

Thomas–Yau conjecture and holomorphic curves

Yang Li

Abstract. The main theme of this paper is the Thomas–Yau conjecture, primarily in the setting of exact, (quantitatively) almost calibrated, unobstructed Lagrangian branes inside Calabi–Yau Stein manifolds. In our interpretation, the conjecture is that Thomas–Yau semistability is equivalent to the existence of special Lagrangian representatives. We clarify how holomorphic curves enter this conjectural picture, through the construction of bordism currents between Lagrangians, and in the definition of the Solomon functional. Under some extra hypotheses, we shall prove Floer theoretic obstructions to the existence of special Lagrangians, using the technique of integration over moduli spaces. In the converse direction, we set up a variational framework with the goal of finding special Lagrangians under the Thomas–Yau semistability assumption, and we shall make sufficient progress to pinpoint the outstanding technical difficulties, both in Floer theory and in geometric measure theory.

Contents

1. Foreword	1
2. Thomas–Yau conjecture backgrounds	5
3. Moduli space of holomorphic curves	35
4. Continuity, LMCF and variational method	91
5. Variational method	101
A. The Fukaya category	131
References	149

1. Foreword

The visionary proposal of Thomas and Yau [82, 83] is the philosophy that *existence and uniqueness questions of certain special Lagrangian branes inside an almost Calabi–Yau manifold (X, ω, Ω) should be governed by stability conditions in the derived Fukaya category* (see Section 2.1). This proposal lies at the intersection of two major mathematical disciplines:

Mathematics Subject Classification 2020: 49Q15 (primary); 53D37 (secondary).

Keywords: special Lagrangian, Fukaya category, stability condition, geometric measure theory, symplectic geometry.

(i) From the viewpoint of *geometric measure theory*, the Thomas–Yau proposal promises a systematic method to prove existence theorems for special Lagrangians, which are certain absolute volume minimizers within prescribed homology classes. Currently, the known construction methods are based on perturbative techniques, symmetry reductions, some special ansatzes, and integrable system techniques specific to low dimensions [44]. In contrast, Thomas and Yau suggests the difficult PDE questions may be reducible to largely topological and algebraic questions in Floer theory.

(ii) From the viewpoint of *mirror symmetry*, the Thomas–Yau proposal is a window beyond homological mirror symmetry. Currently, most of the mathematical works on mirror symmetry are concerned with the duality between pure symplectic topology and pure complex geometry, and as such have a largely topological and algebraic flavour. As soon as one simultaneously considers the full Kähler geometric data on each side of the mirror, then questions of analytic nature become inevitable, and the Thomas–Yau proposal is a central component of this larger picture (see Section 2.4).

Two decades have lapsed since their proposal, but we feel only a small part of the mystery has been unveiled. Part of the problem is that the relevance of holomorphic curves to the Thomas–Yau proposal is insufficiently understood, despite their central role in defining the Fukaya category. One of the main goals of this paper is to propose the following picture, in the more specialized setting of exact Lagrangian branes inside (almost) Calabi–Yau Stein manifolds: Using $(n + 1)$ -dimensional universal families of holomorphic curves associated to certain $(n - 1)$ -dimensional moduli spaces, one can construct bordism currents between two n -dimensional Lagrangians in the same derived category class. More generally, one can also sometimes associate bordism currents between several Lagrangians, the notable example coming from distinguished triangles. This is a special case of the open-closed map well known to symplectic geometers, and some lower brow expositions are given in Section 3.1.

Solomon [77] defined a functional on the universal cover of the space of Lagrangians within the same Hamiltonian deformation class, whose critical points are precisely special Lagrangians. In the exact setting, we will give a more homological formula for the Solomon functional, and propose that the bordism current produced from holomorphic curves allows one to extend the functional to Lagrangian objects in the same derived category class, which is well defined without the universal cover problem (see Section 3.2). A different perspective involving integration over moduli spaces is presented in Section 3.7, and we suggest further how the Solomon functional may generalize to compact Calabi–Yau manifolds.

We further specialize to almost calibrated Lagrangians. We propose that given a distinguished triangle $L_1 \rightarrow L \rightarrow L_2 \rightarrow L_1[1]$, then there are Floer theoretic necessary conditions for L to be a special Lagrangian:

$$\arg \int_{L_1} \Omega \leq \arg \int_L \Omega \leq \arg \int_{L_2} \Omega.$$

Here L_1, L_2 are not required to be special Lagrangians, and the cohomological integral is the only quantitative way they enter into obstruction criteria. In other words, obstructions are of numerical nature. We will prove this assuming further that the bordism current from holomorphic curves satisfies automatic transversality and a *positivity condition* (see Sections 3.3–3.5). The main technique is integration over the $(n - 1)$ -dimensional moduli spaces of holomorphic curves. We also give generalizations to several Lagrangians, to make possible contact with Harder–Narasimhan decomposition. Compared with the proposal of Joyce [45], the Floer theoretic obstructions capture some features of the Bridgeland stability condition, but instead of tackling the entire derived Fukaya category $D^b \text{Fuk}(X)$, we restrict to almost calibrated Lagrangian objects, which morally correspond to the heart of a t -structure.

For a sample application of the moduli space integration technique, see Theorem 3.21, Theorem 3.26, and Proposition 3.40.

Remark 1.1. While Joyce’s prediction of a Bridgeland stability condition is a very important heuristic motivation for this paper, in our proposal we will carefully avoid assuming that the Bridgeland stability condition exists on $D^b \text{Fuk}(X)$, or that $D^b \text{Fuk}(X)$ is idempotent closed, or that the almost calibrated Lagrangians form an abelian subcategory of $D^b \text{Fuk}(X)$, except when we make comparisons with Joyce’s program.

When there are no destabilizing distinguished triangles, we say the derived Fukaya category class of L is *Thomas–Yau semistable* (see Definition 3.32). This is similar to the original viewpoint of Thomas and Yau [82, 83]. Let X be a Calabi–Yau Stein manifold, such that (ω, Ω) satisfies the complex Monge–Ampère equation, and the regularity scale of the Calabi–Yau manifold tends to infinity asymptotically. Let \mathcal{L} be the geometric measure theoretic closure of the class of exact, quantitatively almost calibrated, unobstructed Lagrangians in the $D^b \text{Fuk}(X)$ class. (Beware that we include immersed and singular Lagrangians as in Joyce [45], so the definition of \mathcal{L} is partially conjectural.) According to our interpretation of the Thomas–Yau conjecture, if $L_0 \in \mathcal{L}$ is Thomas–Yau semistable, then there exists a special Lagrangian current in \mathcal{L} (see Section 5.5.1 for the full version). Moreover, it is desirable that the special Lagrangian current carries unobstructed brane structure in some formal sense, to represent the $D^b \text{Fuk}(X)$ class or some weaker equivalence class.

We will make a start on the variational approach to tackle the Thomas–Yau conjecture in the exact and quantitative almost calibrated setting. A summary of the variational strategy is in Section 5.3. We prove a number of uniform estimates independent of the Lagrangians, notably the potential clustering property, which implies a uniform bound on the energy of holomorphic curves, and that the Solomon functional deviates from an elementary functional by a uniformly bounded amount (see Sections 5.1 and 5.2). We are not able to complete the proof of the Thomas–Yau conjecture, but we try to identify the main technical difficulties to be overcome. To make contact with the a priori compactness theorems in geometric measure theory, it is natural to set up the variational program in terms of Lagrangian currents and varifolds. One then encounters both geometric measure theoretic

difficulties, to do with the existence of enough Lagrangian competitors, and the difficulty to make sense of Floer theory for Lagrangian currents. We expect that many aspects of Floer theory will not be robust under passage to such weak limits of Lagrangians, but as a heuristic principle, we hope that the bordism currents relevant to our proposal are robust (see Section 5.4).

Assume there is a suitable compactness theory for Lagrangian objects, then we expect the existence of special Lagrangians to be essentially equivalent to the properness of the Solomon functional. The following heuristic principle would then explain why Thomas–Yau conjecture could be true: for a sequence of Lagrangians whose Solomon functional diverges to infinity, the underlying Lagrangian objects should break up into a bounded number of connected components L_i , such that the Lagrangian potential on each component has uniformly bounded oscillation, and the asymptotic behaviour of the Solomon functional is controlled by Floer theoretic data related to the cohomological integrals $\int_{L_i} \Omega$. We will give evidence for this picture (see Section 5.5), and explain a similar picture for Hermitian Yang–Mills connections (see Section 2.5).

For another summary on the logical structure of the variational program, see Conjecture 5.21 and the subsequent discussions.

Although we restrict to the exact setting in the above, we will leave the room open for more ambitious speculations in the setting of compact almost Calabi–Yau manifolds. The almost calibrated condition, on the other hand, is essential in our proposal, and cannot be dropped within this framework even with substantial efforts.

Aside from the main goal of making the proposal and presenting evidence, this paper also contains a large amount of expository content, which aims to present both a (biased) overview and a critique of the current literature on the Thomas–Yau conjecture, and to explain how our picture fits into this body of works. The most relevant works are the original papers of Thomas and Yau [82, 83], and the major update by Joyce [45]. We will devote substantial attention to their main considerations (see Sections 2.1–2.5, 2.7, and 4.1), and the significance of the Thomas–Yau–Joyce proposal to mirror symmetry (see Section 2.4), and we will constantly draw comparisons between their pictures and ours. We also touch on a number of other relevant works:

- Joyce–Imagi–Santos [41] gave a characterization of Lawlor necks, in which holomorphic curves made an interesting, albeit somewhat technical appearance; see Section 2.3.
- The Hermitian Yang–Mills equation [26, 27, 84] was a major inspiration to Thomas and Yau. The deformed Hermitian Yang–Mills equation [16, 18–20] has been advocated as a mirror version of special Lagrangians; see Sections 2.5 and 2.6.
- Solomon [77–79] introduced a functional and a formal Riemannian metric on the infinite-dimensional space of almost calibrated Lagrangians in a given Hamiltonian isotopy class. This space is too small for our variational purpose, but his functional is fundamental to our proposal; see Section 2.8.

- Lotay and Pacini [57,58] introduced a formal GIT picture for totally real submanifolds. Their notion of ‘geodesic’ was our main inspiration for the bordism currents between Lagrangians; see Section 2.9.
- Neves [62–64] pointed out the inevitability of finite time singularity in the Lagrangian mean curvature flow, which is fundamental to Joyce’s proposal [45]. Some of his technical arguments are adapted to support our variational program in Section 5.2; see Sections 4.1 and 5.2.
- Joyce’s earlier work on wall crossing and special Lagrangian counting problems [42] is in some sense the elliptic cousin to his Lagrangian mean curvature flow proposal [45]; see Section 4.2

Remark 1.2 (Prerequisites). While we have endeavoured to survey most of the previous works directly aimed at the Thomas–Yau conjecture (regrettably with some inevitable omissions), there is a very extensive literature on geometric measure theory and symplectic geometry in the background. We do not assume expertise on these matters, but some previous exposures such as F. Morgan’s introductory book [61], and the excellent surveys of Auroux [9] and Smith [75] would be useful. The most important background facts for our main purpose are recalled in Section 5.1 and the Appendix on the Fukaya category. While the brief summary therein is not completely sufficient for all arguments in this paper, we hope the casual reader could get the main gist, if not some sporadic remarks. The punctilious reader may wish to refer to the Floer degree and sign convention summarized in the Appendix, which is different from, e.g., Seidel’s book [71]. The various allusions to Kähler geometry are mainly for motivational purposes, which can be skipped by readers less interested in these topics.

Remark 1.3 (Rigor). This paper contains a somewhat unconventional mixture of proposals, expositions and proofs. We have attempted to indicate all speculative elements as ‘heuristic’ or ‘conjecture’. Other arguments are either complete proofs, or sketches intended as expositions.

2. Thomas–Yau conjecture backgrounds

2.1. The Thomas–Yau–Joyce picture

An *almost Calabi–Yau* manifold (X, ω, J, Ω) is a complete n -dimensional Kähler manifold with a nowhere vanishing holomorphic volume form. It is called Calabi–Yau if the complex Monge–Ampère equation $\omega^n = \text{const } \Omega \wedge \bar{\Omega}$ holds, whence X is Ricci flat. A real n -dimensional compact submanifold of an almost Calabi–Yau manifold $\iota: L \rightarrow X$ is

called *special Lagrangian* of constant phase angle $\hat{\theta} \in \mathbb{R}$,¹ if

$$\omega|_L = 0, \quad \text{Im}(e^{-i\hat{\theta}}\Omega|_L) = 0.$$

Special Lagrangian submanifolds have unobstructed deformation theory.² Furthermore, if X is Calabi–Yau, then a special Lagrangian is a minimal submanifold, and in fact minimizes the area within its homology class, thanks to the *calibration inequality* for submanifolds, saturated precisely by special Lagrangians:

$$\left| \int_L \text{Re}(e^{-i\hat{\theta}}\Omega) \right| \leq \left| \int_L \Omega \right| \leq \int_L d\text{vol}(L) = \text{Vol}(L).$$

We will always impose some compactness or convexity at infinity on X to ensure the Fukaya category of (X, ω) makes sense.

Thomas–Yau’s proposal. In [82, 83], Thomas and Yau introduced the remarkable philosophy that *existence and uniqueness questions of special Lagrangians inside an (almost) Calabi–Yau manifold (X, ω, Ω) should be related to stability conditions in the Fukaya category*. Turning this intuition into precise mathematical predictions, is however not easy, for at least the following geometric reasons, in addition to the analytic difficulties related to minimal surfaces and mean curvature flows:

(i) The Fukaya category as it currently stands is likely inadequate for the purpose: the special Lagrangian representatives of a given derived Fukaya category class in $D^b\text{Fuk}(X)$, should it exist, is by no means guaranteed to be smooth and embedded. Thus one would like to enlarge the objects of the Fukaya category to include immersed and possibly singular Lagrangians. From the viewpoint of symplectic topology, the lack of Lagrangian objects is a basic difficulty, which is usually hidden in the non-geometric step of taking the twisted complexes in the construction of $D^b\text{Fuk}(X)$, and the idempotent completions in the construction of $D^\pi\text{Fuk}(X)$ (see Section A.2).

(ii) The knowledge of the stability conditions is severely deficient. Thomas–Yau’s paper predates the ingredient of Bridgeland stability,^{3,4} but the basic problem remains open: how to construct a stability condition on the derived Fukaya category $D^b\text{Fuk}(X)$ from the information of a holomorphic volume form Ω on X ? In contrast, in analogous problems such as the existence of Hermitian Yang–Mills connections, the stability condition is known a priori before the more serious endeavour to solve the PDE. One of the

¹Different lifts of the phase angle from $\mathbb{R}/\pi\mathbb{Z}$ to \mathbb{R} shift the grading of the brane structure, so they are different as objects of the Fukaya category.

²However, obstructions may arise if brane structures are taken into account.

³What was available at the time, was the famous μ -stability related to Hermitian Yang–Mills connections, of which Thomas and Yau were of course leading experts.

⁴There are current debates whether Bridgeland stability is the ultimately correct framework for formalizing the physical intuition of stability conditions controlling BPS particle decay [31]. We regard Bridgeland stability as a working definition, to be modified in case future evidence arises.

goals of this paper is to explain some modest progress on this issue, namely that there are non-trivial Floer theoretic obstructions to the existence of special Lagrangians.

Thomas–Yau primarily restricted attention to the case of *almost calibrated* Lagrangians, meaning the Lagrangian angle function θ satisfies⁵

$$-\frac{\pi}{2} < \theta < \frac{\pi}{2}.$$

Notice that this restriction removes any $\pi\mathbb{Z}$ ambiguity of the Lagrangian angle, so the Lagrangian is graded. Since we are focusing on compact Lagrangians, it makes sense to restrict to a more quantitative version, for some fixed small $\varepsilon > 0$:

$$-\frac{\pi}{2} + \varepsilon \leq \theta \leq \frac{\pi}{2} - \varepsilon. \quad (2.1)$$

One immediate consequence is an a priori *volume bound*.

Lemma 2.1. *If L satisfies (2.1), then $\text{Vol}(L) \leq \frac{1}{\sin \varepsilon} \int_L \text{Re } \Omega$.*

Proof. From $e^{i\theta} d\text{vol}_L = \Omega|_L = \text{Re } \Omega|_L + \sqrt{-1} \text{Im } \Omega|_L$, we see $\text{Re } \Omega$ is an *orientation* form on L , and we arrive at

$$\text{Vol}(L) = \int_L \frac{1}{\cos \theta} \text{Re } \Omega \leq \frac{1}{\sin \varepsilon} \int_L \text{Re } \Omega,$$

as required. ■

Remark 2.1. Another immediate consequence of the almost calibrated condition, is that the complex number $Z(L) = \int_L \Omega$ is non-zero, with $\phi(L) = \frac{1}{\pi} \arg Z(L) \in (-\frac{1}{2}, \frac{1}{2})$.

Keeping our narrative closer to the historical development, Thomas and Yau were inspired by μ -stability for Hermitian Yang–Mills connections. They assumed that the principal mechanism which can forbid the Hamiltonian isotopy class of L from admitting a special Lagrangian representative, is related to a distinguished triangle in the Fukaya category, the primary geometric source being that L is Hamiltonian isotopic to a graded Lagrangian connected sum $L = L_1 \# L_2$,⁶ such that $\phi(L_1) \geq \phi(L_2)$. Based on this intuition, Thomas made an attempt to define a notion of stability (see [82, Definition 5.1] and Definition 3.32 below).⁷ Thomas then made the following important prediction.

⁵The important thing is that the upper and lower bounds on θ differ by π . Shifting the interval by a constant is inconsequential.

⁶It is very important for Thomas and Yau that the Lagrangian connected sum is asymmetric in L_1 and L_2 . Our notation for the Lagrangian connected sum agrees with Thomas and Yau, but is opposite to Joyce [43] and a large number of symplectic geometry texts. Our convention is compatible with the distinguished triangle $L_1 \rightarrow L_1 \# L_2 \rightarrow L_2 \rightarrow L_1[1]$.

⁷We will not repeat their definition verbatim here because the author thinks its focus on the Hamiltonian isotopy class, instead of the Fukaya category class, is largely a limitation of its time. It is the spirit rather than the letter of their definition which matters.

Conjecture 2.2 ([82, Conjecture 5.2]). *A graded Lagrangian has a special Lagrangian representative in its Hamiltonian class if and only if it is stable, and this special Lagrangian representative is unique.*

Thomas and Yau [83] analyzed the problem again from the Lagrangian mean curvature flow (LMCF) perspective, and made a somewhat more cautious prediction, which roughly amounts to the following. When a destabilizing decomposition into Lagrangian sums is forbidden, either by a smallness assumption on the oscillation of the phase function θ , or because the volume of L is smaller than $\text{Vol}(L_1) + \text{Vol}(L_2)$ for any putative decomposition, (notice both conditions are preserved under the flow), then they conjectured that the mean curvature flow starting from L will converge into a special Lagrangian [83, Section 7].

Bridgeland stability, Joyce’s proposal. To trace the subsequent development, and move beyond the almost calibrated setting, we need to recall an important piece of homological algebra known as *Bridgeland stability conditions* on a triangulated category [13].

Definition 2.3. A Bridgeland stability condition (Z, \mathcal{P}) on a triangulated category \mathcal{D} consists of a group homomorphism (‘central charge’) Z from the (numerical) Grothendieck group $K(\mathcal{D})$ to \mathbb{C} , and full additive subcategories $\mathcal{P}(\phi) \subset \mathcal{D}$ for each $\phi \in \mathbb{R}$, whose objects are called ‘semistable objects of phase angle $\pi\phi$ ’,⁸ satisfying the following axioms:

Phase. If $L \in \mathcal{P}(\phi)$, then $Z(L) = m(L) \exp(i\pi\phi)$ for some $m(L) > 0$.

Shift. For all $\phi \in \mathbb{R}$, $\mathcal{P}(\phi + 1) = \mathcal{P}(\phi)[1]$.

Monotonicity. If $\phi_1 > \phi_2$ and $L_j \in \mathcal{P}(\phi_j)$, $j = 1, 2$, then $\text{Hom}_{\mathcal{D}}(L_1, L_2) = 0$.

Harder–Narasimhan decomposition. For each non-zero object $L \in \mathcal{D}$, there are a finite sequence of real numbers $\phi_1 > \phi_2 > \cdots > \phi_N$ and a *Harder–Narasimhan decomposition*

$$0 = \mathcal{E}_0 \rightarrow \mathcal{E}_1 \rightarrow \cdots \rightarrow \mathcal{E}_N = L, \quad (2.2)$$

with distinguished triangles

$$\mathcal{E}_{i-1} \rightarrow \mathcal{E}_i \rightarrow L_i \rightarrow \mathcal{E}_{i-1}[1]$$

such that $L_j \in \mathcal{P}(\phi_j)$.

Calibration. For any fixed norm on the finite-dimensional vector space $K(\mathcal{D}) \otimes_{\mathbb{Z}} \mathbb{R}$, we have a uniform constant C , such that any semistable object L satisfies

$$\|L\| \leq C|Z(L)|.$$

Remark 2.2. For the heuristic but naïve geometric meaning, one can imagine that \mathcal{D} is the derived Fukaya category, $K(\mathcal{D})$ is the sublattice of $H_n(X)$ generated by the Lagrangians, the central charge is $Z(L) = \int_L \Omega$, the subcategory $\mathcal{P}(\phi)$ is generated by the special

⁸In the convention of Bridgeland, the phase is ϕ . We have instead opted to call $\pi\phi$ the phase, which is more naturally identified with the phase angle of special Lagrangians.

Lagrangians of constant phase angle $\theta = \pi\phi$, the Harder–Narasimhan decomposition means a multiple Lagrangian connected sum with decreasing phase angles

$$L \simeq L_1 \# L_2 \# \cdots \# L_N,$$

and the calibration property comes from the fact that the total mass of a special Lagrangian is equal to $|Z(L)|$, and the mass bounds any norm of $[L] \in H_n(X)$ using Poincaré duality, assuming $H_n(X)$ is finite-dimensional.

Remark 2.3. We write the subcategory generated by all the $\mathcal{P}(\phi)$ within the interval $\phi \in (\phi_0, \phi_1]$ as $\mathcal{P}(\phi_0 < \phi \leq \phi_1)$. The important property is that $\mathcal{P}(\phi_0 < \phi \leq \phi_0 + 1)$ is an abelian category, known as the heart of a bounded t -structure. The advantage is that while in general triangulated categories only distinguished triangles make sense, for abelian categories we can talk about exact sequences, subobjects and quotients.

With the hindsight of Bridgeland stability condition, and more than one decade of progress on the mean curvature flow as well as symplectic geometry, Joyce [45] produced a major update of the Thomas–Yau picture. We shall discuss more about the LMCF considerations in Section 4.1, but it suffices here to give away its main punchline (see [45, Conjecture 3.34]).

Let (X, ω, Ω) be a Calabi–Yau manifold,⁹ either compact or Stein. There should be an enlarged version of the derived Fukaya category $D^b \text{Fuk}(X)$, including classes of immersed or mildly singular Lagrangians, and a Bridgeland stability condition (Z, \mathcal{P}) on $D^b \text{Fuk}(X)$, whose central charge is

$$Z(L) = \int_L \Omega.$$

Given a Lagrangian brane L (with grading, orientation, relative spin structure, local system and bounding cochain data) such that the Floer cohomology is unobstructed, and suppose L is generic in its Hamiltonian isotopy class. Then the LMCF $(L_t)_{t>0}$ (with brane structure) starting from L exists for all time *with surgeries* at a finite series of singular times. The nature of these surgeries have conjectural descriptions. The Lagrangian can change its Hamiltonian class at these surgeries, but maintains its derived category class. At $t \rightarrow \infty$, the Lagrangians L_t converges in the geometric measure theory sense to a union of graded special Lagrangian currents L_1, \dots, L_N of phase angle $\hat{\theta}_1 > \hat{\theta}_2 > \cdots > \hat{\theta}_N$ with multiplicities counted:

$$\lim_{t \rightarrow \infty} L_t = L_1 + \cdots + L_N. \quad (2.3)$$

If there is only one constant phase angle $\hat{\theta}$ appearing in the infinite time limit, then L defines a semistable object of phase $\hat{\theta}$ in $D^b \text{Fuk}(X)$. Otherwise, L is not semistable with

⁹The complex Monge–Ampère equation is convenient but not indispensable, see Section 4.1.

respect to any phase angle, and the infinite time limit supposedly give rise to the Harder–Narasimhan decomposition.

Joyce’s program is much more ambitious in the sense that it attempts to capture the entire triangulated category $D^b \text{Fuk}(X)$. The original Thomas–Yau proposal, which is concerned only with almost calibrated Lagrangians, fits into the Joyce picture as an *abelian subcategory* $\mathcal{P}(-\pi/2 < \theta \leq \pi/2) \subset D^b \text{Fuk}(X)$,¹⁰ the heart of a certain t -structure. The main appeal of Joyce’s mean curvature flow perspective, is that if the infinite time limiting currents L_1, \dots, L_N can be given suitable brane structures to be admitted as objects of $D^b \text{Fuk}(X)$, then the program gives a conjectural dynamical mechanism to explain the Harder–Narasimhan decomposition, which is the most non-trivial aspect of the Bridgeland stability. Its principal drawback is that Joyce offers no a priori information on the Bridgeland stability beyond the central charge, other than letting off the mean curvature flow to find its own destiny. This limits its practical applicability, for instance to existence questions of special Lagrangians in prescribed classes in $D^b \text{Fuk}(X)$. See Section 3.6 for more discussions.

Why Thomas–Yau has predictive power. At this moment an objection may arise: since the Thomas–Yau–Joyce picture is beset by some vagueness and plenty of technical difficulties, why is it useful as a guiding principle at all? Besides the supportive evidence that we shall soon discuss, the main answer is that the Thomas–Yau philosophy transforms a PDE problem (the existence of special Lagrangians) into a categorical framework, which if better understood is in principle checkable by algebraic means. *A Bridgeland stability is the interplay between a category and a numerical property.* In the analogous problem of Hermitian Yang–Mills connections, the existence criterion is formulated by μ -stability, which compresses the information of the Kähler form into only certain intersection numbers/cohomological integrals (see Section 2.5). Likewise, the stability condition responsible for the existence of special Lagrangians, even though it is not adequately specified, is in principle a compression of the analytic data of a holomorphic volume form, and quite plausibly enters only through cohomological integrals of Ω , as will be discussed more fully in Section 3. As an indication of the possible predictive power of the Thomas–Yau picture, here is a sample question as food for thought:

Question 1. Fix the holomorphic volume form Ω . Let ω_1, ω_2 be two generic Kähler forms, differing only by the differential of a compactly supported 1-form. Can we define a count of special Lagrangian rational homology spheres, such that the numbers agree for ω_1 and ω_2 ?

¹⁰Here we ignore the ‘small’ difference between $<$ and \leq . This would be justified if Ω is suitably generic so that there is no special Lagrangian of phase $\pi/2$, namely that the countably many numbers $\text{Arg} \int_L \Omega$ for $[L] \in H_n(X, \mathbb{Z})$ miss the number $\pi/2$.

Remark 2.4. As we shall see, the main evidence of Thomas–Yau picture (see Section 2.2) does not really use the complex Monge–Ampère equation.¹¹ On the mirror side, as a consequence of the μ -stability characterization, the existence of Hermitian Yang–Mills connections on a compact Kähler manifold does not depend on the choice of the Kähler form within a fixed Kähler class.

2.2. Principal evidence of Thomas–Yau

Thomas and Yau offered a number of arguments in support of their picture. We now describe the evidence from *first principles*, and will later return to the evidence from *analogies*.

The following elementary observation is a first indication about how special Lagrangian objects resemble stable objects of a Bridgeland stability condition.

Proposition 2.4 ([83, Section 5.2]). *Let L, L' be unobstructed Lagrangian branes whose supports are compact special Lagrangian manifolds, with Lagrangian angles $\theta_L > \theta_{L'}$. Then $HF^k(L, L') = 0$ for $k \leq 0$.*

Proof. After C^∞ -small Hamiltonian perturbation we may assume any intersection point p between L and L' is transverse. We can write the tangent planes of L, L' inside $T_p X \simeq \mathbb{C}^n$ in the standard form

$$T_p L = \mathbb{R}^n \subset \mathbb{C}^n, \quad T_p L' = (e^{i\phi_1}, \dots, e^{i\phi_n})\mathbb{R}^n \subset \mathbb{C}^n, \quad 0 < \phi_i < \pi.$$

The Floer degree of the intersection point p is

$$\mu_{L, L'}(p) = \frac{1}{\pi} \left(\theta_L - \theta_{L'} + \sum_1^n \phi_i \right) > 0,$$

whence $HF^k(L, L') = 0$ for $k \leq 0$. ■

Remark 2.5. As is clear from the proof, the special Lagrangian condition can be relaxed to $\inf_L \theta_L > \sup_{L'} \theta_{L'}$.

Remark 2.6. Our Floer degree convention follows Joyce [45], but is opposite to that of Seidel [71] and many other symplectic geometry texts. In our convention, shifting the Lagrangian phase by $k\pi$ for $k \in \mathbb{Z}$ is equivalent to considering the shifted object $L[k]$. Thus a slight extension of the above is that if $\theta_L > \theta_{L'} + l$, then

$$HF^k(L, L') = HF^{k-l}(L, L'[l]) = 0, \quad k \leq l.$$

¹¹The almost Calabi–Yau setting is desirable not only for the sake of generality, but may be essential to achieve suitable genericity. As an additional motivation on the side of physics, the SCFT condition translates into a Kähler condition on the target space metric, which satisfies the Ricci-flatness only approximately [39, Section 14.2.4].

The most compelling evidence due to Thomas and Yau is as follows.

Theorem 2.5 (Thomas–Yau uniqueness [83]). *Let L, L' be unobstructed Lagrangian branes supported on embedded special Lagrangians of the same phase, which define isomorphic objects in $D^b \text{Fuk}(X)$, then their supports coincide.*

Remark 2.7. The most general Thomas–Yau uniqueness, which includes immersed Lagrangians, seems to be due to Imagi [40]. The original Thomas–Yau theorem is phrased in terms of uniqueness in the Hamiltonian isotopy class, even though it can be cast in more general categorical terms. The categorical perspective is preferred, because it is closer to the spirit of homological mirror symmetry, and because one derived Fukaya category class may contain several Hamiltonian isotopy classes. If immersed Lagrangians are allowed, then isomorphism in $D^b \text{Fuk}(X)$ would also identify certain embedded Lagrangian objects with immersed objects of different topologies. When this happens, an interesting corollary of Thomas–Yau uniqueness is that at most one of these Hamiltonian classes contains special Lagrangian branes.

Sketch of the proof. Assume the supports do not coincide. After small Hamiltonian perturbations, we can ensure the perturbed Lagrangians \tilde{L}, \tilde{L}' have transverse intersections, and still define the same isomorphic objects in $D^b \text{Fuk}(X)$. Under the special Lagrangian assumption, through judicious choice of the Hamiltonian via Morse theory as in Thomas–Yau [83, Theorem 4.3], or by using genericity arguments based on real analyticity as in Joyce–Imagi–Santos [41, Section 4.3], one can ensure there is no intersection point $\tilde{L} \cap \tilde{L}'$ of degree 0, n , so in particular

$$HF^0(L, L') \simeq HF^0(\tilde{L}, \tilde{L}') = 0.$$

However, this implies the cohomological unit of $HF^*(L, L)$ is zero, so the Floer cohomology ring of L is zero, namely L is a zero object in $D^b \text{Fuk}(X)$, contradiction. ■

The Thomas–Yau argument reveals the relevance of the Fukaya category to special Lagrangian geometry. The role of holomorphic curves, which are a central ingredient in the Fukaya category, is however rather opaque in this argument; their only appearance is to make the Floer cohomology defined.

2.3. Variant: Uniqueness of the Lawlor neck

A variant of the Thomas–Yau argument¹² appears in the subsequent work of Joyce–Imagi–Santos [41] on the uniqueness of Lawlor necks, where holomorphic curves and the algebraic structures of the Fukaya category appear in a more prominent, albeit somewhat technical way.

¹²Abouzaid and Imagi have mentioned in their talks some other interesting applications on the topology of special Lagrangians inside the cotangent bundle of a special Lagrangian with some fundamental group conditions, using another variant of the Thomas–Yau argument. We look forward to the appearance of their paper.

Lawlor necks. We first recall some basics about Lawlor necks [46, 52], which are non-compact embedded exact special Lagrangians $L_{\phi,A}$ inside the standard Euclidean \mathbb{C}^n , asymptotic at infinity to the union of two planes

$$\Pi_0 = \mathbb{R}^n, \quad \Pi_\phi = (e^{i\phi_1}, \dots, e^{i\phi_n})\mathbb{R}^n, \quad 0 < \phi_i < \pi, \quad \sum \phi_i = \pi.$$

Symplectic topologically, they can be viewed as a realization of the Lagrangian handle $S^{n-1} \times \mathbb{R}$ that appears in the Lagrangian connected sum construction. This motivates the Ansatz

$$L_{\phi,A} = \{(z_1(y)x_1, \dots, z_n(y)x_n) : y \in \mathbb{R}, x_k \in \mathbb{R}, x_1^2 + \dots + x_n^2 = 1\}.$$

The special Lagrangian condition translates into an ODE system on the functions $z_1(y), \dots, z_n(y)$, which can be solved exactly as follows.

Let $n > 2$ and $a_1, \dots, a_n > 0$, and define polynomials p, P by

$$p(x) = (1 + a_1x^2) \cdots (1 + a_nx^2) - 1, \quad P(x) = \frac{p(x)}{x^2}.$$

Define real numbers ϕ_1, \dots, ϕ_n and A by

$$\phi_k = a_k \int_{-\infty}^{\infty} \frac{dx}{(1 + a_kx^2)\sqrt{P(x)}}, \quad A = \int_{-\infty}^{\infty} \frac{dx}{2\sqrt{P(x)}}.$$

Clearly $\phi_k, A > 0$, and elementary integration shows $\sum \phi_i = \pi$. This yields a 1-1 correspondence between n -tuples (a_1, \dots, a_n) with $a_k > 0$, and $(n+1)$ -tuples $(\phi_1, \dots, \phi_n, A)$ with $\phi_k \in (0, \pi)$, $\sum \phi_k = \pi$ and $A > 0$. Setting

$$z_k(y) = e^{i\psi_k(y)} \sqrt{a_k^{-1} + y^2}, \quad \text{where } \psi_k(y) = a_k \int_{-\infty}^y \frac{dx}{(1 + a_kx^2)\sqrt{P(x)}},$$

yields the solution $(z_1(y), \dots, z_n(y))$, hence the *Lawlor necks* $L_{\phi,A}$.

For fixed asymptotic planes Π_0, Π_ϕ , the Lawlor necks arise in a 1-parameter family, related to each other by the rescaling $\vec{z} \rightarrow \lambda \vec{z}$ in \mathbb{C}^n , and A behaves like 2-dimensional area $A \rightarrow \lambda^2 A$ under this scaling. One also observes that asymptotically near infinity, the Lawlor necks are graphs over Π_0 (resp. Π_ϕ) of the differential df , where

$$|f| = O(|\vec{x}|^{2-n}), \quad |\partial^k f| = O(|\vec{x}|^{2-n-k}).$$

We say the Lawlor neck has asymptotic decay rate $2-n$. The upshot is that it approaches $\Pi_0 \cup \Pi_\phi$ sufficiently fast.

Joyce–Imagi–Santos uniqueness theorem.

Theorem 2.6 ([41]). *Let $n \geq 3$.¹³ Assume L is a smooth embedded exact special Lagrangian of phase zero, asymptotic at rate < 0 to the union of the two planes $\Pi_0 \cup \Pi_\phi$ with $\sum \phi_i = \pi$, then L is $L_{\phi,A}$ for some $A > 0$.*

¹³The theorem is also true in complex dimension 2, proved earlier by Lotay and Neves [56] by other means not involving Floer theory.

Here is a sketch of their arguments.

Using the asymptotic assumption on the exact Lagrangian L , one can assign an analytic invariant $A(L)$ to L as follows. Let $f_L: L \rightarrow \mathbb{R}$ be a primitive of the Liouville form λ , namely $df_L = \lambda|_L$, then f_L converges to constants c_0, c_ϕ at the two asymptotic ends along Π_0, Π_ϕ respectively. Then one defines $A(L) = c_\phi - c_0$. If L coincides with the Lawlor neck $L_{\phi,A}$, then $A(L) = A$.

Partially compactify \mathbb{C}^n into a Liouville manifold identified as the plumbing M of two cotangent bundles T^*S^n with T^*S^n . Here the two copies of S^n arise topologically as one-point compactifications of Π_0 and Π_ϕ by adding the points at infinity ∞_0 and ∞_ϕ , and topologically M is the union of \mathbb{C}^n and the two cotangent fibres over ∞_0 and ∞_ϕ respectively. Under suitably fast decay condition at infinity, the unknown special Lagrangian L can be compactified into an exact graded embedded Lagrangian \bar{L} inside M . One would like to compare this to the Lagrangian $\bar{L}_{\phi,A}$ obtained by the compactification of the standard Lawlor necks $L_{\phi,A}$ inside M .

By analyzing the intersection pattern with the two cotangent fibres at infinity, and using the classification results of Abouzaid and Smith [3], one shows that inside $D^b \text{Fuk}(M)$, the Lagrangian object \bar{L} is isomorphic to one of the two Lagrangian connected sums of the two S^n with suitable gradings, and in fact the assumption on Floer degrees

$$\sum \phi_i = \pi$$

singles out $\bar{L} \simeq \bar{L}_{\phi,A} \in D^b \text{Fuk}(M)$, the opposite surgery corresponding to

$$\sum \phi_i = (n-1)\pi.$$

This step needs $n \geq 3$. For contradiction, we assume \bar{L} does not coincide with $\bar{L}_{\phi,A}$ for any choice of parameter $A > 0$.

By modification of the Thomas–Yau argument, one shows that after a small Hamiltonian perturbation \bar{L}'' of $\bar{L}_{\phi,A}$, we can ensure \bar{L}'' is transverse to \bar{L} , there is no degree 0, n intersection points in $\bar{L}'' \cap \bar{L}$ inside \mathbb{C}^n , and there is precisely one intersection point p and q in $\bar{L}'' \cap \bar{L}$ on each of the two cotangent fibres at infinity respectively. Moreover, the $D^b \text{Fuk}(M)$ class and the analytic invariants of \bar{L}'' agree with that of $\bar{L}_{\phi,A}$.

Remark 2.8. The subtlety at infinity prevents one from removing degree 0, n intersections outside the \mathbb{C}^n region, so one does not reach an immediate contradiction as in the Thomas–Yau argument. This technical failure is necessary, because the Lawlor necks with fixed asymptotic planes are not unique, but do arise in a 1-parameter family. It is in overcoming this technical problem that holomorphic curves appear in [41].

Now suppose the Lawlor neck is chosen with the parameter $A = A(L)$, which presumes $A(L) > 0$.

Lemma 2.7 ([41, Theorem 2.15]). *Assume J is a generic almost complex structure on M compatible with the Liouville structure. There exists a J -holomorphic strip Σ with boundary on \bar{L} and \bar{L}'' and two corners at p and q respectively.*

Proof. Consider the Floer cup product with mod 2 coefficients

$$HF^0(\bar{L}'', \bar{L}) \otimes HF^0(\bar{L}, \bar{L}'') \rightarrow HF^0(\bar{L}, \bar{L}),$$

which can be identified as the cup product

$$H^0(\bar{L}, \bar{L}) \otimes H^0(\bar{L}, \bar{L}) \rightarrow H^0(\bar{L}, \bar{L}), \quad 1_{\bar{L}} \cup 1_{\bar{L}} = 1_{\bar{L}},$$

and thus must be non-trivial. However, at chain level this Floer product comes from the A_∞ operation

$$m_2: CF^0(\bar{L}'', \bar{L}) \otimes CF^0(\bar{L}, \bar{L}'') \rightarrow CF^0(\bar{L}, \bar{L}),$$

which must be non-trivial; see the appendix for more exposition on Floer theory, especially Example A.1. The counting interpretation implies there are intersection points

$$p' \in CF^0(\bar{L}, \bar{L}'') \quad \text{and} \quad q' \in CF^0(\bar{L}'', \bar{L}) \simeq CF^n(\bar{L}, \bar{L}'')^\vee$$

and some holomorphic strip in between. Since degree 0, n intersection points cannot occur inside \mathbb{C}^n , they can only occur at infinity, so we must have $\{p, q\} = \{p', q'\}$. ■

Now the area of the J -holomorphic curve can be computed cohomologically. Using the choice of parameter A , we have

$$\begin{aligned} \int_{\Sigma} \omega &= \int_{\partial \Sigma} \lambda = \int_{p' \rightarrow q'} df_{\bar{L}} + \int_{q' \rightarrow p'} df_{\bar{L}''} \\ &= \pm(A(L) - A(L'')) = \pm(A - A) = 0. \end{aligned}$$

This contradicts the positivity of area of the holomorphic curve, which proves L must coincide with $L_{\phi, A}$.

Finally, one needs to a priori justify $A(L) > 0$. This relies on a slightly more complicated holomorphic polygon counting argument, and the main upshot is that one can produce a non-trivial holomorphic triangle from a distinguished triangle in $D^b \text{Fuk}(M)$, with the three edges on \bar{L} , $\Pi_0 \cup \{\infty_0\}$ and $\Pi_\phi \cup \{\infty_\phi\}$. Then one shows $A(L)$ has the interpretation as its area, so must be positive.

Ideal triangles. In [41] the perturbations involved in the partial compactification and the genericity of the almost complex structure makes the holomorphic curves rather difficult to visualize.¹⁴ We now present a heuristic way to see holomorphic triangles with the edges on $\bar{L}_{\phi, A}$, $\Pi_0 \cup \{\infty_0\}$ and $\Pi_\phi \cup \{\infty_\phi\}$, by restricting attention to \mathbb{C}^n with the standard complex structure, and we imagine the two vertices ∞_0, ∞_ϕ as the intersection points at infinity.

¹⁴A typical feature of Floer theory, is that completely realistic examples about holomorphic curves are also non-explicit.

