{\overline{\mathcal{Y}}} \mathcal{E}(\Psi_*(\cdot)) \leq \max_{\overline{\mathcal{Y}}} (\Psi(\cdot)) - c_0^{-1} \epsilon_* \delta_*.$$

This result is the desired nonsense because if maximum of $\mathcal{E}(\Psi(\cdot))$ is less than $\mathcal{E}_{n,m} + \frac{1}{2} c_0^{-1} \epsilon_* \delta_*$, then the maximum of $\mathcal{E}$ on Ψ_* will be less than $\mathcal{E}_{n,m}$ which can't happen by virtue of $\mathcal{E}_{n,m}$ being the infimum of the set of maxima of $\mathcal{E}$ on the maps from $\mathfrak{F}$. q.e.d.

6.4. On the existence of critical points of $\mathcal{E}$ on $\mathcal{C}_{2n}$. The analysis so far allows for the following pathology: There exists a positive integer (call it j) such that all of the critical point of $\mathcal{E}$ that are obtained via Proposition 6.3 using every possible pair of integers n and m (with $n \geq 1$ and $m \geq 0$) lie in $\overline{\mathcal{C}}_{2j}$ (which is a subspace of all $n \geq j$ versions of $\overline{\mathcal{C}}_{2n}$). We can't rule this out at present. In any event, the question directs us to derive a more refined picture of the behavior of $\mathcal{E}$ on $\mathcal{C}_{2n}$ near the lower dimensional strata in its closure $\overline{\mathcal{C}}_{2n}$. The upcoming Lemma 6.6 is an example of what might be needed.

To set the stage for Lemma 6.6, suppose that $j \in \{0, \dots, n-1\}$ and that $\mathfrak{q}$ is the minimal representative of a $\mathbb{Z}/2\mathbb{Z}$ divisor in the $\mathcal{C}_{2j}$ stratum of $\overline{\mathcal{C}}_{2n}$. Let $\{p_1, \dots, p_{2n-2j}\}$ denote distinct points in $S^2 - \mathfrak{q}$ and set $\mathfrak{p} \equiv \{p_1, \dots, p_{2n-2j}\} \cup \mathfrak{q}$. Fix n disjoint, embedded arcs in $S^2 - \mathfrak{q}$ whose endpoints pair the points in the set $\{p_1, \dots, p_{2n-2j}\}$. Let Σ denote

their union. Then choose a point to be denoted by x in $S^2 - (\Sigma \cup \mathbf{q})$. Since the bundle $\mathcal{I}_{\mathbf{q}}$ and $\mathcal{I}_{\mathbf{p}}$ are isomorphic on $S^2 - (\Sigma \cup \mathbf{q})$, the choice of an isomorphism of their fibers at x extends uniquely to an isomorphism between the two line bundle over their common domain $S^2 - (\Sigma \cup \mathbf{q})$. A choice of such a point x and a corresponding isomorphism will be made momentarily.

To continue with the stage setting: Let $f_{\mathbf{q}}$ denote an $\mathcal{I}_{\mathbf{q}}$ -eigensection and let $f_{\mathbf{p}}$ denote an $\mathcal{I}_{\mathbf{p}}$ -eigensection; and let $E_{\mathbf{q}}$ and $E_{\mathbf{p}}$ denote their respective eigenvalues. Take x to be a point where neither $f_{\mathbf{q}}$ nor $f_{\mathbf{p}}$ is zero and then choose the isomorphism between $\mathcal{I}_{\mathbf{q}}$ and $\mathcal{I}_{\mathbf{p}}$ at x so that $f_{\mathbf{q}}f_{\mathbf{p}}$ is positive at x . (The product $f_{\mathbf{q}}f_{\mathbf{p}}$ is a section of the line bundle $\mathcal{I}_{\mathbf{p}'}$ with $\mathbf{p}' \equiv \{p_1, \dots, p_{2n-2j}\}$. The isomorphism between $\mathcal{I}_{\mathbf{p}}$ and $\mathcal{I}_{\mathbf{q}}$ on $S^2 - (\Sigma \cup \mathbf{q})$ gives an isomorphism between $\mathcal{I}_{\mathbf{p}'}$ and the product real line bundle on $S^2 - \Sigma$.)

Now comes a crucial observation: The preceding isomorphism between $\mathcal{I}_{\mathbf{q}}$ and $\mathcal{I}_{\mathbf{p}}$ on $S^2 - (\Sigma \cup \mathbf{q})$ defines two isomorphisms along Σ which differ via multiplication by -1 . These are obtained by following the isomorphism defined in the preceding paragraph along a path from x to Σ that starts at x and ends on Σ (transversely to Σ) whose interior is disjoint from Σ . In particular, a choice of normal vector to Σ 's arcs defines one of these two isomorphisms between $\mathcal{I}_{\mathbf{p}}$ and $\mathcal{I}_{\mathbf{q}}$ along Σ : Take a path from x that hits Σ at its endpoint from the side of the chosen normal vector. Choose this normal vector to define the isomorphism along Σ . With regards to notation, Lemma 6.6 has ∂_n denoting the directional derivative on Σ in the chosen normal direction.

Lemma 6.6. *Suppose that $\mathbf{q}, \mathbf{p}$ and $f_{\mathbf{q}}, f_{\mathbf{p}}$, and $E_{\mathbf{q}}, E_{\mathbf{p}}$, and also the isomorphism between $\mathcal{I}_{\mathbf{p}}$ and $\mathcal{I}_{\mathbf{q}}$ are as described above. Then there is the following identity:*

$$(E_{\mathbf{p}} - E_{\mathbf{q}}) \int_{S^2 - (\mathbf{q} \cup \Sigma)} f_{\mathbf{q}} f_{\mathbf{p}} = -2 \int_{\Sigma} (f_{\mathbf{q}} \partial_n f_{\mathbf{p}} - f_{\mathbf{p}} \partial_n f_{\mathbf{q}}).$$

Proof of Lemma 6.6. Multiply the equation $E_{\mathbf{p}} f_{\mathbf{p}} = -\Delta f_{\mathbf{p}}$ by $f_{\mathbf{q}}$ and likewise multiply the equation $E_{\mathbf{q}} f_{\mathbf{q}} = -\Delta f_{\mathbf{q}}$ by $f_{\mathbf{p}}$ and then subtract the second from the first. Having done that, integrate both sides of the result over $S^2 - (\Sigma \cup \mathbf{q})$ and then integrate by parts. q.e.d.

This lemma is most useful if Σ is part of the zero locus of $f_{\mathbf{p}}$ in which case there is only the $f_{\mathbf{q}} \partial_n f_{\mathbf{p}}$ term in the lemma's right hand side integral whose sign on any component of Σ can be determined by the sign of $f_{\mathbf{q}} f_{\mathbf{p}}$ on the chosen normal's side of Σ . (If this sign is $+$, then the integrand is $-$ and vice versa.) For example, if $j = n - 1$ so that there are just two points in $\mathbf{p}'$ which are very close to each other (and one arc in Σ), and if $f_{\mathbf{q}} f_{\mathbf{p}}$ is positive very near to Σ on the normal direction side of Σ then the lemma's right hand side integral is negative and so the right hand side in the lemma's identity is positive.

The preceding observation implies the claim made in the first sentence of the third bullet of (1.5) to the effect that the supremum of the function $E_{(\cdot)}$ on $\mathcal{C}_{2n}$ is always achieved at some configuration on $\mathcal{C}_{2n}$. Use an induction argument with the integer n starting from $n = 1$ as follows: Supposing that the supremum of $E_{(\cdot)}$ is taken at a configuration in $\mathcal{C}_{2n}$. Take $\mathfrak{q}$ in Lemma 6.6 to be such a configuration and take $f_{\mathfrak{q}}$ to be a $\mathcal{I}_{\mathfrak{q}}$ -eigensection with $\mathcal{I}_{\mathfrak{q}}$ -eigenvalue equal to $E_{\mathfrak{q}}$. Now take $\mathfrak{p} \in \mathcal{C}_{2n+2}$ as follows: Let x denote a point in S^2 where $f_{\mathfrak{q}} \neq 0$ and let p_1 and p_2 denote two points very close to x , equidistant from x on opposite sides of a great circle arc through x . Let r denote the distance between them. Set $\mathfrak{p} = \{p_1, p_2\} \cup \mathfrak{q}$ and let $f_{\mathfrak{p}}$ denote the $\mathcal{I}_{\mathfrak{p}}$ -eigensection with eigenvalue $E_{\mathfrak{p}}$. If r is sufficiently small, then there is a component of the zero locus of $f_{\mathfrak{p}}$ in a small radius ball centered on x that connects p_1 to p_2 . (This can be proved using Lemma 5.2.) Take Σ to be such a component. As explained in the preceding paragraph, the formula in Lemma 6.6 implies in this case that $E_{\mathfrak{p}}$ is greater than $E_{\mathfrak{q}}$.

To elaborate on what is going on here: The set up in the preceding paragraph is designed so that $|f_{\mathfrak{p}}|$ is very close to $|f_{\mathfrak{q}}|$ except in a small radius disk containing the points p_1 and p_2 . Thus, over most of $S^2 - (\mathfrak{q} \cup \{p_1, p_2\})$, the product $f_{\mathfrak{p}}f_{\mathfrak{q}}$ will be positive except very near where either $f_{\mathfrak{q}}$ is zero or $f_{\mathfrak{p}}$ is zero with it understood that the identification between $\mathcal{I}_{\mathfrak{p}} \otimes \mathcal{I}_{\mathfrak{q}}$ and the product line bundle on $S^2 \cup (\mathfrak{q} \cup \{p_1, p_2\})$ is obtained from the identification at x by following paths from x that don't intersect $\mathfrak{q} \cup \Sigma$. In particular, the proportion of $S^2 - (\mathfrak{q} \cup \Sigma)$ where $f_{\mathfrak{p}}f_{\mathfrak{q}} > 0$ approaches 1 as the distance from p_1 to p_2 shrinks to zero. This is why the integral of $f_{\mathfrak{p}}f_{\mathfrak{q}}$ is positive when the distance from p_1 to p_2 is sufficiently small.

7. The case of $\mathcal{C}_2$

The case where $n = 2$ already illustrates much of what is described in the preceding sections. The subsections that follow have some observations on this case.

7.1. Identifying $\mathcal{C}_2$ and $\bar{\mathcal{C}}_2$. The space $\mathcal{C}_2$ can be identified with the unit disk bundle in the tangent bundle to $\mathbb{R}P^2$. To make this explicit, first view S^2 as the radius 1 sphere about the origin in $\mathbb{R}^3$, and then identify the unit disk bundle in its tangent bundle (which is denoted by $T_{<1}S^2$) as

$$(7.1) \quad T_{<1}S^2 = \{(n, v) \in \mathbb{R}^3 \times \mathbb{R}^3 : |n| = 1, |v| < 1 \text{ and } \langle n, v \rangle = 0\}$$

The projection to S^2 sends any given pair (n, v) to n . There is an action of $\{\pm 1\}$ on $T_{<1}S^2$ that covers the $\{\pm 1\}$ action on S^2 whose quotient is $\mathcal{C}_2$, the quotient identification being

$$(7.2) \quad (n, v) \sim (-n, v).$$

The quotient, $\mathcal{C}_2$, is the unit disk bundle in an $\mathbb{R}^2$ bundle over $\mathbb{RP}^2$ that is isometric to the tangent bundle of $\mathbb{RP}^2$. Explicitly, the inverse identification sends the equivalence class $([n], v)$ to the pair of points

$$(7.3) \quad v \pm (1 - |v|^2)^{\frac{1}{2}} n$$

on the unit sphere about the origin in $\mathbb{R}^3$. (Here and below, when n is a unit length vector in $\mathbb{R}^3$, then $[n]$ denotes the corresponding point in $\mathbb{RP}^2$.) Note in particular that the $v = 0$ locus is the copy of $\mathbb{RP}^2$ given by those configurations where the constituent points are antipodal. With $\mathcal{C}_2$ viewed as the disk bundle in an $\mathbb{R}^2$ bundle over $\mathbb{RP}^2$, this $v = 0$ locus is the zero section of that vector bundle.

As for $\bar{\mathcal{C}}_2$, this is the Thom space for this $\mathbb{R}^2$ -bundle over $\mathbb{RP}^2$: It is the space of pairs $(n, v) \in \mathbb{R}^3 \times \mathbb{R}^3$ with $|n| = 1, |v| \leq 1$ and $\langle n, v \rangle = 0$ modulo the equivalence relation in (7.2) and the equivalence

$$(7.4) \quad (n, v) \sim (n', v') \text{ when } |v| = |v'| = 1.$$

Thus, a path in $\mathcal{C}_2$ where $|v| \rightarrow 1$ converges in $\bar{\mathcal{C}}_2$ to the stratum $\mathcal{C}_0$ in $\bar{\mathcal{C}}_2 = \mathcal{C}_2 \cup \mathcal{C}_0$.

In this case, the cohomology of $\mathcal{C}_2$ is that of $\mathbb{RP}^2$ which is $\mathbb{Z}/2\mathbb{Z}$ in degree 0, 1 and 2. Meanwhile, an instance of the Thom isomorphism theorem says that the cohomology of $\bar{\mathcal{C}}_2$ is $\mathbb{Z}/2\mathbb{Z}$ in degree 0, it is zero in degree 1, then $\mathbb{Z}/2\mathbb{Z}$ in degrees 2, 3 and 4. The degree 2 class is the obstruction class ω , and (according to the Thom isomorphism theorem), it's self cup product is non-zero.

7.2. The fiber bundle $\mathcal{F}$. The fiber bundle $\mathcal{F} \rightarrow \mathcal{C}_2$ consists of the set of pairs of the form $(\mathbf{p}, x)$ with $\mathbf{p} \in \mathcal{C}_2$ and with $x \in S^2 - \mathbf{p}$. The fiber of $\mathcal{F}$ over any given configuration is diffeomorphic to an open annulus. This annulus deformation retracts onto a circle bundle whose fiber over a given configuration $\mathbf{p} = ([n], v)$ is the set of unit length vectors in the plane orthogonal to the vector n (this is the great circle whose points are equidistant from the two points in $\mathbf{p}$). Viewed in this light, it is the pull-back from $\mathbb{RP}^2$ of the unit disk bundle in the tangent bundle to $\mathbb{RP}^2$.

The assertion in Section 4.1 to the effect that there is no real line bundle over $\mathcal{F}$ that restricts to the fiber of each configuration $\mathbf{p}$ from $\mathcal{C}_2$ as $\mathcal{I}_{\mathbf{p}}$ is equivalent to the assertion that there is no degree 1 class in the $\mathbb{Z}/2\mathbb{Z}$ cohomology of the unit circle bundle in $T\mathbb{RP}^2$ that restricts non-trivially to each fiber. In turn, the latter is equivalent to the assertion that any given fiber circle represents the zero class in the $\mathbb{Z}/2\mathbb{Z}$ first homology of the unit circle bundle in $T\mathbb{RP}^2$. That this is so can be seen by constructing a section of this bundle over the complement in $\mathbb{RP}^2$ of any given point in $\mathbb{RP}^2$ and then looking at how the section behaves near the omitted point (see below).

For example, let n_* denote a given unit length vector in $\mathbb{R}^3$ and let $[n_*]$ denote the corresponding point in $\mathbb{RP}^2$. A section over $\mathbb{RP}^2 - [n_*]$ is given by the rule

(7.5)
$$[n] \rightarrow \left([n], v = \frac{1}{\sqrt{1 - \langle n, n_* \rangle^2}} (n_* - n \langle n, n_* \rangle) \right).$$

If n is close to n_* , then n can be written as $\frac{1}{\sqrt{1+\epsilon^2}}(n_* + \epsilon u)$ where u has norm 1 and is orthogonal to n_* and where $\epsilon \in (0, 1)$. The section in (7.5) near $\epsilon = 0$ when written in terms of ϵ and u has the form

(7.6)
$$\left[\frac{1}{\sqrt{1 + \epsilon^2}} (n_* + \epsilon u) \right] \rightarrow ([n_*], v = u) + \mathcal{O}(\epsilon)$$

which directly illustrates that the fiber over n_* is a boundary in the unit tangent bundle. Indeed, the preceding construction depicts the fiber over $[n_*]$ as the boundary of an embedded Möbius band. Meanwhile, the boundary of a Möbius band is zero in the $\mathbb{Z}/2\mathbb{Z}$ homology of the Möbius band since it is twice the class of the central circle in the band’s $\mathbb{Z}$ -homology.

7.3. The space $\overline{\mathbb{RP}}$ and the function $\mathcal{E}$. As explained in the appendix of [11], the only critical points of $\mathcal{E}$ on the $\mathcal{C}_2$ version of $\mathbb{RP}$ are over the configurations of antipodal points in $\mathcal{C}_2$. In this case, the set of $\mathcal{I}_p$ –eigenvalues is the set $\{m^2 - \frac{1}{4}\}_{m=1,2,\dots}$; any given integer m eigenvalue has multiplicity $2m$. The eigenspace with eigenvalue $m^2 - \frac{1}{4}$ for a given antipodal configuration is spanned by eigensections that can be described as follows: Take spherical coordinates for S^2 with the longitude being rotation about the axis through the pair of points from the configuration. With these coordinates denoted by (θ, φ) , then the eigenspace consists of the linear combinations of the real and imaginary parts of the eigensection that is depicted below in (7.6) and then the real and imaginary parts of what is obtained from that eigensection by acting by the generators of the $SO(3)$ action no more than $m-1$ times.

(7.7)
$$f_m = (\sin \theta)^{m-\frac{1}{2}} e^{(m-\frac{1}{2})\varphi}$$

Note that the minimal eigenvalue ($m = 1$) has multiplicity two and it does not obey the criteria to be a critical point of $\mathcal{E}$. The next highest eigenvalue is a critical value (the $m = 2$ case in (7.7)). But note also that there is a dimension 4 eigenspace with this value of $\mathcal{E}$. All of the $m \geq 2$ eigenvalues are also critical values of $\mathcal{E}$.

The compactification $\overline{\mathcal{C}}_2$ adds the one point $\mathcal{C}_0$ to $\mathcal{C}_2$. The spectrum of the Laplacian over the $\mathcal{C}_0$ stratum is the set $\{m(m+1)\}_{m=0,1,\dots}$ with any given $m(m+1)$ eigenvalue having multiplicity $2m+1$. This being the case, $\mathcal{I}_{(\cdot)}$ –eigenvalues for configurations in $\mathcal{C}_2$ with the two points very close will be very nearly of the $m(m+1)$ form with m being a

non-negative integer with its values ranging from zero to some positive integer that gets ever larger as the points come together.