We choose any $\vec{x} = (x_1, \dots, x_n) \in S^{n-1}$. Assume first that $x_k \neq 0$. Then coordinate-wise, we have a real curve in \mathbb{C} swept out by $z_k(y)x_k$ as y varies from $-\infty$ to $+\infty$, and two straight rays emanating from the origin defined by $\text{sign}(x_k)\mathbb{R}_+$ and $\text{sign}(x_k)e^{i\phi_k}\mathbb{R}_+$. Inside \mathbb{C} , these three real curves enclose a non-compact holomorphic triangle, with one vertex at the origin, and two idealized intersection points at the infinity of $\text{sign}(x_k)\mathbb{R}_+$ and $\text{sign}(x_k)e^{i\phi_k}\mathbb{R}_+$. In the product space \mathbb{C}^n , this gives rise to a holomorphic triangle $\Sigma_{\vec{x}}$ with boundary on $L_{\phi,A}$, Π_0 , Π_ϕ , and corners at $0, \infty_0, \infty_\phi$. Now in case some $x_k = 0$, there is still a holomorphic triangle in the product space that makes sense; the k -th projection of this triangle is simply the origin. What happens when $x_k \rightarrow 0$, is simply that the k -th projection $\pi_k(\Sigma_{\vec{x}})$ becomes very thinly concentrated near the two rays $\text{sign}(x_k)\mathbb{R}_+$ and $\text{sign}(x_k)e^{i\phi_k}\mathbb{R}_+$, and its area shrinks to zero. Moreover, for any given ε , the subset $\Sigma_{\vec{x}} \cap \pi_k^{-1}(|z| > \varepsilon)$ disappears into the infinity of \mathbb{C}^n as $x_k \rightarrow 0$. Thus when we restrict to any compact subset of \mathbb{C}^n , the holomorphic discs behave continuously as $x_k \rightarrow 0$.

Example 2.8. In the most symmetric case $\phi_1 = \phi_2 = \dots = \phi_n = \frac{\pi}{n}$, the Lawlor neck is invariant under $SO(n, \mathbb{R})$, and these holomorphic triangles are up to $SO(n, \mathbb{R})$ rotation, simply the triangle inside the first coordinate line $\mathbb{C} \subset \mathbb{C}^n$ enclosed by the three Lagrangians.

Using any of the holomorphic triangles $\Sigma_{\vec{x}}$ parametrized by S^{n-1} , we can calculate its area cohomologically by

$$\int_{\Sigma} \omega = \int_{\partial \Sigma} \lambda = \int_{\partial \Sigma \cap L_{\phi,A}} df_{L_{\phi,A}} = A,$$

which is the intuitive explanation of why A must be positive, an important ingredient of [41].

We want to draw attention also to a different aspect not explicit in [41]: that these holomorphic triangles naturally arise in an $(n-1)$ -dimensional *moduli*, rather than as isolated triangles. Consequently, the universal family of such holomorphic triangles is naturally $(n+1)$ -dimensional. A generic point on $\Pi_0 \cup \Pi_\phi \cup L_{\phi,A}$ is swept out precisely once by some $\partial \Sigma_{\vec{x}}$. When the orientations are taken into account, then the total space of this universal family gives rise to an $(n+1)$ -dimensional integration current, which provides a *bordism current* between the integration cycles of $L_{\phi,A}$ and $\Pi_0 \cup \Pi_\phi$. Producing bordism currents via universal families of holomorphic curves will be essential to our proposals concerning the Thomas–Yau conjecture.

2.4. Philosophy of open string mirror symmetry

The Thomas–Yau conjecture is largely motivated by mirror symmetry, based on the analogy between special Lagrangian geometry on the A -side (‘symplectic’), and stability conditions in the derived category of coherent sheaves on the B -side (‘holomorphic’). It is worth emphasizing that mirror symmetry is quantum by nature, and only sometimes admits classical geometric interpretations. The mathematical predictions of open string mirror symmetry chiefly fall into the following three classes.

Homological mirror symmetry (the categorical approach). Given a (compact) Calabi–Yau manifold X viewed as a symplectic manifold, one expects to find a (compact) Calabi–Yau manifold X^\vee viewed as a complex manifold (sometimes over the Novikov field), such that the (derived, idempotent closed) Fukaya category $D^\pi \text{Fuk}(X)$ is equivalent as a triangulated category to the derived category of coherent sheaves $\text{Coh}(X^\vee)$. Frequently, there is a preferred A_∞ enhancement on both sides, such that $D^\pi \text{Fuk}(X) \simeq D^b \text{Coh}(X^\vee)$ holds as an A_∞ -equivalence. The concrete implication is that Lagrangian branes L correspond to objects E of $D^b \text{Coh}(X^\vee)$, such that the Floer groups are identified with the Ext groups, which explains the name ‘homological’:

$$HF^*(L, L') \simeq \text{Ext}^*(E, E').$$

Homological mirror symmetry has many ramifications beyond Calabi–Yau manifolds, involving for instance Fano manifolds and Landau–Ginzburg models. As a general feature, homological mirror symmetry only uses pure symplectic topology (vis-à-vis complex geometry) on either side of the mirror. It is the best understood aspect of mirror symmetry, with many special cases proven.

SYZ mirror symmetry (the moduli space approach). Taken in a somewhat generalized and very optimistic sense,¹⁵ the Strominger–Yau–Zaslow philosophy says that after taking into account the quantum instanton corrections, then the moduli space of (semi)stable Lagrangian branes within fixed homology classes acquires a modified complex structure, and can be identified with certain moduli spaces of (semi)stable objects from the mirror side. The stable Lagrangian branes are believed to be related to special Lagrangians via the Thomas–Yau conjecture. Most of the literature concentrates on the restricted case of (special) Lagrangian torus fibrations arising from the large complex structure limit, in which case the moduli space of Lagrangian branes supported on the fibres inside X should recover the moduli space of points on X^\vee , namely the mirror manifold X^\vee itself.¹⁶ Beyond the torus fibration case, the moduli space of Lagrangian branes defined by the Maurer–Cartan equation typically has singularities, which is compatible with the fact that the moduli space of semistable objects in $D^b \text{Coh}(X^\vee)$ typically is very singular.

DT theory (the counting approach). In the special case of Calabi–Yau 3-folds, the moduli spaces/stacks of semistable objects have virtual dimension zero, so one can hope to extract enumerative invariants, which can be thought of as integrals over the moduli

¹⁵The SYZ paper [81] focuses primarily on special Lagrangian torus fibrations, but hints at more general special Lagrangian branes. Some more discussions can be found in Fukaya’s work [33, Section 8.5].

¹⁶This SYZ mirror construction has seen the intense research effort on the symplectic side by Auroux [8], Abouzaid, Fukaya, among many others. The very recent progress [86] achieves a non-archimedean mirror from Fukaya categorical considerations, but the nature of singular fibres is still not adequately understood. The skeptics may reasonably doubt if special Lagrangians exist at all in the region with large curvature inside the Calabi–Yau manifold. Even if they do exist, there is insufficient evidence that the fibration structure persists. This is maybe the weakest point of the SYZ conjecture.

space, the most basic version being the Euler number. Donaldson–Thomas theory is a highly elaborate framework for extracting such numbers using virtual techniques, on the side of $D^b \text{Coh}(X^\vee)$. It is suggested (e.g., in [33, Section 8.5]) that one can assign DT invariants to the moduli space of semistable objects in the Fukaya category,¹⁷ and these should be equivalent to the DT invariants on the mirror X^\vee . The role of Thomas–Yau conjecture is that in principle it transforms a problem about counting special Lagrangian objects, into a categorical framework involving Calabi–Yau A_∞ -categories and stability conditions, which then supposedly¹⁸ feeds into the grand machinery of Kontsevich and Soibelman [47–50]. This DT perspective may be regarded as the ultimate goal of Thomas–Yau conjecture,¹⁹ and is close to the physical applications involving BPS states counting, although on the A -side it is the most removed from mathematical attempts.

It is worth emphasizing that taking stability questions into account requires one to go beyond homological mirror symmetry. While the Fukaya category (vis-à-vis $D^b \text{Coh}(X^\vee)$) depends only on the symplectic (holomorphic) data, the stability condition remembers information about the holomorphic volume form (polarization class), which now sees the B -side (A -side). As such, neither the above strong version of the SYZ mirror symmetry, nor the DT mirror symmetry are consequences of homological mirror symmetry alone, as indicated already in the SYZ paper [81].²⁰ A very general framework that encompasses all the three aspects of mirror symmetry is in Kontsevich and Soibelman [47], but how to fit the A -side into this picture is wildly conjectural.

Caveats for differential geometers. However, in the differential geometric literature on the Thomas–Yau conjecture, one often attempts to go beyond all these aspects of mirror symmetry, and tries to directly compare the space of Lagrangian submanifolds inside X , with the space of holomorphic vector bundles equipped with Hermitian metrics over X^\vee . This comparison must be treated with caution, because there is no general bijective correspondence between these two infinite-dimensional spaces, and because mirror symmetry is properly a quantum phenomenon which is not fully captured by classical considerations. In fact, mirror symmetry only relates $\text{Fuk}(X)$ and $\text{Coh}(X^\vee)$ after taking the derived category, and there is no general reason for the heart of a preferred t -structure on the derived Fukaya category (e.g., the t -structure defined by some almost calibrated condition) to agree with $\text{Coh}(X^\vee)$. Having vaccinated the reader with these precautions, we shall regard such comparisons as useful formal analogies, and we consider those aspects with quantum interpretations as more reliable.

¹⁷This would presuppose suitable properness and algebraicity on the moduli space of special Lagrangian objects, which is a highly non-trivial claim related to the Thomas–Yau conjecture. For instance, in the SYZ special Lagrangian torus fibration case, compactifying the moduli space requires good understanding in the singular region, which is far beyond current knowledge.

¹⁸modulo heroic efforts

¹⁹As an analogy, the practical calculation of Donaldson’s ASD instanton invariants depends largely on the solution of Hermitian Yang–Mills equation on algebraic surfaces.

²⁰In the words of SYZ, they required the full equivalence of the two type-II string theories, including the full BPS spectrum.

2.5. Hermitian Yang–Mills

Hermitian Yang–Mills (HYM) connections have long been established as the epitome of how *stability conditions* control the existence questions of geometric PDEs, but we shall attempt to say a few new words besides the customary homage. We consider Hermitian metrics H on a holomorphic vector bundle $E \rightarrow X^\vee$ over a compact Kähler manifold X^\vee ,²¹ inducing the Chern connection ∇_H and the curvature $F_H \in \Omega^{1,1}(\text{End}(E))$. The HYM equation can be written as

$$\omega_{X^\vee}^{n-1} \wedge \frac{\sqrt{-1}}{2\pi} F_H = \frac{\mu(E)}{\int_{X^\vee} \omega_{X^\vee}^n} \text{Id}_E \otimes \omega_{X^\vee}^n, \quad \mu(E) = \frac{\int_{X^\vee} c_1(E) \wedge \omega_{X^\vee}^{n-1}}{\text{rk}(E)} = \frac{\deg(E)}{\text{rk}(E)}.$$

This implies the Yang–Mills equation $\nabla_H^* F_H = 0$, and in fact HYM connections are absolute minimizers of the Yang–Mills energy among all unitary connections on the Hermitian vector bundle E . The μ -stability (resp. semistability) means that for all proper coherent subsheaf $E' \subset E$, the slope $\mu(E') < \mu(E)$ (resp. $\mu(E') \leq \mu(E)$). The bundle E is called μ -polystable if it is a direct sum of μ -stable bundles of the same slope.

The famous Donaldson–Uhlenbeck–Yau theorem is as follows.

Theorem 2.9. *On a compact Kähler manifold, the holomorphic bundle E admits an HYM metric if and only if E is μ -polystable.*

Remark 2.9. Certain parallels between HYM and special Lagrangians are already known to Thomas and Yau. The Chern connection ∇_H is analogous to the Lagrangian submanifold L (with $U(1)$ local systems), the HYM equation as a first order equation on ∇_H is analogous to the special Lagrangian condition on L , the second order Yang–Mills equation is analogous to the minimal surface equation on L , and the Yang–Mills energy minimization is analogous to the volume minimization. If the Lagrangian L is represented as the graph of an exact 1-form, then the potential function defining L could be viewed as analogous to the Hermitian metric H . Finding a mirror analogue resembling the μ -stability was one of Thomas and Yau’s principal motivations.

μ -stability and its wider context. We now recall why the HYM equation implies μ -semi-stability. Let $E' \subset E$ be a holomorphic subbundle (or more generally, a proper coherent subsheaf). A basic feature of holomorphic geometry is that pointwise curvature decreases in subbundles: the Chern curvature $F_{E',H}$ for the restricted Hermitian metric $H|_{E'}$ satisfies

$$\sqrt{-1} F_{E',H} \leq \sqrt{-1} F_{E,H}|_{E'}.$$

Wedge both sides with $\omega_{X^\vee}^{n-1}$, taking the trace, and integrating over X^\vee , we get

$$\int_{X^\vee} c_1(E') \wedge \omega_{X^\vee}^{n-1} \leq \mu(E) \text{rk}(E'),$$

²¹If we are interested in non-compact A -model target spaces, the mirror is in fact not a compact Kähler manifold. We hope the reader will excuse us on this issue, since mirror symmetry is only used as motivations in this paper.

namely $\mu(E') \leq \mu(E)$, which is μ -semistability. Although this argument is very transparent, we wish to summarize its key features:

- Even though connections and curvatures make sense in a more general setting, we need the *integrability* of Kähler geometry to obtain pointwise *positivity*.²²
- To derive μ -stability, one integrates over X^\vee , which can be interpreted as the *moduli space* of constant maps into X^\vee .²³
- The input from complex geometry can be interpreted as a short exact sequence

$$0 \rightarrow E' \rightarrow E \rightarrow E/E' \rightarrow 0,$$

which has a categorical meaning in $D^b \text{Coh}(X^\vee)$ as a *distinguished triangle*.

- The role of the Kähler form enters via *cohomological integrals*.
- There is no need for the complex Monge–Ampère equation.

Most of these features are not specific to HYM connections, but similar arguments give rise to obstructions for a large class of PDEs involving holomorphic bundles, such as the deformed Hermitian Yang–Mills equation.

Remark 2.10. The μ -semistability condition can be recast in terms of the central charge $Z(E) = -\deg(E) + \sqrt{-1} \text{rk}(E)$, as saying $\arg Z(E') \leq \arg Z(E)$ for all non-zero proper subsheaves $E' \subset E$. It is worth emphasizing that except for the case of Riemann surfaces, μ -stability does *not* give rise to a Bridgeland stability condition on $D^b \text{Coh}(X^\vee)$, since skyscraper sheaves will generally have zero rank and zero degree, hence zero central charge. A similar but more subtle failure of Bridgeland stability happens in the context of the deformed Hermitian Yang–Mills connections (see [19, Section 4]). Such a failure does not spell doom for the PDE applications, nor for DT theoretic applications.²⁴ Even though Bridgeland stability seems to be a plausible framework for special Lagrangians in the light of Joyce’s proposal, it is probably advisable to maintain a more flexible attitude to stability conditions.

Donaldson functional. The reverse direction, that μ -stability implies the existence of HYM connections, is the hard part of the subject, and a key ingredient is the *Donaldson functional*. We make the not very essential simplification that $c_1(E) = 0$. Donaldson [26, 27] defined a functional \mathcal{M} on the infinite-dimensional space \mathcal{H} of Hermitian metrics on the fixed bundle E , by prescribing its first variation at any point $H \in \mathcal{H}$:

$$\delta \mathcal{M} = \int_{X^\vee} \text{Tr}(\sqrt{-1} F_H \delta H H^{-1}) \wedge \omega_{X^\vee}^{n-1}.$$

²²It would be interesting if the physicists can explain this positivity from supersymmetry, which is closely related to the Kähler condition.

²³Path integrals on the topological B -model typically localizes to the moduli space of constant maps. This suggests a worldsheet interpretation, which will be more apparent on the mirror side.

²⁴R. Thomas defined DT invariants for μ -stability long before the insight of Bridgeland.

Obviously from the definition, the critical points are precisely the HYM metrics. Less obviously, this functional is well defined up to an additive constant fixed by the choice of a reference Hermitian metric H_0 . Different choices are related by

$$\mathcal{M}_{H_0}(H) = \mathcal{M}_{H'_0}(H) + \mathcal{M}_{H_0}(H'_0). \quad (2.4)$$

The space \mathcal{H} can be formally assigned a Riemannian structure with non-positive curvature:

$$|\delta H|^2 = \int_{X^\vee} \text{Tr}(\delta H H^{-1})^2 \wedge \omega_{X^\vee}^n.$$

The geodesics in \mathcal{H} are given by $e^{tA}H$ where A is self adjoint with respect to H , and the Donaldson functional is convex along geodesics.

Proof idea of the Donaldson–Uhlenbeck–Yau theorem. In sketchy terms, one standard proof of the Donaldson–Uhlenbeck–Yau theorem (closest to Simpson’s approach [74]) proceeds via the heat flow. It has two principal steps:

- (i) Consider the HYM heat flow

$$\partial_t H H^{-1} = -\sqrt{-1} \Lambda F_H.$$

Using certain parabolic maximum principles, one proves long time existence by showing that all derivatives of H remain bounded for any given finite time. Furthermore, $\sup_{X^\vee} |\Lambda F|$ and $\|F\|_{L^2}$ are non-increasing in time, so remain uniformly bounded for all time. This step does not use stability.

(ii) The Donaldson functional is non-increasing in time almost by definition, so has a uniform upper bound for all time. Together with the pointwise bound on ΛF , which is like a Laplacian bound, one eventually shows that if $H(t)$ fails to be L^∞ bounded for all time, then there exists an L^2 -subbundle of E with destabilizing properties, which is then interpreted algebraically as a subsheaf. (Roughly, the destabilizing sheaf comes from the eigensubspaces of E corresponding to the small eigenvalues of $H(t)$ with respect to a fixed reference metric; compare the variational viewpoint below.) The stability condition rules out this case; one then shows that $H(t)$ actually converges smoothly at infinite time to a solution of the HYM equation.

Remark 2.11. An alternative approach by Donaldson [27] in the projective manifold case, is also based on the flow method, but uses a dimensional induction in which the stability condition appears indirectly through alternative algebro-geometric characterizations. The approach of Uhlenbeck and Yau [84] uses the continuity method instead, where the stability condition appears in a way similar to the above.

We now discuss the variational perspective to the HYM equation, even though no proofs have been constructed along such lines. One would try to compactify \mathcal{H} in some

weaker topology (which is not known),²⁵ extend the Donaldson functional to this compactification, attempt to find a minimizer of the functional, and then prove its regularity. Since the Donaldson functional is convex, it is natural to expect the existence of minimizer is equivalent to the properness of \mathcal{M} , or roughly equivalently \mathcal{M} should grow at the infinity of \mathcal{H} .

Suppose now that E fits into an extension sequence

$$0 \rightarrow E_1 \rightarrow E \rightarrow E_2 \rightarrow 0.$$

Take arbitrary Hermitian metrics H_E and H_{E_2} on E and E_2 respectively,²⁶ and regard H_{E_2} as a semi-Hermitian metric on E . We can then produce a 1-parameter family of Hermitian metrics on E , via $H(s) = e^{-s} H_E + H_{E_2}$. For $s \gg 1$, the metric $H(s)$ can be understood as equal to $e^{-s} H_E$ when restricted to E_1 , and almost equal to H_{E_2} on the orthogonal complement of E_1 . In terms of bundles with connections, in the limit $s \rightarrow \infty$ we get $E_1 \oplus E_2$ with the Chern connection for $H_E|_{E_1} \oplus H_{E_2}$. As such, it is an easy exercise to show that to leading order

$$\mathcal{M}(H(s)) \sim -2\pi s \int_{X^\vee} c_1(E_1) \wedge \omega_{X^\vee}^{n-1}.$$

Recall we have assumed $c_1(E) = 0$ to simplify the definition of the Donaldson functional. Thus the subbundle E_1 destabilizes E , precisely when the degree of E_1 is positive, so $\mathcal{M}(H(s))$ goes to $-\infty$ along this 1-parameter family.

If one wants to turn the variational approach into an actual proof, one needs to further show that all possible ways to approach the infinity of (the metric completion of) \mathcal{H} can be approximated by such 1-parameter families of algebraic origin. For our motivational purpose, it suffices to emphasize the following conceptual points:

- The interesting limiting behaviour of the Donaldson functional occurs at an *infinite distance* boundary of (the metric completion of) \mathcal{H} . This sits well with the non-positive Riemannian curvature of \mathcal{H} .
- The *stability condition controls the asymptotic behaviour* of the Donaldson functional near the boundary of \mathcal{H} .

Remark 2.12. The non-linear analysis concerning special Lagrangians is more difficult than HYM. For instance, the finite time singularities of the LMCF are inevitable. The HYM equation can be viewed as a toy model which shares some, but by no means all, of the high level features with the special Lagrangian equation.

²⁵What is known is how to compactify the space of Kähler potentials via psh functions, see the work by Boucksom [11] for its fantastic application to Kähler–Einstein metrics. In that context, the boundary at infinity is related to non-archimedean geometry.

²⁶The fact that the choice of Hermitian metrics will not ultimately matter is an expected feature, analogous to the relation between Kähler potentials and non-archimedean potential theory.

Infinite-dimensional GIT picture, and possible lack of mirror analogue. The HYM equation famously fits into a formal geometric invariant theory (GIT) framework. For this, we slightly change viewpoint, and consider the bundle E equipped with a fixed Hermitian structure, while the $(0, 1)$ -connection $\bar{\partial}_E$ encoding the holomorphic structure is allowed to vary. Each $\bar{\partial}_E$ uniquely determines the Chern connection ∇_E . The group of complex gauge transformations $GL(E, \mathbb{C})$ acts on the space of integrable $(0, 1)$ -connections, via

$$\bar{\partial}_E \mapsto g \circ \bar{\partial}_E \circ g^{-1}.$$

This action is analogous to a complex reductive group action on a finite-dimensional Kähler manifold. The subgroup $U(E)$ of unitary gauge transformations is analogous to the maximal compact subgroup. The space $GL(E, \mathbb{C})/U(E)$ can be formally identified with the space of Hermitian metrics on E . The HYM equation arises naturally from considerations of the moment map, and the Donaldson–Uhlenbeck–Yau theorem can be formally motivated from this picture [27].

This kind of infinite-dimensional GIT framework has successfully suggested the answer in many problems within Kähler geometry, so it is only natural that many people have attempted to find an analogue suitable for special Lagrangian geometry. One such attempt is as follows. Thomas [82, Section 3] considered the space

$$\mathcal{Z} = \{(L, A) : L \subset X \text{ is Lagrangian, } A \text{ is a flat } U(1)\text{-connection on } L\}$$

(not up to gauge equivalence!). The tangent space is $Z^1(L) \oplus Z^1(L)$, where $Z^1(L)$ is the space of closed 1-forms on L . This suggests an almost complex structure on \mathcal{Z}

$$J = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}.$$

With some hesitation,²⁷ Thomas attempted to complexify the Hamiltonian group action into a complex infinite-dimensional group action. Unfortunately, there seems to be no natural way to do this in general, and J is not quite an integrable complex structure. On the other hand, the moduli space of special Lagrangians with $U(1)$ -local systems does have a natural complex structure induced from J , which is however naïve in the sense that the complex structure of the moduli space of Lagrangian branes is subject to further quantum corrections due to holomorphic curves, known also as ‘worldsheet instantons’ [8].

There seems to be no agreed interpretation, but in the author’s view, this suggests the infinite-dimensional GIT framework is itself inadequate for the purpose of special Lagrangian geometry.²⁸ As we explained in Section 2.4, the main predictions of mirror symmetry are quantum in nature, and one should be cautious about taking an overly

²⁷Thomas was aware of the possible objections, and did not use the GIT analogy as the principal basis of his proposal.

²⁸Thomas’s suggestion is not the only possible way to achieve a GIT analogy. However, the other proposals [28, 57] do not exhibit the mirror analogy in the same intuitively plausible way.

classical perspective. From our perspective, one underlying reason why the infinite-dimensional GIT framework is successful in Kähler geometry, is that on the B -side one does not see the worldsheet instanton effect directly. On the A -side we have no such luxury.

As a word of console, although GIT is very good at *suggesting* the correct stability conditions for a PDE problem in Kähler geometry, it is almost never involved in the actual *proofs* of existence and uniqueness results for PDEs.

2.6. Deformed Hermitian Yang–Mills

Much recent research concentrates on a more non-linear cousin of HYM, known as the *deformed Hermitian Yang–Mills* equation (dHYM), expertly surveyed in [19]. Although the techniques involved in dHYM are more akin to other areas of Kähler geometry, such as Kähler–Einstein metrics and the J -equation,²⁹ its motivation is in part to find an improved mirror analogue of special Lagrangians, with the distant goal of constructing the Bridgeland stability on the B -side of the mirror.

The equation can be motivated differential geometrically from *semiflat mirror symmetry* [20]. To start, one imposes *toric symmetry*, so that over a local n -dimensional base D , the A and B sides are respectively the T^n -bundles T^*D/Λ and TD/Λ^\vee equipped with the canonical symplectic/complex structure, where Λ and Λ^\vee are dual lattices, so that the torus fibres are naturally dual. A *section* L of T^*D/Λ over D fibrewise determines a flat $U(1)$ -connection on the corresponding fibre of TD/Λ , and the Lagrangian condition on L is equivalent to these fibrewise connections fitting into a connection A on a line bundle E over TD/Λ , with the integrability condition $F_A^{0,2} = 0$. Now a Hessian metric on D induces a Kähler structure on both sides. The condition for L to be a special Lagrangian of phase $\hat{\theta}$, translates into the mirror condition on the holomorphic line bundle:

$$\mathrm{Im}(e^{-i\hat{\theta}}(\omega_{X^\vee} + F_A)^n) = 0.$$

Here the curvature F_A of the $U(1)$ -connection is an imaginary valued $(1, 1)$ -form. Up to rescaling ω_{X^\vee} , the global version is the dHYM equation for a Hermitian metric on a holomorphic line bundle E over X^\vee , thought of as a PDE on the potential function ϕ :

$$\mathrm{Im}(e^{-i\hat{\theta}}(\omega_{X^\vee} + \sqrt{-1}\alpha_\phi)^n) = 0, \quad [\alpha_\phi = \alpha + \sqrt{-1}\partial\bar{\partial}\phi] \in -c_1(E).$$

For an alternative view, we can pointwise simultaneously diagonalize ω_{X^\vee} and α_ϕ , to extract the eigenvalues λ_i ,

$$\omega_{X^\vee} = \frac{\sqrt{-1}}{2} \sum dz_i \wedge d\bar{z}_i, \quad \alpha_\phi = \frac{\sqrt{-1}}{2} \sum \lambda_i dz_i \wedge d\bar{z}_i.$$

²⁹Indeed, the recent breakthrough of Gao Chen [16] on dHYM is largely based on the methods he developed for the J -equation. The work of Collins et al. [18] has strong analogy with the homogeneous complex Monge–Ampère equation important for CSCK metrics.

The dHYM equation is then

$$\Theta_\omega(\alpha) = \sum_i \arctan \lambda_i = \hat{\theta} \bmod \pi \mathbb{Z}.$$

Thus the equation actually separates into several discrete branches depending on the choice of $\hat{\theta}$, and changing $\hat{\theta}$ by $k\pi$ can drastically alter the behaviour of the equations. Replacing $\hat{\theta}$ by $-\hat{\theta}$ (so E is replaced by its dual) does not change the problem, so without loss of generality $0 \leq \hat{\theta} < \frac{n\pi}{2}$. There are two significant limiting cases:

(i) When $\lambda_i \gg 1$ for any i , so $\hat{\theta} \approx \frac{n\pi}{2} - \sum \frac{1}{\lambda_i}$, the equation becomes approximately the J -equation

$$\alpha_\phi^{n-1} \wedge \omega_{X^\vee} = \frac{\int \alpha^{n-1} \wedge \omega_{X^\vee}}{\int \alpha^n} \alpha_\phi^n.$$

Most mathematical works on dHYM assume some lower bound such as $\hat{\theta} > \frac{(n-2)\pi}{2}$, which may be morally interpreted as being near this large phase limit.

(ii) When $|\lambda_i| \ll 1$ for any i , so $\hat{\theta} \approx \sum \lambda_i$, the equation is approximately the line bundle case of the HYM equation

$$\alpha_\phi \wedge \omega_{X^\vee}^{n-1} = \frac{\int \alpha \wedge \omega_{X^\vee}^{n-1}}{\int \omega_{X^\vee}^n} \omega_{X^\vee}^n.$$

This is also known as the large volume limit, in the sense that the Kähler metric length scale is much larger than the curvature scale of the bundle.

Towards a Bridgeland stability condition. The existence of dHYM solutions has algebraic obstructions. Let V be a p -dimensional subvariety of X^\vee , with $1 \leq p < n$. Under the large phase assumption $\hat{\theta} > \frac{(n-2)\pi}{2}$, a pointwise consideration of eigenvalues shows [19, Proposition 3.4]

$$\operatorname{Im} \left(\int_V e^{-\sqrt{-1}(\hat{\theta} - (n-p)\frac{\pi}{2})} (\omega_{X^\vee} + \sqrt{-1}\alpha_\phi)^n \right) > 0.$$

More suggestively, denote

$$Z_V(E) = - \int_V e^{-\sqrt{-1}\omega_{X^\vee}} ch(E) \in \mathbb{R}_{>0} e^{\sqrt{-1}\phi_V(E)}, \quad \phi_V(E) \in (0, \pi),$$

so the algebraic obstruction can be rewritten as

$$\operatorname{Im} \left(\frac{Z_V(E)}{Z_{X^\vee}(E)} \right) > 0, \quad \text{i.e., } \phi_V(E) > \phi_{X^\vee}(E) = \hat{\theta} - \frac{(n-2)\pi}{2}.$$

This bears some resemblance to a Bridgeland stability condition with central charge $Z_{X^\vee}(E)$, even though on a technical level there are some discrepancies [19, Section 3.2]. In the simplest understood examples, such as the blow-up of \mathbb{CP}^2 in a point, the obstruction criterion for dHYM seems to refine Bridgeland stability, in the sense that every dHYM stable object is Bridgeland stable, but not conversely. It is not entirely clear how to interpret this (see Remark 2.10).

Main achievements on dHYM. Some of the most significant results on the dHYM equation are as follows:

- There is a well developed formal GIT picture ([18, 21], [22, Section 2]), including a complexified group action on an infinite-dimensional Kähler manifold, a moment map, a negatively curved symmetric space, and a Donaldson-type functional which is convex along geodesics.
- The geodesic equation with prescribed boundary data have $C^{1,1}$ -solutions and viscosity solutions under a large phase assumption [19, Theorem 2.7]. This leads to non-trivial algebraic obstructions [19, Section 3].
- Assume large phase $\hat{\theta} > \frac{(n-2)\pi}{2}$ and the existence of a subsolution in the suitable sense, then the dHYM equation can be solved [19, Theorem 5.2].
- Assume large phase $\hat{\theta} > \frac{(n-2)\pi}{2}$, then the existence of dHYM solution can be formulated in terms of algebraic obstruction conditions [16].

Will dHYM lead to the mirror Bridgeland condition? Despite substantial progress, many essential difficulties still need to be overcome before the dHYM equation can give rise to a Bridgeland stability condition on $D^b \text{Coh}(X^\vee)$:

- Can one relax the *large phase assumption*?
- Is there a generalization of dHYM equation to *higher rank vector bundles*? (Near the large volume limit, the asymptotic existence criterion for a generalized dHYM equation is understood in [25], but there is no fully non-perturbative result as in the line bundle case.)
- Can this story be extended to *complexes* of vector bundles?
- How can one compare the answer with the *A*-side of the mirror?

Time will tell how far one can push in this program, but we would like to momentarily play the skeptic's advocate.

The differential geometric motivation³⁰ of dHYM assumes semiflat ambient Kähler metrics. Such metrics only arise naturally if the manifold has toric symmetry, or as a good approximate description for degenerating Calabi–Yau metrics near the large complex structure limit/large volume limit. Near this limit, the dHYM equation may be an improvement on the HYM equation as the mirror version of special Lagrangians. But far away from such limits, the instanton corrections cannot be ignored, and indeed a general Kähler metric has no toric symmetry.

The reliance on the large phase assumption is a possible indication of a breakdown once we move too far from a large phase limit. In the closely related special Lagrangian

³⁰There is an independent physics motivation for dHYM from the Dirac–Born–Infeld action. The DBI action is however not an exact result, but depends on the assumption that certain derivative terms of the curvature can be ignored. Unlike the *A*-model side where the central charge $Z(L) = \int_L \Omega$ is believed to hold exactly, on the *B*-side the central charge formula $Z_{X^\vee}(E) = -\int_{X^\vee} e^{-\sqrt{-1}\omega_{X^\vee}} ch(E)$ is believed to be subject to worldsheet instanton corrections.

graph equation, relaxing the phase condition is known to result in rather severe singularities for the viscosity solutions.

Being the section of a torus fibration is a serious assumption on the topology of a Lagrangian.

The difficulty with higher rank vector bundles is a possible indication that $\text{Coh}(X^\vee)$ may be not the natural abelian subcategory of $D^b \text{Coh}(X^\vee)$ suited for dHYM. Furthermore, there is no a priori guarantee for arbitrary stability conditions to correspond to PDEs,³¹ or to have any classical geometrical interpretation at all. In particular, the B -side mirror to the hypothetical special Lagrangian stability condition may well be an abstract stability condition with no particular PDE interpretation.

Due to the lack of imagination, it is hard to see how PDEs can know about the degree shifts of a complex of vector bundles, and how it can detect homotopy equivalence of complexes.

If we take this skeptical view, then we do not expect an exact comparison between the hypothetical Bridgeland stabilities on both sides of the mirror, without adding very substantial assumptions such as toric symmetry. But when HYM shares the same qualitative features as dHYM, which admit a mirror interpretation, it lends more plausibility for these features to appear also on the special Lagrangians.

2.7. Extensions and wall crossing

Part of Thomas and Yau's insights is that one should focus on the *categorical* aspect of mirror symmetry when we look for an analogy between the A -side and the B -side. Their starting observation is that *Lagrangian connected sums are analogous to bundle extensions* [82, Sections 3 and 4].

Extension bundle vs. Lagrangian connection sum. An essential aspect of $\text{Coh}(X^\vee)$ is that new bundles can be constructed from extensions of known bundles E_1, E_2 , namely

$$0 \rightarrow E_1 \rightarrow E \rightarrow E_2 \rightarrow 0.$$

Such extension sequences are classified by the complex vector space $\text{Ext}^1(E_2, E_1)$. Due to the \mathbb{C}^* -scaling, the choice of E is parametrized by the projective space $\mathbb{P}(\text{Ext}^1(E_2, E_1))$. In general, extensions are *not symmetric* in E_1 and E_2 . The extensions

$$0 \rightarrow E_2 \rightarrow E \rightarrow E_1 \rightarrow 0$$

are classified by $\text{Ext}^1(E_1, E_2)$, which is a quite different space. Extensions can also be viewed as distinguished triangles in $D^b \text{Coh}(X^\vee)$.

³¹If a Bridgeland stability condition arises from a PDE, there is no a priori guarantee that its deformations also come from PDEs.

In the mirror picture, exact sequences do not make a priori sense, but one can talk about distinguished triangles, whose geometric sources are the graded Lagrangian connected sums $L_1 \# L_2$, fitting into a distinguished triangle

$$L_1 \rightarrow L_1 \# L_2 \rightarrow L_2 \rightarrow L_1[1].$$

Such distinguished triangles are classified by $HF^1(L_2, L_1)$.³² Again there is a scaling symmetry related to the neck size of the Lagrangian connected sum, and there is an asymmetry between $L_1 \# L_2$ and $L_2 \# L_1$.

This analogy is a prime example of *homological mirror symmetry*. On either side, only pure complex geometry/pure symplectic geometry appears.

Wall crossing. Wall crossing in the categorical context refers to the following phenomenon when the stability condition varies in a 1-parameter family, with central charges Z_t . For $t > 0$, the objects E_1, E_2, E in the extension sequence are all stable, so necessarily

$$\arg Z_t(E_1) < \arg Z_t(E) < \arg Z_t(E_2), \quad Z_t(E) = Z_t(E_1) + Z_t(E_2).$$

At $t = 0$, the phase angles become equal, and for $t < 0$, the phase angle inequality is reversed, and E becomes unstable. This phase alignment occurs on a codimension one locus in the space of stability condition, and thus they are called walls. Every extension sequence potentially gives rise to a wall, and the walls can be dense in general.

A notable special case is μ -stability of bundles for a 1-parameter family of Kähler classes $[\omega_t]$, and the slopes $\mu(E_1), \mu(E_2), \mu(E)$ become equal precisely for $t = 0$. On one side of the wall, the HYM connections exist on E_1, E_2, E , and on the other side E becomes unstable and no longer admits any HYM connection.

Thomas [82] interpreted a gluing construction of Joyce as the mirror analogue of the wall crossing phenomenon for bundles.³³ In the simplest case, let $n \geq 3$, fix a symplectic structure ω , and vary the holomorphic volume form Ω_t in a 1-parameter family while keeping the almost Calabi–Yau condition. Let L_t^1, L_t^2 be smooth special Lagrangians with respect to Ω_t with phase θ_t^1 and θ_t^2 , intersecting transversely at precisely one point p_t , with Floer degree $\mu_{L_t^2, L_t^1}(p_t) = 1$, such that $\theta_t^2 - \theta_t^1$ increases past zero at $t = 0$, and in particular $\theta_0^1 = \theta_0^2$. Thus at $t = 0$, the tangent planes can be put into the standard form inside \mathbb{C}^n .³⁴

$$T_p L_0^1 = (e^{i\phi_1}, \dots, e^{i\phi_n})\mathbb{R}^n, \quad T_p L_0^2 = \mathbb{R}^n, \quad 0 < \phi_k < \pi, \quad \sum \phi_k = \pi,$$

$$\omega = \frac{\sqrt{-1}}{2} \sum dz_k \wedge d\bar{z}_k, \quad e^{-i\theta_0^1} \Omega = a^n dz_1 \wedge \dots \wedge dz_n, \quad a > 0.$$

³²The caveat is that unlike bundles, the neck length of the Lagrangian connected sum cannot be arbitrarily large.

³³It is quite remarkable that Thomas and Yau knew before Bridgeland, that stability conditions make sense categorically beyond μ -stability, and the mirror of the hypothetical special Lagrangian stability condition does not need to be μ -stability.

³⁴The constant a comes from the almost Calabi–Yau structure, and $a = 1$ in the Calabi–Yau case.

This provides the appropriate framing data to topologically glue in a Lawlor neck $L_{\phi,A}$ (see Section 2.3) to desingularize $L_t^1 \cup L_t^2$ for small t , so that the glued Lagrangian has the topology $L_t^1 \# L_t^2$. Joyce [43, Theorem 9.10] shows that for each small $t > 0$, the glued Lagrangian can be perturbed into a special Lagrangian of phase $\theta_t \approx \theta_0^1$. As $t \rightarrow 0$, these special Lagrangians converge as currents to $L_t^1 \cup L_t^2$, while for $t < 0$ this gluing strategy does not produce any new special Lagrangian.

The central charges are

$$Z(L_t^i) = \int_{L_t^i} \Omega_t^i = R_t^i e^{\sqrt{-1}\theta_t^i}, \quad Z(L_t^1 \# L_t^2) = Z(L_t^1) + Z(L_t^2) = R_t e^{\sqrt{-1}\theta_t}.$$

In particular, $R_t^1 \sin(\theta_t^1 - \theta_t) = -R_t^2 \sin(\theta_t^2 - \theta_t)$. The asymmetry between $t > 0$ and $t < 0$ in Joyce’s gluing construction comes from an approximate formula for the Lawlor neck parameter valid for small t ,³⁵

$$R_t^1 \sin(-\theta_t^1 + \theta_t) = a^m A > 0.$$

Thus $\theta_t^1 < \theta_t$ is needed in the gluing construction. This is strongly reminiscent of a Bridgeland stability condition. What happens when t crosses zero can be interpreted as wall crossing.

To summarize, the mirror analogy of wall crossing phenomenon is of categorical nature. However, it goes beyond homological mirror symmetry as soon as it involves the stability conditions, and the new problems involve analytical aspects, beyond purely topological issues. The similarity between the Joyce gluing and the bundle case is a strong motivation for Thomas and Yau. However, the Joyce analysis is only valid in a perturbative regime, and what is missing here is a non-perturbative understanding of when the $D^b \text{Fuk}(X)$ class of the Lagrangian connected sum can admit special Lagrangians.

2.8. Space of almost calibrated Lagrangians

We now return to the A -side of the mirror and describe the work of J. Solomon [77, 78], generally accepted as the canonical picture on the subject. While dHYM is motivated by the dream of a correspondence (semiflat mirror symmetry) between the two sides of the mirror at the level of classical objects, Solomon is interested in the structural similarities between the infinite-dimensional spaces involved in the mirrors, and especially in finding analogues for the HYM equation.³⁶ The reader is thus invited to keep in mind the comparison with Section 2.5.

Let \mathcal{L}^+ be the space of almost calibrated compact immersed Lagrangians in an almost Calabi–Yau manifold (X, ω, Ω) ,

$$\mathcal{L}^+ = \left\{ \iota: L \rightarrow X \mid -\frac{\pi}{2} < \theta < \frac{\pi}{2} \right\}.$$

³⁵For a heuristic short derivation see [42, Section 6]. Beware that Joyce’s Lagrangian connected sum has the opposite convention.

³⁶Solomon’s work predates the substantial works on the dHYM equation.

For the benefit of intuition, we shall loosely identify the Lagrangian immersion with its image. Solomon [78] considers the space \mathcal{O} of Lagrangians which are exact isotopic (aka. local Hamiltonian isotopic)³⁷ within \mathcal{L}^+ to a fixed Lagrangian. It should be borne in mind that unlike the space \mathcal{H} of Hermitian metrics on a bundle, the space \mathcal{O} may have very non-trivial topology; we call its universal cover $\tilde{\mathcal{O}}$. The formal deformations of the Lagrangians are given by the Hamiltonian functions $h: L \rightarrow \mathbb{R}$, up to the ambiguity of an additive constant. Solomon proposes to fix the constant by the normalization condition $\int_L h \operatorname{Re} \Omega = 0$, and assigns a formal Riemannian metric on

$$T_L \mathcal{O} = \left\{ h: L \rightarrow \mathbb{R} \mid \int_L h \operatorname{Re} \Omega = 0 \right\},$$

via the formula

$$\langle h, k \rangle = \int_L h k \operatorname{Re} \Omega.$$

This is positive definite because $\operatorname{Re} \Omega$ is a volume form on L by almost calibratedness. The main result of [78] is a computation on the Riemannian curvature of \mathcal{O} , which is found to be *non-positively curved*, similar to the HYM setting. A further paper [79] computes the geodesic equation with respect to this formal metric, and reinterprets a geodesic between $L_1, L_2 \in \mathcal{O}$ in terms of a 1-parameter family of special Lagrangians (of phase $\frac{\pi}{2}$ instead!) with boundary on $L_0 \cup L_1$.

Another major aspect of Solomon's work is to look for an analogue of the Donaldson functional. The definition of this Solomon functional does not really require the almost calibrated condition, even though some of the main properties do. Choose some appropriate $\hat{\theta} \in (-\frac{\pi}{2}, \frac{\pi}{2})$ so that $\int_L \operatorname{Im}(e^{-i\hat{\theta}} \Omega) = 0$. Now take a 1-parameter family of Lagrangians L_t in \mathcal{O} with associated Hamiltonian functions $h_t: L_t \rightarrow \mathbb{R}$. Solomon defines

$$\mathcal{S} = \int_0^1 dt \int_{L_t} h_t \operatorname{Im}(e^{-i\hat{\theta}} \Omega).$$

His main theorem is the following.

Theorem 2.10 ([77]). *The functional \mathcal{S} is independent of Hamiltonian deformations of the path of Lagrangians fixing the two ends. In particular, by fixing the starting Lagrangian L_0 , we obtain a functional \mathcal{S} of the endpoint Lagrangian, which is well defined on the universal cover $\tilde{\mathcal{O}}$.*

Like the Donaldson functional, this well definition is non-trivial. It is however obvious that the critical points in \mathcal{O} are precisely special Lagrangians of phase $\hat{\theta}$. Furthermore, we have the following theorem.

³⁷Exact isotopies of immersed Lagrangians differ from global Hamiltonian isotopies, in that the Hamiltonian function on the Lagrangians may depend on the local sheets, so may not always extend smoothly to the ambient space.

Theorem 2.11 ([77]). *Assume $\hat{\theta} = 0$. Then the second variation of \mathcal{S} at a critical point is positive semidefinite. Furthermore, along a geodesic with respect to Solomon’s formal Riemannian metric, the functional \mathcal{S} is convex.*

The analogy with Donaldson’s picture in Section 2.5 should be quite clear.

Limitations. Unlike Thomas and Yau who based their bet primarily on the Floer theoretic or categorical aspects, which are closer to the quantum world of topological field theories, Solomon’s picture is predominantly classical, and its chief limitation comes from fixing the topological type of the Lagrangian:

- (i) There is no appearance of the brane structure, or the role of holomorphic curves.
- (ii) Solomon works with exact isotopic Lagrangians, but the Thomas–Yau argument suggests it is more natural to work in a derived Fukaya category class.
- (iii) The Solomon functional is only well defined by passing to a highly non-trivial universal cover. In the very special case where ω and $\text{Im}(e^{i\hat{\theta}}\Omega)$ are exact forms on X , Solomon gave a formula [77, Theorem 1.3] that shows his functional is well defined on \mathcal{O} . We view the exactness on $\text{Im}(e^{i\hat{\theta}}\Omega)$ as too strong an assumption for applications.
- (iv) The infinite-dimensional Riemannian structure is incomplete in a much more severe way compared to the B -side analogues. This means that in non-pathological examples, we can reach the boundary of the exact isotopy class within finite distance in the Solomon metric, such that the Solomon functional remains finite.

This is geometrically very significant. In the LMCF approach, this would strongly suggest the formation of finite time singularity, which is a major difference with the HYM case. In the variational viewpoint, this incompleteness would negate all the favourable arguments from the convexity of the functional and the non-positivity of curvature, and suggest instead that the exact isotopy class is not an adequate framework for finding special Lagrangians. We will discuss later that a more promising variational framework needs to incorporate Lagrangians from the same *derived Fukaya category class*, not just the same exact isotopy class.

Exact isotopy class versus derived category class. The example below is closely related to the most symmetric case of the Lawlor necks (see Section 2.3). It is also morally related to the Lawlor neck pinching singularity in the Joyce program [45, Section 3.5].

Example 2.12. Consider two almost calibrated Lagrangians L^1, L^2 with a unique intersection point p , and there is a Darboux chart around p modelled on $B_1 \subset \mathbb{C}^n$, such that inside the chart L^1, L^2 the local setup is

$$L^1 = e^{i\pi/n}\mathbb{R}^n, \quad L^2 = \mathbb{R}^n, \quad \omega = \frac{\sqrt{-1}}{2} \sum dz_k \wedge d\bar{z}_k, \quad \Omega \approx \prod dz_k.$$

Let $(L_t)_{0 < t \ll 1}$ be a 1-parameter family of Lagrangian connected sums with neck length $O(t)$, which all agree with $L_1 \cup L_2$ except in a compact subset in B_1 . Inside B_1 , we take

the Ansatz

$$L_t = \{(t\gamma(s)x_1, \dots, t\gamma(s)x_n) \mid x_1^2 + \dots + x_n^2 = 1\},$$

where the curve $\gamma(s): \mathbb{R} \rightarrow \mathbb{C}$ can be chosen so that L_t is almost calibrated and agrees with $L_1 \cup L_2$ outside $B_{1/2}$. Clearly, L_t are related by scaling inside B_1 . The Hamiltonian vector field along L_t , which is really a section of $(TX/TL_t)|_{L_t}$, agrees with $(\gamma(s)x_1, \dots, \gamma(s)x_n)$ inside B_1 and is zero outside.³⁸ The corresponding Hamiltonian function h_t is t^2 times a smooth function of one variable s ; a small caveat is that h_t converges to two generally different constants along L_1 and L_2 . Thus it takes finite distance in the Solomon metric to reach the limit $t \rightarrow 0$, and the Solomon functional remains finite, but the topology changes from L_t to $L^1 \cup L^2$.

Remark 2.13. The Hamiltonian functions h_t along L_t can be extended to global functions on X , with

$$\|h_t\|_{C^0} = O(t^2), \quad \|dh_t\|_{C^0} = O(t), \quad \|\nabla^2 h_t\|_{L^\infty} \leq C.$$

But in the $t \rightarrow 0$ limit, the second derivatives fail to be continuous at the origin, and indeed L_t changes topology in the limit.

In this example the essential failure is the breakdown of smoothness. In view of Joyce's program, this suggests that the remedy is to allow for (Floer theoretically unobstructed) almost calibrated Lagrangians connected to each other not just by exact isotopies, but also surgeries such as Lagrangian connected sums. In these transitions the derived category class of the Lagrangian is unchanged, and the Thomas–Yau argument suggests the $D^b \text{Fuk}(X)$ class is a natural framework to look for special Lagrangian representatives. The following fundamental question is thus relevant for the compatibility between the geometric and the categorical perspectives:

Question 2. When are two almost calibrated unobstructed Lagrangian branes isomorphic in $D^b \text{Fuk}(X)$ connected by exact isotopies with surgeries?

Remark 2.14. The Joyce program suggests that running the LMCF would result in a sequence of exact isotopies and surgeries, to connect the initial Lagrangian to its infinite time limit, which one hopes to be the unique representative of the Harder–Narasimhan decomposition. In the almost calibrated case, there is no ‘collapsing zero object’ in this process for homological reasons, so the surgeries should be continuous in the geometric measure theory sense. Since any two such Lagrangians within the same $D^b \text{Fuk}(X)$ class are expected to flow to the same limit, they are supposedly connected to each other through a continuous family of unobstructed Lagrangians.

³⁸This is consistent because near the boundary of B_1 , the position vector is a tangent vector of L , so vanishes in the quotient TX/TL_t .

2.9. Totally real geometry

There is another interesting framework due to Lotay and Pacini [57, 58], which makes the holomorphic curves appear on the forefront, and exhibits good analogy with the classical GIT picture, by enlarging the space of Lagrangians into the space of *totally real submanifolds*.

Let (M, ω, J) be a Kähler manifold. A real n -dimensional submanifold L is called *totally real*, if at every point $T_p L$ is transverse to $JT_p L$. Lotay and Pacini introduce a formal principal bundle $\mathcal{P} \rightarrow \mathcal{T}$, where \mathcal{P} is the space of totally real immersions $\iota: L \rightarrow X$ isotopic to a given immersion, and \mathcal{T} is formally its quotient by the orientation preserving diffeomorphism group $\text{Diff}^+(L)$ of L .³⁹ There is a horizontal distribution, which at each point $\iota \in \mathcal{P}$ assigns the transverse bundle $JT\iota(L)$, so gives a way to lift tangent vectors from \mathcal{T} to \mathcal{P} . They then define a *geodesic* to be a curve ι_t in \mathcal{P} , such that the tangent vector field $\partial_t \iota_t$ is parallel with respect to the horizontal distribution. As a caveat, here the word ‘geodesic’ does not suggest a Riemannian metric. An alternative characterization [57, Lemma 2.2] of a geodesic, is a 1-parameter family of totally real submanifolds $\iota_t: L \rightarrow X$, such that there is a fixed vector field $Y \in \Gamma(L, TL)$, with

$$\frac{d}{dt} \iota_t = J \iota_{t*} Y, \quad [\iota_{t*} Y, J \iota_{t*} Y] = 0.$$

Geometrically, one can imagine $(\iota_t)_{0 \leq t \leq 1}$ sweeps out an $(n+1)$ -dimensional submanifold with boundary, foliated into complexified integral curves of Y , which are *holomorphic curves* inside X .

Lotay and Pacini also define the J -volume functional on \mathcal{T} . Pointwise on a totally real submanifold L , a real cotangent vector of L corresponds to a $(1, 0)$ -form in $\Lambda^{(1,0)} X|_L$, so taking the n -th wedge power, we have a canonical isomorphism between $\Lambda^n L \otimes_{\mathbb{R}} \mathbb{C}$ and $K_X|_L$. Now the Kähler structure on X induces a Hermitian metric on K_X , and pointwise on L an element of K_X with unit Hermitian norm uniquely specifies a volume form $d\text{vol}_J$ on L . Their J -volume functional is

$$\text{Vol}_J(L) = \int_L d\text{vol}_J.$$

It is easy to show Vol_J provides a lower bound $\text{Vol}_J(L) \leq \text{Vol}(L)$ to the Riemannian volume of L , with equality precisely when L is Lagrangian.

In case X is Calabi–Yau, this construction is particularly transparent. The totally real condition is equivalent to the pointwise non-vanishing of Ω when restricted to L , and

$$\left| \int_L \Omega \right| \leq \text{Vol}_J(L) = \int_L e^{-i\theta} \Omega \leq \text{Vol}(L),$$

³⁹Lotay and Pacini did not worry about analytical issues involving the quotient; their picture is entirely formal.

where the phase factor is chosen to make $e^{-i\theta}\Omega$ an orientation form on L . Consequently, any totally real submanifold in a Calabi–Yau manifold with $\theta = \hat{\theta}$ constant, with no need for the Lagrangian condition, is an absolute minimizer of the Vol_J functional. This enormous space of critical points is closely related to the non-ellipticity of the critical point equation. The situation is somewhat better in a negative Kähler–Einstein ambient space, where the critical points coincide with minimal Lagrangians [57, Section 5.5].

One of their main results is the following.

Proposition 2.13 ([57, Theorem 5.10]). *In a Kähler–Einstein ambient manifold with non-positive Ricci curvature, the J -volume functional is convex along the geodesics.*

There is also a formal GIT picture [57, Section 6]: to some extent the space \mathcal{P} can be viewed as the infinitesimal complexification of $\text{Diff}^+(L)$, with the space of orbits \mathcal{T} . The J -volume functional is formally a Kähler potential on \mathcal{P} , and the $\text{Diff}^+(L)$ -moment map gives rise to critical points of the J -functional.

Limitations. From the viewpoint of special Lagrangian geometry, the limitations of the Lotay–Pacini picture are as follows:

(i) The Lagrangians are largely relegated to the back stage. As clear from the above, in the Calabi–Yau case the space of critical points is too enormous. Their framework may be more useful in the negative Kähler–Einstein case, but the lack of ellipticity is a severe obstacle.

(ii) There is no attempt to link up with Floer theory. As such, they lack satisfactory existence criteria for the holomorphic curves, despite making some limited progress in special cases. In fact, Lotay and Pacini’s geodesics seem too oversimplified from the Floer theoretic viewpoint: one needs to address transversality questions in general, and holomorphic disc breaking should occur in moduli spaces of dimension ≥ 1 .

2.10. Analogy with tunnelling effect

Note. *This section is meant to be purely inspirational, and its aim is to present some analogies and comparisons with no claim to physical accuracy.*

Notice the Thomas–Yau argument is a little mysterious from the following classical perspective: how could two Lagrangians in possibly different Hamiltonian isotopy classes communicate with each other? This situation seems conceptually similar to the phenomenon of tunnelling in quantum mechanics: there may be several minima of the classical potential function separated by potential wells, but there is a non-trivial quantum amplitude for the particles to move in between, thereby removing the ground state degeneracy. In the Thomas–Yau setup, we have two putative special Lagrangians, which are analogous to the energy minima points, and the Thomas–Yau uniqueness statement is similar to the removal of degeneracy, resulting in a unique ground state.

The physicists tell us that branes are dynamical objects and can fluctuate. If so, it might make sense to ask about the amplitude for a brane to start with a given configuration and

end with another. Now if Lagrangian branes behave like classical particles moving on an infinite-dimensional space such as Solomon’s space \mathcal{O} , one might expect Solomon’s formal picture to be relevant for describing this amplitude, and the incompleteness of \mathcal{O} would suggest a non-trivial amplitude to tunnel outside \mathcal{O} to another Hamiltonian isotopy class. As a conflicting viewpoint, the use of Floer theory in the Thomas–Yau argument suggests the Lagrangian branes communicate by the strings stretched between them, so one might expect the amplitudes to be computed in terms of worldsheet integrals. Lotay and Pacini’s geodesics fit this viewpoint better, and have the major advantage of making sense outside a given Hamiltonian isotopy class.

Tunnelling effects made a famous appearance in Witten’s interpretation of Morse theory [85] in terms of supersymmetric quantum mechanics, with an eye towards applications in quantum field theory. In Witten’s context, the tunnelling amplitude between two critical points of a Morse function is to leading order proportional to e^{-S} , where S is proportional to the difference of the two critical values. Now Solomon’s functional has some similarity with a Morse function whose only critical points are minima. In Sections 3.2 and 3.7 we will *unify* the Solomon functional with the Lotay–Pacini geodesics, in the special case of exact Lagrangians. Could the Solomon functional have any physical interpretation in terms of the logarithm of the tunnelling amplitudes?⁴⁰

3. Moduli space of holomorphic curves

The theme of this section is the *bordism currents* produced from the $(n + 1)$ -dimensional universal family over the $(n - 1)$ -dimensional moduli spaces of holomorphic curves. This theme unifies the search for the Floer theoretic obstructions, and the problem of extending the Solomon functional to the derived Fukaya category setting. One of our central techniques is *integration over the moduli spaces*, to prove both identities and inequalities. The primary setting of this section is exact graded compact Lagrangians inside Calabi–Yau Stein manifolds, with occasional comments on the compact Calabi–Yau case.

The *Floer theoretic obstructions* are necessary conditions to the existence of special Lagrangians in the exact and almost calibrated setting. We will give a large number of a priori heuristic principles to place very stringent constraints on what kind of obstructions we are looking for, then prove the Floer theoretic obstructions under the extra hypotheses of automatic transversality and a positivity condition (see Section 3.5), and then compare our picture to Joyce’s proposal on Bridgeland stability (see Section 3.6). The converse direction, in search of *sufficient conditions* for the existence of special Lagrangians, will be laid out in Section 5.

⁴⁰As a grain of salt, Witten’s interpretation concerns supersymmetric QFT, while the Thomas–Yau picture concerns the dynamics of Lagrangian branes, which is a target space perspective on string theory, which lacks a fundamental path integral formulation. As such, our suggested physical interpretation is not logically rigorous, but only a guess.

We also present two perspectives on extending the Solomon functional to unobstructed Lagrangians within the same $D^b \text{Fuk}(X)$ class. The more elementary perspective (see Section 3.2) works only in the exact setting, and implies the first variation formula under Hamiltonian isotopies in a rather straightforward manner. The second perspective (see Section 3.7) represents the Solomon functional as an integral over the moduli space of holomorphic curves, which we hope can be generalized to the compact Calabi–Yau setting. Section 3.8 contains more applications of the moduli integral technique.

Remark 3.1. The two *assumptions* of automatic transversality (which roughly means that no perturbation of almost complex structure is needed, see Section 3.3) and the positivity condition (see Section 3.4) will feature prominently in the major results of this section. Correspondingly, we will not dwell on the detail of the perturbation schemes aimed at overcoming the transversality issues, which are painstakingly carried out in [4, 14, 15, 67, 68]. The expert readers may convert to their favourite schemes as they prefer. For more on our clockwise conventions such as gradings and signs, which differ from many symplectic texts, see the appendix.

3.1. Lotay–Pacini picture revisited

From a Floer theoretic perspective, the main insight of Lotay and Pacini (see Section 2.9) is that given two Lagrangians L, L' decorated with suitable brane structures, one should look for a *family* of holomorphic curves with boundary on L and L' , which pass through any generic point of L and L' . In their formal picture, such families are called ‘geodesics’. What Lotay and Pacini did not provide is a good existence criterion for their geodesics. Now, even though we will soon specialize to a much simpler setting, we wish to explain how their geodesics fit into the Thomas–Yau–Joyce picture.

In the setup of Bridgeland stability, the central charge is defined as a homomorphism from the Grothendieck group of a triangulated category to \mathbb{C} , which factorizes through a finitely generated lattice, viewed as a numerical Grothendieck group. In view of the application to special Lagrangians, the triangulated category is $D^b \text{Fuk}(X)$, and the numerical Grothendieck group should be a subgroup of the homology group modulo torsion $H_n(X, \mathbb{Z})/\text{tors}$. In particular,

- if two Lagrangian branes L, L' define the same object in $D^b \text{Fuk}(X)$, then this picture predicts them to lie in the same homology class in $H_n(X)$;
- if $L_1 \rightarrow L \rightarrow L_2 \rightarrow L_1[1]$ is a distinguished triangle, then $[L_1] + [L_2] = [L]$ in $H_n(X)$.

This means Floer theory must provide a bordism current between L, L' (resp. L and $L_1 \cup L_2$), namely an $(n+1)$ -dimensional integration current \mathcal{C} with $\partial\mathcal{C} = L - L'$ (resp. $\partial\mathcal{C} = L - L_1 - L_2$).⁴¹ In Floer theory, cycles are constructed from moduli spaces of

⁴¹The boundary of integration currents do not detect contributions from supports of small enough dimension. In this respect they are similar to pseudocycles, although when we generalize to Lagrangians with very weak regularity later, the language of currents may be more natural.

holomorphic curves and evaluation maps, so this current \mathcal{C} should come from families of holomorphic curves, possibly with highly sophisticated perturbations and virtual techniques. This suggests the Lotay–Pacini geodesics, construed in this homological sense, should be part of the Fukaya category foundation necessary for the Thomas–Yau–Joyce program.

High brow viewpoint. The claim that *the K -theory of the derived Fukaya category of a symplectic Calabi–Yau manifold (exact with suitable convexity at infinity, or compact) factorizes through homology $H_n(X)$ modulo torsion* seems well known to symplectic topology experts, although a precise reference seems rather difficult to find. We now sketch a high brow viewpoint explained to the author by P. Seidel and S. Rezkikov, and will later explain in more detail a more pedestrian approach in the exact setting. The claim is a formal consequence of the existence of maps

$$\begin{aligned} K_0(D^b \text{Fuk}(X)) &\xrightarrow{\text{ch}_0} HH_0(D^b \text{Fuk}(X)), \\ HH_0(D^b \text{Fuk}(X)) &\xrightarrow{OC} QH^n(X; \Lambda_{\text{nov}}) \simeq H_n(X) \otimes \Lambda_{\text{nov}}. \end{aligned}$$

Here $HH_0(D^b \text{Fuk}(X))$ is the Hochschild homology in degree zero. The first map sends the K -theory class of an unobstructed Lagrangian brane L (compact, graded, oriented, with spin and bounding cochain structure) to the unit

$$1_L \in HF^0(L, L) \rightarrow HH_0(D^b \text{Fuk}(X));$$

the well definition of this map is an essentially algebraic fact. Suppose L, L' are isomorphic in $D^b \text{Fuk}(X)$, then there are closed morphisms $\alpha \in CF^0(L, L')$, $\beta \in CF^0(L', L)$ whose derived category compositions are equal to $1_L \in HF^0(L, L)$, $1_{L'} \in HF^0(L', L')$ in cohomology. The Hochschild differential of $\beta \otimes \alpha$ exhibits $1_L - 1_{L'}$ as a coboundary in the Hochschild chain complex, so $1_L = 1_{L'}$ in $HH_0(D^b \text{Fuk}(X))$. Some additional calculation shows the compatibility with distinguished triangles.

The second map is a special case of the open-closed string map, and in general requires working over the Novikov field. One then needs the claim that 1_L is sent to the homology class $[L] \in H_n(X)$, without quantum correction. The intuitive meaning of the open-closed string map is to consider holomorphic discs with boundary on L with an unconstrained boundary marked point, and find the cycle in X traced out by an interior marked point. The claim amounts to saying that the only contribution comes from constant maps. Unfortunately, the author is unable to locate a general reference. Granted this claim, we would get by composition a map from $K_0(D^b \text{Fuk}(X))$ to $H_n(X; \Lambda_{\text{nov}})$ which sends the K -theory class of L to the homology class $[L]$.

Remark 3.2. The claim that the non-constant holomorphic discs do not affect the image of 1_L in $QH^*(X; \Lambda_{\text{nov}})$ under the open-closed map, has the following intuitive explanation, communicated to the author by P. Seidel. The moduli space of holomorphic discs with an interior marked point constrained to a given n -cycle in X , together with an unconstrained boundary marked point, admits an S^1 -action coming from the rotation of the

disc. Unless the holomorphic disc is constant, the S^1 action has finite stabilizer, which suggests that the virtual count contribution from non-constant discs should vanish, hence by Poincaré duality the contribution of non-constant discs to the open-closed map vanishes.

Without assuming this claim, then one has the weaker conclusion that the image of 1_L under the open-closed map is $[L] \in H_n(X, \Lambda_{\text{nov}})$, modulo terms with strictly positive Novikov exponents (due to the positive energy of the non-constant discs). Supposing that L and L' are isomorphic in $D^b \text{Fuk}(X)$, the above discussion would imply that $[L] = [L']$ modulo terms with strictly positive Novikov exponent, and by comparing the leading order term, we would still deduce $L = L'$.

In Sections 3.1.1 and 3.1.2 we will give a more concrete perspective on why L is homologous to L' in the exact setting.

3.1.1. The exact embedded Lagrangian case. We specialize to the setting of Stein manifolds, and all Lagrangians are assumed to be exact, graded and compact, and in particular carry an orientation (see the Appendix for some basic Floer theory). The local systems have holonomy in \mathbb{R} , \mathbb{Q} or \mathbb{Z} . We consider two transverse embedded Lagrangians L, L' in the *same derived Fukaya category class*. By definition, we have closed morphisms $\alpha \in CF^0(L, L')$ and $\beta \in CF^0(L', L)$; we sometimes view the Lagrangian intersections in β as degree n outputs. Moreover, in terms of the product structure on cohomology

$$\begin{cases} HF^0(L, L') \otimes HF^0(L', L) \rightarrow HF^0(L', L'), \\ HF^0(L', L) \otimes HF^0(L, L') \rightarrow HF^0(L, L), \end{cases}$$

the composition $\alpha \circ \beta = 1_{L'}$ and $\beta \circ \alpha = 1_L$. These conditions completely characterize isomorphism in $D^b \text{Fuk}(X)$. Our goal is to explain the following result.

Proposition 3.1. *There is a bordism current \mathcal{C} such that $\partial \mathcal{C} = L - L'$ in the sense of currents.*

The bordism current will be constructed from *universal families of (perturbed) holomorphic strips* with boundary on L and L' , and with ends at α and β (meaning that the ends of the strip converge to intersection points of L, L' of degree 0 and n respectively, and α, β encode the weighting factors to the contribution of these intersection points). We will assume all the usual transversality assumptions in Floer theory are satisfied, so the compactified moduli space $\overline{\mathcal{M}}(\alpha, \beta)/\mathbb{R}$ of perturbed holomorphic strips up to domain translation is a smooth manifold with boundary and corners, of dimension $(n - 1)$. The notation really stands for a formal sum of many moduli spaces, coming from the summands of α, β . The universal family is fibred over this moduli space $\overline{\mathcal{M}}(\alpha, \beta)/\mathbb{R}$, whose fibres are the solutions to the Cauchy–Riemann equation (with domain dependent perturbations of the almost complex structure), which we call perturbed holomorphic curves. The fibres over the boundary of the moduli space are broken holomorphic curves. The orientation on the universal family is induced from the complex orientation on Σ and the

orientation on the moduli space, up to an extra minus sign (see Example A.2 for conventions). Upon evaluation to X we obtain an $(n + 1)$ -dimensional current \mathcal{C} .

Remark 3.3. If the Fukaya category is defined over \mathbb{Z} , then all the weighting factors to the various universal families are all *integers*, and \mathcal{C} is naturally an *integral current*. If we use Fukaya categories over \mathbb{Q} or \mathbb{R} instead, then \mathcal{C} is only guaranteed to be a finite \mathbb{Q} (resp. \mathbb{R}) linear combination of integral currents.

Our main task is to understand the boundary of \mathcal{C} . There are two sources of boundaries as follows:

- The holomorphic curves themselves have boundary along $L \cup L'$. This boundary contribution is always supported on $L \cup L'$.
- The compactified moduli spaces have boundary due to holomorphic strip breaking.

In schematic notation, the boundary of the moduli space is described by

$$\partial(\overline{\mathcal{M}(\alpha, \beta)/\mathbb{R}}) = \bigcup_r \overline{\mathcal{M}(\alpha, r)/\mathbb{R}} \times \overline{\mathcal{M}(r, \beta)/\mathbb{R}}.$$

Here r can range from all intersection points of degree between 1 and $n - 1$. The notation $\overline{\mathcal{M}(\alpha, \beta)/\mathbb{R}}$ stands for a weighted sum of moduli spaces, with weighting coming from the holonomy factors of the local systems.

Next comes a crucial observation. Although $2 \leq \deg r \leq n - 2$ give rise to boundaries of the moduli space $\overline{\mathcal{M}(\alpha, \beta)/\mathbb{R}}$, their contributions to $\partial\mathcal{C}$ are contained in the universal families associated to $\overline{\mathcal{M}(\alpha, r)/\mathbb{R}}$ and $\overline{\mathcal{M}(r, \beta)/\mathbb{R}}$. These smaller moduli spaces have dimension at most $n - 3$, and the corresponding universal families have dimension at most $n - 1$. By rectifiability considerations, the n -dimensional current $\partial\mathcal{C}$ cannot receive contributions from at most $(n - 1)$ -dimensional supports, so such disc breakings do not contribute to $\partial\mathcal{C}$.

Now for $\deg r = 1$, the moduli spaces $\overline{\mathcal{M}(\alpha, r)/\mathbb{R}}$ are zero-dimensional, so their presence merely amounts to some counting factors. The condition for α to be *closed* in $CF^0(L, L')$ is equivalent to the weighted count of $\overline{\mathcal{M}(\alpha, r)/\mathbb{R}}$ being zero. This weighted sum appears as the coefficient of the n -dimensional current defined by the universal family over $\overline{\mathcal{M}(r, \beta)/\mathbb{R}}$. Thus we see that $r \in CF^1(L, L')$ does not contribute to $\partial\mathcal{C}$. Similarly, the condition for β to be closed in $CF^0(L', L)$ implies that $r \in CF^1(L', L)$ (alternatively viewed as degree $n - 1$ intersections from L to L') does not contribute to $\partial\mathcal{C}$. In summary, $\partial\mathcal{C}$ must be an integration cycle supported on $L \cup L'$.

Since $\partial\mathcal{C}$ is itself the boundary of a current, it must be closed. This explains why $\partial\mathcal{C}$ is a constant linear combination of the integration cycle of L and L' , instead of some non-trivial function times these cycles. The constant coefficients can be pinned down by counting the number of holomorphic strips passing through a given generic point on L (resp. L'), and the choice of the generic point does not matter. Such counts are precisely the geometric interpretation of the Floer product $HF^0(L, L') \otimes HF^0(L', L) \rightarrow HF^0(L', L')$, and $HF^0(L', L) \otimes HF^0(L, L') \rightarrow HF^0(L, L)$ (see Example A.1). When

the moduli space orientations are taken into account, we obtain $\partial\mathcal{C} = L - L'$ (see Example A.2 for an exposition on signs).

Remark 3.4 (Homological uniqueness of the bordism current). Some auxiliary perturbation data goes into the construction of \mathcal{C} due to the need to ensure transversality. If we fix L, L' , but change the domain dependent almost complex structures, then the difference of two bordism currents $\mathcal{C} - \mathcal{C}'$ has zero boundary in the sense of currents. Recall that Stein manifolds have the homotopy type of a CW complex of dimension $\leq n$, and thus $H_k(X) = 0$ for $k \geq n + 1$, so $\mathcal{C} - \mathcal{C}'$ must be the boundary of an $(n + 2)$ -dimensional current. For an alternative viewpoint, this $(n + 2)$ -dimensional current can be concretely provided by parametrized families of pseudoholomorphic curves (see Remark 3.6).

3.1.2. Immersed case. Still working in the exact setting, we now allow $(L, b), (L', b')$ to be unobstructed immersed Lagrangians with transverse self intersections. Assume they intersect transversally, and define isomorphic objects in $D^b \text{Fuk}(X)$. We now wish to explain why Proposition 3.1 should continue to hold even in the immersed setting, without delving too deep into the specifics of the perturbation schemes and transversality issues. For some background on the immersed Floer theory, see Section A.2.

The isomorphism condition gives us closed morphisms $\alpha \in CF^0(L, L')$ and $\beta \in CF^0(L', L)$ whose cohomological compositions give the identities. At the chain level,

$$m_2^{b,b'}(\beta, \alpha) = 1_L - m_1^b(\gamma), \quad m_2^{b',b}(\alpha, \beta) = 1_{L'} + m_1^{b'}(\gamma')$$

where $1_L, 1_{L'}$ stand for the geometric units (represented by a sum of local maximum points of Hamiltonian functions on L, L' respectively), and γ, γ' are elements in $CF^{-1}(L, L), CF^{-1}(L', L')$ respectively. Notice in the almost calibrated case, γ, γ' would be both zero, since there are no self intersections of degree -1 .

As before, the bordism current shall be constructed from the universal family of (perturbed) holomorphic curves with boundary on L and L' . But instead of working only with holomorphic strips, we need holomorphic polygons with corners not only at intersection points in α, β , but also at points in b, b' . In addition to the holomorphic strip moduli space $\mathcal{M}(\alpha, \beta)/\mathbb{R}$, we also need the moduli space of polygons $\mathcal{M}(b, \dots, b, \alpha, b', \dots, b', \beta)$, and $\mathcal{M}(b, \dots, b, \gamma), \mathcal{M}(b', \dots, b', \gamma')$. The notation here is a shorthand for a weighted sum of many moduli spaces of polygons. Since the bounding cochain elements have Floer degrees one, these moduli spaces all have dimension $n - 1$. The energy of the polygons satisfies the topological formula (A.4), so by the Novikov positivity requirement of bounding cochains, there is a uniform a priori energy bound once $\alpha, \beta, \gamma, \gamma'$ are given, whence there are in fact only finitely many moduli spaces involved. Each moduli space provides a universal family of holomorphic curves, and the sum of all the contributions defines an $(n + 1)$ -dimensional current \mathcal{C} . For sign conventions, see Section A.2 and Example A.2.

The boundary of \mathcal{C} comes from two sources: the boundary of the individual holomorphic curves which lie on $L \cup L'$, and the boundary of the compactified moduli spaces. In the exact setting, there are no sphere bubbles. As in the embedded case, for support

dimension reasons, the boundaries of the compactified moduli space that can contribute to $\partial\mathcal{C}$, is caused by curve breaking into two pieces arising in 0 and $(n - 2)$ -dimensional moduli spaces. The cancellation of these contributions is very similar to the standard argument for the Floer differential to square to zero (see Section A.2).

For breakings at a nodal point mapping to $CF^1(L, L')$ (resp. $CF^1(L', L)$), the contributions vanish due to the closedness condition $m_1^{b,b'}(\alpha) = 0$ (resp. the closedness of β).

For breakings at degree 2 self intersection point on L (resp. L'), the contributions vanish due to the Maurer–Cartan equation on b (resp. b').

A new way of disc breaking/splitting, is at a degree zero self intersection r on L (the L' case being entirely similar). The discs of $\mathcal{M}(b, \dots, b, \alpha, b', \dots, b', \beta)$ can break into a virtual dimension zero disc with input at $b, \dots, \alpha, b', \dots, \beta, b, \dots, b$ and output at r , and a virtual dimension $n - 1$ disc with input corners at b, \dots, b, r . On the other hand, the discs of $\mathcal{M}(b, \dots, b, \gamma)$ can break into a virtual dimension zero disc with input at $b, \dots, \gamma, b, \dots, b$ and output at r , and a virtual dimension $n - 1$ disc with input corners at b, \dots, b, r . These two effects cancel out.

After the cancellation of all moduli space boundaries, the only contributions $\partial\mathcal{C}$ are supported in $L \cup L'$. As in the embedded case, $\partial\mathcal{C}$ is locally a constant multiple of the underlying cycles of L, L' . The interpretation of the geometric unit pins down $\partial\mathcal{C} = L - L'$ as in the embedded case.

Example 3.2. If L and L' are disjoint Lagrangian branes which both define the zero object in $D^b\text{Fuk}(X)$, then α, β are both zero, and \mathcal{C} comes entirely from the γ, γ' contributions. Of course, zero Lagrangian objects have zero homology class, which cannot happen in the almost calibrated case.

Remark 3.5. If there are degree -2 self intersections of L , then the choice of γ is only unique up to m_1^b of some element in $CF^{-2}(L, L)$. The corresponding choice of \mathcal{C} would be ambiguous by the boundary of an $(n + 2)$ -dimensional integration current. As a closely related issue, our conditions on α, β are merely cohomological, so in general we can adjust α and β by coboundary terms, which would affect \mathcal{C} also by the boundary of an $(n + 2)$ -dimensional current. If we impose L, L' to be almost calibrated, then there are no $CF^{-1}(L, L')$ elements to begin with, and these phenomena do not happen.

On the other hand, \mathcal{C} still depends on the choice of local systems and bounding cochains, which may contribute non-trivial holonomy factors. Gauge equivalent choices affect \mathcal{C} by the boundary of an $(n + 2)$ -dimensional current. One may naturally ask the following questions.

Question 3. Up to gauge equivalence of bounding cochains and local systems, is there an optimal representative of \mathcal{C} ?

Question 4. Given an exact isotopy with surgery between L and L' among unobstructed Lagrangians, is there a preferred choice of \mathcal{C} (see Question 2)?

Distinguished triangles. In our convention, an immersed Lagrangian can be made up of *several connected components*. A prototypical situation is when L' is the union of two immersed Lagrangians L_1 and L_2 , with some degree one intersections in $CF^1(L_2, L_1)$ arising as part of the bounding cochain data of L' . When the brane structure is taken into account, we can view L' as a twisted complex built from L_1, L_2 (with bounding cochains b_1, b_2 suppressed in the notation) and a closed morphism $\gamma \in CF^1(L_2, L_1) = CF^0(L_2[-1], L_1)$. Inside $D^b \text{Fuk}(X)$,

$$L \simeq L' \simeq \begin{pmatrix} (L_2, b_2) & \\ \gamma & (L_1, b_1) \end{pmatrix}.$$

We have a distinguished triangle

$$L_2[-1] \xrightarrow{\gamma} L_1 \rightarrow \text{Cone}(\gamma) \rightarrow L_2,$$

and $L \simeq L' \simeq \text{Cone}(\gamma)$. Rotating the triangles, we get another *distinguished triangle*

$$L_1 \rightarrow L \rightarrow L_2 \xrightarrow{\gamma} L_1[1].$$

The bordism current between L and L' is an $(n+1)$ -dimensional integration current with $\partial\mathcal{C} = L - L_1 - L_2$. In particular, this explains that the Grothendieck group of $D^b \text{Fuk}(X)$ should factorize through $H_n(X)$.

Here is a more geometric perspective on the bordism currents arising from distinguished triangles, which is very close to Thomas and Yau's original viewpoint, where the fundamental phenomenon is *Lagrangian breaking*. In the simplest case, we can imagine L is isomorphic in $D^b \text{Fuk}(X)$ to the Lagrangian connected sum $L_1 \# L_2$ (beware our convention for $L_1 \# L_2$ is the same as Thomas–Yau [82] but different from many symplectic texts), so that we can construct a bordism current between L and $L_1 \# L_2$. Now when $L_1 \# L_2$ deforms, the Lagrangian handle part can shrink, and in the limit $L_1 \# L_2$ can break into two components $L_1 \cup L_2$ (see Example 2.12). The bordism current between L and $L_1 \cup L_2$ should simply be the limit of the sequence of bordism currents. This picture illustrates that even when the topology of the Lagrangians can change under non-smooth convergence, the bordism currents should persist in a continuous way.

One can proceed with the case of *many Lagrangians*, namely we take the immersed Lagrangian L' to be the *twisted complex* (see Section A.2)

$$\begin{pmatrix} (L_N, b_N) & & & & \\ b_{N,N-1} & (L_{N-1}, b_{N-1}) & & & \\ \cdots & & & & \\ b_{N,1} & b_{N-1,1} & \cdots & b_{2,1} & (L_1, b_1) \end{pmatrix}. \quad (3.1)$$

In this case, assuming L is isomorphic to L' in $D^b \text{Fuk}(X)$, the bordism current between L and L' amounts to a bordism current between L and $L_1 \cup L_2 \cup \cdots \cup L_N$.

This multi-Lagrangian situation is built out of many distinguished triangles: for $1 \leq k \leq N$, let \mathcal{E}_k be the immersed Lagrangian $L_1 \cup \cdots \cup L_k$ corresponding to the twisted complex

$$\left(\begin{array}{cccccc} (L_k, b_k) & & & & & \\ b_{k,k-1} & (L_{k-1}, b_{k-1}) & & & & \\ \cdots & & & & & \\ b_{k,1} & b_{k-1,1} & \cdots & b_{2,1} & (L_1, b_1) \end{array} \right). \quad (3.2)$$

Then (suppressing bounding cochains in the notation) we have a sequence in $D^b \text{Fuk}(X)$,

$$0 = \mathcal{E}_0 \rightarrow \mathcal{E}_1 \rightarrow \cdots \rightarrow \mathcal{E}_N \simeq L,$$

with distinguished triangles

$$\mathcal{E}_{i-1} \rightarrow \mathcal{E}_i \rightarrow L_i \rightarrow \mathcal{E}_{i-1}[1].$$

The morphism from L_i to \mathcal{E}_{i-1} comes from $b_{i,j}$ for $j < i$. This setup should be reminiscent of Harder–Narasimhan decompositions (2.2), although at this stage we have not yet brought in stability conditions, which shall be discussed further in Section 3.6.

3.2. Solomon functional revisited

Let (X, ω, Ω) be an almost Calabi–Yau Stein manifold, and L be an exact Lagrangian brane, with $\int_L \text{Im}(e^{-i\hat{\theta}} \Omega) = 0$ for some suitable $\hat{\theta} \in (-\frac{\pi}{2}, \frac{\pi}{2})$. Our goal is to suggest how the Solomon functional may be well defined without the universal cover issue, and extended beyond a given exact isotopy class. Both issues are essential for the variational approach to the Thomas–Yau conjecture, see Section 5.

Homological nature of the Solomon functional. Write the Liouville 1-form as λ , so $d\lambda = \omega$, and the potential of the immersed Lagrangian L as f_L , so $df_L = \lambda|_L$. We consider the potential as part of the brane data, so adding a constant to f_L is viewed as a different Lagrangian brane. We shall consider a path of such Lagrangians L_t , with associated Hamiltonian functions h_t , so there is a preferred way to parallel transport f_L , as recalled below.

Lemma 3.3. *We have*

$$\int_0^1 dt \int_{L_t} h_t \text{Im}(e^{-i\hat{\theta}} \Omega) = \int_{L_t} f_{L_t} \text{Im}(e^{-i\hat{\theta}} \Omega)|_{t=0}^t - \int_{\cup_t L_t} \lambda \wedge \text{Im}(e^{-i\hat{\theta}} \Omega). \quad (3.3)$$

Proof. Let X_t be the Hamiltonian vector field along L_t associated to h_t , namely $dh_t = \omega(X_t, \cdot)$. We calculate the time derivative of f_{L_t} along L_t

$$\mathcal{L}_X \lambda = \iota_X d\lambda + d(\iota_X \lambda) = \iota_X \omega + d(\iota_X \lambda) = d(h_t + \iota_X \lambda),$$

so there is a preferred parallel transport of f_L along the path L_t ,

$$\partial_t f_{L_t} = h_t + \iota_X \lambda.$$

Hence,

$$\partial_t \int_{L_t} f_{L_t} \operatorname{Im}(e^{-i\hat{\theta}} \Omega) = \int_{L_t} (h_t + \iota_X \lambda) \operatorname{Im}(e^{-i\hat{\theta}} \Omega) + \int_{L_t} f_{L_t} \mathcal{L}_X \operatorname{Im}(e^{-i\hat{\theta}} \Omega).$$

Now by the Cartan formula and the closedness of Ω ,

$$\mathcal{L}_X \operatorname{Im}(e^{-i\hat{\theta}} \Omega) = d\iota_X \operatorname{Im}(e^{-i\hat{\theta}} \Omega),$$

so after integration by part,

$$\int_{L_t} f_{L_t} \mathcal{L}_X \operatorname{Im}(e^{-i\hat{\theta}} \Omega) = - \int_{L_t} df_t \wedge \iota_X \operatorname{Im}(e^{-i\hat{\theta}} \Omega) = - \int_{L_t} \lambda \wedge \iota_X \operatorname{Im}(e^{-i\hat{\theta}} \Omega).$$

Combining the above,

$$\partial_t \int_{L_t} f_{L_t} \operatorname{Im}(e^{-i\hat{\theta}} \Omega) = \int_{L_t} h_t \operatorname{Im}(e^{-i\hat{\theta}} \Omega) + \int_{L_t} \iota_X (\lambda \wedge \operatorname{Im}(e^{-i\hat{\theta}} \Omega)).$$

Integrating in t gives the result. ■

We observe that by the Kähler condition $\omega \wedge \Omega = 0$, so

$$d(\lambda \wedge \operatorname{Im}(e^{-i\hat{\theta}} \Omega)) = \omega \wedge \operatorname{Im}(e^{-i\hat{\theta}} \Omega) = 0.$$

This means the term $\int_{\cup_t L_t} \lambda \wedge \operatorname{Im}(e^{-i\hat{\theta}} \Omega)$ is a homological quantity, in the sense that we can replace $\cup_t L_t$ by any compactly supported $(n+1)$ -current \mathcal{C} with $\partial \mathcal{C} = L_1 - L_0$, which would automatically satisfy

$$[\mathcal{C} - \cup_t L_t] = 0 \in H_{n+1}(X)$$

since $H_{n+1}(X) = 0$ for Stein manifolds. In particular, this explains Solomon's theorem that his functional is invariant under Hamiltonian deformations of the path of Lagrangians. The advantage of our homological interpretation is to allow more general currents \mathcal{C} , which in particular can come from families of holomorphic curves.