To elaborate: If the two points are moved from being almost antipodal to being almost coincident along a great circle arc, then the eigenvalues from the antipodal set $\{m^2 - \frac{1}{4}\}_{m=1,2,\dots}$ must move in a continuous fashion to generate the coincident set $\{m(m+1)\}_{m=0,1,\dots}$ and likewise for the eigensections. Keeping in mind that the antipodal eigenvalues are alternating with the coincident eigenvalues on the non-negative real axis, here is how this works: When the points are moved together along a great circle arc, m of $2m$ -linearly independent eigensections for a given antipodal eigenvalue $m^2 - \frac{1}{4}$ have decreasing eigenvalue which limit to $m(m-1)$, and m of them have increasing eigenvalue which limit to $m(m+1)$.

With regards to the preceding: The precise eigenspaces that give decreasing or increasing eigenvalues depends on the choice of the great circle arc that is used to bring the two points together. This can be seen most clearly for the $m = 1$ case of the antipodal configuration eigenspace (eigenvalue $\frac{3}{4}$) whose normalized eigensections have the form

$$(7.8) \quad z \sin^{\frac{1}{2}} \theta \sin \left(\frac{1}{2}(\varphi - \alpha) \right)$$

with $\alpha \in \mathbb{R}/2\pi\mathbb{Z}$ and with z being independent of α . Consider a path in $\mathcal{C}_2$ that starts with an antipodal configuration and then changes the configuration so as to bring the two points together along a constant φ half great circle of longitude (say $\varphi = \varphi_*$). It follows from Proposition 2.6 that the $\alpha = \varphi_*$ version of (7.6) deforms to a corresponding family of $\mathcal{I}_{(\cdot)}$ -eigensections along the configuration path with the corresponding path of eigenvalues decreasing to 0 as the points converge. Meanwhile, the $\alpha = \varphi_* + \pi$ version of (7.6) deforms to a family of $\mathcal{I}_{(\cdot)}$ -eigensections along the configuration path with the corresponding path of eigenvalues increasing to 2 along the configuration path (keep in mind that 2 is the $m = 1$ version of $m(m+1)$).

One can imagine that there is a topological explanation for the interesting spectral flow over $\mathcal{C}_2$ via an analysis of the effect on $\overline{\mathbb{RP}}$ of the rotation group's action on S^2 .

Appendix A. The cohomology of a weak $\mathbb{RP}^\infty$ bundle

This appendix states and then proves a proposition that describes the $\mathbb{Z}/2\mathbb{Z}$ cohomology of a weak $\mathbb{RP}^\infty$ bundle over a space X . An additional lemma at the end of the appendix (Section A.6) points out a useful corollary from the proof. (A proposition much like the one below might well be in the literature in which case we apologize to its authors for failing to find it and hence reference it.)

To set the stage for the proposition, suppose that X is a compact CW complex and that $\pi : \mathcal{R} \rightarrow X$ is a weak $\mathbb{R}\mathbb{P}^\infty$ bundle. (The assumption holds for the case $X = \overline{\mathcal{C}}_{2n}$.)

There is an obstruction class in $H^2(X)$ which is defined just as in Section 5.4. Let ω denote this obstruction class. Now introduce $\mathcal{A}^*$ to be the kernel in $H^*(X)$ of the homomorphism $\omega \wedge (\cdot)$ to $H^{*+2}(X)$, and let $\mathcal{B}^*$ denote the cokernel. By way of an example, let a denote the highest non-zero self-cup product of ω . This is to say that $\omega^a \neq 0$ but $\omega^{a+1} = 0$. Then ω^a is in $\mathcal{A}^{2a}$.

Proposition A.1. *The $\mathbb{Z}/2\mathbb{Z}$ cohomology of $\mathcal{R}$ is isomorphic (non-canonically) to the vector space of polynomials in a degree two class τ with coefficient ring $\mathcal{C}$ as described in the subsequent bullets. Thus,*

$$H^*(\mathcal{R}) \approx C^* \oplus (C^{*-2} \wedge \tau) \oplus (C^{*-4} \wedge \tau^2) \oplus \dots$$

- *The definition of τ : Let $\hat{\omega}$ denote a given representative cocycle for ω . Then $\pi^*\hat{\omega}$ can be written as $d\nu$ with ν being a 1-cochain whose restriction to each fiber algebraically generates the cohomology of the fiber. The cochain $\nu \wedge \nu$ is closed and non-zero in $H^*(\mathcal{R})$. The class of $\nu \wedge \nu$ is the class τ . This class can depend on the choice for $\hat{\omega}$ and ν but any such change has the form $\pi^*\lambda$ with λ a non-zero degree 2-class in $H^*(X)$ that can be represented as the square of a 1-cochain (as $\mu \wedge \mu$).*
- *The definition of $\mathcal{C}$: The degree $*$ summand of $\mathcal{C}$ is zero if $*$ is negative, it is isomorphic to $\mathbb{Z}/2\mathbb{Z}$ if $*$ = 0, and for positive degrees, it is isomorphic (non-canonically) to $\mathcal{B}^* \oplus \mathcal{A}^{*-1}$ with an isomorphism defined in the three parts that follow:*
 - *Fix a basis for $\mathcal{B}^*$ and then a lift of each basis element as a closed cocycle that represents a class in $H^*(X)$ which projects back to the basis element (the lift for a given basis element B is denoted $\hat{B}$).*
 - *Fix a basis for $\mathcal{A}^{*-1}$, and for each basis element, fix a cocycle that represents that element (the cocycle for the basis element A is denoted by $\hat{A}$. Having fixed $\hat{A}$, then fix degree $*$ cochain $\hat{i}_A$ with the property that $d\hat{i}_A = \hat{\omega} \wedge \hat{A}$.*
 - *The isomorphism $C^* \approx \mathcal{B}^* \oplus \mathcal{A}^{*-1}$ sends any given pair of basis elements B and A to the class in $H^*(\mathcal{R})$ of the cocycle $\pi^*\hat{B} + \hat{\nu} \wedge \pi^*(\hat{A}) + \pi^*\hat{i}_A$.*

The proof of this proposition is given in Section A.5 of this appendix. The proof employs the auxiliary $\mathbb{R}\mathbb{P}^{2m+1}$ bundles over X that are constructed in the next subsection. As explained in Section A.6 of this appendix, these finite dimensional fiber bundles can also be used to represent the homology of $\mathcal{R}$.

A.1. The $\mathbb{RP}^{2m+1}$ bundles. Fix a non-negative integer m . The two parts of this subsection use the obstruction cochain to construct a useful $\mathbb{RP}^{2m+1}$ bundle over X . This $\mathbb{RP}^{2m+1}$ bundle is denoted by $\mathcal{Y}_m$.

Part 1. The class ω on $\bar{\mathcal{C}}_{2n}$ can be represented by a Čech 2-cocycle as defined with respect to a finite cover. Let $\mathcal{U}$ denote such a cover. The corresponding (multiplicative) cocycle assigns either -1 or 1 to each component of each mutual intersection of three sets from $\mathcal{U}$. If U'' , U' and U are three sets from $\mathcal{U}$ with points in common, then the associated number (-1 or 1) is denoted by $z_{U''U'U}$.

Now let $\{\varphi_U\}_{U \in \mathcal{U}}$ denote a partition of unity subordinate to the cover $\mathcal{U}$. This can be used to produce maps from pairwise intersections of sets from $\mathcal{U}$ to the Lie group $SO(2m+2)$ as follows: Supposing that U and U' are sets from $\mathcal{U}$ that share points, set

$$(A.1) \quad w_{U'U} = \frac{1}{2} \sum_{U'' \in \mathcal{U}} \varphi_{U''} (1 - z_{U''U'U}),$$

which is a map from $U \cap U'$ to $[0, 1]$. The latter map is used to construct a map from $U \cap U'$ to the group $SO(2)$:

$$(A.2) \quad A_{U,U'} \equiv \begin{pmatrix} \cos(\pi w_{U'U}) & -\sin(\pi w_{U'U}) \\ \sin(\pi w_{U'U}) & \cos(\pi w_{U'U}) \end{pmatrix}.$$

This is the desired map in the case when $m = 0$. If $m > 0$, the desired map is the tridiagonal, $(2m+2) \times (2m+2)$ matrix that is depicted below in (A.3) as a diagonal $(m+1) \times (m+1)$ matrix whose entries are 2×2 matrices with zero off of the diagonal and $A_{U'U}$ on the diagonal:

$$(A.3) \quad \begin{pmatrix} A_{U'U} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & A_{U'U} \end{pmatrix}$$

The latter matrix is denoted also by $A_{U'U}$. The important point now is that if U'' , U' and U have points in common, then

$$(A.4) \quad A_{U''U'} A_{U'U} A_{UU'} = Z_{U''U'U} \mathbb{I},$$

where $\mathbb{I}$ denotes the $(2m+2) \times (2m+2)$ identity matrix.

Part 2. The plan for this part of the proof is to use the constructions from Part 1 to build an $\mathbb{RP}^{2m+1}$ fiber bundle over $\bar{\mathcal{C}}_{2n}$ that has two key properties: It is not the $\{1, -1\}$ quotient of a sphere bundle in a $2m+2$ dimensional vector bundle; and the corresponding 2-dimensional obstruction class is the class ω . (The statement that this $\mathbb{RP}^{2m+1}$ bundle is not the quotient of a sphere bundle is saying in effect that there is no class in the first cohomology of the total space that restricts to the fiber over any given point in X as the generator of first cohomology of that fiber. Equivalently: There is no real line bundle over the total space

that restricts to any given fiber as the one non-trivial line bundle over the fiber.)