Proposed extension of the Solomon functional. Taking the homological interpretation of (3.3) as starting point, a natural way to extend the Solomon functional is to make use of the bordism current \mathcal{C} between an unobstructed Lagrangian L and a fixed unobstructed reference Lagrangian L_0 . We have $\partial \mathcal{C} = L - L_0$ as currents, and \mathcal{C} comes from the universal family of holomorphic curves. Our proposed formula is

$$\mathcal{S}(L) = \int_L f_L \operatorname{Im}(e^{-i\hat{\theta}} \Omega) - \int_{L_0} f_{L_0} \operatorname{Im}(e^{-i\hat{\theta}} \Omega) - \operatorname{Im} \int_{\mathcal{C}} \lambda \wedge e^{-i\hat{\theta}} \Omega. \quad (3.4)$$

A few conceptual points are in order.

(i) We emphasize that this depends not only on the underlying Lagrangian, but also on the potential f_L .

(ii) There is no need to pass to any universal cover in the space of Lagrangians, as in Solomon’s work (see Section 2.8).

(iii) The topology of L is no longer fixed, and in particular the Hamiltonian isotopy class may change.

(iv) We view (3.4) as a unification of the very different viewpoints of Solomon and Lotay–Pacini. In Section 2.10 we suggested that this extended Solomon functional may be relevant for quantum tunnelling amplitudes between the branes L_0 and L .

(v) Suppose we vary the Lagrangian L within a 1-parameter exact isotopy family of unobstructed Lagrangians L_t . The bordism currents \mathcal{C}_t between L_t and L_0 satisfy

$$\mathcal{C}_{t_2} = \mathcal{C}_{t_1} + \cup_{t_1 \leq t \leq t_2} L_t \text{ modulo exact } (n+1)\text{-dim currents,}$$

then the computation in Lem 3.3 proves the *first variation formula for the Solomon functional*

$$\frac{d}{dt} \mathcal{S}(L_t) = \int_{L_t} h_t \operatorname{Im}(e^{-i\hat{\theta}} \Omega), \quad (3.5)$$

which is of course the defining feature of the Solomon functional. Consequently, the formula (3.4) extends Solomon’s definition in our exact setting, and fixes the multi-valuedness problem (i.e., the need to pass to universal covers) in Solomon’s work.

Change of reference Lagrangian. The definition of the Solomon functional depends on the reference Lagrangian L_0 , and we write $\mathcal{S}_{L_0}(L)$ when we wish to emphasize this dependence. The following feature of the Solomon functional resembles the Donaldson functional in the HYM context (see (2.4)).

Proposition 3.4. *Under the change of reference Lagrangians,*

$$\mathcal{S}_{L_0}(L) = \mathcal{S}_{L'_0}(L) + \mathcal{S}_{L_0}(L'_0). \quad (3.6)$$

Proof. We shall use the homological nature of the Solomon functional and the fact that $H_{n+1}(X) = 0$. We pick $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$ such that

$$\partial \mathcal{C}_1 = L - L'_0, \quad \partial \mathcal{C}_2 = L'_0 - L_0, \quad \partial \mathcal{C}_3 = L - L_0.$$

Then $\mathcal{C}_1 + \mathcal{C}_2$ is homologous to \mathcal{C}_3 , so we can replace \mathcal{C}_3 by $\mathcal{C}_1 + \mathcal{C}_2$ to compute $\mathcal{S}_{L_0}(L)$, whence (3.6) follows. ■

Remark 3.6. A more Floer theoretic argument that $\mathcal{C}_1 + \mathcal{C}_2$ is homologous to \mathcal{C}_3 , which does not appeal to $H_{n+1}(X) = 0$ directly, can be sketched as follows. We assume L_0, L'_0, L are three unobstructed Lagrangians mutually isomorphic in $D^b \operatorname{Fuk}(X)$, and

$$HF^{-1}(L_0, L_0) = 0.$$

Of course, the self Floer cohomologies of L_0, L'_0, L are all isomorphic, and $HF^{-1} = 0$ is a necessary condition if the $D^b \operatorname{Fuk}(X)$ class admits any almost calibrated representative

at all. We consider $\alpha \in CF^0(L_0, L'_0)$, $\beta \in CF^0(L'_0, L)$, $\gamma \in CF^0(L, L_0)$ representing the generators of HF^0 , such that at the level of Floer cohomology

$$\gamma \circ \beta \circ \alpha = 1_{L_0}, \quad \alpha \circ \gamma \circ \beta = 1_{L'_0}, \quad \beta \circ \alpha \circ \gamma = 1_L.$$

For simplicity we first assume almost calibratedness, so that $CF^{-1} = 0$, and there is no ambiguity for these generators. Notice the compositions $\beta \circ \alpha$, $\gamma \circ \beta$, $\alpha \circ \gamma$ provide generators of $HF^0(L_0, L)$, $HF^0(L'_0, L_0)$, $HF^0(L, L'_0)$. Consider the n -dimensional moduli spaces $\tilde{\mathcal{M}}$ of holomorphic discs with corners at α, β, γ and the self intersection points corresponding to the bounding cochains. The corresponding universal family $\tilde{\mathcal{C}}$ provides an $(n+2)$ -dimensional current, whose boundary comes from disc bubbling and disc breaking. Most of the boundary contributions are eliminated by the Maurer–Cartan equation of the bounding cochains, the closedness of α, β, γ , and support dimension reasons, and only three boundary contributions survive. These are the $(n+1)$ -dimensional bordism currents between L_0, L'_0 (resp. L'_0, L and L, L_0) constructed from the universal family of holomorphic curves associated to the generators $\alpha, -\gamma \circ \beta$ (resp. $\beta, -\alpha \circ \gamma$ and $\gamma, \beta \circ \alpha$). We can identify these as $\mathcal{C}_2, \mathcal{C}_1, -\mathcal{C}_3$. The upshot is that Floer theory explicitly provides the $(n+2)$ -dimensional current that exhibits the homological relation between $\mathcal{C}_1 + \mathcal{C}_2$ and \mathcal{C}_3 .

In general without assuming almost calibratedness, then CF^{-1} can be non-zero. Then we need some extra n -dimensional moduli spaces to account for the non-uniqueness of cohomological representatives of HF^0 , an issue quite similar to Section 3.1.2. A subtle new issue is that the moduli space $\tilde{\mathcal{M}}$ receives new boundary contributions involving the m_3^b products (this shorthand notation indicates the presence of bounding cochain elements, see (A.8)) of α, β, γ . The three cyclic permutations of α, β, γ produce three m_3^b products, which are elements in $CF^{-1}(L_0, L_0)$, $CF^{-1}(L'_0, L'_0)$ and $CF^{-1}(L, L)$ respectively, and the $(n-1)$ -dimensional moduli of polygons with one corner at the CF^{-1} intersections and the other corners at bounding cochain elements contribute to $\partial\tilde{\mathcal{C}}$. Now by the A_∞ relation, and the closedness of α, β, γ ,

$$m_1^b(m_3^b(\gamma, \beta, \alpha)) + m_2^b(\gamma, m_2^b(\beta, \alpha)) - m_2^b(m_2^b(\gamma, \beta), \alpha) = 0.$$

Writing

$$m_2^b(\gamma, m_2^b(\beta, \alpha)) = 1_{L_0} + m_1^b(\delta_1) \quad \text{and} \quad m_2^b(m_2^b(\gamma, \beta), \alpha) = 1_{L_0} + m_1^b(\delta_2),$$

we see $m_3^b(\gamma, \beta, \alpha) + \delta_1 - \delta_2$ is m_1^b -closed, so by the assumption that $HF^{-1}(L_0, L_0) = 0$, it is in fact $-m_1^b(\varepsilon_1)$ for some $\varepsilon_1 \in CF^{-2}(L_0, L_0)$. We can then produce an n -dimensional moduli space, from polygons with a corner at ε_1 , and other corners at the bounding cochain elements. Completely analogously, one can produce two other n -dimensional moduli spaces from $\varepsilon_2 \in CF^{-2}(L'_0, L'_0)$ and $\varepsilon_3 \in CF^{-2}(L, L)$. Combining the $(n+2)$ -dimensional universal families over the n -dimensional moduli spaces, results in an explicit bordism current between $\mathcal{C}_1 + \mathcal{C}_2$ and \mathcal{C}_3 .

3.3. Automatic transversality

In our later applications, it is not enough to just have a bordism current between Lagrangians L, L' constructed from perturbed pseudoholomorphic curves. Two additional conditions are desirable: *automatic transversality* and *positivity condition*. These are natural properties of the highly idealized picture of Lotay and Pacini (see Section 2.9), but may seem rather strong for Floer theorists.

In this section we discuss various sufficient conditions for automatic transversality, which intuitively means that the bordism current \mathcal{C} is constructed without perturbing the integrable complex structure. This requires that the (extended) linearized Cauchy–Riemann operator is surjective, namely the moduli space is regular. The next section will discuss the positivity condition. Complex integrability and the existence of holomorphic volume form Ω will be assumed throughout. All holomorphic curves are assumed to be non-constant.

Automatic transversality. There exist a finite collection of $(n - 1)$ -dimensional *smooth* moduli spaces of holomorphic curves $u: \Sigma \rightarrow X$ with respect to the *integrable complex structure*, constructed from the inputs in $CF^0(L, L')$, $CF^0(L', L) \simeq CF^n(L, L')^\vee$ and the bounding cochain data as in Section 3.1, such that by taking the weighted sum of the $(n + 1)$ -dimensional universal families of holomorphic curves, we obtain a current \mathcal{C} with $\partial\mathcal{C} = L - L'$. This bordism current \mathcal{C} agrees with the bordism currents constructed from generically perturbed almost complex structures, up to the boundary of an $(n + 2)$ -dimensional current.

Moreover, considering the boundary of holomorphic curves $\partial\Sigma$ varying in the $(n - 1)$ -dimensional regular moduli spaces, we obtain n -dimensional universal families, sweeping out the cycle $L - L'$; we require the evaluation map from these n -dimensional universal family to $L \cup L'$ to be *immersions*, except at the corner points of $\partial\Sigma$ mapping to the Lagrangian intersections, where the failure of immersion is ‘minimal’ (see below for details). We say that the bordism current \mathcal{C} consists purely of ‘automatically transverse curves’.

Automatic transversality, weak version. We can allow certain holomorphic curves

$$u: \Sigma \rightarrow X$$

arising in $(n - 1)$ -virtual dimensional moduli spaces, which are not automatically transverse, subject to the following requirements on these extra bad curves:

- (1) When virtual perturbation theory is taken into account, $\partial\mathcal{C} = L - L'$ still holds.
- (2) At any such bad curve $u: \Sigma \rightarrow X$, given any $n - 1$ first order deformation vector fields, the 1-form $\Omega(\cdot, v_1, \dots, v_n)$ restricted to Σ vanishes identically. Intuitively, this means Ω vanishes identically on the Zariski tangent space of the universal family at $u: \Sigma \rightarrow X$. As a caveat, these Zariski tangent spaces may be higher-dimensional.

- (3) The boundary evaluation $u: \partial\Sigma \rightarrow L \cup L'$ for all such bad holomorphic curves is contained in some subset of $L \cup L'$ of Hausdorff dimension $\leq n - 1$. As such, at almost every point on $L \cup L'$, only automatically transverse curves pass through it.
- (4) The Solomon functional can be computed by integrating only on the part of \mathcal{C} consisting of automatically transverse curves.

The automatic transversality assumption should be viewed as a higher-dimensional generalization of the fact that on Riemann surfaces, the non-trivial holomorphic polygons are immersions up to the boundary (see [71, Section 13 (b)]). The intuition for the weak version is that we sometimes need extra holomorphic curves to maintain $\partial\mathcal{C} = L - L'$, but for questions related to the Solomon functional and the boundary evaluation to the Lagrangians, these extra curves do not contribute.

Index theory preliminary. For a pseudoholomorphic polygon $u: \Sigma \rightarrow X$ with inputs at p_1, \dots, p_k and an output at q , arranged in clockwise order, the index is

$$\deg q - \sum_{k=1}^k \deg p_k,$$

where the degree convention is (A.1). The index amounts to a Maslov number computation, and an alternative topological description is as follows: take a section s of the complex line bundle $\Lambda^n TM \rightarrow \Sigma$, which restricts on $\partial\Sigma$ to a section of the real line bundle $\Lambda^n TL$. (When several Lagrangians are involved, it is understood that TL refers to the appropriate Lagrangian on the portion of $\partial\Sigma$.) We assume s has isolated zeros up to the boundary and the corner (aka. strip like ends). We may also regard s as a function $\Omega(s)$ on the polygon, by contraction with Ω . Then

$$\begin{aligned} \text{Index} = 2 \sum (\text{interior zeros}) + \sum (\text{boundary zeros}) \\ + \sum (\text{excess corner zeros}) + n. \end{aligned} \quad (3.7)$$

Here the order of zeros is computed from winding numbers, and for general sections s may take positive and negative values. At a corner where $\partial\Sigma$ passes from L_+ to L_- in the clockwise direction, we can put the tangent spaces $TL_{\pm} \subset TX$ into the standard form respecting the complex structure

$$L_+ = \mathbb{R}^n, \quad L_- = (e^{i\phi_1}, \dots, e^{i\phi_n})\mathbb{R}^n, \quad 0 < \phi_i < \pi, \quad TX = \mathbb{R}^n \otimes \mathbb{C},$$

so if $\Omega(s) \sim z^\alpha$ in the complex coordinate of the upper half plane model, the excess vanishing order at the corner is $\alpha - \frac{1}{\pi} \sum_1^n \phi_i$. Formula (3.7) is equivalent to the standard index formula by a topological version of the Cauchy residue formula.

3.3.1. Automatically transverse cases.

Holomorphic strip. We first consider the holomorphic strip case with input p and output q , and the *integrability* of the complex structure will be important. The *first order deformations* of the holomorphic strips Σ are given by *holomorphic sections* of $TX|_{\Sigma}$, which takes boundary value in TL over $\partial\Sigma$, and decays at the corners.

Lemma 3.5. *If v_1, \dots, v_n are first order deformation vector fields, then either*

$$\Omega(v_1, \dots, v_n) = 0$$

everywhere on Σ , or we must have $\deg q - \deg p \geq n$, and when the equality holds then $\Omega(v_1, \dots, v_n)$ only vanishes at the ends with excess vanishing order zero.

Proof. We have a section of $\Lambda^n TX|_{\Sigma}$ given by $v_1 \wedge \dots \wedge v_n$, which takes boundary value in $\Lambda^n TL$ on $\partial\Sigma$. Since v_1, \dots, v_n are all *holomorphic*, so must be the function $\Omega(v_1, \dots, v_n)$. Assume this function is not identically zero. By holomorphicity, the zeros are isolated. We claim that the order of zeros must be non-negative everywhere. This is clear for the interior and the boundary points. We analyze the ends of the strip as the origin in the upper half plane model with holomorphic coordinate z , putting TL_{\pm} in the standard form at the corner point. The deformation vector field has the leading asymptote

$$v_k = (a_{k1}z^{\phi_1/\pi}, \dots, a_{kn}z^{\phi_n/\pi}) + O(z), \quad k = 1, 2, \dots, n,$$

hence

$$\Omega(v_1, \dots, v_n) = z^{(\sum \phi_k)/\pi} (\det(a_{kj}) + o(1)),$$

and the excess vanishing order is non-negative. By the index formula (3.7), the index $\deg q - \deg p \geq n$, and when equality is achieved all vanishing orders must be zero. In particular, $\det(a_{kj}) \neq 0$ at the corners. ■

Corollary 3.6 (Automatic transversality, strip case). *Suppose that $\deg q - \deg p = n$. If v_1, \dots, v_n are \mathbb{R} -linearly independent at some point on $\partial\Sigma$ away from the two corners, then v_1, \dots, v_n span the space of all first order deformations, the obstruction space vanishes, and the moduli space is smooth at $u: \Sigma \rightarrow X$. Moreover, the holomorphic strip is an immersion up to the boundary.*

Proof. Since v_1, \dots, v_n are \mathbb{R} -linearly independent at a point on $\partial\Sigma$, they span TL at the point, so $\Omega(v_1, \dots, v_n) \neq 0$. By the lemma above v_1, \dots, v_n are pointwise complex linearly independent as sections of the holomorphic vector bundle TX over Σ , so any holomorphic first order deformation can be written as

$$v = f_1 v_1 + \dots + f_n v_n.$$

The functions f_1, \dots, f_n are holomorphic on Σ up to boundary, and even up to corners due to $\det(a_{kj}) \neq 0$. Now subtracting a constant linear combination of v_1, \dots, v_n , we

can ensure v vanishes at any chosen point on $\partial\Sigma$. Then $\Omega(v, v_2, \dots, v_n)$ has a zero, so must be identically zero by the above Lemma, whence $f_1 = 0$ identically. Similar all $f_k = 0$, so $v = 0$. This proves that v_1, \dots, v_n span all first order deformations. Since the index is n , and the first order deformation space is n -dimensional, we must have *vanishing obstruction space*.

There is a special deformation vector field from \mathbb{R} translation. The non-vanishing result then implies that the holomorphic strip is an immersion up to boundary. At the corners, the holomorphic strip is to leading order

$$(a_1 z^{\phi_1/\pi} + O(z), \dots, a_n z^{\phi_n/\pi} + O(z)), \quad a_k \neq 0 \quad \forall k, \quad |z| \ll 1.$$

By $\det(a_{kj}) \neq 0$, this translation vector field cannot be $O(z)$ at the corner, so for at least one choice of k , we have $a_k \neq 0$. We say *the failure of immersion at the corner is 'minimal'*. ■

Holomorphic polygon. We now move on to holomorphic polygons with $k + 1$ corner points for $k \geq 2$. The extended linearized Cauchy–Riemann equation (see [71, Chapter 9]) involves a vector field $v \in C^\infty(\Sigma, u^*TX)$ decaying at the ends, and $\rho \in \Omega^{0,1}(\Sigma, T\Sigma)$ representing a tangent vector of the Stasheff associahedron (i.e., the deformation of Riemann surface structure on the domain Σ), satisfying

$$\bar{\partial}v + \frac{1}{2}J_X \circ du \circ \rho = 0,$$

where J_X is the complex structure on X . Here ρ can be taken to be compactly supported, so $\bar{\partial}v = 0$ near the corners. The next lemma immediately follows.

Lemma 3.7. *Given first order deformation vector fields v_1, \dots, v_{n-1} , then the $(1, 0)$ -form $\Omega(\cdot, v_1, \dots, v_{n-1})$ on Σ is holomorphic.*

Remark 3.7. Adding vector fields on Σ valued in $T\Sigma$ to v_1, \dots, v_{n-1} does not affect $\Omega(\cdot, v_1, \dots, v_{n-1})$ as a 1-form on Σ . Thus this 1-form is insensitive to how one represents the Riemann surface structures on the abstract polygon.

We impose that the input at one of the k corners maps to an intersection point in $CF^0(L, L')$, and the other inputs map to degree one self intersections of L or L' . The output maps to $q \in CF^*(L, L')$.

Proposition 3.8 (Automatic transversality, polygon case). *Suppose v_1, \dots, v_{n-1} are linearly independent first order deformation vector fields. Either $\Omega(\cdot, v_1, \dots, v_{n-1})$ vanishes identically as a 1-form on Σ , or we must have $\deg q \geq n$, and when the equality holds then $\Omega(\cdot, v_1, \dots, v_{n-1})$ only vanishes at corners. In this case, all first order deformation vector fields are spanned by v_1, \dots, v_{n-1} , the holomorphic polygon $u: \Sigma \rightarrow X$ is an immersion up to the boundary, the obstruction of the extended linearized operator vanishes, and the moduli space is smooth at $u: \Sigma \rightarrow X$.*

Proof. By viewing the domain of the polygon as a strip with extra boundary punctures, we produce a holomorphic vector field v_n as the \mathbb{R} -translation vector field. However, unlike in the strip case, at the degree one self intersection corners v_n does not typically have the required decay to be admitted as a deformation vector field. Indeed, by thinking about such a corner point as the origin in the upper half plane model of Σ with local coordinate z , then zv_n decays at the corner, but not necessarily v_n itself.

Now $v_1 \wedge \cdots \wedge v_n$ is a section of $\Lambda^n TX$ with boundary value on $\Lambda^n TL$, and in addition $\Omega(v_1, \dots, v_n)$ is a holomorphic function on Σ . We assume from now on that it is not identically zero. The index of the ordinary Cauchy–Riemann operator is

$$\deg q - \sum_1^k \deg p_k = \deg q - k + 1.$$

Invoking (3.7) this is computable from the vanishing orders of $\Omega(v_1, \dots, v_n)$:

$$\deg q - k + 1 = 2 \sum (\text{interior zeros}) + \sum (\text{boundary zeros}) + \sum (\text{corner zeros}) + n.$$

The interior and boundary vanishing orders are non-negative. Since the v_k are holomorphic near the corners without correction, the proof of Lemma 3.5 shows that the excess vanishing order at the $CF^0(L, L')$ corner and the q corner are both non-negative. At the degree one self intersection corners, the excess vanishing order of $\Omega(v_1, \dots, v_{n-1}, zv_n)$ is non-negative by the same previous arguments, so $\Omega(v_1, \dots, v_{n-1}, v_n)$ itself has excess vanishing order ≥ -1 . Hence,

$$\deg q - k + 1 \geq n - k + 1,$$

namely $\deg q \geq n$.

When the equality is achieved, all bounds are saturated. In particular, $\Omega(v_1, \dots, v_n)$ can only vanish at the corners, so $u: \Sigma \rightarrow X$ is an *immersion up to boundary*. At the corners, the same arguments in Corollary 3.6 shows the failure of immersion is minimal.

If v is the deformation vector field corresponding to an arbitrary kernel element of the extended linearized operator, then after subtracting off a constant linear combination of v_1, \dots, v_{n-1} , we may assume v is tangent to Σ at any chosen point on $\partial\Sigma$. The same argument in Corollary 3.6 shows v is tangent to the image of Σ . The immersion property allows us to lift v to the domain Σ . There is no room to deform the complex structure of Σ , nor is there any automorphism of Σ , so in fact v vanishes identically. This shows that v_1, \dots, v_{n-1} span all first order deformations. But $\deg q = n$ implies that the index of the extended linearized operator is

$$\deg q - \sum_1^k p_k + k - 2 = n - 1.$$

Thus the cokernel dimension is zero, namely the obstruction space vanishes. Consequently, the moduli space of such holomorphic polygons is smooth. ■

Remark 3.8. Using the Floer degree formula given in (A.1), the asymptotic behaviour of $\Omega(v_1, \dots, v_n)$ at the corners can be extracted from the above proof: at the $CF^0(L, L')$ corner point p ,

$$\Omega(v_1, \dots, v_n) = a_p z^{(\theta_{L'} - \theta_L)(p)/\pi} (1 + O(z)), \quad a_p \neq 0, \quad \arg a_p = \theta_L(p) \bmod \pi \mathbb{Z}.$$

At the degree one self intersections $p_l \in CF^1(L_+, L_-)$ on L or L' ,

$$\Omega(v_1, \dots, v_n) = a_l z^{(\theta_{L_-} - \theta_{L_+})(p_l)/\pi} (1 + O(z)), \quad a_l \neq 0, \quad \arg a_l = \theta_{L_+}(p_l) \bmod \pi \mathbb{Z}.$$

At the degree n output q ,

$$\Omega(v_1, \dots, v_n) = a_q z^{(\theta_L - \theta_{L'})(q)/\pi} (1 + O(z)), \quad a_q \neq 0, \quad \arg a_q = \theta_{L'}(q) \bmod \pi \mathbb{Z}.$$

Weighted Sobolev space with exponential growth. We now discuss solutions to linearized Cauchy–Riemann equations in weighted Sobolev spaces $W^{1,2;\mu}$ (see [72, Section 2]). These spaces agree with their unweighted counterparts along the strip like input ends, but at the strip like output end $s \gg 0$, a vector field $v \in W^{1,2;\mu}$ means that $\exp(-\mu s)v$ lies in $W^{1,2}$. Generally we choose μ to avoid a discrete set of indicial values. The main point of these weighted Sobolev spaces is that they allow for holomorphic vector fields with prescribed exponential growth along the output end, which is conceptually similar to allowing for meromorphic functions in Riemann surface theory. If we think of the strip like end q as the infinity (resp. the origin) in the upper half plane model of Σ , then the natural coordinate is $z = e^{\pi(s+it)}$ (resp. $z = e^{-\pi(s+it)}$), and the exponential growth $o(e^{\mu s})$ becomes $o(|z|^{\mu/\pi})$ (resp. $o(|z|^{-\mu/\pi})$).

For larger μ , more vector fields are included in the Sobolev space and the index increases by one each time μ crosses an indicial value (counted with multiplicity). In our problem, the indicial values are

$$\phi_1 + \pi \mathbb{Z}, \quad \phi_2 + \pi \mathbb{Z}, \dots, \phi_n + \pi \mathbb{Z},$$

where ϕ_1, \dots, ϕ_n are the characterizing angles at the Lagrangian intersection point q at the output end. Then the index for the linearized Cauchy–Riemann operator $W^{1,2;\mu} \rightarrow L^{2,\mu}$ is

$$\deg q - \sum_1^k \deg p_i + \text{number of indicial values between } 0 \text{ and } \mu, \quad (3.8)$$

where k is the number of input ends. In particular, for holomorphic strips with $\deg p = \deg q$ (resp. $\deg q - \deg p = 1$), then the index for $\mu = \pi$ is equal to n (resp. $n + 1$). In contrast, the ordinary index (for the $\mu = 0$ case) is zero, and the moduli space obtained by taking \mathbb{R} -quotient has virtual dimension -1 (resp. zero). There are in fact sufficient conditions to rule out the negative dimension moduli spaces, and constrain the zero-dimensional moduli spaces.

Lemma 3.9. *In the holomorphic strip case, assume v_1, \dots, v_n are in the kernel of the linearized Cauchy–Riemann operator on $W^{1,2;\mu=\pi}$, such that $\Omega(v_1, \dots, v_n)$ does not vanish identically. Then $\deg q - \deg p \geq 1$. When the equality is achieved, the holomorphic strip is an immersion up to the boundary with minimal vanishing at the corner, and the zero-dimensional moduli space is regular.*

Proof. We modify the proof of Lemma 3.5 and Corollary 3.6. We think of the corner q as the origin in the upper half plane model. Without loss of generality v_n is the \mathbb{R} -translation vector field of the holomorphic strip. Then the leading order asymptotic is

$$v_k = (a_{k1}z^{\phi_1/\pi-1}, \dots, a_{kn}z^{\phi_n/\pi-1}) + O(1), \quad k = 1, 2, \dots, n-1,$$

and

$$v_n = (a_{n1}z^{\phi_1/\pi}, \dots, a_{nn}z^{\phi_n/\pi}) + O(z),$$

hence

$$\Omega(v_1, \dots, v_n) = z^{(\sum \phi_k)/\pi - n + 1} (\det(a_{kj}) + o(1)).$$

The excess vanishing order is $\geq 1 - n$, where negative order stands for poles. By the index formula (3.7) for the ordinary linearized Cauchy–Riemann equation, we have

$$\deg q - \deg p \geq 1 - n + n = 1,$$

and equality forces $\Omega(v_1, \dots, v_n)$ to have no interior zero, no boundary zero, minimal zero at p , and $\det(a_{kj}) \neq 0$ at q . The argument in Corollary 3.6 shows $v_1, \dots, v_n, \frac{v_n}{z}$ span the real vector space of first order deformations in $W^{1,2;\pi}$. In particular, the only first order deformation which decays at q is the \mathbb{R} -translation vector field. Thus the cokernel to the ordinary linearized Cauchy–Riemann operator vanishes, and the moduli space is regular. ■

A very analogous statement holds in the polygon case, and is left to the reader.

Lemma 3.10. *In the holomorphic polygon case, we assume that v_1, \dots, v_{n-1} are in the kernel of the extended linearized Cauchy–Riemann operator on $W^{1,2;\mu=\pi}$ such that $\Omega(\cdot, v_1, \dots, v_{n-1})$ does not vanish identically as a 1-form on Σ . Then*

$$\deg q - \sum_1^k \deg p_k + k - 2 \geq 0.$$

When the equality is achieved, the holomorphic polygon is an immersion up to the boundary with minimal vanishing at the corner, and the zero-dimensional moduli space of holomorphic polygons is regular at $u: \Sigma \rightarrow X$.

A similar statement applies to teardrop curves.

Lemma 3.11 (Regularity of teardrops). *Let $u: \Sigma \rightarrow X$ be a teardrop curve with a unique output q and no input ends. Assume v_1, \dots, v_{n-1} are in the kernel of the linearized Cauchy–Riemann operator on $W^{1,2;\mu=\pi}$ such that $\Omega(\cdot, v_1, \dots, v_{n-1})$ does not vanish identically as a 1-form on Σ . Then $\deg q \geq 2$. When the equality is achieved, the teardrop curve is an immersion up to the boundary with minimal vanishing at the corner, and the kernel of the ordinary Cauchy–Riemann operator is spanned as a real vector space by the Möbius vector fields on Σ fixing the q corner, and the cokernel vanishes.*

Proof. We modify the proof of Lemma 3.9. We think of the corner q as the origin in the upper half plane model, and take v_n instead to be the Möbius vector field $z^2 \partial_z$ on Σ . This has one higher order of vanishing:

$$v_n = (a_{n1} z^{\phi_1/\pi+1}, \dots, a_{nn} z^{\phi_n/\pi+1}) + O(z^2).$$

This leads to

$$\Omega(v_1, \dots, v_n) = z^{(\sum \phi_k)/\pi - n + 2} (\det(a_{kj}) + o(1)),$$

so the excess vanishing order at q is $\geq 2 - n$. The index of the ordinary linearized Cauchy–Riemann operator is

$$\deg q = 2 \sum (\text{interior zeros}) + \sum (\text{boundary zeros}) + \sum (\text{excess corner zeros}) + n,$$

whence $\deg q \geq 2$.

When the equality is achieved, then there is no interior or boundary zero and

$$\det(a_{kj}) \neq 0$$

at the corner, hence the immersion claim. The argument in Corollary 3.6 shows that $v_1, \dots, v_n, \frac{v_n}{z}, \frac{v_n}{z^2}$ span the real vector space of first order deformations in $W^{1,2;\pi}$. In particular, the only first order deformation which decays at q are spanned by v_n and $z^{-1}v_n$, namely the Möbius generators. Since the index of the ordinary Cauchy–Riemann operator is two, the cokernel must have dimension zero, namely the obstruction vanishes. ■

The above lemma describes the optimal case for teardrop curves. Such $\deg q = 2$ teardrop curves arise in isolated zero-dimensional moduli spaces after taking the $\text{Aut}(D^2, q)$ quotient, and the counting contribution to m_0 are ± 1 depending on the spin structure and the orientation issues.

Structure of linearized Cauchy–Riemann equation. Let Σ be a holomorphic polygon with $k \geq 0$ input ends p_i , and one output end at q . The case $k = 0$ corresponds to teardrops, and $k = 1$ corresponds to strips. We consider the ordinary linearized Cauchy–Riemann operator in weighted Sobolev spaces $W^{1,2;\mu}$, to classify the structure of the first order deformation theory. As usual, the complex structure is integrable. Since $\bar{\partial}$ is elliptic, its cokernel in L^2 is finite-dimensional, represented by holomorphic 1-forms on Σ , which must have finite order of vanishing at q . For large enough μ , the dual space $L^{2;- \mu}$ for $L^{2;\mu}$

imposes an exponential decay condition $O(e^{-\mu s})$ at q , so the cokernel eventually vanishes for $\mu \gg 1$. Then the kernel dimension in $W^{1,2;\mu}$ is equal to the index, computed by (3.8). For convenience, we use $\mu \in \pi\mathbb{N}$, which avoids the indicial values. Then

$$\dim(\ker \bar{\partial} \subset W^{1,2;\mu}) = \deg q - \sum_1^k \deg p_i + \frac{n\mu}{\pi}. \quad (3.9)$$

It is convenient to view the domain Σ of the holomorphic polygon as the upper half-plane with coordinate z , with corners p_i on the real line and q at infinity.

Lemma 3.12. *If $v \in \ker \bar{\partial} \subset W^{1,2;\mu}$, then $v = fw$ for some real coefficient polynomial function f on the upper half plane, such that w is non-vanishing on $\mathbb{R} \setminus \{p_1, \dots, p_k\}$, and vanishes minimally at p_i (meaning w is indivisible by $z - p_i$).*

Proof. If v vanishes at any boundary point a on $\mathbb{R} \setminus \{p_1, \dots, p_k\}$, or if v vanishes at p_i beyond minimal order, then $\frac{v}{z-a}$ (resp. $\frac{v}{z-p_i}$) is also a first order deformation with the same TL boundary condition, subject to the growth constraints at infinity. Since the kernel dimension is finite, the divisions can only happen a finite number of times, producing the polynomial f . ■

Let $p \in \partial\Sigma \simeq \partial\mathbb{H}$, and let K be the maximal number depending on p , such that there exist \mathbb{R} -linearly independent $v_1, \dots, v_K \in \ker \bar{\partial}$ satisfying the following conditions:

- If p is not a corner point, then $v_1(p), \dots, v_K(p)$ are \mathbb{R} -linearly independent vectors.
- If $p = p_i$ is a corner point, then the non-zero elements in the \mathbb{R} -span of v_1, \dots, v_K are vector fields vanishing minimally at p .

Lemma 3.13. *Any $v \in \ker \bar{\partial} \subset W^{1,2;\mu}$ is of the form $g_1 v_1 + \dots + g_K v_K$ for some real coefficient rational functions g_1, \dots, g_K , non-singular at p .*

Proof. Without loss of generality $p = 0$. Let w_0 be any first order deformation. By the maximality of K , we can choose real numbers a_i , such that the first order deformation $w_0 - \sum_1^K a_i v_i$ vanishes at zero, so $w_0 - \sum_1^K a_i v_i = z^{k_1} w_1$ for some first order deformation w_1 which is non-zero at the origin. Finite dimensionality means this process can be repeated for only a finite number of times:

$$\left\{ \begin{array}{l} w_0 = \sum a_{i1} v_i + z^{k_1} w_1, \\ w_1 = \sum a_{i2} v_i + z^{k_2} w_2, \\ \vdots \\ w_{N-1} = \sum a_{iN} v_i + z^{k_N} w_N. \end{array} \right.$$

We choose the smallest N such that $v_1, \dots, v_K, w_0, \dots, w_N$ are \mathbb{R} -linearly dependent as vector fields; notice v_1, \dots, v_K are linearly independent, so $N \geq 0$. We then get a linear

relation

$$f_0(z)w_N = \sum_1^K f_i(z)v_i,$$

where f_0, \dots, f_K are polynomials, and $f_0(0) \neq 0$. This implies the claim. \blacksquare

We can also apply a Möbius transform to make the output end q lie at the origin. The growth condition translates to $o(|z|^{-\mu/\pi})$ at zero. Let K be maximal, such that there are \mathbb{R} -linearly independent vector fields $z^{-\mu/\pi}v_1, \dots, z^{-\mu/\pi}v_K \in \ker \bar{\partial} \subset W^{1,2;\mu}$, and any non-zero element in the \mathbb{R} -span of v_1, \dots, v_K vanishes minimally at $q = 0$. As a caveat, this does not assume v_1, \dots, v_K satisfy the growth constraints at infinity to lie in $\ker \bar{\partial} \subset W^{1,2;\mu}$. Minor adaptations give the following lemma.

Lemma 3.14. *Any $v \in \ker \bar{\partial} \subset W^{1,2;\mu}$ is of the form $z^{-\mu/\pi}(g_1v_1 + \dots + g_Kv_K)$ for some real coefficient rational functions g_1, \dots, g_K , non-singular at q .*

Corollary 3.15. *The number K is independent of the boundary and corner points on $\partial\Sigma$.*

We view the boundary $\partial\Sigma \simeq \mathbb{P}^1(\mathbb{R})$. By the above lemmas, there is a real algebraic vector bundle \mathcal{E} of rank K over $\mathbb{P}^1(\mathbb{R})$ such that v_1, \dots, v_K provide the basis of local sections. By Grothendieck's classification of vector bundles, we obtain the following result.

Proposition 3.16. *The vector bundle $\mathcal{E} \simeq \oplus_1^K \mathcal{O}(n_i)$ for some $n_i \in \mathbb{Z}$.*

Since the rank K is non-decreasing in μ , it eventually stabilizes for $\mu \gg 0$. Since around any given point, the same choice of v_1, \dots, v_K is valid for all large μ , the algebraic vector bundle \mathcal{E} is independent of $\mu \gg 0$. The elements of $\ker \bar{\partial} \subset W^{1,2;\mu}$ can be interpreted as global sections of $\mathcal{E} \otimes \mathcal{O}(\frac{\mu}{\pi})$. Thus for all large μ ,

$$\dim(\ker \bar{\partial} \subset W^{1,2;\mu}) = \dim \Gamma\left(\mathbb{P}^1(\mathbb{R}), \mathcal{E} \otimes \mathcal{O}\left(\frac{\mu}{\pi}\right)\right) = K\left(\frac{\mu}{\pi} + 1\right) + \sum_1^K n_i.$$

Contrasting with the index formula (3.9), we have the following.

Corollary 3.17. *The rank $K = n$, and the degree $\sum_1^n n_i = \deg q - \sum_1^k \deg p_i - n$.*

The structure of $\mathcal{E} \simeq \oplus_1^K \mathcal{O}(n_i)$ provides meromorphic sections v_1, \dots, v_n which are a basis of local sections on $\mathbb{R} \subset \mathbb{P}^1(\mathbb{R})$, and have excess vanishing orders n_1, \dots, n_n at q . Consider the function $\Omega(v_1, \dots, v_n)$. By construction, it has no boundary zero, and its excess corner vanishing order is $\sum_1^n n_i$. Comparing with the index formula (3.7),

$$\deg q - \sum_1^k p_i = 2 \sum (\text{interior zeros}) + \sum n_i + n.$$

Since all interior vanishing orders are non-negative by holomorphicity, we then obtain the following corollary.

Corollary 3.18. *We have $\Omega(v_1, \dots, v_n) \neq 0$ in the interior of Σ .*

The significance is that the *algebraic vector bundle* structure on $\mathcal{E} \rightarrow \mathbb{P}^1(\mathbb{R})$ now extends over the entire Σ . The v_1, \dots, v_n now provide the basis of local sections for the vector bundle $u^*TX|_\Sigma$. One upshot is that an algebraic structure arises on $u^*TX \rightarrow \Sigma$ from solving the Cauchy–Riemann equation with Lagrangian boundary:

$$(u^*TX, u^*TL) \simeq (\oplus_1^n \mathcal{O}(n_i), \text{natural real structure}). \quad (3.10)$$

In contrast, the Lagrangians are only assumed to be smooth, not necessarily real analytic.

To analyze obstructions, Serre duality motivates us to consider the dualized cokernel to the ordinary (unweighted, unextended) linearized Cauchy–Riemann operator. A dualized cokernel element η is represented by a holomorphic 1-forms in $\Omega^{1,0}(\Sigma, u^*T^*X)$ with L^2 integrability, and its T^*X factor lies in the annihilator of the TL boundary condition. Equivalently, for all test vector fields $v \in W^{1,2}(\Sigma, TX, TL)$,

$$\int_\Sigma \langle \bar{\partial}v \wedge \eta \rangle = 0,$$

where $\langle \cdot, \cdot \rangle$ is the pairing of TX with T^*X , and the wedge takes care of the forms on Σ . In the canonical form (3.10), this dualized cokernel is isomorphic to

$$\Gamma(\mathbb{R}\mathbb{P}^1, \oplus_1^n \mathcal{O}(-2 - n_i)).$$

In particular, we have the following result.

Corollary 3.19. *In the teardrop curve case $k = 0$, the strip case $k = 1$ and the triangle case $k = 2$, the vanishing of cokernel is equivalent to $n_i \geq -1$ for all i .*

For $k \geq 3$ the deformation of the holomorphic polygons is governed instead by the extended Cauchy–Riemann equation, since the punctured Riemann surface structure on Σ is allowed to vary. The dualized cokernel of the extended Cauchy–Riemann operator, is the subspace of the dualized cokernel of the ordinary Cauchy–Riemann operator, which pairs trivially with $J_X \circ du \circ \rho$ for any ρ representing some tangent vector of the Stasheff associahedron.

Hamiltonian deformations and transversality. We now consider the parametrized moduli space of holomorphic curves over the infinite-dimensional space of Hamiltonian deformations for the Lagrangian L . Infinitesimally around a holomorphic curve Σ , we have a Hamiltonian vector field X_H defined by $\omega(X_H, \cdot) = dH$, viewed as a TX -valued vector field over Σ . We are interested in whether the Hamiltonian deformation kills the cokernel of the ordinary Cauchy–Riemann operator. This question was first addressed by Oh [65]. The following account follows a similar strategy but differs in details.

Recall that the ordinary Cauchy–Riemann operator maps $W^{1,2}(\Sigma, u^*TX, u^*TL)$ to $L^2(\Sigma, u^*TX \otimes T^{*(1,0)}\Sigma)$. The effect of Hamiltonian deformation is to enlarge the domain of the $\bar{\partial}$ operator, by including the vector fields X_H for all the allowed Hamiltonians H . The question is to analyze the pairing of $\bar{\partial}X_H$ with the dualized cokernel elements.

Proposition 3.20. *Let $u: \Sigma \rightarrow X$ be a holomorphic disc which is immersed near some point $z_0 \in \partial\Sigma$ with the boundary injectivity property $u|_{\partial\Sigma}^{-1}(u(z_0)) = \{z_0\}$. Let η be a non-zero dualized cokernel element for the ordinary linearized Cauchy–Riemann operator. Then there is a Hamiltonian H supported in any prescribed small ball on X containing $u(z_0)$, such that*

$$\int_{\Sigma} \langle \bar{\partial} X_H \wedge \eta \rangle \neq 0.$$

Proof. Since η is a holomorphic 1-form valued in u^*T^*X , Stokes' theorem gives

$$\int_{\Sigma} \langle \bar{\partial} X_H \wedge \eta \rangle = \int_{\partial\Sigma} \langle X_H, \eta \rangle,$$

where $\langle \cdot, \cdot \rangle$ stands for the pairing between TX and T^*X . On $\partial\Sigma$, we can write

$$\eta = \omega(\cdot, Y) ds$$

for some vector field Y valued in u^*TX , and s is any local coordinate on $\partial\Sigma$. The cokernel element condition implies $\omega(v, Y) = 0$ for any $v \in u^*TL$, so Y must in fact be valued in the Lagrangian subbundle u^*TL . Thus,

$$\int_{\partial\Sigma} \langle X_H, \eta \rangle = \int_{\partial\Sigma} \omega(X_H, Y) ds = \int_{\partial\Sigma} dH(Y) ds.$$

We suppose for contradiction, that this pairing vanishes identically for any H supported in the prescribed ball.

By the holomorphicity of η , its zeros are isolated, so without loss of generality Y does not vanish in the local portion of $\partial\Sigma$ where u is injective and immersed. Suppose first that Y is not tangent to the image of Σ . Then we find some local function h on a small ball in X with $dh(Y) = 1$ and $h = 0$ on the local portion of $\partial\Sigma$, and another cutoff function $h_2 \geq 0$ with $dh_2(Y) = 0$ along $\partial\Sigma$, supported in a small ball. Taking $H = hh_2$, then

$$\int_{\partial\Sigma} dH(Y) ds = \int_{\partial\Sigma} h_2 ds \neq 0.$$

This contradiction shows Y is tangent to the image of Σ in the local portion of $\partial\Sigma$. We can write $Y = f \partial_s$ for some local function f . Then requiring

$$\int_{\partial\Sigma} dH(Y) ds = \int_{\partial\Sigma} f \partial_s H ds = - \int_{\partial\Sigma} H \partial_s f ds$$

for any compactly supported local function H , implies that f is constant in the local portion of $\partial\Sigma$. Thus up to multiplying by a non-zero constant, locally

$$Y = \frac{\partial u}{\partial s} ds, \quad \eta = \omega\left(\cdot, \frac{\partial u}{\partial s}\right) ds. \quad (3.11)$$

We now produce holomorphic vector fields on Σ . For holomorphic strips or polygons with $k + 1 \geq 3$ corners, we select one input end as p , and call the output q as usual, and

represent Σ as a strip with $k - 1$ boundary punctures. This perspective provides a natural translation vector field $\frac{\partial u}{\partial s}$, which have exponential decay along the p, q ends, but may not be L^2 near the other $k - 1$ ends. Instead, by thinking about the $k - 1$ ends as the origin in the upper half plane model, we see

$$\frac{\partial u}{\partial s} = O(|z|^{\alpha-1}), \quad \alpha = \min\{\phi_1/\pi, \dots, \phi_n/\pi\}$$

for the characterizing angles ϕ_1, \dots, ϕ_n at the Lagrangian intersection point. The $T^{(1,0)}X$ part of $2J_X \frac{\partial u}{\partial s}$ is

$$J_X \frac{\partial u}{\partial s} + \sqrt{-1} \frac{\partial u}{\partial s}.$$

Contracting this with the $T^{*(1,0)}X \otimes T^{*(1,0)}\Sigma$ part of η yields a 1-form on Σ :

$$\zeta = \eta \left(J_X \frac{\partial u}{\partial s} + \sqrt{-1} \frac{\partial u}{\partial s} \right),$$

which is also holomorphic with boundary value along Σ :

$$\zeta = \omega \left(J_X \frac{\partial u}{\partial s} + \sqrt{-1} \frac{\partial u}{\partial s}, Y \right) ds = \omega \left(J_X \frac{\partial u}{\partial s}, Y \right) ds. \quad (3.12)$$

Here $\omega(\frac{\partial u}{\partial s}, Y) = 0$ since both vectors satisfy the TL boundary condition. Notably, the boundary condition of ζ is real valued. In the upper half plane model, the Schwartz reflection principle allows us to extend ζ meromorphically over \mathbb{CP}^1 .

At any of the $k - 1$ ends, since $\eta \in L^2$, we know by holomorphicity $|\eta| = O(|z|^\alpha)$, so $\zeta = O(|z|^{2\alpha-1})$ in the upper half plane model, hence has no pole. At the p, q ends, by the decay of the holomorphic $\frac{\partial u}{\partial s}$ and η , we likewise infer that ζ has no pole in the upper half plane model. In conclusion, the extension of ζ over \mathbb{CP}^1 has no pole, so must in fact vanish. However, by (3.11) and (3.12), on a local portion of $\partial\Sigma$, we have

$$\zeta = \omega \left(J_X \frac{\partial u}{\partial s}, \frac{\partial u}{\partial s} \right) ds \neq 0.$$

This contradiction proves the Proposition in the $k \geq 1$ case.

Finally, for the teardrop curve case $k = 0$, we replace the holomorphic vector field $\frac{\partial u}{\partial s}$ by the Möbius vector fields vanishing at the corner, and the rest of the arguments are entirely similar. \blacksquare

The upshot is that by the Sard–Smale theorem, provided we can always ensure ‘somewhere boundary injectivity’ for any holomorphic disc in a given moduli space, then generic Hamiltonian perturbation would be able to achieve regularity for the moduli space.

Remark 3.9. In the exact setting, there is no closed holomorphic curve. The failure of ‘somewhere boundary injectivity’ is often associated with multiple cover issues, namely $u: \Sigma \rightarrow X$ may decompose into several domain components, each of which factorizes

through a somewhere boundary injective holomorphic disc (see the paper [51] for the case of Lagrangian boundary with no corners).

In the simplest case, if u factorizes through another disc, then the corner points would be repeated several times on $\partial\Sigma$. This phenomenon does not happen for the curves appearing in the bordism current \mathcal{C} , which involve only one corner at $CF^0(L, L')$ and one corner at $CF^0(L', L)$. Nor does this occur for teardrop curves, which have only one corner at a degree two self intersection point. This raises hope that the failure of ‘somewhere boundary injectivity’ may be *highly non-generic*, or in certain situations can be ruled out altogether.

Further comments on automatic transversality. We now comment on the gap between what we have proved and the (weak version of) automatic transversality that we will later assume.

(1) Proposition 3.8, Lemma 3.5 and Corollary 3.6 establish the dichotomy for holomorphic discs $u: \Sigma \rightarrow X$ arising in virtual dimension $n - 1$ moduli spaces, that either u is automatically transverse, or $\Omega(\cdot, v_1, \dots, v_{n-1})$ vanishes for any $(n - 1)$ first order deformation vectors. This argument does not establish unperturbed regularity for the lower-dimensional moduli spaces, so it is not completely clear if complex structure perturbations can be removed in the arguments for $\partial\mathcal{C} = L - L'$ in Section 3.1.2.

(2) For the bad curves, $\Omega(\cdot, v_1, \dots, v_{n-1})$ vanishes identically as a 1-form on Σ , so at any point on the boundary, v_1, \dots, v_{n-1} and the tangent vector to $\partial\Sigma$ are \mathbb{R} -linearly dependent. Suppose for the moment that the moduli spaces are regular, then the boundary evaluation to $L \cup L'$ for the bad curves arise in Hausdorff dimension at most $n - 1$. Moreover, since the Solomon functional is defined through

$$\int_{\mathcal{C}} \lambda \wedge \Omega,$$

and Ω vanishes around the bad curves, smoothness assumptions imply that the bad curves cannot contribute.

When regularity assumptions are dropped, one needs to appeal to virtual techniques, so these conclusions require further justification. One problem is that the standard virtual perturbation techniques based on Kuranishi structures do not necessarily produce virtual cycles inside the original moduli spaces, but only inside their small neighbourhoods. This perturbation step destroys the identical vanishing of Ω , by a small amount corresponding to the size of the perturbation. As one shrinks the size of the perturbations, one needs uniform mass bound on the virtual chains to justify that the integral contribution to $\int_{\mathcal{C}} \lambda \wedge \Omega$ from the bad curves actually converges to zero.

(3) Alternatively, one can hope to replace Lagrangians by arbitrarily small Hamiltonian perturbations to achieve transversality. This is mostly adequate for our purpose, except that one needs to justify the ‘somewhere boundary injectivity’ property (see Remark 3.9).

3.4. Positivity condition

The $(n - 1)$ -dimensional moduli spaces contributing to the bordism current come with orientation signs and weighting factors. The *positivity condition* (i.e., *no cancellation of signs*) means that around any automatically transverse curves, if the first order deformations v_1, \dots, v_{n-1} form an oriented basis of the moduli space, and v_0 be a clockwise ordered vector field on $\partial\Sigma$, then upon boundary evaluation $v_0 \wedge \dots \wedge v_{n-1}$ agrees with the orientation of $L - L'$. For the other holomorphic curves, the question of orientation does not arise, because the boundary evaluation maps have degenerate differentials everywhere on $\partial\Sigma$.

The positivity condition forbids two curves passing through a generic point with the evaluation maps contributing opposite signs. Such a requirement is *geometric rather than homological*, and if we go beyond the almost calibrated case, it also *depends on the choice of the generators* $\alpha \in CF^0(L, L')$ and $\beta \in CF^0(L', L)$, rather than only their classes in HF^0 .

Question 5. Let L, L' be two unobstructed Lagrangian objects which are isomorphic in $D^b \text{Fuk}(X)$. When can we make gauge choices for the local systems and the bounding cochain data, and choices of the Floer cohomology group generators, such that the bordism current \mathcal{C} produced from the universal family of holomorphic curves satisfies the positivity condition?

The positivity condition will arise in the applications as follows. We will write various quantities as integrals over the $(n - 1)$ -dimensional moduli spaces of holomorphic curves, and the positivity condition would in each case imply the *pointwise positivity of the integrand*. In the mirror analogy, this corresponds to the pointwise positivity of curvature integrands, which features for instance in the proof that the Hermitian Yang–Mills equation implies the semistability of bundles (see Section 2.5).

Morse theory analogy. The intuition of the positivity condition can be explained through the following *analogy with Morse theory*. Given a compact oriented manifold M with a Morse–Smale function H , then

(i) the degree zero (resp. n) elements in the Morse cochain complex are generated by the local maxima (resp. local minima) of H whose unstable submanifolds (resp. stable submanifolds) have *preferred orientations*;

(ii) the *fundamental class* $1 \in H^0(M)$ is represented by the sum of the local maxima with the preferred orientation. Similarly with the generator of $H^n(M)$;

(iii) a generic point on M lies on exactly one Morse flowline which starts with one local maximum point and ends on one local minimum point. More formally, if we construct the universal family of Morse flowlines that start with some local maximum and ends on some local minimum, the evaluation map would *sweep out the fundamental cycle* of M as an n -dimensional current, *without any cancellation effect*.

For simplicity, if we start with an embedded Lagrangian brane L , and perform a small generic Hamiltonian deformation $\phi_H(L)$, then the Floer cohomology $HF^*(L, \phi_H(L))$ is computed as the Morse cohomology, so the Morse theory statements above would imply the positivity condition at least in such special cases. The intuition is that if the Lagrangian L' (together with its brane structure) is a sufficiently small deformation of L , then we expect the positivity condition to hold for the bordism current between L and L' .

Positivity for individual moduli spaces. In general the bordism current \mathcal{C} receives contributions from many $(n - 1)$ -dimensional moduli spaces of holomorphic curves. We now focus on one moduli space by fixing the choice of the Lagrangian intersections p_1, \dots, p_k, q and the homotopy type of $u: \Sigma \rightarrow X$, and consider a connected open subset of the moduli space which contains only automatically transverse curves. We make the following observations:

(i) For a fixed automatically transverse holomorphic curve $u: \Sigma \rightarrow X$, along $\partial\Sigma$ between any two successive corners, by the nowhere vanishing of $\Omega(\cdot, v_1, \dots, v_n)$, the Jacobian of the boundary evaluation map cannot change sign. That is, either the orientation of the universal family agrees with the orientation of L (resp. $-L'$) along $u: \partial\Sigma \rightarrow L$ at every point along the boundary portion of $\partial\Sigma$, or the two orientations disagree at every point.

(ii) The Lagrangians are graded by assumption, and the orientations are canonically determined by $e^{-i\theta}\Omega$. The corner behaviour (see Remark 3.8) implies that at the degree one self intersections, $e^{-i\theta}\Omega(\cdot, v_1, \dots, v_{n-1})$ does not change sign. On the other hand, at the $CF^0(L, L')$ and the $CF^0(L', L)$ ends along $\partial\Sigma$, the 1-form $e^{-i\theta}\Omega(\cdot, v_1, \dots, v_{n-1})$ changes orientation sign. Thus at a fixed automatically transverse curve, the orientation of the universal family and $L - L'$ either completely agree along every point of $\partial\Sigma$, or completely disagree.

(iii) As we deform among automatically transverse curves, the orientation signs cannot change. Thus either all these holomorphic curves contribute positively to $\partial\mathcal{C}$, or they all contribute negatively.

The above discussion also suggests the limitation of the positivity condition: if we encounter a holomorphic curve in the moduli space, which is *not automatically transverse*, then it is possible to switch orientation signs. For arbitrary exact immersed Lagrangians, it seems unreasonable to expect the positivity condition, and it is conceivable that counterexamples may arise from h -principle constructions. Whether counterexamples occur for more restrictive Lagrangians seems less clear, and we leave the following sample questions as food for thought.

Question 6. How does the positivity condition behave under exact isotopy with surgery?

Question 7. Are there examples of exact Calabi–Yau manifolds such that the positivity condition is satisfied for bordism currents between all exact, almost calibrated, unobstruc-

ted immersed Lagrangians equipped with suitable brane structures? What if the Lagrangians are quantitatively almost calibrated (see (2.1))?

Remark 3.10. If the Fukaya category is defined over \mathbb{Z} , we can require all holonomy factors to be integer valued. The positivity condition requires all the contributions to $\partial\mathcal{C} = L - L'$ to have the same orientation sign. This has the amusing consequence that all holonomy factors associated with $(n - 1)$ -dimensional moduli spaces contributing to \mathcal{C} , must in fact all be $+1$. Intuitively, this means there is a unique such holomorphic curve through any generic point of L , and the boundary evaluation of universal family to L is transverse.

3.5. Floer theoretic obstructions

General features of obstruction conditions. Our goal is to look for obstructions to the existence of special Lagrangians within given $D^b \text{Fuk}(X)$ classes, which is the ‘easy direction’ of the conjectural stability condition. Before specializing to a technically oversimplified setup, we first explain the features we expect from these obstructions, which may hold in much more general contexts. The mirror analogy (see our discussion on the μ -stability in Section 2.5) suggests that:

- the obstructions are associated to certain *positivity* of signs, which essentially depend on the *integrability* of Kähler geometry;
- the quantity involved in the obstruction can be expressed as an *integral over a moduli space* of worldsheet instantons (i.e., holomorphic curves), and its sign comes from a *pointwise positivity* of the integrand on the moduli space;
- the input from Floer theory is associated to a *distinguished triangle* in $D^b \text{Fuk}(X)$, or possible generalizations to several Lagrangians;
- the role of the holomorphic volume form enters via *cohomological integrals*;
- there is no need for the complex Monge–Ampère equation. Only the *almost Calabi–Yau* condition is needed.

Furthermore, out of the many moduli spaces that may arise in Floer theory, we will only make use of certain $(n - 1)$ -dimensional moduli spaces of holomorphic curves, whose associated $(n + 1)$ -dimensional universal family provides bordism currents between the n -dimensional Lagrangians. Here are some a priori reasons why we restrict attention to these:

(i) The holomorphic volume form is naturally integrated over n -cycles. This explains the dimension.

(ii) We need bordism currents canonically associated to the distinguished triangles. In Floer theory, the A_∞ -structure only becomes an invariant when considered as a whole, and individual A_∞ products are not invariants, so invariance constrains how moduli spaces can enter into stability conditions. As mentioned in Section 3.1, the existence of the bor-

dism current is the geometric manifestation of linear relations in the zeroth Hochschild homology HH_0 of the Fukaya category, which contains important invariant information.

(iii) From a variational viewpoint which will be discussed more fully in Section 5, it is desirable to extend Floer theory to Lagrangians with much weaker regularity, in the varifold and current sense. We shall explain there that most of Floer cohomologies and A_∞ products cannot be expected to pass to the limit when the Lagrangians degenerate in such weak topologies, and we hope that the bordism currents we use are among the few pieces of Floer theory that *may* be well behaved under rather severe degenerations of Lagrangians.

These requirements are very stringent. We notice two other features:

(iv) We shall crucially rely on the almost calibrated condition.

(v) When we test the stability of L via the distinguished triangle

$$L_1 \rightarrow L \rightarrow L_2 \rightarrow L_1[1],$$

we do not wish to assume L_1 or L_2 is special Lagrangian. In our view, stability conditions should be expressed in Floer theoretic terms, without a priori knowledge of what special Lagrangians there are inside a given almost Calabi–Yau manifold.

The Floer theoretic obstruction condition. The following Floer theoretic obstruction will crucially require *complex integrability* and the *almost calibrated condition*. Assume

$$L_1 \rightarrow L \rightarrow L_2 \xrightarrow{\gamma} L_1[1]$$

to be a distinguished triangle of unobstructed exact immersed Lagrangian branes with bounding cochains, such that L_1, L, L_2 are all almost calibrated, and all intersections are transverse. In other words, the Lagrangian brane L is isomorphic in $D^b \text{Fuk}(X)$ to the immersed Lagrangian corresponding to the twisted complex (see Sections 3.1 and A.2)

$$L' \simeq \begin{pmatrix} (L_2, b_2) \\ \gamma & (L_1, b_1) \end{pmatrix}.$$

We obtain a bordism current \mathcal{C} with $\partial\mathcal{C} = L - L' = L - L_1 - L_2$. In our generality, the domains of L_1, L_2, L may have many connected components.

Theorem 3.21 (Floer theoretic obstruction). *Assume the automatic transversality and the positivity condition hold for the bordism current \mathcal{C} . Assume the destabilizing condition*

$$\hat{\theta}_1 = \arg \int_{L_1} \Omega > \hat{\theta}_2 = \arg \int_{L_2} \Omega.$$

Then the Lagrangian phase angle of L has a lower bound on its oscillation:

$$\sup_L \theta_L \geq \hat{\theta}_1, \quad \inf_L \theta_L \leq \hat{\theta}_2, \quad (3.13)$$

and moreover the J -volume of L (see Section 2.9) has a non-trivial lower bound

$$\text{Vol}_J(L) = \int_L e^{-i\theta} \Omega \geq \left| \int_{L_1} \Omega \right| + \left| \int_{L_2} \Omega \right|. \quad (3.14)$$

Proof. At a holomorphic polygon Σ in the universal family \mathcal{C} , denote v_1, \dots, v_{n-1} as the first order deformation vector fields representing an oriented basis of tangent vectors to the moduli space. In the special case of holomorphic strips, the moduli space refers to the \mathbb{R} -translation quotient. We noted in Section 3.3 that $\Omega(\cdot, v_1, \dots, v_{n-1})$ restricts to a holomorphic 1-form on Σ , so can be written as the differential of a *holomorphic function* F by the simply connectedness of Σ :

$$dF = \Omega(\cdot, v_1, \dots, v_{n-1}). \quad (3.15)$$

The corners on Σ are arranged in *clockwise* order with the following possibilities:

(i) In the primary case, we encounter some degree one self intersections on L from bounding cochains, a corner $p \in CF^0(L, L_2)$, some degree one self intersections on L_2 , a corner r from $\gamma \in CF^1(L_2, L_1)$, some degree one self intersections on L_1 , and a corner at $q \in CF^0(L_1, L)$. Notice the Lagrangian boundary follows L, L_2, L_1 in clockwise order, and we cannot go reversely from L_1 to L_2 instead.

(ii) In the secondary cases, the boundary data may miss either L_1 or L_2 . For instance, we may encounter some degree one intersections on L , a corner $p \in CF^0(L, L_2)$, some degree one intersections on L_2 and a corner at $q \in CF^0(L_2, L)$. The Lagrangian boundary follows L, L_2 in clockwise order. The alternative possibility of Lagrangian boundary along L and L_1 is entirely similar.

In all cases, there is precisely one corner p at $CF^0(L, L')$ and a corner q at $CF^0(L', L)$. We can normalize $F(q) = 0$ to fix the constant. In the primary case, there is a corner $r \in CF^1(L_2, L_1)$, which is absent in the secondary cases. In general, the bordism current \mathcal{C} receives contributions from many moduli spaces, and all three cases may arise depending on the generators of $HF^0(L, L')$ and $HF^0(L', L)$.

We can now define *complex valued volume forms* on the $(n-1)$ -dimensional moduli spaces of holomorphic curves. Recall v_1, \dots, v_{n-1} represent the tangent vectors to the moduli spaces, and the holomorphic function F depends on $v_1 \wedge \dots \wedge v_{n-1}$. In the primary case, we define

$$\begin{cases} \tilde{\Omega}_L(v_1, \dots, v_{n-1}) = F(p), \\ \tilde{\Omega}_{L_1}(v_1, \dots, v_{n-1}) = F(r), \\ \tilde{\Omega}_{L_2}(v_1, \dots, v_{n-1}) = F(p) - F(r). \end{cases}$$

In the secondary cases, if the Lagrangian boundary lies on L and L_1 , then

$$\begin{cases} \tilde{\Omega}_L(v_1, \dots, v_{n-1}) = F(p), \\ \tilde{\Omega}_{L_1}(v_1, \dots, v_{n-1}) = F(p), \\ \tilde{\Omega}_{L_2}(v_1, \dots, v_{n-1}) = 0. \end{cases}$$

If the Lagrangian boundary lies on L and L_2 , then

$$\begin{cases} \tilde{\Omega}_L(v_1, \dots, v_{n-1}) = F(p), \\ \tilde{\Omega}_{L_1}(v_1, \dots, v_{n-1}) = 0, \\ \tilde{\Omega}_{L_2}(v_1, \dots, v_{n-1}) = F(p). \end{cases}$$

The values of F should be understood as the integral of dF on the appropriate portions of $\partial\Sigma$. The key point is that since $\partial\mathcal{C}$ sweeps out the cycle $L - L_1 - L_2$, we can write the *period integrals* as *integrals on the $(n-1)$ -dimensional moduli spaces of holomorphic curves*:

$$\int_L \Omega = \int_{\mathcal{M}} \tilde{\Omega}_L, \quad \int_{L_i} \Omega = \int_{\mathcal{M}} \tilde{\Omega}_{L_i}, \quad i = 1, 2,$$

where \mathcal{M} is a shorthand for the weighted sum over contributions from all the $(n-1)$ -dimensional moduli spaces involved in the construction of \mathcal{C} , see Section A.2.

Recall the *positivity condition* means that if v_0 stands for a *clockwise oriented* tangent vector on $\partial\Sigma$, then $v_0 \wedge v_1 \wedge \dots \wedge v_{n-1}$ agrees with the orientation on L , and is *opposite* to the orientation on L' . The non-vanishing of $v_0 \wedge v_1 \wedge \dots \wedge v_{n-1}$ is a consequence of the *immersion property* from the automatic transversality (see Corollary 3.6, Proposition 3.8). The *almost calibrated* condition implies that $\text{Re } \Omega > 0$ on the Lagrangians with respect to the orientation on L and L' . Thus we make the following claim.

Claim 3.22 (Monotonicity). Clockwise along $\partial\Sigma$, the function $\text{Re } F$ is *increasing* on the L boundary portion, but *decreasing* on the $L' = L_1 \cup L_2$ boundary portion. In particular,

$$0 = \text{Re } F(q) \leq \text{Re } F \leq \text{Re } F(p).$$

More intrinsically, the real part of the complex volume forms on the moduli spaces are *non-negative*.

The holomorphic function F maps Σ into a bounded region in the complex plane. The behaviour at the corners is specified in Remark 3.8. Since each vertical line intersects $\partial F(\Sigma) \subset \mathbb{C}$ at ≤ 2 points by the monotonicity claim above, the boundary and corner local behaviours imply the following.

Claim 3.23 (Image curve). The image $F(\Sigma) \subset \mathbb{C}$ lies *above* its L' boundary portion, and *below* its L boundary portion.

We turn to the proof of the *Lagrangian phase angle inequality* (3.13). For each curve that contributes non-trivially to $\tilde{\Omega}_L$, by the monotonicity claim we can find a unique point r' on the L boundary of $\partial\Sigma$, such that

$$\begin{cases} \text{Re } F(r') = \text{Re } F(r), & \text{primary case,} \\ r' = q, & \text{secondary case, boundary on } L \text{ and } L_2, \\ r' = p, & \text{secondary case, boundary on } L \text{ and } L_1. \end{cases}$$

From the image curve claim, we always have $\operatorname{Im} F(r') \geq \operatorname{Im} F(r)$ in the primary case. Integrating over the moduli space of holomorphic curves, we obtain

$$\operatorname{Re} \int_{\mathcal{M}} F(r') = \operatorname{Re} \int_{\mathcal{M}} \tilde{\Omega}_{L_1} = \operatorname{Re} \int_{L_1} \Omega, \quad \operatorname{Im} \int_{\mathcal{M}} F(r') \geq \operatorname{Im} \int_{\mathcal{M}} \tilde{\Omega}_{L_1} = \operatorname{Im} \int_{L_1} \Omega.$$

We now introduce two almost everywhere defined functions χ_{A_1}, χ_{A_2} on L . The recipe is that at any generic point $P \in L$, if an automatically transverse holomorphic curve in the universal family passes through P on the boundary portion of $\partial\Sigma$ joining q to r' (resp. r' to p), then it gives an additive contribution to $\chi_{A_1}(P)$ (resp. $\chi_{A_2}(P)$) equal to the weighting factor of the curve. Intuitively χ_{A_1}, χ_{A_2} should be understood as the characteristic functions of weighted subsets $A_1, A_2 \subset L$. The positivity condition gives $\chi_{A_i} \geq 0$, and $\partial\mathcal{C} = L - L'$ gives $\chi_{A_1} + \chi_{A_2} = 1$. Intuitively A_1, A_2 give a (weighted) partition of L .

The moduli space integrals now have target space interpretations:

$$\int_{\mathcal{M}} F(r') = \int_L \chi_{A_1} \Omega =: \int_{A_1} \Omega, \quad \int_{\mathcal{M}} F(p) - F(r') = \int_L \chi_{A_2} \Omega =: \int_{A_2} \Omega.$$

Since L is homologous to $L_1 + L_2$, we have $\int_L \Omega = \int_{L_1} \Omega + \int_{L_2} \Omega$. Whence, we make the following claim.

Claim 3.24. There is a weighted partition $L = A_1 + A_2$ such that

$$\operatorname{Re} \int_{A_i} \Omega = \operatorname{Re} \int_{L_i} \Omega > 0, \quad \operatorname{Im} \int_{A_2} \Omega \leq \operatorname{Im} \int_{L_2} \Omega, \quad \operatorname{Im} \int_{A_1} \Omega \geq \operatorname{Im} \int_{L_1} \Omega.$$

Consequently,

$$\arg \int_{A_2} \Omega \leq \arg \int_{L_2} \Omega = \hat{\theta}_2 \quad \text{and} \quad \arg \int_{A_1} \Omega \geq \arg \int_{L_1} \Omega = \hat{\theta}_1,$$

so in particular $\inf_L \theta_L \leq \hat{\theta}_2$ and $\sup_L \theta_L \geq \hat{\theta}_1$.

Finally, we deal with the J -volume lower bound (3.14). By the triangle inequality,

$$\int_L e^{-i\theta} \Omega = \int_L |\Omega| = \int_{A_1} |\Omega| + \int_{A_2} |\Omega| \geq \left| \int_{A_1} \Omega \right| + \left| \int_{A_2} \Omega \right|.$$

The right-hand side is at least $|\int_{L_1} \Omega| + |\int_{L_2} \Omega|$, due to an elementary numerical fact.

Lemma 3.25. Let z, w be complex numbers, with fixed real parts $0 < \operatorname{Re}(z) < \operatorname{Re}(w)$. Then as a function of $\operatorname{Im}(z)$, the function $|z| + |w - z|$ is decreasing when $\arg z \leq \arg w$, and increasing when $\arg z \geq \arg w$.

This concludes the proof of (3.14). ■

A few remarks are in order to clarify the relevance to special Lagrangian geometry.

Remark 3.11. Recall that for L to be a special Lagrangian, then its phase angle is constant, and its J -volume is

$$\text{Vol}_J(L) = \left| \int_L \Omega \right| < \left| \int_{L_1} \Omega \right| + \left| \int_{L_2} \Omega \right|,$$

using the triangle inequality, the homological relation $[L] = [L_1 + L_2] \in H_n(X)$ and the assumption that $\hat{\theta}_1 > \hat{\theta}_2$. Thus the conclusion of the theorem is a *quantitative obstruction* for L to be special Lagrangian. In Section 3.6 we will discuss the relation to Joyce's LMCF program and the Bridgeland stability condition.

Remark 3.12. If L_1, L_2 are actually special Lagrangians, then the phase angle bounds given in (3.13) would be evident from the Floer degree formula (A.1) applied to the intersection points $L \cap L'$. One main feature of the theorem is that we do not need a priori knowledge on the existence of special Lagrangians, and the holomorphic volume form enters the obstruction criterion only through cohomological information.

Variant: Twisted complex case. The Floer theoretic obstruction for distinguished triangles can be easily generalized to involve many Lagrangians. Let L' be an exact immersed Lagrangian with bounding cochain built from the data of a twisted complex (3.1). We assume L is isomorphic to L' in $D^b \text{Fuk}(X)$, so we obtain a bordism current \mathcal{C} with $\partial \mathcal{C} = L - L' = L - \sum_1^N L_i$. As before, all Lagrangians are assumed to be almost calibrated, and all intersections are transverse.

Theorem 3.26 (Floer theoretic obstruction, multiple Lagrangian case). *Assume the automatic transversality and the positivity condition hold for the bordism current \mathcal{C} . Assume the destabilizing condition*

$$\hat{\theta}_1 > \hat{\theta}_2 > \cdots > \hat{\theta}_N, \quad \hat{\theta}_i = \arg \int_{L_i} \Omega.$$

Then the Lagrangian phase angle of L has a lower bound on its oscillation:

$$\sup_L \theta_L \geq \hat{\theta}_1, \quad \inf_L \theta_L \leq \hat{\theta}_N, \quad (3.16)$$

and moreover the J -volume of L has a non-trivial lower bound

$$\text{Vol}_J(L) = \int_L e^{-i\theta} \Omega \geq \sum_1^N \left| \int_{L_i} \Omega \right|. \quad (3.17)$$

Proof. Since most parts of the proof are identical to the distinguished triangle case, we will only sketch the main difference.

The Lagrangian boundaries on $\partial \Sigma$ are arranged in the *clockwise* order as $L, L_N, L_{N-1}, \dots, L_1$. We construct the holomorphic function F as in (3.15), and use it to produce complex valued volume forms on the $(n-1)$ -dimensional moduli spaces, such that

for any $m = 1, \dots, N$, we have

$$\int_L \Omega = \int_{\mathcal{M}} \tilde{\Omega}_L, \quad \int_{L_i} \Omega = \int_{\mathcal{M}} \tilde{\Omega}_{L_i}, \quad i = 1, 2, \dots, N.$$

As before, the real part of these complex volume forms are all non-negative, as a consequence of the positivity condition. The claim on the image curve $F(\Sigma)$ holds verbatim. Similarly to the distinguished triangle case, we produce non-negatively weighted subsets $A_1, \dots, A_N \subset L$ with $L = \sum A_i$, such that

$$\operatorname{Re} \int_{A_m} \Omega = \operatorname{Re} \int_{L_m} \Omega > 0, \quad \operatorname{Im} \sum_1^m \int_{A_i} \Omega \geq \operatorname{Im} \sum_1^m \int_{L_i} \Omega, \quad \int_L \Omega = \sum_1^N \int_{L_i} \Omega.$$

This implies

$$\arg \int_{A_1} \Omega \geq \arg \int_{L_1} \Omega, \quad \arg \int_{A_N} \Omega \leq \arg \int_{L_N} \Omega,$$

whence the phase inequality (3.16).

The J -volume can be bounded below by

$$\operatorname{Vol}_J(L) = \int_L |\Omega| = \sum_1^N \int_{A_i} |\Omega| \geq \sum_1^N \left| \int_{A_i} \Omega \right| \geq \sum_1^N \left| \int_{L_i} \Omega \right|.$$

The last step uses the purely numerical Lemma 3.27 below. ■

Lemma 3.27. *Let z_1, \dots, z_N be complex numbers, and a_1, \dots, a_N be fixed complex numbers with positive real parts, such that $\arg a_1 > \arg a_2 > \dots > \arg a_N$. Assume*

$$\operatorname{Re}(z_i) = \operatorname{Re}(a_i), \quad \operatorname{Im} \sum_1^m z_i \geq \operatorname{Im} \sum_1^m a_i, \quad \sum_1^N z_i = \sum_1^N a_i.$$

Then $\sum_1^N |z_i| \geq \sum_1^N |a_i|$.

Proof. We argue by induction. The $N = 2$ case is implied by Lemma 3.25. In general, we view $\sum_1^N |z_i|$ as a function of the imaginary parts of z_1, \dots, z_N subject to the constraints. Clearly, this function achieves its minimum for some (z_i) . If $z_1 = a_1$, then we can conclude by induction. Otherwise, $\operatorname{Im} z_1 > \operatorname{Im} a_1$. If $\arg z_2 < \arg z_1$, then we can fix $z_1 + z_2$ and decrease $\sum_1^2 |z_i|$ by Lemma 3.25, which would contradict minimality. Proceeding with this argument, we are forced to have

$$\arg z_1 \leq \arg z_2 \leq \dots \leq \arg z_N,$$

whence

$$\arg \sum_1^N z_i \geq \arg z_1 > \arg a_1 > \arg \sum_1^N a_i,$$

which contradicts $\sum_1^N z_i = \sum_1^N a_i$. ■

What if we relax the positivity condition? We now discuss a weaker version that does not require the positivity condition on holomorphic curves. This amounts to dropping pointwise positivity of the integrand for the moduli space integral, which breaks some parts of our mirror analogy.

As in Theorem 3.21, we consider L' built from two immersed Lagrangians L_1, L_2 , and L fits into the distinguished triangle. All Lagrangians are almost calibrated, and all intersections are transverse. We classify the automatically transverse holomorphic curves (see Corollary 3.6, Proposition 3.8) into \pm types according to whether the boundary evaluation to L agrees with the orientation on L or its opposite. The complex volume forms $\tilde{\Omega}_{L_i}$ on the moduli space can be split into the sum of two parts according to whether $u: \Sigma \rightarrow X$ is of \pm types:

$$\tilde{\Omega}_{L_i} = \tilde{\Omega}_{L_i}^+ + \tilde{\Omega}_{L_i}^-, \quad \int_{\mathcal{M}} \tilde{\Omega}_{L_i} = \int_{L_i} \Omega.$$

In particular, the signed measure $\text{Re } \tilde{\Omega}_{L_i}$ is decomposed into its positive and negative parts, and $\text{Re } \int_{\mathcal{M}} \tilde{\Omega}_{L_i}^- \leq 0$. We define

$$\hat{\theta}_i^+ = \arg \int_{\mathcal{M}} \tilde{\Omega}_{L_i}^+, \quad \hat{\theta}_i^- = \arg \left(- \int_{\mathcal{M}} \tilde{\Omega}_{L_i}^- \right), \quad i = 1, 2.$$

Here $\hat{\theta}_i^-$ is only defined when $\int_{\mathcal{M}} \tilde{\Omega}_{L_i}^-$ is non-zero, namely the case not covered already by the positivity condition.

Theorem 3.28 (Floer theoretic obstruction, relaxing positivity condition). *Assume the automatic transversality holds for the bordism current \mathcal{C} between L and L' . Then the Lagrangian phase angle of L has a lower bound on its oscillation:*

$$\sup_L \theta_L \geq \max\{\hat{\theta}_1^+, \hat{\theta}_1^-\}, \quad \inf_L \theta_L \leq \min\{\hat{\theta}_2^+, \hat{\theta}_2^-\}.$$

Proof. We will only sketch the modifications. The positivity condition enters through the Monotonicity Claim 3.22. Once we drop this, we would allow holomorphic discs $u: \Sigma \rightarrow X$ that sweep out parts of $L \cup L'$ with the reversed orientation. For such curves, Claim 3.22 is modified to the following.

Claim 3.29. Clockwise along $\partial\Sigma$, the function $\text{Re } F$ is *decreasing* on the L boundary portion, but *increasing* on the $L' = L_1 \cup L_2$ boundary portion. In particular,

$$0 = \text{Re } F(q) \geq \text{Re } F \geq \text{Re } F(p).$$

More intrinsically, the real part of the complex volume forms on the moduli spaces at such $u: \Sigma \rightarrow X$ are *non-positive*.

The corresponding Claim 3.23 is then modified as follows.

Claim 3.30. The image $F(\Sigma) \subset \mathbb{C}$ lies *below* its L' boundary portion, and *above* its L boundary portion.

At almost every point on $L \cup L'$, only automatically transverse holomorphic curves pass through it, since by assumption the boundary evaluation of the other holomorphic curves is contained in some subset of $L \cup L'$ with Hausdorff dimension $\leq n - 1$. According to the \pm types of the automatically transverse curves, we decompose the weighted characteristic functions χ_{A_i} on L into its positive and negative parts $\chi_{A_i^+} \geq 0$ and $\chi_{A_i^-} \leq 0$.

Then upon integration over the moduli space,

$$\begin{aligned} \operatorname{Re} \int_L \chi_{A_i^+} \Omega &= \operatorname{Re} \int_{\mathcal{M}} \tilde{\Omega}_{L_i}^+ \geq 0, & \operatorname{Re} \int_L \chi_{A_i^-} \Omega &= \operatorname{Re} \int_{\mathcal{M}} \tilde{\Omega}_{L_i}^- \leq 0, \\ \operatorname{Im} \int_L \chi_{A_2^+} \Omega &\leq \operatorname{Im} \int_{\mathcal{M}} \tilde{\Omega}_{L_2}^+, & \operatorname{Im} \int_L \chi_{A_1^+} \Omega &\geq \operatorname{Im} \int_{\mathcal{M}} \tilde{\Omega}_{L_1}^+, \\ \operatorname{Im} \int_L \chi_{A_2^-} \Omega &\geq \operatorname{Im} \int_{\mathcal{M}} \tilde{\Omega}_{L_2}^-, & \operatorname{Im} \int_L \chi_{A_1^-} \Omega &\leq \operatorname{Im} \int_{\mathcal{M}} \tilde{\Omega}_{L_1}^-. \end{aligned}$$

In particular,

$$\arg \int_L \chi_{A_2^+} \Omega \leq \hat{\theta}_2^+, \quad \arg \left(- \int_L \chi_{A_2^-} \Omega \right) \leq \hat{\theta}_2^-,$$

and

$$\arg \int_L \chi_{A_1^+} \Omega \geq \hat{\theta}_1^+, \quad \arg \left(- \int_L \chi_{A_1^-} \Omega \right) \geq \hat{\theta}_1^-.$$

Now $\chi_{A_i^+} \geq 0$ and $\chi_{A_i^-} \leq 0$, and the special case where $\chi_{A_i^-} = 0$ almost everywhere is already covered by the positivity condition. The theorem follows. \blacksquare

Remark 3.13. In the special case where L_1, L_2 are special Lagrangians of phase $\hat{\theta}_1, \hat{\theta}_2$, then clearly $\hat{\theta}_i^\pm = \hat{\theta}_i$. The conclusion in this case can be deduced easily from Floer degree considerations at Lagrangian intersections, similar to Section 2.2.

Remark 3.14. Recall $\hat{\theta}_i = \arg \int_{L_i} \Omega$. The caveat is that $\max\{\hat{\theta}_i^\pm\}$ and $\min\{\hat{\theta}_i^\pm\}$ do not quite control $\hat{\theta}_i$, so the above Theorem 3.28 *does not* imply the phase angle inequality (3.13). In this sense the conclusion of Theorem 3.28 is weaker than Theorem 3.21, illustrating the power of the positivity condition.

On the other hand, we will heuristically argue in Section 3.6 that in the Thomas–Yau–Joyce picture, once we assume the existence of Joyce’s Bridgeland stability condition, the phase angle inequality (3.13) can be deduced without the positivity condition in Theorem 3.21.

We think it is very interesting to either prove the positivity condition as a consequence of the other assumptions, or to find another Floer theoretic argument for (3.13) that requires neither the positivity condition, nor the a priori knowledge of special Lagrangian representatives.