The desired fiber bundle is obtained from the disjoint union of the sets $\{U \times \mathbb{RP}^{2m+1}\}_{U \in \mathcal{U}}$ by the equivalence relation given momentarily. To set the notation: Depict $\mathbb{RP}^{2m+1}$ as the quotient of the radius 1 sphere centered at the origin in $\mathbb{R}^{2m+2}$ by the action of multiplication by $\{1, -1\}$. When a point in this sphere is denoted by x , then its image in $\mathbb{RP}^{2m+1}$ is denoted by $[x]$. Now for the equivalence relation: The relation identifies a point $(\mathbf{p}, [x])$ in the space $U \times \mathbb{RP}^{2m+1}$ with a point $(\mathbf{p}', [x'])$ in $U' \times \mathbb{RP}^{2m+1}$ if and only if $\mathbf{p} = \mathbf{p}'$ and $[x] = [A_{U'U}x']$. This is a valid equivalence relation (and hence defines a topological space) because each instance of $z_{U''U'U}$ in (A.4) is either 1 or -1 .

It also follows from (A.4) that the bundle $\mathcal{Y}_m$ cannot be lifted to a $(2m+1)$ dimensional sphere bundle because (A.4) identifies its obstruction class in $H^2(X)$ as the class ω .

A.2. The pull-back of the obstruction class ω to $\mathcal{Y}_m$. Let π_m denote the projection map from $\mathcal{Y}_m$ to X . Fix a 2-cocycle to be denoted by $\hat{\omega}$ on X that gives the class ω . Part 1 of this section proves that $\pi_m^* \hat{\omega}$ is exact. Part 2 proves that if $\hat{\nu}$ is a 1-cochain with $d\hat{\nu} = \pi^*$, then $\hat{\nu}$ must restrict to each fiber of π so as to algebraically generate the cohomology of the fiber (which is $\mathbb{RP}^{2m+1}$).

Part 1. To prove that $\pi_m^* \hat{\omega}$ is exact: View $\pi_m^* \mathcal{Y}_0$ as an $\mathbb{RP}^1$ bundle over $\mathcal{Y}_m$. Its obstruction class is $\pi_m^* \omega$. On the other hand, $\pi_m^* \mathcal{Y}_0$ can be viewed as the fiber product $\mathcal{Y}_m \times_X \mathcal{Y}_0$ with the convention being that the left hand $\mathcal{Y}_m$ is the ‘base’ manifold for the fiber bundle $\pi_m^* \mathcal{Y}_0$. Let π_L denote the projection to this left hand $\mathcal{Y}_m$ and let π_R denote the projection to the right hand $\mathcal{Y}_0$. Viewed as an $\mathbb{RP}^{2m+1} \times \mathbb{RP}^1$ bundle over X , there are two line bundles over any given fiber, $\mathcal{I}_L$ and $\mathcal{I}_R$ with the former being the tautological line bundle over $\mathbb{RP}^{2m+1}$ and the latter being the tautological line bundle over $\mathbb{RP}^1$. The obstruction cocycle for the tensor product bundle $\mathcal{I}_L \otimes \mathcal{I}_R$ to extend over $\mathcal{Y}_m \times_X \mathcal{Y}_0$ is (multiplicatively) given on triple intersections of sets U'', U' and U from the cover $\mathcal{U}$ by the Čech 2-cocycle $z_{U''U'U}^2$, which is 1. As a consequence of this, there is a line bundle (denoted by $\mathcal{I}_{LR}$) on $\mathcal{Y}_m \times_X \mathcal{Y}_0$ that restricts to each fiber over X as $\mathcal{I}_L \otimes \mathcal{I}_R$. With $\mathcal{Y}_m \times_X \mathcal{Y}_0$ now viewed as $\pi_m^* \mathcal{Y}_0$, this line bundle $\mathcal{I}_{L,R}$ restricts to each fiber $\mathbb{RP}^1$ of the projection π_L to $\mathcal{Y}_m$ as the tautological line bundle over the fiber. Since the obstruction for this is $\pi_m^* \omega$, the latter class is zero in $H^2(\mathcal{Y}_m)$. This implies that $\pi_m^* \hat{\omega}$ can be written as $d\hat{\nu}$.

Part 2. This part of the proof explains why $\hat{\nu}$ restricts to each fiber $\mathbb{RP}^{2m+1}$ as the generator of the first cohomology of the fiber. To do this, consider first the $m = 0$ case. This case is simpler because $\mathbb{RP}^1$ is diffeomorphic to S^1 and thus $\mathcal{Y}_0$ is a circle bundle over X . As such,

there is the Gysin sequence (see e.g. Hatcher's book [3]):

(A.5)

$$0 \rightarrow H^1(X) \xrightarrow{\pi_0^*} H^1(\mathcal{Y}_0) \xrightarrow{\iota_0} H^0(X) \xrightarrow{\hat{e} \wedge (\cdot)} H^2(X) \xrightarrow{\pi_0^*} H^2(\mathcal{Y}_0) \rightarrow \dots$$

where $\hat{e}$ denotes the $\mathbb{Z}/2\mathbb{Z}$ Euler class and $\hat{e} \wedge (\cdot)$ denotes the cup product with this class. According to (A.5), the homomorphism $\iota_0 : H^1(\mathcal{Y}_0) \rightarrow H^0(X)$ is surjective if and only if there is a class in $H^1(\mathcal{Y}_0)$ whose pairing with the generator of the first homology of each fiber is non-zero. In this case, there is no such generator because the obstruction class ω is zero. Therefore, $\hat{e} \wedge (\cdot)$ maps $H^0(X)$ isomorphically to $H^2(X)$ and the class $\hat{e}$ is the unique class in $H^2(X)$ whose pull-back to $\mathcal{Y}_0$ is zero. Since the pull-back of the class ω to $\mathcal{Y}_0$ is zero, it follows that ω must be the Euler class $\hat{e}$.

With $\hat{e}$ understood to be ω , fix $x \in X$ and consider the following diagram:

(A.6)

$$\begin{array}{ccccccc} H^1(\mathcal{Y}_0|_x) & \longrightarrow & H^2(\mathcal{Y}_0, \mathcal{Y}_0|_x) & \longrightarrow & H^2(\mathcal{Y}_0) & \longrightarrow & \\ \uparrow & & \uparrow & & \uparrow & & \\ 0 & \longrightarrow & H^2(X, x) & \longrightarrow & H^2(X) & \longrightarrow & 0 \\ \uparrow & & \uparrow & & \uparrow & & \\ 0 & \longrightarrow & H^0(X, x) & \longrightarrow & H^0(X) & \longrightarrow & H^0(x) \end{array}$$

The top row is from the exact sequence of the pair $(\mathcal{Y}_0, \mathcal{Y}_0|_x)$ and the second and third rows are from the exact sequence of the pair (X, x) . The vertical arrows are the Gysin sequence homomorphisms; the top row of vertical arrows are π_0^* ; and the lower vertical arrows are cup product with the Euler class (which is ω in our case) or the relative Euler class as the case may be.

Now observe that the middle row says that ω which is in $H^2(X)$ comes from a unique class in $H^2(X, x)$ which will be denoted by ω_x . Since $\pi_0^* \omega = 0$, it follows that $\pi_0^* \omega_x$ is in the image of the connecting homomorphism from $H^1(\mathcal{Y}_0|_x)$. In particular, this implies the claim about $\hat{\nu}$'s restriction to the fiber if it is the case that $\pi_0^* \omega_x$ is non-zero. But the fact that $\pi_0^* \omega_x$ is not zero follows from the fact that $H^0(X, x)$ is zero since the vertical arrows make for an exact sequence.

The assertion about $\hat{\nu}$ for the case when $m > 0$ follows directly from the $m = 0$ case because of the block diagonal form of the matrix in (A.3). Indeed, by virtue of (A.3) being block diagonal, the fiber bundle $\mathcal{Y}_0$ sits as a subbundle in $\mathcal{Y}_m$ whose fiber over any given point is a standard $\mathbb{RP}^1$ in the $\mathbb{RP}^{2m+1}$ fiber of $\mathcal{Y}_m$.

A.3. The cohomology of $\mathcal{Y}_m$. This section first states and then proves a lemma that describes the $\mathbb{Z}/2\mathbb{Z}$ cohomology of $\mathcal{Y}_m$. The notation uses

$\mathcal{A}^*$ to denote the kernel in $H^*(X)$ of the homomorphism $\omega \wedge$ and $\mathcal{B}^*$ denotes the cokernel of this homomorphism. The lemma that follows uses $\mathcal{C}^*$ to denote $\mathcal{B}^* \oplus \mathcal{A}^{*-1}$. A convention in what follows is that all modules with negative degrees are equal to zero.