3.6. Towards a Bridgeland stability condition

Joyce’s proposal and Bridgeland stability condition revisited. We now seek a better appreciation of the Bridgeland stability aspect of Joyce’s proposal (see Section 2.1). To

specify the Bridgeland stability condition on $D^b \text{Fuk}(X)$, we need the central charge

$$Z(L) = \int_L \Omega,$$

and all the subcategories $\mathcal{P}(\phi_0 < \phi < \phi_1)$ for any interval (ϕ_0, ϕ_1) . Joyce's proposal [45] strongly suggests two claims:

(i) If an unobstructed Lagrangian brane L has phase angle function $\theta \in (\pi\phi_0, \pi\phi_1)$, then L defines an object in the subcategory $\mathcal{P}(\phi_0 < \phi < \phi_1)$ generated by all stable objects with $\phi_0 < \phi < \phi_1$. This claim is because the infinite time limit of the LMCF should provide the stable objects which generate L , and by the monotonicity of the Lagrangian angle (see Section 4.1 below), we can predict a priori $\theta \in (\pi\phi_0, \pi\phi_1)$ for all these stable objects.

(ii) Any object in $\mathcal{P}(\phi_0 < \phi < \phi_1)$ can be generated by unobstructed Lagrangian branes with phase angle function $\theta \in (\pi\phi_0, \pi\phi_1)$.

Thus one can simply *define* $\mathcal{P}(\phi_0 < \phi < \phi_1)$ to be the subcategory of the derived Fukaya category (suitably enlarged to allow for immersed and singular objects) generated by all unobstructed Lagrangian branes L with $\theta \in (\pi\phi_0, \pi\phi_1)$, and then reconstruct $\mathcal{P}(\phi')$ as the intersection of all $\mathcal{P}(\phi_0 < \phi < \phi_1)$ for all $\phi_0 < \phi' < \phi_1$. Such a definition would make the Thomas–Yau proposal nearly tautological, and the difficult part of Joyce's proposal is to verify this indeed defines a Bridgeland stability. In fact, by the discussions in Sections 2.1 and 2.2, the only formidable part is the *Harder–Narasimhan decomposition*, for which Joyce's LMCF provides the conjectural mechanism.

There are two primary applications of the Thomas–Yau–Joyce proposal to keep in mind:

(i) The existence of special Lagrangians is important for geometric measure theory. A definition of stability conditions along the above lines is too tautological to be useful.

(ii) Defining special Lagrangian DT invariants is important for mirror symmetry (see Section 2.4). Knowing the existence of a Bridgeland stability condition on $D^b \text{Fuk}(X)$ is of great theoretic significance in view of Kontsevich and Soibelman's framework [47, 49], but without a more Floer theoretic characterization it would lack computability.

Thus even if Joyce's conjectures can be proved along the lines in [45], it is still desirable to have a Floer theoretic characterization of the Bridgeland stability condition. We first revisit Theorem 3.21 in the light of the Thomas–Yau–Joyce conjectural picture, but without assuming the automatic transversality and positivity condition.

Conjecture 3.31. *Suppose we have almost calibrated exact Lagrangian objects L_1 , L_2 , and L , fitting into a distinguished triangle $L_1 \rightarrow L \rightarrow L_2 \rightarrow L_1[1]$, and satisfies the destabilizing condition*

$$\hat{\theta}_1 = \arg \int_{L_1} \Omega > \hat{\theta}_2 = \arg \int_{L_2} \Omega.$$

Then the phase angle inequality (3.13) follows. In particular, the derived category class of L admits no special Lagrangian representative.

Proof (heuristic). Consider the Harder–Narasimhan decomposition (2.2) of L_1 :

$$0 = \mathcal{E}_0 \rightarrow \mathcal{E}_1 \rightarrow \cdots \rightarrow \mathcal{E}_N = L_1,$$

fitting into the distinguished triangles

$$\mathcal{E}_{i-1} \rightarrow \mathcal{E}_i \rightarrow L'_i \rightarrow \mathcal{E}_{i-1}[1],$$

where L'_i represents an object in $\mathcal{P}(\phi_i)$, with $\phi_1 > \phi_2 > \cdots > \phi_N$. Since, by assumption, L_1 is almost calibrated, we have $L_1 \in \mathcal{P}(-\frac{1}{2} < \phi < \frac{1}{2})$, hence $-\frac{1}{2} < \phi_i < \frac{1}{2}$. Since the central charges satisfy

$$Z(L_1) = \sum_1^N Z(L'_i), \quad \arg Z(L_i) = \pi\phi_i,$$

we must have $\hat{\theta}_1 \leq \pi\phi_1$. The conjectural description of the Bridgeland stability condition requires that L'_i has a special Lagrangian representative with constant Lagrangian phase $\pi\phi_i$. A weaker requirement which suffices for us is that there exists a representative L'_i with Lagrangian angle function $\theta_{L'_i}$ satisfying the oscillation bound $|\theta_{L'_i} - \pi\phi_i| < \varepsilon$ for any given $\varepsilon > 0$. It is expected that this flexibility allows one to assume sufficient smoothness on the Lagrangian.

By combining the distinguished triangles, we obtain a new distinguished triangle

$$L'_1 \rightarrow L \rightarrow L'' \rightarrow L'_1[1].$$

Here L'_1 , L , L'' are all almost calibrated. Suppose for contradiction that $\sup_L \theta_L < \hat{\theta}_1$. Then $\sup_L \theta_L < \pi\phi_1$, and we can arrange $\sup_L \theta_L < \inf_{L'_1} \theta_{L'_1}$. The Floer degree formula (A.1) implies $CF^0(L'_1, L) = 0$, and in particular $HF^0(L'_1, L) = 0$. The distinguished triangle splits:

$$L'' \simeq L \oplus L'_1[1].$$

Since L'' is almost calibrated, it lies in $\mathcal{P}(-\frac{1}{2} < \phi < \frac{1}{2})$, and so must $L'_1[1]$. But $L'_1 \in \mathcal{P}(\phi_1)$ implies $L'_1[1] \in \mathcal{P}(\phi_1 + 1)$. Since $\phi_1 + 1 > \frac{1}{2}$, we know

$$\mathcal{P}(\phi_1 + 1) \cap \mathcal{P}\left(-\frac{1}{2} < \phi < \frac{1}{2}\right) = \emptyset,$$

a contradiction. This proves $\sup_L \theta_L \geq \hat{\theta}_1$, subject to the conjectural existence of the Bridgeland stability condition.

A similar argument, beginning with the Harder–Narasimhan decomposition of L_2 , would show $\inf_L \theta_L \leq \hat{\theta}_2$. ■

Remark 3.15. In the above argument, once we achieved $CF^0(L'_1, L) = 0$, there is a different way to proceed. We reinterpret the distinguished triangle $L'_1 \rightarrow L \rightarrow L'' \rightarrow L'_1[1]$ as an isomorphism in $D^b \text{Fuk}(X)$ between L and a twisted complex built from $L'_1 \cup L''$. This

would give rise to a bordism current constructed from the universal family of holomorphic curves, with $\partial\mathcal{C} = L - L'_1 - L''$. However, $CF^0(L'_1, L) = 0$ implies that no holomorphic curve contributing to \mathcal{C} passes from L'_1 to L in the clockwise direction of $\partial\Sigma$. The twisted complex structure on $L'_1 \cup L''$ forbids the passage from L'_1 to L'' in the clockwise direction of $\partial\Sigma$. Thus if the holomorphic curve has any boundary portion on L'_1 , its entire boundary would lie on L'_1 , which cannot happen in the almost calibrated setting.

This motivates the following definition, whose precise meaning depends on the conjectural enlargement of the derived Fukaya category by incorporating singular Lagrangian objects.

Definition 3.32. Let L be an almost calibrated exact Lagrangian brane representing a class in the (suitably enlarged) derived Fukaya category. Suppose for any almost calibrated exact Lagrangian objects L_1, L_2 fitting into a distinguished triangle

$$L_1 \rightarrow L \rightarrow L_2 \rightarrow L_1[1],$$

we always have

$$\hat{\theta}_1 = \arg \int_{L_1} \Omega \leq \hat{\theta}_2 = \arg \int_{L_2} \Omega, \quad \text{resp. } \hat{\theta}_1 = \arg \int_{L_1} \Omega < \hat{\theta}_2 = \arg \int_{L_2} \Omega,$$

then we say L is *Thomas–Yau semistable* (resp. *strictly stable*). If L fails to be Thomas–Yau semistable, we say it is Thomas–Yau unstable.

We have attributed this definition to Thomas–Yau [82, 83], since it is in their spirit that stability conditions should be Floer theoretic conditions to be tested on the distinguished triangles, and that one should restrict attention only to almost calibrated Lagrangians. We now argue that if Joyce’s conjectural Bridgeland stability exists with its expected properties, then its semistable objects should agree with Thomas–Yau semistability.

Conjecture 3.33. *An almost calibrated exact Lagrangian brane L defines a semistable object in $D^b \text{Fuk}(X)$ under Joyce’s Bridgeland stability, if and only if it is Thomas–Yau semistable.*

Proof (heuristic). If the derived category class of L is semistable in Joyce’s sense, then we can choose an optimal representative which is a special Lagrangian, or at least has phase oscillation arbitrarily small. By Conjecture 3.31, we cannot have any destabilizing distinguished triangle, i.e., L is Thomas–Yau semistable.

Conversely, if L is not a semistable object in Joyce’s sense, then from its Harder–Narasimhan decomposition we can produce a destabilizing distinguished triangle

$$L_1 \rightarrow L \rightarrow L_2 \rightarrow L_1[1]$$

with almost calibrated L_1, L_2 , which violates Thomas–Yau semistability. ■

Having discussed the semistable objects, it is interesting to see what Joyce’s LMCF picture suggests about the Harder–Narasimhan decomposition.

Conjecture 3.34. *Assume further that the Kähler metric on X is Calabi–Yau. Suppose L is an almost calibrated exact Lagrangian brane in $D^b \text{Fuk}(X)$, with Harder–Narasimhan decomposition (2.2)*

$$0 = \mathcal{E}_0 \rightarrow \mathcal{E}_1 \rightarrow \cdots \rightarrow \mathcal{E}_N = L,$$

fitting into the distinguished triangles

$$\mathcal{E}_{i-1} \rightarrow \mathcal{E}_i \rightarrow L_i \rightarrow \mathcal{E}_{i-1}[1],$$

such that $L_i \in \mathcal{P}(\phi_i)$ with $\phi_1 > \cdots > \phi_N$. We have $\hat{\theta}_i = \pi\phi_i = \arg \int_{L_i} \Omega$. Then the phase angle inequality (3.16) and the volume lower bound (3.17) hold.

Proof (heuristic). In Joyce’s conjectural program, the Harder–Narasimhan decomposition is constructed by running the LMCF (L_t) starting from the unobstructed Lagrangian L , and take the infinite time limit (2.3) to obtain the limiting special Lagrangians L_1, \dots, L_N with angles $\hat{\theta}_1 > \hat{\theta}_2 > \cdots > \hat{\theta}_N$, assuming L_1, \dots, L_N have enough regularity to be admitted as objects of $D^b \text{Fuk}(X)$. It is expected that L_1, \dots, L_N generate L in $D^b \text{Fuk}(X)$ via (2.2).

A basic feature of LMCF in Calabi–Yau manifolds is that the Lagrangian angle satisfies a heat equation (see Section 4.1), so $\sup_{L_t} \theta$ is non-increasing in time (resp. $\inf_{L_t} \theta$ is non-decreasing). Comparing the initial time with the infinite time limit, this suggests $\sup \theta_L \geq \hat{\theta}_1$ and $\inf \theta_L \leq \hat{\theta}_N$.

Moreover, if the ambient metric is Calabi–Yau, then LMCF is a special case of mean curvature flow, so the volume functional decreases in time. This monotonicity is not affected by the surgeries in Joyce’s LMCF. Thus,

$$\text{Vol}(L) \geq \sum_1^N \text{Vol}(L_i) = \sum_1^N |Z(L_i)| = \sum_1^N \left| \int_{L_i} \Omega \right|.$$

Under the Calabi–Yau metric, the volume of the Lagrangian L agrees with the J -volume:

$$\text{Vol}(L) = \int_L e^{-i\theta} \Omega = \int_L |\Omega|,$$

so (3.17) follows. ■

Remark 3.16. Analogously in the context of HYM connections, if a holomorphic bundle E is unstable, then its Harder–Narasimhan decomposition provides a lower bound on the Yang–Mills energy of any Chern connection on E compatible with $\bar{\partial}_E$, which improves the topological energy bound. This type of phenomenon is common in Kähler geometry, for instance it also happens in the context of K-stability. These topics are covered in the introduction of [29].

In conclusion, Joyce’s conjectural picture suggests that if a Lagrangian object is unstable, then it satisfies certain angle and volume inequalities which quantitatively forbids it to be a special Lagrangian, and these obstructions detect the features of the Harder–Narasimhan decomposition. This should be compared with Theorem 3.26, which contains the main features of the obstructions, but makes no a priori reference to the LMCF or special Lagrangian representatives. The cost is that Theorem 3.21 and 3.26 require the positivity condition as an extra hypothesis.

Almost calibrated case: Categorical predictions of the Joyce picture. A complete characterization of Bridgeland stability conditions on a triangulated category, known since the inception of the subject [13, Proposition 5.3], is that $\mathcal{P}(\phi_0 < \phi \leq \phi_0 + 1)$ defines an abelian subcategory (‘the heart of a bounded t -structure’), and the central charge function on this abelian subcategory satisfies the Harder–Narasimhan condition.

In the Thomas–Yau–Joyce picture, $\mathcal{P}(-\frac{1}{2} < \phi \leq \frac{1}{2})$ essentially is the same as the subcategory of almost calibrated Lagrangians, if we assume there is no special Lagrangian of phase $\frac{\pi}{2}$, which holds as long as the discrete set of values of Ω -periods on $H_n(X, \mathbb{Z})$ miss the phase angle $\frac{\pi}{2}$. Then this picture would predict almost calibrated Lagrangians to form an *abelian category*, which moreover *generate* the entire derived Fukaya category using the shift operator. Morally, this is asserting that there are sufficiently many almost calibrated Lagrangians, which is evidently very deep since constructing geometric Lagrangian objects is known to be a difficult problem in symplectic topology. Another deep prediction of the existence of Bridgeland stability condition [45, Conjecture 3.6], is that the derived Fukaya category $D^b \text{Fuk}(X)$ (after incorporating immersed and singular Lagrangians with local systems) is automatically *idempotent complete*, so agrees with $D^\pi \text{Fuk}(X)$. These predictions, if correct, are very interesting structural results on the Fukaya category, but at the moment they are controversial.

In Section 5, we will set up a variational framework to find special Lagrangian representatives of $D^b \text{Fuk}(X)$ classes under the assumption of Thomas–Yau semistability. By restricting only to the subcategory of almost calibrated Lagrangians, our program evades these difficult structural claims on the entire $D^b \text{Fuk}(X)$. It would thus not have the same strength as the Joyce program, nor is it subject to the same falsification criteria.

3.7. Moduli integral formula for the Solomon functional

3.7.1. Moduli integral formula for the Solomon functional. Assuming automatic transversality, we can rewrite the Solomon functional (3.4) as a moduli space integral in terms of the notations introduced in Section 3.5. Let $u: \Sigma \rightarrow X$ be a holomorphic polygon, with first order deformation vector fields v_1, \dots, v_{n-1} , so we can define a holomorphic function F via (3.15). In clockwise order on $\partial\Sigma$, we encounter the degree one intersections on L , an intersection $p \in CF^0(L, L_0)$, the degree one self intersections on L_0 , and an intersection $q \in CF^0(L_0, L)$. As before, we fix the additive constant by $F(q) = 0$.

In the following calculation, we will use the complex orientation on Σ , and the *counterclockwise orientation* on $\partial\Sigma$. The Solomon functional contains a term

$$-\int_{\mathcal{C}} \lambda \wedge \text{Im}(e^{-i\hat{\theta}} \Omega).$$

Now $-\int_{\mathcal{C}} \lambda \wedge \Omega$ can be expressed as an integral of the following $(n-1)$ -form over the $(n-1)$ -dimensional moduli spaces \mathcal{M} of holomorphic curves:

$$\int_{\Sigma} \lambda \wedge \Omega(\cdot, v_1, \dots, v_{n-1}) = \int_{\Sigma} \lambda \wedge dF.$$

Notice that since $u: \Sigma \rightarrow X$ is a holomorphic curve and Ω is an $(n, 0)$ -form, adjusting v_i by a vector field tangent to Σ does not change this integrand, and all v_i must hit Ω instead of the 1-form λ . After integration by parts, we have

$$\int_{\Sigma} \lambda \wedge dF = \int_{\Sigma} F d\lambda - \int_{\partial\Sigma} F \lambda = \int_{\Sigma} F \omega - \int_{\partial\Sigma} F \lambda.$$

The Solomon functional contains another two terms

$$\int_L f_L \text{Im}(e^{-i\hat{\theta}} \Omega) \quad \text{and} \quad - \int_{L_0} f_{L_0} \text{Im}(e^{-i\hat{\theta}} \Omega).$$

Now $\int_L f_L \Omega$ can be expressed as an integral of the following $(n-1)$ -form over the moduli spaces \mathcal{M} :

$$-\int_{\partial\Sigma \cap L} f_L \Omega(\cdot, v_1, \dots, v_{n-1}) = - \int_{\partial\Sigma \cap L} f_L dF.$$

The abused notation $\partial\Sigma \cap L$ means the part of $\partial\Sigma$ mapping to L instead of L_0 . Similarly, $-\int_{L_0} f_{L_0} \Omega$ is the moduli space integral of the $(n-1)$ -form

$$-\int_{\partial\Sigma \cap L_0} f_{L_0} \Omega(\cdot, v_1, \dots, v_{n-1}) = - \int_{\partial\Sigma \cap L_0} f_{L_0} dF.$$

The extra minus sign comes from the fact that $\partial\mathcal{C}$ sweeps out the cycle $-L_0$ instead of L_0 .

Combining all the three contributions, the Solomon functional is the moduli space integral with integrand

$$\mathcal{I} = \text{Im} \int_{\Sigma} e^{-i\hat{\theta}} F \omega - \text{Im} \int_{\partial\Sigma \cap L} e^{-i\hat{\theta}} d(f_L F) - \text{Im} \int_{\partial\Sigma \cap L_0} e^{-i\hat{\theta}} d(f_{L_0} F).$$

The last two terms involve total derivatives, so can be integrated along boundary segments between the corner points, to yield

$$-\text{Im} \int_{\partial\Sigma \cap L} e^{-i\hat{\theta}} d(f_L F) - \text{Im} \int_{\partial\Sigma \cap L_0} e^{-i\hat{\theta}} d(f_{L_0} F) = \text{Im} \sum_{\text{all corners}} e^{-i\hat{\theta}} F f|_{-}^{+}, \quad (3.18)$$

where $f|_{-}^{+}$ stands for the difference of the potentials $f_{L_+} - f_{L_-}$ at a Lagrangian intersection point, such that $\partial\Sigma$ moves from L_+ to L_- in the clockwise direction. In the more

general framework of Floer theory with Novikov coefficients, $f|_{\pm}^{\pm}$ have the interpretation as the Novikov exponents of these intersection points. The bounding cochain elements have $f|_{\pm}^{\pm} \geq 0$, while $p \in CF^0(L, L_0)$, $q \in CF^0(L_0, L)$ may have negative Novikov exponents.

Proposition 3.35 (Moduli space integral formula). *The Solomon functional $\mathcal{S}(L)$ is the integral of the following complex valued volume form over the $(n-1)$ -dimensional moduli spaces of holomorphic curves:*

$$\mathcal{S}(L) = \int_{\mathcal{M}} \mathcal{I}, \quad \mathcal{I} = \operatorname{Im} \int_{\Sigma} e^{-i\hat{\theta}} F \omega + \operatorname{Im} \sum_{\text{all corners}} e^{-i\hat{\theta}} F f|_{\pm}^{\pm}. \quad (3.19)$$

Remark 3.17. The normalization $F(q) = 0$ is convenient, but changing F by a constant along Σ does not affect \mathcal{I} , due to the energy identity

$$\int_{\Sigma} \omega + \sum_{\text{all corners}} f|_{\pm}^{\pm} = 0.$$

Remark 3.18. We have focused the discussion on the holomorphic curves with boundary on both L and L_0 , which are the only curves relevant for the bordism current \mathcal{C} in the almost calibrated case. In general we need also curves involving corners at $CF^{-1}(L, L)$ or $CF^{-1}(L_0, L_0)$, and the formula (3.19) takes into account all these contributions.

3.7.2. Change of reference Lagrangians formula revisited. We now revisit Proposition 3.4 from the moduli space integral perspective, which we expect is better suited for generalization to compact Calabi–Yau settings. All transversality requirements of moduli spaces will be assumed, and in this sense the calculations below are formal.

In Remark 3.6 we sketched that under the extra assumption $HF^{-1}(L_0, L_0) = 0$, there is an $(n+2)$ -dimensional universal family $\tilde{\mathcal{C}}$ over some n -dimensional moduli $\tilde{\mathcal{M}}$, such that the boundary of $\tilde{\mathcal{C}}$ has three $(n+1)$ -dimensional contributions, corresponding up to sign to the three bordism currents $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$ between L_0, L'_0, L , which are in turn the universal families over the $(n-1)$ -dimensional moduli spaces \mathcal{M}_i for $i = 1, 2, 3$. Here \mathcal{M}_i can be viewed as certain boundary strata of the compactification of $\tilde{\mathcal{M}}$. The integrand \mathcal{I} naturally makes sense as an $(n-1)$ -form on $\tilde{\mathcal{M}}$, and restricts naturally to \mathcal{M}_i . The change of reference formula (3.6) amounts to

$$\int_{\mathcal{M}_3} \mathcal{I} = \int_{\mathcal{M}_1} \mathcal{I} + \int_{\mathcal{M}_2} \mathcal{I}. \quad (3.20)$$

Our strategy is to use *Stokes' formula on the moduli spaces*. As usual, the holonomy weighting factors will be suppressed in the moduli integral notations. Then (3.20) reduces to the following two claims.

Claim 3.36. The n -form $d\mathcal{I} = 0$ over the moduli space $\tilde{\mathcal{M}}$.

Claim 3.37. The Stokes boundary term is

$$\int_{\tilde{\mathcal{M}}} d\mathcal{I} = \int_{\mathcal{M}_1} \mathcal{I} + \int_{\mathcal{M}_2} \mathcal{I} - \int_{\mathcal{M}_3} \mathcal{I}.$$

We first explain Claim 3.36. First, we calculate the derivatives of F . Let y_1, \dots, y_n be local coordinates on $\tilde{\mathcal{M}}$, so $\frac{\partial}{\partial y_i}$ can be identified as first order deformations of holomorphic curves. The local coordinates on Σ are denoted as s, t . We can write

$$F = \sum F_i (-1)^{i-1} dy_1 \wedge \cdots \wedge \widehat{dy_i} \wedge \cdots \wedge dy_n,$$

such that along Σ ,

$$\begin{aligned} \partial_s F_i &= (-1)^{i-1} \Omega\left(\frac{\partial}{\partial s}, \frac{\partial}{\partial y_1}, \dots, \widehat{\frac{\partial}{\partial y_i}}, \dots, \frac{\partial}{\partial y_n}\right), \\ \partial_t F_i &= (-1)^{i-1} \Omega\left(\frac{\partial}{\partial t}, \frac{\partial}{\partial y_1}, \dots, \widehat{\frac{\partial}{\partial y_i}}, \dots\right). \end{aligned}$$

The holomorphic volume form satisfies $d\Omega = 0$, whence

$$\sum_{i=1}^n \partial_s \partial_i F_i = \partial_s \left(\Omega\left(\frac{\partial}{\partial y_1}, \dots, \frac{\partial}{\partial y_n}\right) \right), \quad \sum_{i=1}^n \partial_t \partial_i F_i = \partial_t \left(\Omega\left(\frac{\partial}{\partial y_1}, \dots, \frac{\partial}{\partial y_n}\right) \right).$$

Notice $\Omega\left(\frac{\partial}{\partial y_1}, \dots, \frac{\partial}{\partial y_n}\right)$ vanishes at the corners due to the decay of the first order deformation vector fields, and comparing with the additive normalization convention on F , we find

$$\sum_{i=1}^n \partial_i F_i = \Omega\left(\frac{\partial}{\partial y_1}, \dots, \frac{\partial}{\partial y_n}\right). \quad (3.21)$$

In particular, at all the corner points $\sum_i \partial_i F_i = 0$. The term $f|^\pm$ at the corners are independent of the moduli space parameters, so the only contribution to $d\mathcal{I}$ comes from the $\text{Im} \int_{\Sigma} e^{-i\hat{\theta}} F\omega$ term in formula (3.19).

We calculate

$$\sum_{i=1}^n \partial_i \int_{\Sigma} F_i \omega = \int_{\Sigma} \sum_i \partial_i F_i \omega + \int_{\Sigma} \sum_i F_i \partial_i \omega.$$

Here $\partial_i \omega$ is the Lie derivative of the symplectic form ω with respect to the vector field $\frac{\partial}{\partial y_i}$, which by Cartan's formula is

$$\partial_i \omega = d\left(\omega\left(\frac{\partial}{\partial y_i}, \cdot\right)\right) + \iota_{\frac{\partial}{\partial y_i}} d\omega = d\left(\omega\left(\frac{\partial}{\partial y_i}, \cdot\right)\right).$$

Thus,

$$\int_{\Sigma} \sum_i F_i \partial_i \omega = \int_{\partial\Sigma} \sum_i F_i \omega\left(\frac{\partial}{\partial y_i}, \cdot\right) - \int_{\Sigma} \sum_i dF_i \wedge \omega\left(\frac{\partial}{\partial y_i}, \cdot\right).$$

Notice that $\partial\Sigma$ and $\frac{\partial}{\partial y_i}$ are both tangent to the Lagrangian boundary, so the $\partial\Sigma$ integrand vanishes. We are left with

$$\begin{aligned} \sum_{i=1}^n \partial_i \int_{\Sigma} F_i \omega &= \int_{\Sigma} \sum_i \partial_i F_i \omega - \int_{\Sigma} \sum_i dF_i \wedge \omega \left(\frac{\partial}{\partial y_i}, \cdot \right) \\ &= \int_{\Sigma} \Omega \left(\frac{\partial}{\partial y_1}, \dots, \frac{\partial}{\partial y_n} \right) \omega \\ &\quad + \sum_i (-1)^{i-1} \omega \left(\frac{\partial}{\partial y_i}, \cdot \right) \wedge \Omega \left(\cdot, \frac{\partial}{\partial y_1}, \dots, \widehat{\frac{\partial}{\partial y_i}}, \dots, \frac{\partial}{\partial y_n} \right). \end{aligned}$$

Here we used the definition of F_i via $\partial_s F_i$, $\partial_t F_i$, and the formula (3.21) for $\sum_i \partial_i F_i$. We contract the identity $\omega \wedge \Omega = 0$ with $\frac{\partial}{\partial y_1}, \dots, \frac{\partial}{\partial y_n}$. When two $\frac{\partial}{\partial y_i}$ hit ω , the Ω term will be contracted only $(n-2)$ times, which produces a $(2, 0)$ -form vanishing identically on the holomorphic curve Σ . When at most one $\frac{\partial}{\partial y_i}$ hits ω , we obtain the above integrand. In effect, the integrand vanishes identically:

$$\sum_{i=1}^n \partial_i \int_{\Sigma} F_i \omega = 0,$$

which then implies $d\mathcal{I} = 0$.

We next explain Claim 3.37. In general, the compactified moduli space has many boundary strata corresponding to disc bubbling and disc splitting.

Claim 3.38. Only the boundary strata corresponding to gluing holomorphic curves with virtual dimension 0 and $n-1$, can have non-zero contributions to the Stokes boundary term.

To see this, we need to understand how \mathcal{I} (and notably F) behaves near the boundary of the moduli space. Recall that when the holomorphic disc is degenerating to several disc components, then under transversality conditions, the cokernel of the extended linearized Cauchy–Riemann operator vanishes, and for small fixed gluing parameters, the kernel elements (i.e., first order deformations) are up to small perturbation obtained by gluing the kernel elements from the degenerate disc components. The perturbation effect tends to zero as we approach the moduli space boundary. Now the kernel elements from different disc components have essentially disjoint supports, so unless we have at least $(n-1)$ kernel elements supported on one disc component such that $\Omega(\cdot, v_1, \dots, v_{n-1})$ does not vanish identically, we will have $dF = 0$ for the moduli boundary strata, so that $F = \text{const}$ along Σ , whence $\mathcal{I} = 0$ by Remark 3.17. This shows Claim 3.38. We comment that this phenomenon is closely related to the fact that many moduli boundary strata do not contribute to the boundary of the bordism current \mathcal{C} due to support reasons (see Section 3.1).

On the boundary strata, the only contributions to the integral \mathcal{I} come from the $(n-1)$ -dimensional moduli spaces. The role of the holomorphic curves of virtual dimension zero,

is to provide the counting factors, in a manner entirely analogous to Section 3.1.2. Most contributions cancel out due to the Maurer–Cartan equation on the bounding cochains, and the closedness of the HF^0 generators. The remaining contributions produce the right-hand side in Claim 3.37.

3.7.3. First variation formula revisited. In this subsection we explain how to semiheuristically understand the first variation formula (3.5) as a consequence of Proposition 3.4, from the perspective of the moduli space integral formula (3.19). We hope this viewpoint is better suited for generalization to compact Calabi–Yau settings.

Suppose we are given a 1-parameter exact isotopy of unobstructed exact immersed Lagrangians L_t , and we wish to calculate $\frac{d}{dt} \mathcal{S}_{L_0}(L_t)$ at $t = 0$. The change of reference Lagrangian formula (see Proposition 3.4) allows us to replace $\mathcal{S}_{L_0'}(L_t)$ by $\mathcal{S}_{L_0}(L_t)$. The Lagrangian L_t for $|t| \ll 1$ is approximately the graph of $t dh$ in T^*L_0 (understood in an immersed sense), for the Hamiltonian function $h = h_t|_{t=0}$ on L_0 . The holomorphic discs between L_t and L_0 for $|t| \ll 1$ have *small energy* of order $O(|t|)$, and are locally approximated by Morse trajectories of h . Write X_h as the Hamiltonian vector field, namely $\omega(X_h, \cdot) = dh$, then

$$\left. \frac{d}{dt} \right|_{t=0} \int_{\Sigma} F \omega = \int_{\partial \Sigma \cap L_0} F \omega(\cdot, X_h) = - \int_{\partial \Sigma \cap L_0} F dh.$$

Now we examine $\sum_{\text{all corners}} F f|_{\pm}^{\pm}$ for very small t . The Lagrangian intersections come in two types:

- (i) The corners $p \in CF^0(L_t, L_0)$ and $q \in CF^0(L_0, L_t)$ correspond to the local extrema of the Hamiltonian h .
- (ii) Any self intersection p_i between two local sheets L_+, L_- of L_0 can be paired with a very nearby self intersection of p_i^t between two sheets L_+^t, L_-^t of L_t . The bounding cochain on L_t is thus induced from the bounding cochain on L_0 .

At the intersection points p, q , we have

$$\begin{cases} f|_{\pm}^{\pm}(q) = f_{L_0}(q) - f_{L_t}(q) = -th(q) + O(t^2), \\ f|_{\pm}^{\pm}(p) = f_{L_t}(p) - f_{L_0}(p) = th(p) + O(t^2). \end{cases}$$

The self intersections are usually not important here, because the smallness of energy prevents their appearance on $\partial \Sigma$, unless $f|_{\pm}^{\pm}(p_i^t) = O(|t|)$, and $f|_{\pm}^{\pm}(p_i) = 0$, which is a rather non-generic situation. When the self intersections do appear, the evolution of the potential under exact isotopy gives

$$\begin{aligned} (f_{L_+^t} - f_{L_-^t})(p_i^t) &= (f_{L_+} - f_{L_-})(p_i) + \int_0^t (h_{+, \tau} - h_{-, \tau})(p_i^{\tau}) d\tau \\ &= f|_{\pm}^{\pm}(p_i) + th|_{\pm}^{\pm}(p_i) + O(t^2), \end{aligned}$$

where h_{\pm} keeps track of the Hamiltonian on the different sheets of L_0 . Thus

$$f|_{\pm}^{\pm}(p_i^t) = f_{L_-^t}(p_i^t) - f_{L_+^t}(p_i^t) = -f|_{\pm}^{\pm}(p_i) - th|_{\pm}^{\pm}(p_i) + O(t^2).$$

Here we have a tricky sign reversal, because if L_+ and L_- are clockwise ordered on $\partial\Sigma$, then L_+^t and L_-^t are counterclockwise ordered. In summary,

$$\frac{d}{dt}\Big|_{t=0} \sum_{\text{all corners}} Ff|_{-}^+(t) = \lim_{t \rightarrow 0} \left\{ -Fh(q) + Fh(p) - \sum_{L_0 - \text{ self intersection corners}} Fh|_{-}^+(p_i) \right\}.$$

Combining the above, and integrating by parts, we obtain

$$\begin{aligned} \frac{d}{dt}\Big|_{t=0} & \left(\int_{\Sigma} F\omega + \sum_{\text{all corners}} Ff|_{-}^+(t) \right) \\ &= \lim_{t \rightarrow 0} \left\{ -Fh(q) + Fh(p) - \sum_{L_0 - \text{corners}} Fh|_{-}^+(p_i) - \int_{\partial\Sigma \cap L_0} F dh \right\} \\ &= \lim_{t \rightarrow 0} \left\{ \int_{\partial\Sigma \cap L_0} h dF \right\}. \end{aligned}$$

Observe that for very small t , as the holomorphic curves vary in the $(n-1)$ -dimensional moduli spaces \mathcal{M} , under the *counterclockwise* sign convention for $\partial\Sigma$, the boundary evaluation of $\partial\Sigma \cap L_0$ sweeps out the cycle L_0 (beware of the sign!), and any generic point on L_0 is swept out precisely once due to the Morse theory limiting description. Consequently, we have the moduli space integral

$$\lim_{t \rightarrow 0} \int_{\mathcal{M}} \left\{ \int_{\partial\Sigma \cap L_0} h dF \right\} = \int_{L_0} h\Omega,$$

hence

$$\frac{d}{dt}\Big|_{t=0} \int_{\mathcal{M}} \left(\int_{\Sigma} F\omega + \sum_{\text{all corners}} Ff|_{-}^+(t) \right) = \int_{L_0} h\Omega.$$

By the moduli integral formula (3.19) of the Solomon functional,

$$\frac{d}{dt}\Big|_{t=0} \mathcal{S}(L_t) = \frac{d}{dt}\Big|_{t=0} \int_{\mathcal{M}} \mathcal{I} = \int_{L_0} h \operatorname{Im}(e^{-i\hat{\theta}} \Omega).$$

This recovers the first variation formula (3.5).

3.7.4. Speculations on compact Calabi–Yau manifolds. Floer theoretic foundations are much more complicated beyond the exact setting, and the foundations concerning the open-closed string map in the immersed Fukaya category setting are not fully written out in the literature. Nonetheless, due to the interest of the topic, we shall offer some speculations about how the Solomon functional formula (3.19) generalizes to the compact almost Calabi–Yau setting. Our local systems will have coefficients in \mathbb{R}, \mathbb{Q} , i.e., the parallel transport in the local system have only Novikov exponent zero components. This convention is somewhat more restrictive than [4, 67, 68].

Remark 3.19. In Joyce’s LMCF, bounding cochains and local systems can be created ex nihilo during the flow, but in all the mechanisms the author is aware of, the flow preserves the above class of local systems.

First, we recall the role of *Novikov coefficients* in Floer theory. All Floer cochain spaces CF^* are modules over the Novikov field

$$\Lambda = \left\{ \sum_i a_i T^{\lambda_i} : \lambda_i \in \mathbb{R}, \lambda_1 < \lambda_2 < \dots \rightarrow +\infty \right\},$$

and $a_i \in \mathbb{Q}$ or \mathbb{R} depending on the coefficient field choice.⁴² There are a few conventions to define A_∞ -operations. Let L be a compact immersed Lagrangian with transverse self intersections. In the Morse model [67], the self Floer cochain space $CF^*(L, L)$ is generated by the Morse critical points on L , and the ordered self intersections (twisted by local system hom and orientation factors as usual). The Fukaya A_∞ -algebra is a collection of Novikov-multilinear operations

$$m_k: CF^*(L, L) \otimes \dots \otimes CF^*(L, L) \rightarrow CF^*(L, L)[2 - k]$$

defined by counting *holomorphic treed discs* $u: \Sigma \rightarrow X$ (see [67, Definition 3.1]), weighted by the holonomy and orientation factors, and an *energy factor* $T^{E(u)}$. Very roughly, the domain Σ have surface parts (which consist of discs, and spheres attached to them), and tree parts connecting the disc boundaries. Then u is a holomorphic map with Lagrangian boundary on the surface parts, and Morse gradient flowlines on the tree parts. The role of CF^* elements is to specify the limiting behaviour of the Morse flowlines, and the Lagrangian self intersections on $\partial\Sigma$. The energy $E(u) = \int_\Sigma \omega$ is the sum of $\int u^* \omega$ on all the surface parts of Σ .

A non-trivial fact is that (after complicated perturbation schemes, or virtual techniques) this gives rise to a *curved A_∞ -algebra* structure [67]. The most important new feature, absent in the exact case, is that the disc bubbling can occur at points of L , which are not necessarily self intersection points. The domain disc splits into two discs, attached at a boundary node. This phenomenon is compensated by considering two discs joined by a gradient flowline segment, whose length shrinks to zero, producing the same nodal discs in the degeneration limit. With the appropriate weights and orientations taken into account, these two effects would cancel algebraically. On the other hand, the length parameter of the tree parts can tend to infinity, causing the Morse gradient flow line to break, a phenomenon which contributes to the boundary of the 1-dimensional moduli spaces, reflected algebraically in the A_∞ -relations.

Similar to the exact immersed case (see Appendix A.2), the bounding cochains are $b \in CF^1(L, L)$ elements satisfying the non-negative Novikov exponent requirement, and the Maurer–Cartan equation

$$m_0 + m_1(b) + m_2(b, b) + \dots = 0.$$

⁴²We do not know if the Fukaya category can be defined over integers in general. One should not confuse the coefficients a_i with the Novikov exponents λ_i . Typically a_i are rational numbers related to counting, while λ_i are real numbers related to the energy.

Generally speaking, the sum is infinite, but after truncating the Novikov series at any given high energy, only finitely many terms appear due to Gromov compactness, so the sum makes formal sense. As usual, the Lagrangian with bounding cochain structures are called unobstructed.

The framework for setting up the Fukaya algebra of a single immersed Lagrangian, also assigns meanings to Floer cohomologies between two immersed Lagrangians with bounding cochain structures. Suppose $\alpha \in CF^0(L, L')$ and $\beta \in CF^0(L', L)$ represent elements in $HF^0(L, L')$ and $HF^0(L', L)$ whose cohomological compositions are the identities. The α, β are in generally represented by infinite series in the Novikov variable T , where some Novikov exponents may be negative, and may not be bounded above, but at least they are bounded from below depending on α, β . We now speculate that there is a bordism current \mathcal{C} with $\partial\mathcal{C} = L - L'$, constructed from the universal families of treed holomorphic discs over the $(n-1)$ -dimensional moduli spaces \mathcal{M} . The monomial summands of α, β and the bounding cochain elements prescribe the corners of the treed holomorphic discs, and monomials with different Novikov exponents are viewed as independent contributions to \mathcal{M} and \mathcal{C} . Beyond the almost calibrated case, one would also need to incorporate degree -1 self intersections as usual. We think the moduli spaces that contribute to \mathcal{C} would satisfy the Novikov exponent condition

$$\int_{\Sigma} \omega + \sum_{\text{all corners}} \text{Novikov exponent} = 0. \quad (3.22)$$

Here the corners include the monomial summands of α, β (or the degree -1 self intersections as appropriate), and the bounding cochains at the degree one self intersections/Morse critical points of L, L' . Since all Novikov exponents at the bounding cochains are non-negative (not so at α, β , and the degree -1 self intersections!), this condition would impose an energy upper bound on the holomorphic treed discs depending on α, β , whence *only finitely many moduli spaces contribute to the bordism current*.

Remark 3.20. The Novikov exponents correspond to $f|_{\pm}^{\pm}$ in the exact case. While in the exact case (3.22) is an automatic consequence of the energy identity (A.4), in general it is an extra condition on the moduli spaces. It sits well with the fact that the geometric unit has zero Novikov exponent.

Now the moduli space integral formula (3.19) *formally* makes sense almost verbatim, ignoring all virtual perturbation nuances. For first order deformations v_1, \dots, v_{n-1} of the holomorphic treed discs, we can define F on the domain Σ via the 1-form $dF = \Omega(\cdot, v_1, \dots, v_{n-1})$. On the surface parts of Σ , we would obtain a holomorphic function F by complex integrability as usual (which must be constant on the holomorphic sphere components by the Liouville theorem), while on the tree parts, there is no obstruction for the 1-form to be exact. Next, we replace the appearance of $f|_{\pm}^{\pm}$ in (3.19) by the Novikov exponents of the monomial summands at the corners, to define the moduli integrand \mathcal{I} . The term $\int_{\Sigma} F\omega$ is understood to only involve integration on the surface parts of Σ . The

Solomon functional $\mathcal{S}(L)$ still has the form $\int_{\mathcal{M}} \mathcal{I}$. Notice that adding a constant to F would not change the moduli integrand \mathcal{I} , thanks to (3.22).

In the absence of the Lagrangian potential, the Novikov exponents of α, β are no longer canonically fixed. Suppose we replace α by $T^\mu \alpha$, and β by $T^{-\mu} \beta$, for some $\mu \in \mathbb{R}$. This would affect the moduli integrand \mathcal{I} , by the amount

$$\mu \operatorname{Im}\{e^{-i\hat{\theta}}(F(p) - F(q))\},$$

where p, q stand for the components of α, β . By analogy with the exact case, we expect

$$\operatorname{Im}\left(e^{-i\hat{\theta}} \int_{\mathcal{M}} F(p) - F(q)\right) = \operatorname{Im}\left(e^{-i\hat{\theta}} \int_L \Omega\right) = 0,$$

whence $\mathcal{S}(L)$ is independent of μ .

Once the foundations are in place, we expect the following result.