Lemma A.2. *If $m = 0$: There is a non-canonical isomorphism between $H^*(\mathcal{Y}_0)$ and $\mathcal{C}^*$. If $m \geq 1$: There is a class $\tau \in H^2(\mathcal{Y}_m)$ with τ^{2k} non-zero when $k \in \{1, \dots, m\}$ (its pairing is 1 with the generator of the $2k$ 'th homology of any given fiber); and there is a (non-canonical) direct sum decomposition*

$$H^*(\mathcal{Y}_m) \approx \mathcal{C}^* \oplus (C^{*-2} \wedge \tau) \oplus (C^{*-4} \wedge \tau^2) \oplus \dots$$

with it understood that $\mathcal{C}^* \equiv 0$ when $*$ is negative, that $\mathcal{C}^0 = \mathcal{B}^0$ and that the highest power of τ that can appear is τ^m .

Proof of Lemma A.2. The proof has four parts.

Part 1. The assertion of the lemma for the case of $\mathcal{Y}_0$ follows from the Gysin sequence in (A.5) which leads to the exact sequence

$$(A.7) \quad 0 \rightarrow \mathcal{B}^* \rightarrow H^*(\mathcal{Y}_0) \rightarrow \mathcal{A}^{*-1} \rightarrow 0.$$

This sequence splits because the coefficients are $\mathbb{Z}/2\mathbb{Z}$ and $H^*(\mathcal{Y}_0) \approx \mathcal{B}^* \oplus \mathcal{A}^{*-1}$, but the splitting is not canonical. Here is one isomorphism: Fix a cocycle $\hat{\omega}$ to represent ω and then fix $\hat{\nu}$ with $d\hat{\nu} = \pi_0^* \hat{\omega}$. Now fix a basis for $\mathcal{B}^*$ and then a lift of each basis element as a closed cocycle that represents a class in $H^*(X)$ which projects back to the basis element (the lift for a given basis element B is denoted $\hat{B}$). Meanwhile, fix a basis for $\mathcal{A}^{*-1}$, and for each basis element, fix a cocycle that represents that element (the cocycle for the basis element A is denoted by $\hat{A}$). Having fixed $\hat{A}$, then fix degree $*$ cochains $\{\hat{i}_A\}$ with the property that $d\hat{i}_A = \hat{\omega} \wedge \hat{A}$. An isomorphism to split (A.7) sends any given pair of basis elements B and A to the classes in $H^*(\mathcal{Y}_0)$ of the respective cocycles

$$(A.8) \quad (B, A) \rightarrow \pi_0^* \hat{B} + \hat{\nu} \wedge \pi_0^* (\hat{A}) + \pi_0^* \hat{i}_A.$$

By way of a parenthetical remark: The splitting in (A.7) is non-canonical because $\hat{i}_A$ can be changed to $\hat{i}_A + \hat{C}$ with $\hat{C}$ any closed cochain. If $\hat{C}$ is not exact, and if it projects to a non-zero class in $\mathcal{B}^*$, then this change in $\hat{i}_A$ is accounted for by suitable changes to the basis for $\mathcal{B}^*$. In addition, the cochain $\hat{\nu}$ can be changed by adding $\pi_0^* \hat{C}$ with $\hat{C}$ being a closed degree 1 cochain. If $\hat{C}$ is not exact and if $\hat{C} \wedge \hat{A}$ represents a non-zero element in $\mathcal{B}^*$, then this change is also accommodated by changing the basis $\mathcal{B}^*$.

By way of a relevant example: Let a denote the positive integer such that ω^a is non-zero and ω^{a+1} is zero. Fix a cocycle representative for ω to be denoted by $\hat{\omega}$ and then fix a degree $2a+1$ cochain $\hat{i}$ such that $d\hat{i} = \hat{\omega}^{a+1}$. The cocycle $\hat{\nu} \wedge \pi^* \hat{\omega}^a + \pi^* \hat{i}$ represents a non-zero class in

$H^{2a+1}(\mathcal{Y}_0)$. Note in this regard that when $a = \frac{1}{2} \dim(X)$, then the resulting class in $H^{\dim(X)+1}(\mathcal{Y}_0)$ is independent of the choice for the representative cocycles, $\hat{\omega}$, $\hat{\nu}$ and $\hat{i}$ because $H^*(X) = 0$ when $*$ is greater than $\dim(X)$.

By way of a second example: The degree 2 cochain $\hat{\nu} \wedge \hat{\nu}$ is closed; so what is this class? In general, the cohomology class of $\hat{\nu} \wedge \hat{\nu}$ has the form

$$(A.9) \quad \pi_0^* \hat{B}_0 + \hat{\nu} \wedge \pi_0^* \hat{A}_0 + \pi_0^* \hat{i}_{A_0}$$

where $\hat{B}_0$ is a closed form that represents a non-zero class in $\mathcal{B}^2$, where $\hat{A}_0$ represents a degree 1 class on X that is annihilated by wedge product with ω ; and where $\hat{i}_{A_0}$ is a degree 2-cochain with $d\hat{i}_{A_0} = \hat{\omega} \wedge \hat{A}_0$.

Part 2. To prove the lemma when $m > 0$, some preliminary observations are needed. The first is this: If A denotes a non-zero class from $\mathcal{A}^{*-1}$ (which is the kernel in $H^{*-1}(X)$ of cup product with ω) and $\hat{A}$ is a representative cocycle and $\hat{i}_A$ is a cochain that obeys $d\hat{i}_A = \hat{\omega} \wedge \hat{A}$, then $\hat{\nu}_R \wedge \pi_m^* \hat{A} + \pi_m^* \hat{i}_A$ is a closed cocycle on $\mathcal{Y}_m$ that represents a non-zero class; it is non-zero because it is already non-zero when restricted to the sub-fiber bundle $\mathcal{Y}_0$. For the same reason: If B is from $\mathcal{B}^*$ and $\hat{B}$ is a representative cocycle, then $\pi_m^* \hat{B}$ represents a non-zero class in $H^*(\mathcal{Y}_m)$. The second observation is that the preceding constructions define a summand in $H^*(\mathcal{Y}_m)$ that is isomorphic to $\mathcal{B}^* \oplus \mathcal{A}^{*-1}$ (although not canonically). This summand is denoted by $\mathcal{C}^*$.

Here is a third observation: If k is any integer from 1 through m , then $\hat{\nu}_R^{2k}$ is a closed cocycle which is non-zero in $H^{2k}(\mathcal{Y}_m)$. It is closed because 2 equals 0 in $\mathbb{Z}/2\mathbb{Z}$; and it is a non-zero class because it has non-zero pairing with the non-zero class in the degree $2k$ homology of any given fiber of the projection to X . This class is not in $\mathcal{C}^*$. Let τ henceforth denote the class that is represented by $\hat{\nu}_R^2$ in $H^2(\mathcal{Y}_0)$.

Part 3. To prove the lemma when $m > 0$, return to the S^1 bundle $\pi_m^* : \mathcal{Y}_0 \rightarrow \mathcal{Y}_m$. Write $\pi_m^* \hat{\omega}$ as $d\hat{\nu}_m$ and write $\pi_0^* \hat{\omega}$ as $d\hat{\nu}_0$. Then write $\pi_L^* \hat{\nu}_m$ as $\hat{\nu}_L$ and write $\pi_R^* \hat{\nu}_0$ as $\hat{\nu}_R$. Writing cocycles additively now, note that $d(\hat{\nu}_L - \hat{\nu}_R) = 0$ because $\pi_L^* \pi_m^* \hat{\omega}$ is the same cocycle as $\pi_R^* \pi_0^* \hat{\omega}$. This implies that the 1-cochain $\hat{\nu}_L - \hat{\nu}_R$ is a closed cochain on $\mathcal{Y}_m \times_X \mathcal{Y}_0$. In this regard: The cohomology class of $\hat{\nu}_L - \hat{\nu}_R$ represents the first Stiefel-Whitney class of the line bundle $\mathcal{I}_{L,R}$ from Part 1 of Section A.2. (It follows from what is said in Part 2 of Section A.2 about the $\hat{\nu}_m$ and $\hat{\nu}_0$ versions of what is denoted there by $\hat{\nu}$ that $\hat{u}$ restricts to $\mathbb{RP}^{2m+1} \times \mathbb{RP}^1$ fibers of the projection map from $\mathcal{Y}_m \times_X \mathcal{Y}_0$ to X as the first Stiefel-Whitney class of the tensor product of the respective $\mathbb{RP}^{2m+1}$ and $\mathbb{RP}^1$ tautological line bundles.)

The cochain $\hat{\nu}_L - \hat{\nu}_R$ is denoted by $\hat{u}$ and u is used to denote its cohomology class.

Part 4. The key observation is that $\mathcal{Y}_m \times_{\mathbb{R}} \mathcal{Y}_0$ can be viewed on the one hand as an $\mathbb{R}P^{2m+1}$ bundle over $\mathcal{Y}_0$ and on the other, as a $\mathbb{R}P^1$ bundle of $\mathcal{Y}_m$. In either case, the pull-backs of the powers of the class $\hat{u}$ generate the cohomology of the fiber. As a consequence, the Leray-Hirsch theorem (see [3]) can be used to depict the cohomology of $\mathcal{Y}_m \times_X \mathcal{Y}_0$ in two ways which is what is done in (A.10) below.

(A.10)

- $H^*(\mathcal{Y}_m \times_X \mathcal{Y}_0) \approx H^*(\mathcal{Y}_0) \oplus H^{*-1}(\mathcal{Y}_0) \oplus H^{*-2}(\mathcal{Y}_0) \oplus \dots \oplus H^{*-2m-1}.$
- $H^*(\mathcal{Y}_m \times_X \mathcal{Y}_0) \approx H^*(\mathcal{Y}_m) \oplus H^{*-1}(\mathcal{Y}_m).$

(The notation is such that modules with negative degree are equal to 0.) These isomorphism work as follows: In the case of the top bullet, a class $C_k \in H^{*-k}(\mathcal{Y}_0)$ is sent to the class $\pi_R^* C_k \wedge u^k$; and in the case of the lower bullet, a class $C_0 \in H^*(\mathcal{Y}_0)$ or $C_1 \in H^{*-1}(\mathcal{Y}_m)$ is sent to $\pi_L^* C_0$ or $\pi_L^* C_1 \wedge u$ as the case may be.

To exploit the two direct sum representations in (A.10), note first that

(A.11)
$$\hat{u}^2 = \pi^* \hat{B}_0 + \hat{\nu}_L \wedge \pi^* \hat{A}_0 + \hat{\nu}_R^2$$

where π denotes the projection from $\mathcal{Y}_m \times_X \mathcal{Y}_0$ to X , and where $\hat{B}_0$ and $\hat{A}_0$ are as depicted in (A.9). Note that the identity in (A.11) can also be written as

(A.12)
$$\hat{u}^2 = \pi^* \hat{B}_0 + \hat{\nu}_R \wedge \pi^* \hat{A}_0 + \hat{\nu}_R^2 + \hat{u} \wedge \pi^* \hat{A}_0$$

The preceding identity can be used to write the summands that appear in the top bullet of (A.10) using the direct sum depiction in the lower bullet of (A.10) and vice-versa. Doing that (and remembering what is said by Part 2) leads in a direct line to Lemma A.2’s assertion.
q.e.d.

A.4. Fiber preserving maps from $\mathcal{Y}_m$ to $\mathcal{R}$. The purpose of this subsection is to construct a fiber preserving map from $\mathcal{Y}_m$ to $\mathcal{R}$ with certain desirable properties. The lemma that follows summarizes. (The lemma refers to notions from (5.5).)

Lemma A.3. *Fix a non-negative integer for m . There exists a fiber preserving, fiber-wise embedding of $\mathcal{Y}_m$ into $\mathcal{R}$ with the following property: If k is a sufficiently large integer, then this embedding over any given open set U from the cover $\mathfrak{U}_k$ maps $\mathcal{Y}_m|_U$ into the Ψ_U -image of $(V_U - 0)/\mathbb{R}^*$. Moreover, the latter map is the Ψ_U image of a fiber-wise fiber bundle embedding from $\mathcal{Y}_m|_U$ to $(V_U - 0)/\mathbb{R}^*$ that identifies $\mathcal{Y}_m|_U$ with the $\mathbb{R}^*$ -quotient of the complement of the zero section in an m -dimensional sub-vector bundle in V_U .*

The rest of this subsection has the proof of this lemma. It is used in the next subsection to prove Proposition A.1.

Proof of Lemma A.3. There are two parts to the proof

Part 1. Start by fixing a large integer (to be denoted by k) that is in any event greater than m . Let $\mathfrak{U}_k$ denote the corresponding open cover of X that is described in (5.5). Supposing that U is a set from $\mathfrak{U}_k$, the first observation is that the obstruction cocycle ω is zero upon restriction to $H^2(U)$. This is because (5.5)'s embedding Ψ_U induces an isomorphism on the first homology of the fibers of $\mathcal{R}$ over U . With this understood, it follows that $\mathcal{Y}_m|_U$ can also be written as the $\mathbb{R}^*$ quotient of the complement of the zero section in a $2m + 2$ dimensional vector bundle over U . Denote that bundle by V_{mU} . An instance of the upcoming Lemma A.4 implies this: If k is sufficiently large (a lower bound is determined by m and X), then there is an injective vector bundle homomorphism $\varphi_{mU} : V_{mU} \rightarrow V_U$. Such a map induces then a fiberwise embedding of $\mathcal{Y}_m|_U$ into $(V_{mU} - 0)/\mathbb{R}^*$ which composes with Ψ_U to give a fiberwise injective embedding of $\mathcal{Y}_m|_U$ into $\mathcal{R}|_U$.

Lemma A.4. *Fix a positive integer to be denoted by m . There exists a positive integer (denoted by k) with the following significance: If $V \rightarrow \mathcal{X}$ is a vector bundle with fiber dimension m and if $E \rightarrow \mathcal{X}$ is a vector bundle with fiber dimension k , then there is an injective vector bundle homomorphism from V to E . Moreover, any two injective vector bundle homomorphisms are homotopic through injective bundle homomorphism.*

Proof of Lemma A.4. Let $\text{HOM}_{m,k}$ denote the vector space of linear maps from $\mathbb{R}^m$ to $\mathbb{R}^k$ (the space of matrices with m rows and k columns, thus $\mathbb{R}^{mk}$). Inside $\text{HOM}_{m,k}$ sits the subspace of injective maps which is denoted by $\text{INJ}_{m,k}$. The complement is a closed subvariety whose codimension is $k - m - 1$. It follows as a consequence that the homotopy groups of $\text{INJ}_{m,k}$ vanish in dimensions less than $k - m$.

With the preceding as background, the vector bundle $\text{HOM}(V_{mU}, V_U)$ is a vector bundle over U whose fiber is $\text{HOM}_{m,k}$; and inside this bundle sits the fiber bundle $\text{INJ}(V_{mU}, V_U)$ whose fiber over any given point $x \in U$ is the space of linear, injective maps from $V_{mU}|_x$ to $V_U|_x$ which is a copy of $\text{INJ}_{m,k}$. Remembering that X is assumed to be a CW complex, there is an upper bound to the dimension of its cells. Denote that by z . If $k - m$ is greater than z , then an inductive construction starting from the 0-cells, then the 1-cells and so on will construct a section of $\text{HOM}(V_{mU}, V_U)$ that sits entirely in $\text{INJ}(V_{mU}, V_U)$, which is to say that it is injective on each fiber. If $k - m > z + 1$, the same sort of inductive construction can be used to prove that any two injective vector bundle homomorphisms are homotopic through a path of fiber-wise injective homomorphisms. q.e.d.

Part 2. To continue with $\mathcal{Y}_m$ and $\mathcal{R}$: For any choice of k , the obstruction cocycle restricts as zero to any given set from $\mathfrak{U}_k$. As a consequence, the fiber bundle $\mathcal{Y}_m$ over any given such set (call it U) can be written as $(V_{mU} - 0)/\mathbb{R}^*$ with $V_{mU} \rightarrow U$ denoting a vector bundle with fiber dimension m .

To choose k : Take k' sufficiently large (given m and the dimension of X as a CW complex) so that if $k \geq k'$, then Lemma A.4 can be used to find an injective vector bundle homomorphism over each $U \in \mathfrak{U}_k$ from V_{mU} to V_U . With $k \gg k'$, choose such an injective vector bundle homomorphism for each such U and denote it by η_U .

To patch these together: Label the sets in $\mathfrak{U}_k$ by consecutive integers starting from 1 so that when U_j and $U_{j'}$ are from $\mathfrak{U}_k$ with $j' > j$, then $\dim(V_{U_{j'}}) \geq \dim(V_{U_j})$. Supposing that U_1 and U_2 share points, then over $U_1 \cap U_2$ sits the homomorphism η_{U_2} that sends V_{mU_2} to V_{U_2} ; and also sitting there is the homomorphism $T_{U_2U_1}\eta_{U_1}$ that sends V_{mU_1} to V_{U_2} . Meanwhile, V_{mU_2} and V_{mU_1} are canonically isomorphic over $U_1 \cap U_2$ up to the action of $\mathbb{R}^*$. Having fixed such an isomorphism, then $T_{U_2U_1}\eta_{U_2}$ and η_{U_2} are two injective bundle homomorphisms over $U_1 \cap U_2$ from V_{mU_2} to V_{U_2} . With this understood, Lemma A.4 can be invoked over $U_1 \cap U_2$ to construct (with the help of cut-off functions) a modification of η_2 so that the new version and $T_{U_2U_1}\eta_{U_1}$ agree on $U_1 \cap U_2$ up to the action of $\mathbb{R}^*$. Having done this, much the same construction can be used with η_3 and Lemma A.4 to extend η_1 and η_2 from $U_1 \cup U_2$ to $U_1 \cup U_2 \cup U_3$. Continuing in this same vein with U_4 and then U_5 and so on will construct the fiber preserving map for Lemma A.3. q.e.d.

A.5. Proof of Proposition A.1. The proof of Proposition A.1 has four parts. By way of a look ahead, the basic strategy is this: The third bullet of (5.5) says in effect that any given cohomology class on $\mathcal{R}$ is visible in sufficiently large k versions of $\bigcup_{U \in \mathfrak{U}_k} \Psi_U((V_U - 0)/\mathbb{R}^*)$. Mean-

while, the upcoming Lemma A.5 says in effect that if m is sufficiently large, then the pull-back by one of Lemma A.3's embeddings (with k sufficiently large) will make this class visible in $\mathcal{Y}_m$. With that understood, it then follows from Lemma A.2's description of the cohomology of $\mathcal{Y}_m$ that this class has a cocycle representative that can be written as a sum with summands of the form $\pi^*\mathfrak{q} \wedge \tau^b$ with $\mathfrak{q}$ a cocycle that represents an element from $\mathcal{C}^*$, with ν as described in the second bullet of Proposition A.1 and with b a non-negative integer. Conversely, any cochain on $\mathcal{R}$ of the form just described is closed and, it follows from Lemmas A.2, A.3 and A.5 that such a cochain is not exact.