Conjecture 3.39. *Fix a compact almost Calabi–Yau manifold X . The Solomon functional is well defined for graded immersed unobstructed Lagrangians in the same $D^b \operatorname{Fuk}(X)$ class of a reference Lagrangian L_0 , satisfying that*

- *the change of reference Lagrangian formula (3.6) holds;*
- *gauge equivalent bounding cochains give rise to the same functional;*
- *cohomologous choices of HF^0 generators α, β give rise to the same functional;*
- *the first variation formula (3.5) holds for any 1-parameter exact isotopy of unobstructed Lagrangians.*

The slogan is that Floer theory should fix the multi-valuedness problem of the Solomon functional (see Section 2.8). As a more technical observation, once the change of reference Lagrangian formula (3.6) is established, one can remove the assumption for L to be transverse to L_0 , using a perturbation L'_0 of L_0 .

3.8. More applications of moduli space integrals

We collect a number of further topics involving the moduli space integral technique. Sections 3.8.2 and 3.8.3 are applications of the moduli integral formula (3.19) for the Solomon functional.

3.8.1. Lotay–Pacini convexity. Lotay and Pacini proved the convexity of their J -functional (see Proposition 2.13) through rather heavy calculations, so it is instructive to see that in the Calabi–Yau case, this result has a much simpler conceptual argument.

We interpret their ‘geodesic’ as a bordism current \mathcal{C} between two Lagrangians L, L' , constructed from universal families of holomorphic curves, such that automatic transversality and the positivity condition hold. In their highly idealized setting, only holomorphic strips $\Sigma \simeq \mathbb{R}_s \times [0, 1]_t$ appear in the construction of \mathcal{C} . We define the holomorphic function F as usual. The 1-parameter family of totally real submanifolds is given by the

constant t -coordinate slices \mathcal{C}_t of \mathcal{C} , whose J -volume functional is expressible through moduli space integrals

$$\text{Vol}_J(\mathcal{C}_t) = \int_{\mathcal{C}_t} |\Omega| = \int_{\mathcal{M}} \int_{\mathbb{R}_s \times \{t\}} \left| \frac{\partial F}{\partial s} \right| ds.$$

Since F is holomorphic, so is $\frac{\partial F}{\partial s}$, whence $\left| \frac{\partial F}{\partial s} \right|$ is subharmonic, which combined with the exponential decay at $s \rightarrow \pm\infty$ implies the convexity of the function in t :

$$\int_{\mathbb{R}_s \times \{t\}} \left| \frac{\partial F}{\partial s} \right| ds.$$

Thus $\text{Vol}_J(\mathcal{C}_t)$ is convex as a function of t , as Lotay and Pacini observed.

3.8.2. Lower bound of the Solomon functional. The theme of Section 5 will be on the variational approach to find special Lagrangians by minimizing the Solomon functional in a fixed derived category class. As an important motivation, special Lagrangians are *formal local minimizers* of the Solomon functional under Hamiltonian deformations (see Section 2.8). In fact, we can do better *under the automatic transversality and the positivity condition*.

Proposition 3.40 (Special Lagrangians are minimizers). *Suppose that L_0 is an exact immersed special Lagrangian of phase $\hat{\theta} \in (-\frac{\pi}{2}, \frac{\pi}{2})$ with unobstructed bounding cochain structure. We let L be an almost calibrated, exact, immersed Lagrangian in the same $D^b \text{Fuk}(X)$ class, which intersects L_0 transversely. Suppose the bordism current \mathcal{C} with $\partial\mathcal{C} = L - L_0$ satisfies automatic transversality and the positivity condition. Then, we have that $\mathcal{S}(L) \geq \mathcal{S}(L_0)$.*

Proof. The incline angle of the tangent vector to $F(\partial\Sigma) \subset \mathbb{C}$ is equal to the Lagrangian angle modulo $\pi\mathbb{Z}$. Since L_0 is a special Lagrangian, along the L_0 boundary portion $\arg F = \hat{\theta}$. Thus $\text{Im}(e^{-i\hat{\theta}} F) = 0$ at p, q and the self intersections on L_0 . The Solomon functional integrand simplifies to

$$\text{Im} \int_{\Sigma} e^{-i\hat{\theta}} F \omega + \sum_{L\text{-self intersections on } \partial\Sigma} \text{Im}(e^{-i\hat{\theta}} F) f_L|_{\pm}^{\pm}.$$

By the almost calibrated assumption on L, L_0 , and the positivity condition, we obtain Claim 3.23, namely $F(\Sigma)$ lies above its L_0 boundary,

$$\text{Im}(e^{-i\hat{\theta}} F) \geq 0 \quad \text{on } \Sigma.$$

Moreover, the Novikov positivity requirement for the bounding cochain on L says that $f_L|_{\pm}^{\pm} \geq 0$ at the degree one self intersections on $\partial\Sigma \cap L$. Thus the Solomon functional integrand is non-negative, which implies $\mathcal{S}(L) \geq 0 = \mathcal{S}(L_0)$. \blacksquare

Remark 3.21. Suppose we drop the positivity condition, then the key step $\text{Im}(e^{-i\hat{\theta}} F) \leq 0$ would break down, so the above proof of Proposition 3.40 would be invalidated. However, the conclusion may still be true (see Section 5.7.1).

3.8.3. Bounded part of the Solomon functional. Let L, L_0 be both exact, immersed Lagrangians with unobstructed bounding cochain structures, lying in the same $D^b \text{Fuk}(X)$ class, such that all intersections are transverse. Assume the bordism current \mathcal{C} with $\partial \mathcal{C} = L - L_0$ satisfies automatic transversality and the positivity condition. We consider L_0 as a fixed reference Lagrangian, while L can vary. We wish to find uniform a priori bound on certain parts of the Solomon functional, under natural conditions on L .

We shall assume the following conditions.

Quantitative almost calibratedness. Both L and L_0 have Lagrangian phase angles within $[-\frac{\pi}{2} + \varepsilon, \frac{\pi}{2} - \varepsilon]$ for some fixed small constant ε .

Potential clustering. (See Lemma A.3.) The immersed Lagrangian L can be represented by a twisted complex (3.1) built from the immersed Lagrangians L_1, \dots, L_N , such that the oscillation of the Lagrangian potentials have uniform bounds

$$\sup_{L_i} f_{L_i} - \inf_{L_i} f_{L_i} \leq A,$$

while for any $i > j$,

$$\sup_{L_j} f_{L_j} \leq \inf_{L_i} f_{L_i}.$$

Without loss of generality, we also assume

$$\sup_{L_0} f_{L_0} - \inf_{L_0} f_{L_0} \leq A$$

for the fixed Lagrangian L_0 .

Proposition 3.41 (Uniform energy bound). *Under the potential clustering assumption, all holomorphic polygons $u: \Sigma \rightarrow X$ with boundary on L and L_0 contributing to \mathcal{C} have uniformly bounded energy independent of L :*

$$E(u) = \int_{\Sigma} u^* \omega \leq A(N + 1),$$

and along $\partial \Sigma$ the degree one self intersections of L_0, L_1, \dots, L_N arising from the bounding cochains satisfy a uniform bound

$$\sum_i \sum_{b_i} f_{L_i}|_+^+(b_i) \leq A(N + 1).$$

Proof. We consider holomorphic polygons whose boundary $\partial \Sigma$ encounters in the clockwise order intersections in

$$\begin{aligned} p_{N-1} &\in CF^1(L_N, L_{N-1}), \dots, p_1 \in CF^1(L_2, L_1), \\ p_0 &\in CF^0(L_1, L_0), \quad p_N \in CF^0(L_0, L_N), \end{aligned}$$

juxtaposed possibly by more degree one self intersections b_i of L_i . The notation here does not constrain the number of self intersections of L_i that can occur on $\partial\Sigma$. The topological energy formula (A.4) expresses $E(u)$ in terms of the Lagrangian potentials at the intersections

$$\begin{aligned} E(u) + \sum_i \sum_{b_i} f_{L_i}|_{-}^+(b_i) &= f_{L_N}(p_N) - f_{L_0}(p_N) + \sum_{i=0}^{N-1} (f_{L_i} - f_{L_{i+1}})(p_i) \\ &= f_{L_0}(p_0) - f_{L_0}(p_N) + \sum_{i=1}^N (f_{L_i}(p_i) - f_{L_i}(p_{i-1})) \\ &\leq (N+1)A. \end{aligned}$$

By the Novikov positivity requirement of the bounding cochains $f_{L_i}|_{-}^+(b_i) \geq 0$, and the energy of the holomorphic curve is also positive, so they are individually bounded.

More generally, the polygons may miss some of the Lagrangians in L_1, \dots, L_N , but cannot reverse the order of the Lagrangians. This amounts to using a smaller effective value N , and the same argument implies the energy bound. ■

We now consider the holomorphic function F as before. Recall by Claim 3.22 we have $0 \leq \operatorname{Re}(F) \leq \operatorname{Re} F(p)$, where p is the intersection point in $CF^0(L, L_0)$. In fact, $F(\Sigma)$ must be contained in a triangular region determined by $\operatorname{Re} F(p)$.

Lemma 3.42 (Wedge region bound). *Under the quantitative almost calibrated hypothesis, we have $|\arg F| \leq \frac{\pi}{2} - \varepsilon$, or equivalently $|\operatorname{Im}(F)| \leq (\cot \varepsilon) \operatorname{Re}(F)$. In particular,*

$$|F| \leq \frac{1}{\sin \varepsilon} \operatorname{Re} F \leq \frac{1}{\sin \varepsilon} \operatorname{Re} F(p).$$

Proof. The incline angle of the tangent vector of $F(\partial\Sigma)$ is equal to the Lagrangian angle mod $\pi\mathbb{Z}$. Together with Claim 3.22 this implies $|\arg F| \leq \frac{\pi}{2} - \varepsilon$ on the $\partial\Sigma$, whence the same bound holds on Σ by the maximum principle for holomorphic functions. ■

We now introduce an *elementary functional*

$$\bar{\mathcal{S}}(L) = \operatorname{Im} \left(\sum_1^N \left(\sup_{L_i} f_{L_i} \right) e^{-i\hat{\theta}} \int_{L_i} \Omega \right) - \left(\sup_{L_0} f_{L_0} \right) \operatorname{Im} \left(e^{-i\hat{\theta}} \int_{L_0} \Omega \right).$$

As in Section 3.5, we introduce complex valued volume form $\tilde{\Omega}_{L_i}$ on the $(n-1)$ -dimensional moduli spaces of holomorphic curves, whose core properties are

$$\operatorname{Re} \tilde{\Omega}_{L_i} \geq 0, \quad \int_{\mathcal{M}} \tilde{\Omega}_{L_i} = \int_{L_i} \Omega, \quad i = 0, 1, \dots, N.$$

Thus the elementary functional is also a moduli space integral, with integrand

$$\operatorname{Im} \sum_1^N (e^{-i\hat{\theta}} \tilde{\Omega}_{L_j}) \sup f_{L_j} - \operatorname{Im} (e^{-i\hat{\theta}} \tilde{\Omega}_{L_0}) \sup f_{L_0}. \quad (3.23)$$

We decompose the Solomon functional into $\bar{\mathcal{S}}(L)$ and $\mathcal{S}(L) - \bar{\mathcal{S}}(L)$.

Theorem 3.43 (Bounded part of the Solomon functional). *Under the quantitative almost calibratedness and the potential clustering assumption, and all the standing assumptions of this section, there is a uniform a priori bound independent of L ,*

$$|\mathcal{S}(L) - \bar{\mathcal{S}}(L)| \leq \frac{A(4N+2)}{\sin \varepsilon} \int_{L_0} \operatorname{Re} \Omega.$$

Proof. We analyze the moduli space integrand (3.19) of the Solomon functional. Applying the uniform energy bound and the wedge region bound, the first term is bounded by

$$\left| \operatorname{Im} \int_{\Sigma} e^{-i\hat{\theta}} F \omega \right| \leq \int_{\Sigma} |F| \omega \leq \frac{\operatorname{Re}(F(p))}{\sin \varepsilon} \int_{\Sigma} \omega \leq \frac{\operatorname{Re}(F(p))}{\sin \varepsilon} A(N+1).$$

More intrinsically $F(p)$ defines the complex valued volume form $\tilde{\Omega}_{L_0}$ on the moduli space, hence

$$\left| \operatorname{Im} \int_{\Sigma} e^{-i\hat{\theta}} F \omega \right| \leq \frac{1}{\sin \varepsilon} A(N+1) \operatorname{Re}(\tilde{\Omega}_{L_0}). \quad (3.24)$$

The other two terms in (3.19) are rewritten as a sum of contributions from intersection points in (3.18). As in Proposition 3.41, we consider holomorphic polygons whose boundary $\partial \Sigma$ encounters in the clockwise order

$$\begin{aligned} p_{N-1} &\in CF^1(L_N, L_{N-1}), \dots, p_1 \in CF^1(L_2, L_1), \\ p &= p_0 \in CF^0(L_1, L_0), \quad q = p_N \in CF^0(L_0, L_N) \end{aligned}$$

juxtaposed possibly by more degree one self intersections b_i of L_i . (The other cases, where $\partial \Sigma$ misses some Lagrangians, can be handled completely similarly.) We first deal with these extra self intersections. Using the wedge region bound and the Novikov positivity requirement, we obtain

$$\begin{aligned} \left| \operatorname{Im} \sum_i \sum_{b_i} e^{-i\hat{\theta}} F f_{L_i}|_+^+(b_i) \right| &\leq \sum_i \sum_{b_i} |F(b_i)| f_{L_i}|_+^+(b_i) \\ &\leq \frac{\operatorname{Re}(F(p))}{\sin \varepsilon} \sum_i \sum_{b_i} f_{L_i}|_+^+(b_i). \end{aligned}$$

By Lemma 3.41, we have

$$\left| \operatorname{Im} \sum_i \sum_{b_i} e^{-i\hat{\theta}} F f_{L_i}|_+^+(b_i) \right| \leq \frac{1}{\sin \varepsilon} A(N+1) \operatorname{Re}(\tilde{\Omega}_{L_0}). \quad (3.25)$$

We are left with the contributions of p_0, p_1, \dots, p_N to (3.18):

$$\operatorname{Im} \sum_{j=0}^{N-1} e^{-i\hat{\theta}} F(f_{L_{j+1}} - f_{L_j})(p_j).$$

If we replace f_{L_i} by its supremum value $\sup_{L_i} f_{L_i}$ for all $i = 0, 1, \dots, N$, the new expression would be

$$\begin{aligned} & \operatorname{Im} \sum_{j=0}^{N-1} e^{-i\hat{\theta}} F(p_j) (\sup f_{L_{j+1}} - \sup f_{L_j}) \\ &= \operatorname{Im} \sum_{j=1}^N e^{-i\hat{\theta}} \sup f_{L_j} (F(p_{j-1}) - F(p_j)) - \operatorname{Im} (e^{-i\hat{\theta}} F(p) \sup f_{L_0}), \end{aligned}$$

which is more intrinsically the integrand (3.23) of the elementary functional. Using the potential clustering assumption and the wedge region bound lemma, the error of replacing the potentials by $\sup_{L_j} f_{L_j}$ can be bounded by

$$2A \sum_{j=1}^N |F(p_j)| \leq 2AN \frac{\operatorname{Re}(F(p))}{\sin \varepsilon} = \frac{2AN}{\sin \varepsilon} \operatorname{Re}(\tilde{\Omega}_{L_0}). \quad (3.26)$$

Now (3.24), (3.25), (3.26) are upper bounds on the three contributions to the difference between the Solomon functional integrand (3.19) and the elementary functional integrand. Their sum is bounded by

$$\frac{A(4N+2)}{\sin \varepsilon} \operatorname{Re}(\tilde{\Omega}_{L_0}),$$

so after integration on the moduli space,

$$|\mathcal{S}(L) - \bar{\mathcal{S}}(L)| \leq \frac{A(4N+2)}{\sin \varepsilon} \int_{L_0} \operatorname{Re} \Omega,$$

as required. ■

Remark 3.22. In Section 5.2 below we will deduce the potential clustering and an upper bound on N as consequences of almost quantitative calibratedness, and very mild conditions on the ambient manifold X . In Section 5.5 the boundedness of $|\mathcal{S} - \bar{\mathcal{S}}|$ will be essential for relating the asymptote of the Solomon functional to stability conditions.

Remark 3.23. In Kähler geometry, it is often useful to decompose natural functionals into two parts. For instance, the K-energy functional can be decomposed into an entropy part and a pluripotential part [17, Section 2.4], which is important in the study of constant scalar curvature Kähler metrics.

Remark 3.24. We suggested in Section 2.10 that the Solomon functional is essentially the logarithm of the tunnelling amplitude between Lagrangian branes. Pushing forth with this physics analogy, we may regard the elementary functional as a semiclassical approximation,⁴³ and $\mathcal{S} - \bar{\mathcal{S}}$ as quantum fluctuation effects. Our main assertion then becomes that

⁴³The elementary functional is proportional to the period integrals over the cycles L_i , which may be regarded as coming from integration over the moduli of constant maps. Such integrals are regarded as more classical than those involving non-trivial holomorphic curves.

quantitative almost calibratedness with some extra hypotheses imply the a priori bound on the quantum fluctuation effects. The author is not aware of previous suggestions in the physics literature, but Jake Solomon’s formal Riemannian picture in Section 2.8 may offer partial explanations for the relevance of the almost calibrated condition.

What if we relax the positivity condition? Suppose we drop the positivity condition on the bordism current, but keep all the other assumptions. Then the key difference is that for holomorphic curves contributing negatively to $\partial\mathcal{C}$, we need to replace Claim 3.22 by Claim 3.29, and Claim 3.23 by Claim 3.30. Correspondingly, all appearance of $\operatorname{Re} F(p)$ is replaced by its absolute value. Then the conclusion in Theorem 3.43 is replaced by

$$|\mathcal{S}(L) - \bar{\mathcal{S}}(L)| \leq \frac{A(4N+2)}{\sin \varepsilon} \int_{\mathcal{M}} |\operatorname{Re} \tilde{\Omega}_{L_0}|.$$

The problem is that the right-hand side is no longer a manifestly a priori bounded quantity.

4. Continuity, LMCF and variational method

We now proceed to the more analytic aspects of the problem of finding special Lagrangians. There are three principal methods for existence theorems in geometric analysis: continuity method, parabolic flows, and the calculus of variations. Generally speaking, continuity or flow methods are often closely related as the elliptic and parabolic cousins of each other, and allow one to work with a priori reasonably smooth objects, but finding a good continuity path or proving the long time existence of the flow may be difficult in a particular problem; the variational approach, on the other hand, operates with the space of a priori less regular objects to achieve some weak compactness, and then attempt to improve the smoothness via regularity theorems. Each approach contains substantial outstanding difficulties. We will try to maintain some equipoise, and compare the main difficulties in each approach. The sections on the flow and the continuity method borrow largely from various writings of Joyce with some new contents; the variational approach is essentially original, and will be presented in Section 5.

4.1. Lagrangian mean curvature flow

LMCF basics. We now return to some analytic aspects of Joyce’s proposal [45] related to the *Lagrangian mean curvature flow* (LMCF) inside a Calabi–Yau manifold. Recall a smooth mean curvature flow means a family of immersions $\iota_t: L \rightarrow M$ parametrized by time $t \in [0, T)$, such that the velocity is equal to the mean curvature

$$\partial_t \vec{x} = \vec{H}.$$

The starting point of LMCF is an early observation of Smoczyk, which justifies the name.

Proposition 4.1 ([76, Section 4.2]). *Let (L_t) be a smooth and compact mean curvature flow inside a Kähler–Einstein manifold, then the Lagrangian condition is preserved by the flow.*

Remark 4.1. In more general Kähler settings, the Lagrangian condition is still preserved provided one couples the mean curvature flow to the Kähler–Ricci flow (see [58]). Joyce’s program may have natural extensions to the almost Calabi–Yau setting. Indeed the generalization of the flow may even be advantageous for achieving certain genericity conditions, as in the work of Woodward and Palmer [67, 68].

Assuming (L_t) is a smooth and compact LMCF inside a Calabi–Yau ambient manifold, then if the initial Lagrangian is graded, i.e., the Lagrangian angle is well defined as a real valued function on the Lagrangian, then so is L_t . The grading is highly desirable, because of the foundational fact that the Lagrangian angle function $L_t \rightarrow \mathbb{R}$ satisfies the heat equation

$$(\partial_t - \Delta_{L_t})\theta_t = 0, \quad (4.1)$$

which among many other things, implies that $\sup_{L_t} \theta_t$ can only decrease in time, and $\inf_{L_t} \theta_t$ can only increase in time, and in particular the almost calibrated condition would be preserved by the flow. Inside a Calabi–Yau manifold, the mean curvature of L_t is related to the Lagrangian angle θ_t by an appealing formula

$$\bar{H} = J\nabla\theta_t,$$

where $\nabla\theta$ stands for the gradient of θ along L_t . Along the LMCF

$$\omega(\partial_t \vec{x}, \cdot) = \omega(\bar{H}, \cdot) = \omega(J\nabla\theta_t, \cdot) = -d\theta_t,$$

so the Lagrangians evolve by the local Hamiltonian function $-\theta_t$ up to an additive constant.

Mean curvature flow in codimension greater than one does not satisfy the avoidance principle. As such embedded Lagrangians can become immersed during the flow, so the program should at least include immersed Lagrangians. Joyce further suggests that certain ‘stable Lagrangian singularities’ should be admitted. For instance, inside Calabi–Yau 3-folds one should allow Lagrangians with local conical singularity modelled on the Harvey–Lawson T^2 -cone [45, Example 2.7]. The adjective ‘stable’ here means that the flow should preserve this class of singularities at least for a short amount of time, even if one makes a generic perturbation of the initial data.

Finite time singularity, and prototypical bad behaviours. The central difficulty of the subject is that *finite time singularities are in general inevitable*, starting from complex dimension two. Indeed, a theorem of Neves [64, Theorem 6.1] says that for any embedded Lagrangian submanifold inside a Calabi–Yau surface, there exists a Lagrangian within the same Hamiltonian isotopy class, such that the LMCF with this initial data forms finite

time singularity.⁴⁴ There is also a good geometric reason why singularities must occur in Joyce’s program: the Thomas–Yau uniqueness theorem applies to Lagrangians within the same derived Fukaya category class, which may include several Hamiltonian isotopy classes, at most one of which can have special Lagrangian representatives. In order for an initial Lagrangian in the wrong Hamiltonian isotopy class to find its way back to the right class along the LMCF, it must undergo a sequence of surgeries.

Now there is a substantial theory of weak solutions of mean curvature flows in the context of varifolds and currents, known as ‘Brakke flows’ [12], which exist under very general conditions. The problem is that such solutions are too weak to guarantee uniqueness of the flow, and the total mass of the varifold may jump down at discrete time. Even more fatally for our purpose, once the smoothness of the flow is dropped, the Lagrangian condition may not be preserved any more. It is instructive to look at the prototypical bad behaviours.

Example 4.2. Schoen and Wolfson [69] found area minimizers within certain Lagrangian isotopy classes, which are not minimal surfaces.⁴⁵ The Brakke flow with such initial data further decreases mass in time, so must cease to be Lagrangian. However, these examples are not graded, so do not contradict Joyce’s program. A possible lesson is that *non-graded Lagrangians are bad*.

Example 4.3. Consider a figure eight curve inside \mathbb{R}^2 ,⁴⁶ whose two looms have unequal areas. Along the mean curvature flow (known as the ‘curve shortening flow’ in this context) one loom shrinks first to zero size. At the moment of singularity, the Lagrangian angle at the self intersection point has a jump. From a more generalizable perspective, one notices that each loom encloses a holomorphic disc, and this singularity is associated with one holomorphic disc shrinking to zero size and disappearing. The general lesson is that *the shrinking down of small area holomorphic discs messes up the grading*, so it is desirable to exclude them if possible.⁴⁷

Example 4.4. Consider any compact Lagrangian inside the unit ball of \mathbb{C}^n . By an easy maximum principle argument, during the flow L_t remains inside the shrinking ball

$$\left\{ \sum |z_i|^2 \leq 1 - 2nt \right\},$$

so must develop a finite time singularity at some $t \leq \frac{1}{2n}$. From the Floer theoretic perspective, since such Lagrangians can always be displaced off itself by the Hamiltonian

⁴⁴Whether the same holds for almost calibrated initial data is an interesting open problem.

⁴⁵There is no contradiction: area minimization among Lagrangians by no means guarantees area stationarity among submanifolds.

⁴⁶Recall that curves in \mathbb{R}^2 are automatically Lagrangian.

⁴⁷Indeed, one important ingredient in Neves’s proof of singularity formation [64] is the destruction of grading related to shrinking enclosed 2-dimensional areas. Although it is not explicit in Neves’s work, these areas seem related to holomorphic discs.

isotopy corresponding to translations in \mathbb{C}^n , its Floer cohomology is either obstructed or zero. As such, a compact Lagrangian supported in a small coordinate ball is invisible to the derived Fukaya category. From a different perspective, since such Lagrangians have zero homology class, they are excluded in the almost calibrated case.

Joyce’s LMCF proposal. With this background in mind, one may better appreciate the upshot of Joyce’s perspective: *LMCF should be better behaved if the Lagrangians support unobstructed brane structures*, namely the bad singularities that would spell ruin in a more general context do not actually occur in his program. The moral reasons are as follows:

- The Lagrangian branes are always assumed to be graded.
- Unobstructed Lagrangian branes cannot bound holomorphic curves with very small areas unless their Floer theoretic contributions exactly balance out, for otherwise certain positivity requirements in the Novikov ring will be violated. (See the figure eight example [45, Example 3.24] for a concrete instance.)
- Assume the Lagrangian decomposes into two pieces, one of which is contained inside a small coordinate ball. Since this piece only contributes a zero object in $D^b \text{Fuk}(X)$, discarding this piece does not affect the $D^b \text{Fuk}(X)$ class of the Lagrangian.

This gain comes at the burdensome cost of carrying the brane structure along the flow, which leads to somewhat counterintuitive prescriptions such as surgeries before the Lagrangian itself reaches a singularity, so that the unobstructed brane structure may not be lost prematurely. Joyce describes a number of singularities and surgeries that are expected to occur generically in his program.

Opening up the neck. The Lagrangian brane can develop new self intersection points, and may flow from unobstructed to obstructed at $t = t_0$ through the shrinking of certain holomorphic curves with boundary on L_t , even through the underlying Lagrangian remains smooth. The Floer theoretic mechanism causing the obstruction (to do with positivity conditions in the Novikov ring) precisely ensures an angle condition at certain self intersection points of L_{t_0} , so that a Joyce–Lee–Tsui Lagrangian expander [46] can be glued into L_{t_0} to continue the LMCF. Woodward and Palmer [67, 68] have performed substantial checks that suitable brane structures can be assigned to such surgeries so that L_t remains unobstructed, and the Floer cohomology remains continuous throughout the surgery.

Neck pinching. In some sense converse to the above process, the LMCF L_t may contain a local region modelled on a Lawlor neck with small length scale parameters $\varepsilon(t)$, which shrinks ‘slowly’ in time, and $\varepsilon(t) \rightarrow 0$ at time $t \rightarrow t_0$.⁴⁸ In some cases this may cause the domain of the immersed Lagrangian to become disconnected.

⁴⁸A gluing construction of neck pinching examples is in working progress with T. Collins. There is however a crucial sign difference concerning Lagrangian angles, between our work and the Joyce prediction, which may have rather disconcerting consequences for Joyce’s program.

Collapsing zero objects. After a number of surgeries, the Lagrangian L_t may be decomposed into several disconnected pieces, some of which are zero objects in $D^b \text{Fuk}(X)$, so in particular have homology class zero. A typical situation is that the zero objects are contained in small coordinate balls.⁴⁹ We simply discard these pieces and continue the flow for the remaining pieces.

Stable singularities. As mentioned above, one may need to include Lagrangians with certain local singularities, since generic perturbations cannot remove such singularities. How such singularities can form dynamically starting with smooth initial data is less clear, but Joyce offers some analogy with the setting of $U(1)$ -invariant special Lagrangians in \mathbb{C}^3 (see [45, Example 2.8]), where Harvey–Lawson T^2 -cone singularities can appear and disappear in pairs within a 1-parameter family of deformations, and in particular smooth objects can be continuously deformed to such singular objects.

The Joyce program of LMCF with surgery contains a number of potentially counter-intuitive phenomena.

Example 4.5 ([45, Example 3.15]). After incorporating the surgery of the brane structures, Joyce’s LMCF is no longer identical to the LMCF of the underlying Lagrangian. The most extreme case is to start with the union L of two unobstructed special Lagrangians L_1, L_2 with phase angles $\theta_{L_1} < \theta_{L_2}$, with an intersection point $b \in CF^1(L_2, L_1)$ defining a closed morphism. We regard L as an immersed Lagrangian with bounding cochain b . This fits into the distinguished triangle

$$L_1 \rightarrow L \rightarrow L_2 \rightarrow L_1[1],$$

which is *not* destabilizing for L . This configuration is stationary in ordinary LMCF. However, under Joyce’s LMCF, the bounding cochain b evolves in time, and loses positivity in the Novikov ring in finite time, after which one is supposed to ‘open up the neck’ to continue the flow in a non-trivial fashion.

Example 4.6 ([45, Section 3.4]). Joyce’s LMCF in general needs to incorporate non-trivial rank one local systems. Even if the initial brane structure has trivial local system, it is possible for surgeries to create non-trivial local systems from the bounding cochain data at self intersection points. One may imagine such bounding cochain data to be a holonomy contribution concentrated at points, which can be converted into a smeared out holonomy contribution from a non-trivial local system.

Example 4.7. The ‘opening-up the neck’ surgery is governed by the Novikov positivity requirement of the bounding cochain, which depends on the choice of the bounding cochain, not just the underlying Lagrangian submanifold. The same underlying Lagrangian with different bounding cochains may therefore flow to different infinite time limits.

⁴⁹Joyce suggests plausibly that Neves’s example [64] exhibits this behaviour by splitting off a small Whitney sphere, although this is not proven.

In summary, the main difficulty of Joyce's LMCF program is that there is a huge gap between the general Brakke flow framework, and the kind of regularity control required for the long time existence of the LMCF. It would represent very substantial progress⁵⁰ if one can classify possible singularity types under suitable genericity assumptions, say for almost calibrated Lagrangians inside Calabi–Yau 3-folds. A large list of problems, from routine level up to the impossible, can be found in Joyce's excellent original paper [45].

Infinite time limit and its difficulties. Provided one can prove long time existence of LMCF, the total mass of L_t will be uniformly bounded since it decreases during the flow. Under mild conditions to ensure L_t does not escape to spatial infinity (e.g., if the ambient Calabi–Yau manifold is compact), one can extract the infinite time subsequential limits of L_t as currents. From the heat equation (4.1) on the Lagrangian angle θ , we obtain

$$(\partial_t - \Delta_{L_t})|\theta|^2 = -2|\nabla\theta|^2 = -2|\vec{H}|^2.$$

If the Lagrangians remain sufficiently smooth, then an integration by part calculation shows

$$\int_0^T \int_{L_t} |\vec{H}|^2 d\text{vol}_{L_t} dt \leq \frac{1}{2} \left(\int_{L_0} |\theta|^2 d\text{vol}_{L_0} - \int_{L_T} |\theta|^2 d\text{vol}_{L_T} \right).$$

Even if the volume mass can jump down at discrete time, such as during the collapsing of zero objects, we still expect

$$\int_0^T \int_{L_t} |\vec{H}|^2 d\text{vol}_{L_t} dt \leq \frac{1}{2} \int_{L_0} |\theta|^2 d\text{vol}_{L_0} < \infty, \quad \forall T > 0. \quad (4.2)$$

In particular, we can find a sequence of time $t_i \rightarrow \infty$, with

$$\int_{L_{t_i}} |\nabla\theta|^2 d\text{vol} \rightarrow 0.$$

This strongly suggests that the subsequential limit is a union of special Lagrangian currents with multiplicities.

In the heuristic logic of Thomas–Yau–Joyce program, the infinite time limit supposedly provides the Harder–Narasimhan decomposition. In general one cannot expect the special Lagrangian currents to be smooth, so this raises the question how to make sense of singular Lagrangians as representatives of $D^b \text{Fuk}(X)$ classes, or whether we should use some weaker equivalence class. Another interesting open problem is whether the limiting current is unique. In order to run the Thomas–Yau argument, one presumably also needs Floer theory for singular Lagrangians.

⁵⁰Joyce [45] assesses the difficulty of his program in the Calabi–Yau 3-fold case to be comparable to Perelman's breakthrough on the Poincaré conjecture. Indeed, ruling out the cigar solution in the context of the Ricci flow is in itself already a major achievement of Perelman.

Joyce [45] already observed that it is not obvious how singular Lagrangians can carry brane structures, and it is logically possible for some Floer theoretic information to be lost in the infinite time limit. The suggestion is that hopefully the Lagrangian L_t at large but finite time $t \gg 1$, can serve as a substitute for the infinite time limit, which presumably has better smoothness properties [45]. There is however no known justification (and probably false) that the surgeries terminate after some finite time, and L_t decomposes into the union of several Lagrangian objects, in order to provide a Harder–Narasimhan decomposition.

We think Floer theory for singular Lagrangians is one of the foundational open questions necessary for an adequate solution of the Thomas–Yau conjecture. See Section 5.4 for further discussion.

4.2. Continuity method

The continuity path. The general idea of the continuity method is to work with a 1-parameter family of PDEs, and attempt to deform from an initial given solution, to a solution of the final PDE, provided the deformation encounters no obstruction, and satisfies suitable compactness properties. The hope that the continuity method *may* be useful here, is based on the foundational fact that compact special Lagrangian submanifolds inside almost Calabi–Yau manifolds have unobstructed deformation theory, before taking brane structures into account.

However, problems immediately ramp up once one attempts to set up a continuity path. The most naïve suggestion, based on the analogy with the HYM equation, is to prescribe the Lagrangian angle as a function on the domain of L . This however breaks the domain parametrization invariance of $\iota: L \rightarrow X$, and the author knows no satisfactory way to make general sense of this approach beyond graphical Lagrangians. Instead, we fix ω and allow Ω to vary in an infinite-dimensional parameter space subject to the almost Calabi–Yau condition. In non-compact almost Calabi–Yau manifolds, we need to also keep the metric asymptote fixed at infinity. The continuity path is a generic 1-parameter family of Ω . This setup strongly resemble the *wall crossing phenomenon* studied by Joyce [42] in the context of special Lagrangian enumerative invariants, and indeed the rest of this section liberally borrows from the ideas therein.⁵¹

Two main obstacles. There are two fundamental obstacles:

- How can one find the *initial* special Lagrangian?
- How can one guarantee compactness?

How to find an initial special Lagrangian. The question about finding the initial special Lagrangian is specific to the continuity method, and does not appear in the LMCF

⁵¹To the author’s awareness, the continuity method for the Thomas–Yau conjecture has not appeared in the literature, except for the wall crossing ideas sketched in [42]. An unpublished program to develop the Thomas–Yau conjecture in complex dimension two via the continuity method is mentioned in the talks of Joyce.

approach. A natural suggestion is to look for special Lagrangians near certain degenerate limits, and our relaxation of the complex Monge–Ampère equation ought to give much more flexibility. For instance, conifold degenerations are known to give rise to special Lagrangian spheres appearing as vanishing cycles [38].⁵² Another general source is to work near a suitable large complex structure limit, so that the Kähler metric can be made almost toric outside a small region, such that the torus fibres are much smaller compared to the characteristic length scale of the base. We can then attempt to find special Lagrangians via adiabatic limits, in close analogy with the standard procedure to find holomorphic curves via tropical degenerations [55].⁵³ The most accessible special Lagrangians in this approach, should be obtainable by small perturbations of the torus fibres.⁵⁴ The next candidate suggested by the Leray filtration of the torus fibration is already much harder.

Question 8. Construct special Lagrangians whose toric projection to the base are small thickenings of certain 1-dimensional graphs.

One would expect that locally along an edge these Lagrangians are perturbations of $T^{n-1} \times \mathbb{R}$, with T^{n-1} contained in the torus fibre direction, so that we obtain $(n-1)$ locally defined closed 1-forms $\int_{S^1} \omega$ on the base corresponding to the cycles $S^1 \subset T^{n-1}$, and the edge is to leading approximation given by requiring these 1-forms to vanish. The local model for the junction where three edges meet,⁵⁵ may have the following topological description. In the $n=2$ case, we have a ‘pair of pants’ inside $T^2 \times \mathbb{R}^2$ with three asymptotic ends $S^1 \times \mathbb{R}$; topologically this is the same as algebraic surface

$$\{z_1 + z_2 = 1\} \subset \mathbb{C}^* \times \mathbb{C}^*.$$

In higher dimensions, we take a product of the pair of pants with T^{n-2} .

In the next order of perturbation, we expect the deformation of the Lagrangian in the base direction to be at least comparable to the length scale of the fibre, and presumably is fixed by the ‘special condition’ $\text{Im } \Omega|_L = 0$.

Remark 4.2. From the viewpoint of the Thomas–Yau–Joyce picture, there is an additional problem to assign unobstructed brane structures to the initial special Lagrangian.

Compactness and genericity. The question about compactness largely reflects problems we already encountered in the LMCF approach. The essential issue is that without any further condition on the Kähler metric, special Lagrangians may be too singular, so

⁵²While Hein and Sun’s result is highly non-trivial, the entire difficulty goes into understanding the Calabi–Yau metric near the conifold point. If we are given the license to prescribe arbitrary Kähler metrics, the problem of finding special Lagrangian vanishing spheres near the conifold point becomes easy.

⁵³The large complex structure limit is supposed to correspond to the large volume limit in the mirror, which is related to the μ -stability, thus offering the hope of a mirror calculation of counting invariants.

⁵⁴The difficulty in [54] to construct SYZ special Lagrangian fibrations again comes from the Calabi–Yau metrics. If one can freely prescribe Kähler metrics, then finding a special Lagrangian torus is not difficult.

⁵⁵This is conceptually related to Matessi’s ‘Lagrangian pair of pants’ [59].

that the McLean deformation theory [60] for special Lagrangians may fail. The natural answer, closely related to the LMCF viewpoint, is that we should only work with *generic* Kähler structures, and 1-parameter families thereof. According to the philosophy advocated by Joyce [42], the importance of singularities are ranked according to their genericity. For special Lagrangians with first Betti number k , the moduli space of deformations is k -dimensional, so in a generic 1-parameter family of Kähler structures, one expects to encounter singularities with genericity index up to $k + 1$, and those singularities of index 0, 1 are the most important.⁵⁶ Before one can seriously pursue the rest of this strategy, it is necessary to have a classification of index 0, 1 special Lagrangian singularities in complex dimension n .

Question 9. In complex dimension 3, classify all special Lagrangian singularities of index 0 and 1, namely all singularities that can occur in a generic 1-parameter family of special Lagrangians when ω is fixed and Ω varies.

Example 4.8. The *Harvey–Lawson T^2 -cone* is a special Lagrangian cone inside \mathbb{C}^3 with link T^2 , invariant under the diagonal $T^2 \subset SU(3)$. Explicitly,

$$L_{HL} = \{(z_1, z_2, z_3) \in \mathbb{C}^3 : |z_1| = |z_2| = |z_3|, \operatorname{Im}(z_1 z_2 z_3) = 0, \operatorname{Re}(z_1 z_2 z_3) \geq 0\}.$$

Haskins [36, Theorem 1] proved that up to unitary transformations, this is the only strictly stable⁵⁷ special Lagrangian cone with smooth embedded link diffeomorphic to T^2 .

The Harvey–Lawson cone admits three distinct 1-parameter deformations into smooth embedded special Lagrangians L_s^1, L_s^2, L_s^3 for $s > 0$. Here

$$L_s^1 = \{(z_1, z_2, z_3) \in \mathbb{C}^3 : |z_1|^2 - s = |z_2|^2 = |z_3|^2, \operatorname{Im}(z_1 z_2 z_3) = 0, \operatorname{Re}(z_1 z_2 z_3) \geq 0\},$$

and L_s^2, L_s^3 arise via cyclic permutations of z_1, z_2, z_3 . Notably, there is a holomorphic disc D_t^1 of area πs with boundary on L_s^1 (and similarly for L_s^2, L_s^3),

$$D_t^1 = \{(z_1, 0, 0) : |z_1|^2 \leq s\}.$$

In particular, L_s^a for $a = 1, 2, 3$ cannot be exact Lagrangians, but have non-zero Lagrangian flux. The $s \rightarrow 0$ limit corresponds to the holomorphic discs shrinking to zero area, or equivalently the Lagrangian flux tends to zero.

Now on a compact special Lagrangian inside an almost Calabi–Yau manifold, the Harvey–Lawson cone can arise as a local model for conical singularities. The gluing results of Joyce [43, Section 10] shows that when certain homological conditions are satisfied, then there exist desingularization families of special Lagrangians locally modelled

⁵⁶Understanding the moduli space of special Lagrangians requires singularities up to index $k + 1$, but if we restrict to Lagrangian deformations with zero Lagrangian flux, then index ≤ 1 may suffice in the optimistic view.

⁵⁷Strict stability here is a condition on the Laplacian spectrum of the link. Unfortunately, the word ‘stable’ is overloaded with many standard meanings in the literature.

on L_S^a , such that the singular special Lagrangians carrying the T^2 -cone singularity arise in codimension one, so in this case the T^2 -cone is an index one singularity in Joyce's sense.⁵⁸ This gluing result is not sensitive to varying Ω . On the other hand, if one restricts to deformations with Lagrangian flux zero, which can be regarded as the analogue of exact isotopies in the mildly singular case, then an isolated local T^2 -cone singularity cannot be desingularized, but instead keeps the singularity as it deforms.

Wall crossing. Some of the generic singularities in the LMCF are expected to have elliptic analogues in the continuity method approach. We fix ω and consider a generic 1-parameter family of Ω , and we follow the Lagrangian flux zero deformations of a given special Lagrangian.

Under exact isotopy, immersed Lagrangians may lose the unobstructed condition. One expects the surgery of the brane structure suggested by Joyce has an elliptic analogue, involving the same ingredients as the ‘Maslov flow’ studied by Woodward and Palmer [67, 68].

As already discussed in Section 2.7, the Lawlor neck is responsible for the gluing of two immersed special Lagrangians. This corresponds to the ‘Lawlor neck pinching’ singularity, as well as the ‘opening the neck’ surgery in the LMCF.