Part 1. This part of the proof introduces and then proves the aforementioned Lemma A.5.

Lemma A.5. *Fix a non-negative integer for m and then a sufficiently large, positive integer k so that there is a fiberwise embedding $\varphi : \mathcal{Y}_m \rightarrow \mathcal{R}$ that factors through $\bigcup_{U \in \mathfrak{U}_k} \Psi_U((V_U - 0)/\mathbb{R}^*)$ as described by*

Lemma A.3. The homomorphism of pull-back by φ is an isomorphism from $H^*\left(\bigcup_{U \in \mathfrak{U}_k} \Psi_U((V_U - 0)/\mathbb{R}^*)\right)$ to $H^*(\mathcal{Y}_m)$ if the degree $*$ is $2m + 1$ or less.

Proof of Lemma A.5. The proof goes by way of an inductive Mayer-Vietoris argument that has three steps to it.

Step 1: Supposing that U is a set from $\mathfrak{U}_k$, let φ_U denote the fiberwise embedding of $\mathcal{Y}_m|_U$ into $(V_U - 0)/\mathbb{R}^*$. Instances of the Leray-Hirsch theorem (see [3]) compute the cohomology of both $(V_U - 0)/\mathbb{R}^*$ and $\mathcal{Y}_m|_U$ to have the (non-canonical) form

$$(A.13) \quad H^*(U) \oplus (H^{*-1}(U) \wedge \nu_U) \oplus (H^{*-2}(U) \wedge \nu_U^2) \oplus \cdots \oplus \nu_U^b$$

where the notation is as follows: First, $H^*(U)$ is identified with its pull-back via the projection to U , which is an injective homomorphism. Second, ν_U is a degree 1 class that restricts to each fiber as the generator of the degree 1 cohomology of that fiber, and where b is the smaller of $*$ and the dimension of the fiber of $(V_U - 0)/\mathbb{R}^*$ or $\mathcal{Y}_m$ as the case may be.

The respective isomorphisms in (A.13) imply in particular that the homomorphism of pull-back via φ_U identifies the respective $\mathcal{Y}_m|_U$ and $(V_U - 0)/\mathbb{R}^*$ summands of $H^*(U) \wedge \nu_U^h$ in (A.13) for values of h starting at 0 and ending at the minimum of $*$ and $(2m + 1)$. This is because the pull-back of the $(V_U - 0)/\mathbb{R}^*$ version of ν_U can be used for the $\mathcal{Y}_m|_U$ version due to how φ_U embeds the fibers of $\mathcal{Y}_m|_U$. In particular, the pull-back by φ is an isomorphism on cohomology of degree $2m + 1$ or less.

Step 2: Label the sets in $\mathfrak{U}_k$ consecutively from 1 in the manner of Part 2 from the proof of Lemma A.4. For any labeling integer j , let X_j denote the union of the sets from $\mathfrak{U}_k$ with label j or greater, thus $U_j \cup U_{j+1} \cup \cdots \cup U_N$ where N is the largest label of the sets from $\mathfrak{U}_k$. Supposing that U_{j-1} shares points with X_j , let $\hat{U}$ denote their intersection. The intersection of $\Psi_{U_{j-1}}((V_{U_{j-1}} - 0)/\mathbb{R}^*)$ with $\bigcup_{j \leq i \leq N} \Psi_{U_i}((V_{U_i} - 0)/\mathbb{R}^*)$ is

the restriction to $\hat{U}$ of the fiber bundle $\Psi_{U_{j-1}}((V_{U_{j-1}} - 0)/\mathbb{R}^*)$ because the dimension of the fiber of $V_{U_{j-1}}$ is not greater than that of any V_U when U 's label is greater than j . Thus, this intersection is homeomorphic to $((V_{U_{j-1}} - 0)/\mathbb{R}^*)|_{\hat{U}}$. As a consequence, another instance of the

Leray-Hirsch theorem describes the cohomology of the intersection between the sets $\Psi_{U_{j-1}}((V_{U_{j-1}}-0)/\mathbb{R}^*)$ and $\bigcup_{j \leq i \leq N} \Psi_{U_j}((V_{U_{j-1}}-0)/\mathbb{R}^*)$ via

a version of (A.13) that has U replaced by $\hat{U}$ and has b being the minimum of $*$ and the dimension of the fiber of $(V_{U_{j-1}}-0)/\mathbb{R}^*$. Meanwhile the cohomology of $\mathcal{Y}_m|_{\hat{U}}$ also has the form of (A.13) with U replaced with $\hat{U}$ and with b being the minimum of $*$ and $2m+1$.

As was the case in Step 1, this implies that pull-back via φ identifies the respective $\mathcal{Y}_m|_{\hat{U}}$ and $((V_{U_{j-1}}-0)/\mathbb{R}^*)|_{\hat{U}}$ summands of $H^*(\hat{U}) \wedge \nu_{\hat{U}}^b$ in the respective $\hat{U}$ versions of (A.13) for values of b starting at 0 and ending at the minimum of $*$ and $(2m+1)$. This last conclusion implies in particular that pull-back by φ is an isomorphism on cohomology of degree $2m+1$ or less.

Step 3: Consider now an induction on the integer j starting from $j=N$ and working down to $j=1$. The induction assumption is this:

(A.14)
For a given integer $j \in \{2, \dots, N\}$, the embedding φ as a map from $\mathcal{Y}_m|_{X_j}$ to $\bigcup_{j \leq i \leq N} \Psi_{U_i}((V_{U_i}-0)/\mathbb{R}^*)$ induces via pull-back an isomorphism between the respective cohomology up to degree $2m+1$.

Note that it follows as an instance of what is said in Step 1 that this assumption holds when $j=N$. As explained directly, it follows from Steps 1 and 2 that if (A.14) holds for a given integer j , then it also holds with j replaced by $j-1$. The assertion of the lemma follows if this claim is true for all j .

The preceding claim is proved by comparing the respective Meyer-Vietoris exact sequences for $\mathcal{Y}_m|_{X_{j-1}}$ and $\bigcup_{j-1 \leq i \leq N} \Psi_{U_i}((V_{U_i}-0)/\mathbb{R}^*)$ which come from the decomposition of X_{j-1} as $\bar{U}_{j-1} \cup X_j$. This comparison has the form depicted below in (A.15). The notation has $\mathcal{Z}_j$ denoting $\bigcup_{j \leq i \leq N} \Psi_{U_i}((V_{U_i}-0)/\mathbb{R}^*)$. The upcoming (A.15) also writes the intersection between $\mathcal{Z}_{j-1}$ and $(V_{U_{j-1}}-0)/\mathbb{R}^*$ as $((V_{U_{j-1}}-0)/\mathbb{R}^*)|_{\hat{U}}$ with $\hat{U} \equiv U_{j-1} \cap X_j$.

(A.15)

$$\begin{array}{ccccccc} \cdots & \longrightarrow & H^{*-1}(\mathcal{Y}_m|_{\hat{U}}) & \longrightarrow & H^*(\mathcal{Y}_m|_{X_{j-1}}) & \rightarrow & H^*(\mathcal{Y}_m|_{X_j}) \oplus H^*(\mathcal{Y}_m|_{U_{j-1}}) \rightarrow \cdots \\ & & \uparrow & & \uparrow & & \uparrow \\ \cdots & \rightarrow & H^*((V_{U_{j-1}}-0)/\mathbb{R}^*)|_{\hat{U}} & \rightarrow & H^*(\mathcal{Z}_{j-1}) & \rightarrow & H^*(\mathcal{Z}_j) \oplus H^*((V_{U_{j-1}}-0)/\mathbb{R}^*) \rightarrow \cdots \end{array}$$

To explain the arrows: The horizontal sequences are the two Mayer-Vietoris sequences, whereas the vertical arrows represent the homomorphism of pull-back by φ .

It follows from what is said in Step 2 that the left most vertical arrow is an isomorphism on degrees $2m + 1$ or less. Meanwhile, the induction hypothesis and what is said in Step 1 imply the same for the right most pair of arrows. This implies that the middle arrow is also an isomorphism for cohomology degrees $2m + 1$ or less. q.e.d.

Part 2. This part of the proof establishes the first bullet of Proposition A.1. The first claim to be proved in this regard is that $\pi^*\omega$ is zero in the cohomology of $\mathcal{R}$. To prove this, assume to the contrary that $\pi^*\omega$ is non-zero so as to generate nonsense. To start: If $\pi^*\omega \neq 0$, then the class $\pi^*\omega$ would restrict to all sufficiently large k versions of $\bigcup_{U \in \mathcal{U}_k} \Psi_U((V_U - 0)/\mathbb{R}^*)$ as a non-zero class (this is the third bullet assumption in (5.5)). Supposing that $m > 1$, and supposing that k is sufficiently large, then the pull-back of this non-zero class to $\mathcal{Y}_m$ by one of Lemma A.3's fiber preserving embeddings would be non-zero (see Lemma A.5). But that event is nonsense because the latter pull-back is the same as ω 's pull-back by the projection π_m from $\mathcal{Y}_m$ to X which is zero (see Lemma A.2).

With the preceding understood, now fix a degree 2-cocycle on X to be denoted by $\hat{\omega}$ to represent the class ω . Having chosen $\hat{\omega}$, then fix a degree 1 cochain on $\mathcal{R}$ (to be denoted by $\hat{\nu}$) with $d\hat{\nu} = \hat{\omega}$. The 2-cochain $\hat{\nu} \wedge \hat{\nu}$ is closed and it is not exact because it evaluates to 1 on the generator of the second homology of each fiber of π .

The class of $\hat{\nu} \wedge \hat{\nu}$ is denoted by τ . If a coboundary, say $d\mu$, is added to $\hat{\omega}$, then this is accommodated by changing $\hat{\nu}$ to $\hat{\nu}' = \hat{\nu} + \pi^*\mu$. This in turn changes τ when $\mu \wedge \mu$ is a non-zero class in $H^2(X)$ because $\hat{\nu}' \wedge \hat{\nu}' = \hat{\nu} \wedge \hat{\nu} + \pi^*(\mu \wedge \mu)$. In general, $\hat{\nu}$ can be changed by adding a closed 1-cocycle (call it σ) which changes $\hat{\nu} \wedge \hat{\nu}$ to $\hat{\nu} \wedge \hat{\nu} + \sigma \wedge \sigma$. But, as explained in the next paragraph, the cohomology class of σ must be the pull-back via π of a class from $H^1(X)$ because these are the only 1 dimensional classes in $H^1(\mathcal{R})$; and thus the class of $\sigma \wedge \sigma$ is also pulled back by π .

To prove that the class of σ is pulled back by π : Any class in $H^1(\mathcal{R})$ is visible in $H^1\left(\bigcup_{U \in \mathcal{U}_k} \Psi_U((V_U - 0)/\mathbb{R}^*)\right)$ when k is sufficiently large by virtue of the assumption in the third bullet of (5.5). Therefore, such a class is visible in $H^1(\mathcal{Y}_m)$ for $m \geq 1$ when k is sufficiently large (given m) by virtue of Lemma A.5. And, according to Lemma A.2, all such classes are the pull-backs from $H^1(X)$ via the projection map from $\mathcal{Y}_m$ to X .

Part 3. This part proves that any class in $H^*(\mathcal{R})$ can be written as a sum of those that can be represented by cocycles that have the form

(A.16)
$$(\pi^*\hat{B} + \hat{\nu} \wedge \pi^*(\hat{A}) + \pi^*\hat{i}_A) \wedge \hat{\nu}^b$$

where the notation is as follows: What is denoted here by $\hat{\nu}$ is the 1-cochain from Part 1 that obeys $d\hat{\nu} = \hat{\omega}$. Meanwhile, $\hat{B}$, $\hat{A}$ and $\hat{i}_A$ are cochains on X . In particular, what is denoted by $\hat{B}$ is closed and represents a class in $H^*(X)$ that projects to a non-zero class in $\mathcal{B}^*$; what is denoted by $\hat{A}$ is a cocycle on X that represents a class in $H^{*-1}(X)$ that is annihilated by cup product with ω ; and what is denoted by $\hat{i}_A$ is a degree $*$ cochain obeying $d\hat{i}_A = \hat{A} \wedge \hat{\omega}$.

To prove the preceding claim, let H denote a given cohomology class on $\mathcal{R}$. If k is sufficiently large, then by virtue of the third bullet assumption in (5.5), this class is non-zero in all sufficiently large k versions of $\bigcup_{U \in \mathcal{U}_k} \Psi_U((V_U - 0)/\mathbb{R}^*)$. Thus, its pull-back by one of Lemma A.3's maps is non-zero on $\mathcal{Y}_m$ for m sufficiently (by virtue of Lemma A.5). It then follows from Lemma A.2 that this class can be represented as a sum of classes that have the form depicted in (A.16) with $\hat{\nu}$ denoting the pull-back by the Lemma A.3 embedding of its namesake on $\mathcal{R}$. (Note in this regard that the pull-back of a cochain on X by the projection from $\mathcal{Y}_m$ to X is identical to the pull-back of that cochain by the composition of first the projection from $\mathcal{R}$ to X and then Lemma A.3's embedding.)

Part 4. This part proves that any class on $\mathcal{R}$ that is represented by a cocycle that can be written in the form of (A.16) with either $\hat{B}$'s class in $\mathcal{B}^*$ non-zero or $\hat{A}$'s class in $\mathcal{A}^{*-1}$ non-zero (or both) represents a non-zero class in $H^*(\mathcal{R})$. To do this, fix a pair (B, A) in $\mathcal{C}^* = \mathcal{B}^* \oplus \mathcal{A}^{*-1}$ and construct the cocycle depicted in (A.16). The corresponding cohomology class is non-zero in $H^*(\mathcal{R})$ because its pull-back to $H^*(\mathcal{Y}_m)$ for m sufficiently large via one of Lemma A.3's fiber preserving embeddings is non-zero (take k very large in Lemma A.3 and then see Lemma A.5).

A.6. Representing homology in $\mathcal{R}$. The lemma that follows states some implications from the preceding proof of Proposition A.1.

Lemma A.6. *Fix a positive integer to be denoted by m . There is a fiber preserving embedding from $\mathcal{Y}_m$ to $\mathcal{R}$ that induces (via pull-back) an isomorphism between $H^*(\mathcal{Y}_m)$ and the (finite) summand in $H^*(\mathcal{R})$ consisting of the polynomials in τ of degree at most m with $\mathcal{C}^*$ coefficients.*

Turning this around, this lemma asserts in effect that the homology of $\mathcal{R}$ that is dual to the summand in question is obtained via push-forward from the homology of $\mathcal{Y}_m$ via a fiber preserving embedding of $\mathcal{Y}_m$ into $\mathcal{R}$.

Proof of Lemma A.6. Let $\mathcal{K}_m \subset H^*(\mathcal{R})$ denote the summand consisting of the polynomials in τ of degree at most m with $\mathcal{C}^*$ coefficients. Note in particular (from Lemma A.2) that this summand is the cohomology of $\mathcal{Y}_m$. By virtue of the third bullet assumption in (5.5), this summand $\mathcal{K}_m$ is visible in all sufficiently large k versions of the subspace $\bigcup_{U \in \mathcal{U}_k} \Psi_U((V_U - 0)/\mathbb{R}^*)$. It then follows from Lemmas A.5 that

if m' is sufficiently large, then pull-back to $\mathcal{Y}_{m'}$ of $\mathcal{K}_m$ via sufficiently large k versions of Lemma A.3's fiber preserving embeddings of $\mathcal{Y}_{m'}$ into $\mathcal{R}$ defines an injective homomorphism from $\mathcal{K}_m$ into the cohomology of $\mathcal{Y}_{m'}$. Because these maps are fiber preserving, the image of $\mathcal{K}_m$ is the summand in the cohomology of $\mathcal{Y}_{m'}$ given by the polynomials in τ of degree at most m with $\mathcal{C}^*$ coefficients. With this understood, note that if $m' > m$, then there is a canonical fiber preserving embedding of $\mathcal{Y}_m$ into $\mathcal{Y}_{m'}$ which is fiberwise the standard embedding of $\mathbb{R}P^{2m+1}$ in $\mathbb{R}P^{2m'+1}$. This follows from the fact that the matrix in (A.3) is block diagonal. As a consequence, this embedding pulls back the $\mathcal{K}_m$ summand of the cohomology of $\mathcal{Y}_{m'}$ isomorphically onto the cohomology of $\mathcal{Y}_m$ to give Lemma A.2's identification of this cohomology with $\mathcal{K}_m$. q.e.d.

References

- [1] L. Faddeev, Operator anomaly for Gauss's law. *Phys. Lett. B*, 145 81–84., 1984.
- [2] D. B. Fuks, *Cohomology of the braid group mod 2*; Funkcional. Anal. i Prilozhen 4 (1970) 62–173.
- [3] A. Hatcher, *Algebraic Topology*; Cambridge University Press 2001.
- [4] T. Kato, *Perturbation theory for linear operators*; Springer 1995, MR4306622.
- [5] L. Müller. Extended functorial field theories and anomalies in quantum field theory. *arXiv:2003.08217*, 2003.
- [6] J. Mawhin and M. Willem, *Origin and evolution of the Palais-Smale condition in critical point theory*; Journ. Fixed Point Theory and Appl. 7 (2010) 265–290, MR2729392, Zbl 1205.58009.
- [7] R. Palais and S. Smale, *A generalized Morse theory*; Bull. Amer. Math. Soc. 70 (1964) 165–172, MR0158411.
- [8] C. Schiessel, *Integral cohomology of configuration spaces of the sphere*; Homology, Homotopy and Appl. 21 (2019) 283–302, MR3866779, Zbl 1404.55018.
- [9] G. Segal, Faddeev's anomaly in Gauss' law. *preprint circa*, 1984.
- [10] C. H. Taubes, The zero loci of $\mathbb{Z}/2$ harmonic spinors in dimension 2, 3 and 4. *arXiv:1407.6206*, 2014.
- [11] C. H. Taubes and Y-Y. Wu, Examples of singularity models for $\mathbb{Z}/2$ harmonic 1-forms and spinors in dimension three *J. Gökova Geom. Topol.*, 2021, MR4251085, Zbl 1472.53060.

DEPARTMENT OF MATHEMATICS
HARVARD UNIVERSITY
CAMBRIDGE, MA
USA

E-mail address: chtaubes@math.harvard.edu

DEPARTMENT OF MATHEMATICS
UNIVERSITY OF HOUSTON
HOUSTON, TX
USA

E-mail address: theresewu@gmail.com