(Stable singularity.) Joyce suggests from his work on $U(1)$ -invariant special Lagrangians in \mathbb{C}^3 ([45, Example 2.8]), that in a continuous 1-parameter family, isolated singular points of special Lagrangian 3-folds with local T^2 -cone singularities can appear and disappear in pairs, by making the two T^2 -cone singularities collide with each other and then smooth out. This is the main motivation for admitting the T^2 -cone singularity in the LMCF, and it seems likely to be a generic singularity in the continuity method as well.

Remark 4.3. The phenomenon of ‘collapsing zero object’ in Joyce’s LMCF has no analogue in the continuity approach, since the special Lagrangian condition forbids any homologically trivial component.

What’s the role of the brane structure? As Bridgeland observed [13, Theorem 1.2], the central charge map gives a local homeomorphism between the space of Bridgeland stability conditions, and the hom space from the numerical Grothendieck group to \mathbb{C} . The Thomas–Yau–Joyce philosophy then suggests that for fixed ω and small deformations of Ω , the stability condition only depends on the cohomology class $[\Omega]$.

A possible geometric interpretation consistent with the wall crossing picture, is that in a generic 1-parameter family of deformations of the almost Calabi–Yau structure fixing ω and $[\Omega]$, the special Lagrangian may undergo surgeries, such that the topology and the Hamiltonian isotopy class may change, but when equipped with the brane structures, the

⁵⁸Joyce’s gluing result is quite subtle. Under certain homological conditions, the smoothing can be forbidden, in which case the T^2 -cone is an index zero singularity. In other cases, due to some linear dependence of certain homology classes, two T^2 -cone singularity may not behave independently, but together behave like an index one singularity. See [43, Section 10].

$D^b \text{Fuk}(X)$ class (or possibly some weaker equivalence class) remains constant. As in the LMCF approach, the role of the unobstructed brane structure is morally to prevent holomorphic discs from having very small area. However, the troubles caused by shrinking discs may be less severe in the continuity method than in the LMCF method, since the continuity method only deals with special Lagrangians, and cannot lose grading in a process like Example 4.3. Instead, a significant amount of difficulty in the continuity method is absorbed into the problem of finding the initial special Lagrangians.

Comparison with LMCF. To summarize the pros and cons compared to the LMCF approach, in the continuity method finding an initial special Lagrangian is a significant new difficulty. However, we no longer need to confront the extremely difficult task of long time existence for the flow. On a slightly more technical level, provided one has a classification for low index singularities, the *genericity assumption* is likely to be easier to use in the continuity method than it is in the LMCF framework.

5. Variational method

We begin by significantly narrowing the scope as follows.

(i) We only consider exact Lagrangians inside Calabi–Yau Stein manifolds. The Stein assumption is meant to simplify Floer theory, at the cost of non-compactness of the ambient space. In particular, $H_{n+1}(X, \mathbb{Z}) = 0$ and $H_n(X, \mathbb{Z})$ has no torsion, since Stein manifolds have the topological type of CW complexes of dimension $\leq n$. The complex Monge–Ampère equation of the Calabi–Yau metric ensures special Lagrangians are area minimizing currents, at the cost of losing genericity. The author thinks the complex Monge–Ampère equation is convenient but not completely indispensable to the Thomas–Yau conjecture.

(ii) We impose another mild condition on the Calabi–Yau metric, namely that the regularity scale grows to infinity asymptotically (see Section 5.2). Its main goal is to prevent the Lagrangians from escaping to infinity.

(iii) We only consider graded Lagrangians whose phase angle function satisfies a quantitative almost calibrated condition $-\pi/2 + \varepsilon \leq \theta \leq \pi/2 - \varepsilon$. An alternative characterization is that $\pm \text{Im } \Omega \leq (\cot \varepsilon) \text{Re } \Omega$ when restricted to L , whence the quantitative almost calibrated condition is preserved under weak limits of currents.

The main question we wish to examine from a variational viewpoint is the following.

Question 10. Fix a $D^b \text{Fuk}(X)$ class in the exact Calabi–Yau manifold (X, ω, Ω) , represented by some non-trivial unobstructed exact Lagrangian brane L_0 with phase function $-\pi/2 + \varepsilon < \theta < \pi/2 - \varepsilon$. Denote

$$\hat{\theta} = \arg \int_{L_0} \Omega \in \left(-\frac{\pi}{2} + \varepsilon, \frac{\pi}{2} - \varepsilon \right).$$

When does there exist a possibly singular special Lagrangian representative L of phase $\hat{\theta}$, in the same derived Fukaya category class (or under some weaker equivalence relation)?

Remark 5.1. Making sense of the derived Fukaya category for Lagrangians with weak regularity is part of the question, which seems highly non-trivial. As a basic meta-principle, any two Lagrangian branes in the same derived category class must lie in the same homology class in $H_n(X)$, so the homology class is fixed. This still leaves open some ambiguity on the choice of local system (see Remark A.8), and the question of which singular Lagrangians to include (see Sections 5.3 and 5.4).

One possible interpretation of the *Thomas–Yau conjecture* is as follows.

Conjecture 5.1 (Thomas–Yau existence). *Assume the derived Fukaya class of the quantitatively almost calibrated Lagrangian brane L_0 is Thomas–Yau semistable (see Definition 3.32). Then there is a special Lagrangian representative in the same $D^b \text{Fuk}(X)$ class (or some weaker S -equivalence relation, see Remark 5.29).*

Our limited goal is not to prove this conjecture completely, but to clear up enough easier obstacles in order to pinpoint the deeper issues that need to be resolved. Whenever we make difficult claims that we are yet unable to prove, we will try to at least provide some heuristic reasons.

5.1. Compactness and regularity

5.1.1. Standard geometric measure theory. We will use the standard language of geometric measure theory, see Federer [32] or Morgan [61] for the terminologies. The starting point of the variational approach is that there are foundational *compactness theorems* in geometric measure theory.

Theorem 5.2 (Federer–Fleming compactness theorem [32]). *Let L_i be a sequence of m -dimensional integral currents in a complete Riemannian manifold X , all supported in a fixed bounded subset, with uniform bounds*

$$\text{Mass}(L_i) \leq C \quad \text{and} \quad \text{Mass}(\partial L_i) \leq C.$$

Then up to subsequence L_i converges weakly in the current topology to an m -dimensional integral current L with the same bounds.

Remark 5.2. While compactness in the current topology is elementary, the claim that the limit is also an integral current is non-trivial, and can be viewed as a regularity result. The same holds with the Allard compactness theorem below. For our applications, we will always work with closed integral currents, namely $\partial L_i = 0$, which implies $\partial L = 0$ in the limit. To such currents one can associate a homology class.

Remark 5.3. A more technical version of Federer–Fleming compactness replaces the current topology by the flat norm topology, which is a slightly stronger topology. The flat

norm of an integral current T is

$$\|T\|_{\text{flat}} = \inf\{\text{Mass}(A) + \text{Mass}(B) \mid T = A + \partial B, A, B \text{ are integral currents}\},$$

and the convergence $T_i \rightarrow T$ in this topology simply means $\|T - T_i\|_{\text{flat}} \rightarrow 0$.

Theorem 5.3 (Allard compactness [5]). *Let L_i be a sequence of m -dimensional integer rectifiable varifolds in a complete Riemannian manifold X , all supported in a fixed bounded subset, with a uniform volume upper bound $\text{Mass}(L_i) \leq C$ and a uniform bound on the first variation $\int_{L_i} |\vec{H}| \leq C$. Then up to subsequence, L_i converges to an m -dimensional integer rectifiable varifold L with the same bounds.*

Remark 5.4. Federer–Fleming and Allard are somewhat complementary. Integral currents are a special kind of distribution valued forms, while varifolds are a special kind of measures on the real Grassmannian bundle $\text{Gr}(TX, m)$ over X whose fibres parametrize m -dimensional planes in the tangent spaces of X . One key advantage of currents is that they know about orientations, while varifolds do not. The integral current L recovers the underlying rectifiable subset $\text{supp}(L)$ with multiplicity, so can be canonically associated with a varifold L^{var} . On the other hand, the natural topology on varifolds (i.e., the topology as measures on $\text{Gr}(TX, m)$) remembers tangent plane information, which can be lost under the flat norm convergence of integral currents. Moreover, assuming all the varifolds in the sequence are contained in a bounded region, then the total volume mass converges under varifold convergence, but not necessarily so under flat norm convergence. The intuition is that *morally the varifold topology detects one more derivative than the flat norm topology*. This explains why Allard requires some integral control on the mean curvature, while Federer–Fleming does not.

We shall later use the informal terminology of ‘*varifold/current topology*’ to refer to convergence simultaneously in the varifold topology and the flat norm topology on integral currents.

Example 5.4. Inside $S^1_{2\pi} \times \mathbb{R}$ with the standard Euclidean metric, take L_k as the graph over S^1 of the function $\frac{1}{k} \sin(kx)$. Then L_k are Lagrangian currents, which converge to S^1 as currents, but due to the high oscillation,

$$\liminf \text{Mass}(L_k) > \text{Mass}(S^1),$$

and L_k do not converge to S^1 in the varifold sense. The Lagrangian angle of L_k is prescribed by $\tan \theta = \cos(kx)$, which converges to zero in the current sense, but not strongly in L^1 .

One of the best *regularity theorems* in geometric measure theory is the following.

Theorem 5.5 (Almgren’s big regularity theorem [6]). *Let L be a compactly supported m -dimensional closed integral current inside a complete Riemannian manifold, which minimizes the volume among all closed integral currents in the same homology class, then*

away from a closed subset $S \subset \text{supp}(L)$ with Hausdorff dimension at most $m - 2$, the rectifiable subset $\text{supp}(L) \setminus S$ is a smooth submanifold.

Remark 5.5. Real codimension two singularity is the optimal result, as easily seen from the examples of singular algebraic curves in \mathbb{CP}^2 , which are automatically area minimizers in their homology classes.

Remark 5.6. Almgren's big regularity theorem is well known for its monumental size of around 1000 pages. The recent works of Delellis et al. have somewhat simplified the proof, which still remains very non-trivial (see [23] for some introduction).

Remark 5.7. For almost Calabi–Yau ambient structures, with

$$\frac{\omega^n}{n!} = (-1)^{n(n-1)} \left(\frac{\sqrt{-1}}{2} \right)^n e^{2\rho} \Omega \wedge \bar{\Omega},$$

the n -form $\text{Re}(e^{-i\hat{\theta}} e^{\rho} \Omega)$ has pointwise norm ≤ 1 , situated precisely by special Lagrangian tangent planes of phase $\hat{\theta}$. Spolaor [80] calls this a semicalibration form (which needs not be closed), and proved that the semicalibrated integral currents (i.e., special Lagrangians in our case) satisfy the same regularity as in Almgren's big theorem.

A standard way to apply these theorems, for instance inside a compact ambient space, is to fix the homology class, and minimize the volume among all the integral currents therein. The compactness theorem guarantees the existence of an absolute volume minimizer, and the regularity theorem then improves its regularity to be more like submanifolds. This strategy is highly effective in producing minimal surfaces, but there is no useful criterion⁵⁹ to guarantee the volume minimizers to be special Lagrangians, which is why producing special Lagrangians is a highly non-trivial problem in geometric measure theory.

5.1.2. Exact Lagrangians under weak regularity. We need to ensure the class of Lagrangians in the variational setup is closed under the varifold/current topology. We shall say an integral current L is Lagrangian, if its tangent space is a Lagrangian subspace a.e. on L , or equivalently

$$\int_L \omega \wedge \eta = 0,$$

for any smooth and compactly supported $(n - 2)$ -form η .

A trivial observation is the following.

Lemma 5.6. *Let L_i be closed Lagrangian integral currents, and suppose $L_i \rightarrow L$ in the current topology, then the Lagrangian/quantitative almost calibratedness conditions pass to the limit.*

⁵⁹If there is at least one special Lagrangian within the given homology class, then all absolute minimizers must be special Lagrangians, by an easy calibration argument. This however does not answer how to find the special Lagrangian in the first place.

Let L be a closed Lagrangian integral current, and f_L be an L^∞ function on L . We say the exact condition $\lambda = df_L$ holds in the weak sense if for any compactly supported test $(n-1)$ -form χ , we have

$$\int_L \lambda \wedge \chi = - \int_L f_L d\chi. \quad (5.1)$$

To make sense of the right-hand-side, notice the rectifiability of L allows the integration of the L^∞ -valued n -form $f_L d\chi$. Equivalently, the normal current $f_L L$ has distributional derivative $\chi \rightarrow \int_L \chi \wedge \lambda$.

Remark 5.8. The examples of immersed Lagrangians show that we cannot require f_L to have a continuous extension to X , so L^∞ -regularity is the best we can impose on f_L .

Lemma 5.7. *All Lagrangians are assumed to be contained in a fixed bounded region of X , homologous to L_0 , and are quantitatively almost calibrated. If L_i is a sequence of exact Lagrangians with potential f_{L_i} , such that f_{L_i} are uniformly bounded in L^∞ . Then up to subsequence, there is a Lagrangian L with potential f_L , such that $L_i \rightarrow L$ and $f_{L_i} L_i \rightarrow f_L L$ as currents.*

Proof. By Lemma 2.1 the volume mass is uniformly upper bounded. By Federer–Fleming compactness, subsequentially $L_i \rightarrow L$ in the flat topology for some Lagrangian integral current L homologous to L_0 . This implies that

$$\int_{L_k} g \operatorname{Re} \Omega \rightarrow \int_L g \operatorname{Re} \Omega$$

for any C^∞ test function g , even though $\liminf_i \operatorname{Mass}(L_i)$ may be strictly greater than $\operatorname{Mass}(L)$, as we do not assume varifold convergence.

We focus on a coordinate ball. The n -currents $f_{L_k} L_k$ can be viewed as a collection of $\binom{2n}{n}$ signed measures

$$g \mapsto \int_{L_k} f_{L_k} g dx_{i_1} \wedge \cdots \wedge dx_{i_n}.$$

Each of these measures are bounded by the measure

$$g \mapsto (\sup \|f_{L_i}\|_{L^\infty}) \frac{C}{\sin \varepsilon} \int_{L_k} g \operatorname{Re} \Omega,$$

whose total mass is uniformly bounded for all k . By the weak compactness of measures, subsequentially these signed measures converge, and the limiting signed measures have L^∞ Radon–Nykodim derivatives $a_{i_1, \dots, i_n}(x)$ with respect to the measure $g \mapsto \int_L g \operatorname{Re} \Omega$:

$$\int_{L_k} f_{L_k} g dx_{i_1} \wedge \cdots \wedge dx_{i_n} \rightarrow \int_L g a_{i_1, \dots, i_n} \operatorname{Re} \Omega.$$

Thus inside the coordinate ball, the currents $f_{L_k} L_k$ converge to $\lim_k f_{L_k} L_k$:

$$\eta \mapsto \int_L \sum_{i_1 < i_2 < \dots < i_n} \eta(\partial_{i_1} \wedge \dots \wedge \partial_{i_n}) a_{i_1, \dots, i_n} \operatorname{Re} \Omega,$$

where η is any test n -form.

Now L is an integral current, so \mathcal{H}^n -a.e. $y \in \operatorname{supp}(L)$ there is a well defined tangent space $T_y L$ and a local integer multiplicity $\Theta(y)$. Recall a blow-up limit of an n -current N at a point $y \in X$ refers to a subsequential limit of the currents on $T_y X$ as $r \rightarrow 0$:

$$\eta \mapsto \int_N \operatorname{rescale}_{y,r}^* \eta, \quad \operatorname{rescale}_{y,r}: x \mapsto \frac{x}{r} \text{ in the geodesic coordinates around } y.$$

For a.e. $y \in \operatorname{supp}(L)$, there is a unique blow-up limit for the current $\lim_k f_{L_k} L_k$, which is

$$\eta \mapsto \int_{T_y L} \sum \eta(\partial_{i_1} \wedge \dots \wedge \partial_{i_n}) a_{i_1, \dots, i_n}(y) \Theta(y) \operatorname{Re} \Omega,$$

whose $\binom{2n}{n}$ component signed measures are just constant multiples of the Lebesgue measure on $T_y X$.

Observe that the weak formulation (5.1) passes to the limit

$$\int_L \lambda \wedge \chi = -(\lim_k f_{L_k} L_k)(d\chi).$$

Thus the blow-up limit of $\lim_k f_{L_k} L_k$ at a.e. $y \in \operatorname{supp}(L)$ is in fact a closed current. Consequently, the polyvector

$$\sum_{i_1 < i_2 < \dots < i_n} a_{i_1, \dots, i_n} \partial_{i_1} \wedge \dots \wedge \partial_{i_n}$$

must be a pure tensor lying in $\Lambda^n T_y L \subset \Lambda^n T_y X$. Hence,

$$\lim_k f_{L_k} L_k = f_L L$$

for some L^∞ -function f_L . ■

Continuity of the Solomon functional. Inside Stein manifolds, the Solomon functional can be defined for any Lagrangian L with potential which is homologous to L_0 , without further Floer theoretic inputs: the formula (3.4) makes sense after choosing any bordism current \mathcal{C} with $\partial \mathcal{C} = L - L_0$ in the sense of currents, and the choice does not matter.

Lemma 5.8 (Continuity of the Solomon functional). *All Lagrangians are assumed to be contained in a fixed bounded region of X , homologous to L_0 , and are quantitatively almost calibrated. Suppose L_i is a sequence of Lagrangian integral currents with potential f_{L_i} , such that $L_i \rightarrow L$ in the flat norm, and $f_{L_i} L_i$ converge to $f_L L$ as currents, then the Solomon functionals converge: $\mathcal{S}(L_i) \rightarrow \mathcal{S}(L)$.*

Proof. Since $f_{L_i} L_i$ converges to $f_L L$ as currents, then

$$\int_{L_i} f_{L_i} \Omega \rightarrow \int_L f_L \Omega.$$

It suffices to justify

$$\int_{\mathcal{C}_i} \lambda \wedge \Omega \rightarrow \int_{\mathcal{C}} \lambda \wedge \Omega,$$

where $\partial \mathcal{C}_i = L_i - L_0$, and $\partial \mathcal{C} = L - L_0$.

Now the flat norm convergence gives $L_i - L = A_i + \partial B_i$ for some integral currents A_i, B_i , with $\text{Mass}(A_i) + \text{Mass}(B_i) \rightarrow 0$. Since $[L_i] = [L] = [L_0] \in H_n(X)$, the homology class of A_i is zero. A version of the isoperimetric theorem (see Proposition 5.11 below) then says $A_i = \partial B'_i$ with $\text{Mass}(B'_i) \leq C \text{Mass}(A_i)^{(n+1)/n} \rightarrow 0$. Without loss of generality we absorb B'_i into B_i . Then we can simply choose $\mathcal{C}_i = \mathcal{C} + B_i$, which is legitimate since it satisfies $\partial \mathcal{C}_i = L_i - L_0$. The claim follows by

$$\left| \int_{B_i} \lambda \wedge \Omega \right| \leq C \text{Mass}(B_i) \rightarrow 0. \quad \blacksquare$$

Robustness of potential clustering. We revisit the potential clustering property (see Section 3.8.3) from the geometric measure theory perspective. Assume as always that the Lagrangian integral current L is quantitatively almost calibrated, homologous to L_0 , and equipped with Lagrangian potential f_L . Given constants $N \in \mathbb{N}, A \in \mathbb{R}_+$, we say L satisfies (N, A) -potential clustering, if

$$L = \sum_1^N L_i, \quad f_L L = \sum_1^N f_{L_i} L_i,$$

for quantitatively almost calibrated, closed Lagrangian integral currents L_i with potential f_{L_i} , contained inside the support of L , such that the oscillation of the Lagrangian potentials have uniform bounds

$$\sup_{L_i} f_{L_i} - \inf_{L_i} f_{L_i} \leq A,$$

while for any $i > j$,

$$\sup_{L_j} f_{L_j} \leq \inf_{L_i} f_{L_i}.$$

Without loss of generality, we assume $\sup_{L_0} f_{L_0} - \inf_{L_0} f_{L_0} \leq A$ for the fixed Lagrangian L_0 .

Remark 5.9. Here we allow L_i to have overlapping supports. For instance, it is possible for $L_1 = L_2$ as currents, but f_{L_2} and f_{L_1} differ by a constant.

We will later be interested in uniform upper bounds on N, A . For now, we observe the robustness under limits.

Corollary 5.9. *Fix the choice of N, A . Suppose we are given a sequence of Lagrangians $L^{(j)}$ with potential satisfying (N, A) -potential clustering, and assume $L_i^{(j)} \rightarrow L_i$ as currents for $1 \leq i \leq N$, all $f_{L_i}^{(j)}$ have uniform L^∞ bounds, and $f_{L_i^{(j)}} L_i^{(j)} \rightarrow f_{L_i} L_i$ as currents. Then the limit*

$$L = \sum_{i=1}^N L_i$$

with its potential f_L also satisfies (N, A) -potential clustering.

Proof. Notice a potential bound such as $f_L \geq c$ can be characterized by the positivity of the measure

$$g \mapsto \int_L (f_L - c) g \operatorname{Re} \Omega.$$

This characterization is robust under current convergence, so

$$\sup_{L_i} f_{L_i} \leq \limsup_j \sup_{L_i^{(j)}} f_{L_i^{(j)}}, \quad \inf_{L_i} f_{L_i} \geq \liminf_j \inf_{L_i^{(j)}} f_{L_i^{(j)}},$$

hence the potential clustering bounds pass to the limit. ■

5.2. Quantitative almost calibratedness

One major advantage of the quantitative almost calibrated condition is that within a fixed homology class of Lagrangians, it guarantees an a priori volume upper bound (see Lemma 2.1). We shall explain that, under very mild asymptotic conditions on the ambient Calabi–Yau manifolds, it also guarantees that the Lagrangian remains within a bounded region. As such, the Federer–Fleming compactness applies automatically, and Allard compactness applies under the additional hypothesis of a uniform bound on $\int_L |\vec{H}|$. We also discuss a number of instructive but not particularly difficult consequences of quantitative almost calibratedness.

Remark 5.10. The arguments in this section are adaptations of Neves [63]. They can also be easily adapted to the almost Calabi–Yau setting under mild conditions on the volume density.

As a preliminary, we will say the regularity scale near a given point P on the Calabi–Yau manifold is at least $O(R)$, if P is contained in a complex coordinate ball with Euclidean radius at least R , on which $\Omega = f dz_1 \wedge \cdots \wedge dz_n$, and⁶⁰

$$C^{-1} \delta_{ij} \leq g_{i\bar{j}} \leq C \delta_{ij}, \quad |\partial^k g_{i\bar{j}}| \leq C(k) R^{-k}, \quad |f - 1| \leq \frac{1}{2}, \quad |\partial^k f| \leq C R^{-k}. \quad (5.2)$$

⁶⁰We will actually only use the metric uniform equivalence. But for Calabi–Yau metrics, the higher derivative estimates are in any event implied by metric equivalence, after shrinking the balls slightly, by Evans–Krylov theory.

From now on we will *assume the regularity scale tends to infinity for $P \rightarrow \infty$* . This is a very mild condition on the Calabi–Yau manifold, for instance satisfied by asymptotically conical Calabi–Yaus.

Remark 5.11. This asymptotic condition is essentially the weakest that can prevent the almost calibrated Lagrangian in a given homology class from escaping to infinity. For instance, if L is a special Lagrangian in X , then $L \times S^1$ is a special Lagrangian in the product $X \times T^*S^1$, which can obviously be translated in the \mathbb{R} direction to escape to infinity, albeit not preserving the $D^b \text{Fuk}(X)$ class. One can still hope to obstruct escaping to infinity using Floer theory, but then incorporating singular Lagrangians requires more foundational work.

Isoperimetric inequality.

Lemma 5.10 (Isoperimetric inequality [63, Lemma 3.10]). *Let L be a closed Lagrangian integral current in (X, ω, Ω) , and consider a Euclidean coordinate ball B_{2R} on which the regularity scale is at least $O(R)$. Assume the quantitative almost calibrated condition $\cos \theta \geq \sin \varepsilon$. Then there is a universal constant C depending only on ε, n and the metric uniform equivalence constant in (5.2), so that*

$$\text{Mass}(A)^{(n-1)/n} \leq C \text{Mass}(\partial A)$$

for all closed subsets A of $\text{supp}(L) \cap B_R$ with rectifiable boundary. Here the Hausdorff measures are computed using the Calabi–Yau metric.

Proof. The isoperimetric theorem [73, Theorem 6.1] guarantees the existence of an integral current A' supported in B_R such that $\partial A' = \partial A$ and for which

$$\text{Mass}(A')^{(n-1)/n} \leq C \text{Mass}(\partial A).$$

Notice the metric uniform equivalence means we do not need to be careful to distinguish the Hausdorff measure for the Euclidean metric and the Calabi–Yau metric. Let T denote the cone over the current $A - A'$, then $\partial T = A - A'$, and thus by the quantitative calibrated condition, we have

$$\begin{aligned} \text{Mass}(A) &\leq \frac{1}{\sin \varepsilon} \int_A \text{Re } \Omega = \frac{1}{\sin \varepsilon} \int_{A'} \text{Re } \Omega + \frac{1}{\sin \varepsilon} \int_{\partial T} \text{Re } \Omega \\ &\leq \frac{1}{\sin \varepsilon} \text{Mass}(A') + \frac{1}{\sin \varepsilon} \int_T d \text{Re } \Omega \leq \frac{1}{\sin \varepsilon} C \text{Mass}(\partial A)^{n/(n-1)}, \end{aligned}$$

which is the isoperimetric inequality. ■

Remark 5.12. The standard isoperimetric theorem inside the Euclidean space [73, Theorem 6.1] does not state $\text{supp}(A') \subset B_R$. Once we have found some A' with $\partial A' = \partial A$ with mass control, we can pushforward by a Lipschitz retraction map $F: \mathbb{R}^{2n} \rightarrow B_1$:

$$F(x) = \begin{cases} x, & x \in B_R, \\ \frac{Rx}{|x|}, & x \in \mathbb{R}^{2n} \setminus B_R. \end{cases}$$

Replacing A' by F_*A' , the mass cannot increase, and $\partial F_*(A') = F_*(\partial A') = F_*(\partial A) = \partial A$, and we have ensured the support is contained in B_R .

The following version of the isoperimetric theorem should be well known to experts, but for lack of a reference we include a proof below.

Proposition 5.11 (Isoperimetric theorem on complete manifolds). *Let X be a complete Riemannian manifold, and ∂A be an m -dimensional exact integral current supported in a fixed bounded open subset $U \subset X$. Then there is an integral current A' supported in a fixed large bounded subset of X , with $\partial A = \partial A'$ and*

$$\text{Mass}(A') \leq \text{const} \cdot \min\{\text{Mass}(\partial A)^{(m+1)/m}, \text{Mass}(\partial A)\}.$$

Proof. We first isometrically embed X into an ambient Euclidean space \mathbb{R}^N , so A can be regarded as an integral current compactly supported in X . Fix a small number $\rho_0 > 0$ such that over U the ρ_0 -neighbourhood in \mathbb{R}^N is isomorphic to the normal bundle, so there is a smooth retraction map F back to U . The Lipschitz norm of F is approximately one.

Applying the deformation theorem for \mathbb{R}^N [73, Section 5.3] to the current ∂A , with a parameter $\rho \leq \rho_0$ to be fixed, we can write

$$\partial A = P + \partial R,$$

where P, R are integral currents inside \mathbb{R}^N , supported in the $O(\rho)$ neighbourhood of $\text{supp}(\partial A)$, with

$$\text{Mass}(R) \leq C\rho \text{Mass}(\partial A), \quad \text{Mass}(P) \leq C \text{Mass}(\partial A),$$

where the constant C depends only on N, m . Moreover, P is an integral linear sum of m -dimensional faces in the standard grid decomposition of \mathbb{R}^N with cube size ρ . We now push forward via F :

$$F_*P + \partial F_*R = F_*(\partial A) = \partial A,$$

since $\partial A \subset U \subset X$ is fixed by F . Note that F_*P, F_*R both live inside U , and their mass bounds are essentially the same as P, R respectively.

Suppose first that $\text{Mass}(\partial A) \ll 1$. If P is non-zero, then by the grid description of P ,

$$\rho^m \leq \text{Mass}(P) \leq C \text{Mass}(\partial A).$$

So by choosing $\rho = 2(C \text{Mass}(\partial A))^{1/m}$ in the above, we force $P = 0$, so $\partial A = \partial F_*R$, with mass bound $\text{Mass}(F_*R) \leq \text{const} \text{Mass}(\partial A)^{(m+1)/m}$, so it suffices to take $A' = F_*R$.

Now suppose $\text{Mass}(\partial A) \gtrsim 1$, then we choose $\rho = \rho_0$. Without loss of generality, we can replace ∂A by F_*P , and pretend $R = 0$. We know that

- P is an integral linear combination of grid cube faces, where all the cubes lie in a bounded region of a fixed grid;
- F_*P is an exact current on X .

The set of all such P form a finitely generated abelian group, which by classification is isomorphic to the direct sum of $\bigoplus_1^r \mathbb{Z}e_i$ and a finite abelian group. For any given element

$$P = \sum a_i e_i + \text{finite group part},$$

the linear coefficients a_i of e_i for $i = 1, \dots, r$ are bounded by

$$|a_i| \lesssim \text{Mass}(P) \lesssim \text{Mass}(\partial A).$$

Each e_i gives rise to an exact simplicial chain F_*e_i inside X , which is the boundary of a finite mass integral current Q_i . Thus

$$\text{Mass}\left(\sum_1^r a_i Q_i\right) \leq \sum_1^r |a_i| \text{Mass}(Q_i) \leq \text{const} \cdot \text{Mass}(\partial A).$$

The finite group part gives rise to another simplicial chain inside X which is the boundary of some finite mass integral current. Thus we have produced an integral current A' with $\partial A' = \partial A$, and mass bound

$$\text{Mass}(A') \leq C(\text{Mass}(\partial A)) + C \leq \text{const} \text{Mass}(\partial A)$$

since we are in the $\text{Mass}(\partial A) \gtrsim 1$ case. ■

Volume monotonicity and lower bound.

Corollary 5.12 (Volume lower bound). *If P is in the support of L , then there is a uniform lower bound on the volume of L inside Euclidean coordinate balls of radius less than R :*

$$\text{Vol}_g(L \cap B(P, r)) \geq C^{-1} r^n \quad \forall r \leq R. \quad (5.3)$$

Proof. Let $f(r) = \mathcal{H}^n(L \cap B(r)) > 0$, then f is increasing in r , and for a.e. $0 < r < R$, by the coarea formula,

$$f'(r) = \int_{\partial B(r) \cap L} \frac{1}{|\nabla r|} d\mathcal{H}^{n-1} \geq C^{-1} \mathcal{H}^{n-1}(\partial B(r) \cap L) \geq C^{-1} f(r)^{(n-1)/n}.$$

The last inequality is the isoperimetric inequality. Thus $\frac{d}{dr} f^{1/n} \geq C^{-1}$, whence we have the volume lower bound $f(r)^{1/n} \geq C^{-1} r$. ■

Remark 5.13. Volume lower bounds like (5.3) are familiar in minimal surface theory, but usually require some integral bound on the mean curvature. Notably, here we need no such assumption; the quantitative almost calibrated condition only concerns the antiderivative θ of $\tilde{H} = J\nabla\theta$.

No escape to spatial infinity.

Corollary 5.13. *Assume near the infinity of X , the regularity scale grows to infinity. Fix the homology class of the quantitatively almost calibrated Lagrangian L . Then L is contained in a fixed compact subset of X .*

Proof. From Lemma 2.1 there is an a priori volume bound $\text{Vol}(L) \leq C$. But if P is in the support of L , then the regularity scale of X near P is bounded by

$$R^n \leq C \text{Vol}(L \cap B(R)) \leq C,$$

whence P must remain in a fixed compact subset. ■

Non-triviality of homology classes.

Corollary 5.14 (Non-triviality of homology classes). *There is a lower bound*

$$\int_L \text{Re } \Omega \geq C^{-1}$$

depending only on the ambient Calabi–Yau and the almost calibratedness constant ε . Here we do not a priori specify the homology class of L .

Proof. The asymptotic hypothesis on the regularity scale implies in particular that the regularity scale has a global lower bound on X . This gives a uniform lower bound on $\text{Vol}(L)$, which by the quantitative almost calibratedness gives a lower bound on the homological integral $\int_L \text{Re } \Omega$. ■

The significance is that in a variational setup to find special Lagrangians among certain Lagrangian currents in the fixed homology class $[L_0]$, it may happen that the Lagrangian breaks into several connected components, each given by a closed Lagrangian current L_i with

$$\sum_i \int_{L_i} \text{Re } \Omega = \int_L \text{Re } \Omega.$$

The non-triviality of each homology class then puts an upper bound on the number of connected components. Moreover, each L_i would inherit a volume upper bound from L . Since all Lagrangians are contained in a fixed compact region, Poincaré duality easily gives an upper bound on the homology classes of L_i :

$$\|[L_i]\|_{H_n(X)} \leq C \text{Vol}(L_i) \leq C.$$

Combining the above, *only finitely many homology classes, multiplicities, and number of components can appear in a given variational problem.*

5.2.1. Intrinsic distance bound and potential clustering. Suppose L is a *smooth* immersed Lagrangian, then it inherits a Riemannian metric from the restriction of the Calabi–Yau metric, so we can speak of the intrinsic distance function on L . Instead of extrinsic balls such as $B(P, r)$, we can talk about intrinsic balls $\hat{B}(P, r)$. A key distinction is that intrinsic distance does not need to extend to a continuous function on $X \times X$, the prototypical example being the union of two embedded Lagrangians, whose domains are disjoint, but whose images in X intersect. Clearly, the intrinsic distance bounds extrinsic geodesic distance, so $\hat{B}(P, r)$ is always contained in an extrinsic geodesic ball of radius r , but the converse is far from true. The intrinsic distance between two distinct connected components would simply be infinity. One can think of intrinsic distance as a *quantitative measurement of connectedness*.

As usual, the regularity scale on X grows to infinity asymptotically by assumption, so the regularity scale has a global lower bound. The following lemma has the same proof as Corollary 5.12. (The essence of this argument also appears in Neves [63, Lemma 3.9].)

Lemma 5.15 (Intrinsic ball volume lower bound). *Let L be a smooth immersed compact Lagrangian in X , satisfying the quantitative almost calibrated condition. For any P in the support of L , there is a uniform bound*

$$\text{Vol}(\hat{B}(P, r) \cap L) \geq C^{-1} r^n, \quad r \leq 1.$$

Corollary 5.16 (Intrinsic diameter bound). *Assume further that the smooth, quantitatively almost calibrated compact Lagrangian L has connected domain. Then within a fixed homology class, the intrinsic distance of L has a uniform upper bound.*

Proof. Let $d = \text{dist}(P, P') \geq N = [d]$ be the intrinsic diameter of L . By connectedness, we can find P_1, \dots, P_N with $\text{dist}(P, P_i) = i$. Now the intrinsic balls $\hat{B}(P_i, \frac{1}{2})$ are disjoint, but each takes up a non-trivial amount of volume $\geq C^{-1}$. Thus

$$C^{-1} N \leq \text{Vol}(L) \leq \frac{1}{\sin \varepsilon} \int_L \text{Re}(\Omega) \leq C,$$

so there is an a priori bound on N , hence on d . ■

Corollary 5.17 (Potential oscillation bound). *Assume L is an exact, immersed, compact Lagrangian, satisfying the almost calibrated condition, such that the Lagrangian potential f_L has connected range. Then the potential f_L has an a priori bound*

$$\sup_L f_L - \inf_L f_L \leq C.$$

Proof. Given any P, P' on L , we write the potential as a line integral of the Liouville 1-form

$$f_L(P) - f_L(P') = \int_{P'}^P \lambda \leq \|\lambda\|_{C^0} \text{dist}(P, P').$$

Thus the intrinsic ball volume lower bound implies that if a lies in the range of f_L , then

$$\text{Vol}\left(\left\{|f_L - a| < \frac{1}{2}\right\}\right) \geq C^{-1}.$$

The range of f_L is by assumption a closed interval. If the interval has length $\geq N$, then we can find N distinct values of a with disjoint $\{|f_L - a| < \frac{1}{2}\}$, so

$$C^{-1}N \leq \text{Vol}(L) \leq \frac{1}{\sin \varepsilon} \int_L \text{Re}(\Omega) \leq C.$$

This provides an a priori bound on N , hence on the potential oscillation. \blacksquare

Corollary 5.18. *Assume L is an exact, immersed, compact Lagrangian, satisfying the almost calibrated condition, then it satisfies (N, A) -potential clustering (see Section 5.1.2) for some N, A with uniform bounds. Moreover, if L carries an unobstructed brane structure, then the potential clustering property is consistent with the twisted complex interpretation.*

Proof. By Lemma A.3, a general immersed Lagrangian can be decomposed into a union of Lagrangians L_i such that the ranges of the potentials f_{L_i} are connected, and any bounding cochain structure naturally produces a twisted complex. The oscillation of each L_i is bounded in terms of the quantitative almost calibration condition, and the ambient features of X . Moreover, the number of Lagrangian components is also a priori bounded. \blacksquare

As a notable consequence, we obtain the *uniform energy bound* on the holomorphic curves (see Proposition 3.41).

Remark 5.14. Although it is unclear how to make sense of the intrinsic distance on a general Lagrangian integral current, mildly singular Lagrangians (for instance with local conical singularities) do have a sensible notion of intrinsic distance, and the arguments in this section extend practically to all non-pathological examples, covering all Lagrangians that appear in Joyce's LMCF program. Furthermore, the potential clustering is robust under limits (see Corollary 5.9). As such, we believe it holds for all Lagrangians relevant to our variational program (see the class \mathcal{L} in Section 5.3 below).

5.2.2. Bounded part of the Solomon functional revisited. In Section 3.8.3, we discussed a Floer theoretic method to bound the difference between the Solomon functional and the elementary functional. The following alternative method is based on the homological nature of the Solomon functional (3.4), and has a more geometric measure theory flavour. Holomorphic curves do not feature in this approach, so the automatic transversality and the positivity conditions are not needed, and in fact the discussion works in the weak setting of varifolds and currents. It is vital to this approach that $H_{n+1}(X) = 0$ for Stein manifolds, so no homological ambiguity can arise for the bordism current.

Proposition 5.19. *Assume the Lagrangian L with potential f_L is quantitatively almost calibrated, homologous to L_0 , and satisfies (N, A) -potential clustering. The elementary functional (3.23) makes sense verbatim, with no smoothness assumptions. Then*

$$|\mathcal{S}(L) - \bar{\mathcal{S}}(L)|$$

has a uniform upper bound independent of L .

Proof. We know $L - L_0$ is homologous to zero in X , and contained in a bounded subset of X by Corollary 5.13. By a version of the isoperimetric theorem (see Proposition 5.11), we can find some compactly supported integral current \mathcal{C} with $\partial\mathcal{C} = L - L_0$, with mass bound

$$\text{Mass}(\mathcal{C}) \leq \text{const} \cdot \min\{\text{Mass}(L - L_0)^{(n+1)/n}, \text{Mass}(L - L_0)\}.$$

This \mathcal{C} has no relation to holomorphic curves. The quantitative almost calibratedness implies a volume bound on L (see Lemma 2.1), hence $\text{Mass}(\mathcal{C})$ is a priori bounded. The homological nature of the Solomon functional (3.4) gives

$$\mathcal{S}(L) = \int_L f_L \text{Im}(e^{-i\hat{\theta}} \Omega) - \int_{L_0} f_{L_0} \text{Im}(e^{-i\hat{\theta}} \Omega) - \text{Im} \int_{\mathcal{C}} \lambda \wedge e^{-i\hat{\theta}} \Omega.$$

The mass bound then implies

$$\left| \int_{\mathcal{C}} \lambda \wedge e^{-i\hat{\theta}} \Omega \right| \leq C \|\lambda\|_{C^0} \|\Omega\|_{C^0} \text{Mass}(\mathcal{C}) \leq \text{const}.$$

Here since L is contained in a bounded region, the terms λ and Ω are bounded. Finally, using the potential clustering bound,

$$\begin{aligned} & \left| \int_L f_L \text{Im}(e^{-i\hat{\theta}} \Omega) - \int_{L_0} f_{L_0} \text{Im}(e^{-i\hat{\theta}} \Omega) - \bar{\mathcal{S}}(L) \right| \\ & \leq A \left(\int_L |\text{Im}(e^{-i\hat{\theta}} \Omega)| + \int_{L_0} |\text{Im}(e^{-i\hat{\theta}} \Omega)| \right) \\ & \leq A(\text{Mass}(L) + \text{Mass}(L_0)) \leq \frac{2A}{\sin \varepsilon} \int_{L_0} \text{Re} \Omega. \end{aligned}$$

Combining the above shows the a priori bound on $|\mathcal{S}(L) - \bar{\mathcal{S}}(L)|$. ■

5.3. Variational strategy

The variational strategy to find special Lagrangians is the following:

(i) Find a suitable subset \mathcal{L} among all the quantitatively almost calibrated, exact Lagrangian integral currents homologous to L_0 . The class \mathcal{L} is closed in the varifold/current topology. It is very desirable to ensure Allard compactness and Federer–Fleming compactness both apply to \mathcal{L} .

(ii) Extend enough of Floer theory from the smooth setting to Lagrangian currents. Morally, the class \mathcal{L} consists of those Lagrangians that can be equipped with unobstructed brane structures in some weak sense, all isomorphic to L_0 in $D^b \text{Fuk}(X)$.

(iii) When the Lagrangian is equipped with the potential f_L , the additive constant freedom of f_L is a source of non-compactness, which affects \mathcal{S} . We need to ultimately match up the asymptotic behaviour of \mathcal{S} with the Floer theoretic obstructions. In other words, the role of stability conditions is to ensure the properness of the Solomon functional.

(iv) Once the Solomon functional is proper, we will follow the direct minimization strategy to find its minimum. We need to justify that the minimum L must be a special Lagrangian closed integral current, and then Almgren regularity will be able to ensure smoothness away from codimension two. Furthermore, we need a sufficiently robust version of the Thomas–Yau uniqueness argument to prove that the special Lagrangian representative is unique.

The class \mathcal{L} is a balance between two requirements: the *approximability by sufficiently smooth objects*, and the existence of *sufficiently many competitors*. A moral definition of \mathcal{L} is as follows:

- Among all the quantitatively almost calibrated, exact Lagrangian integral currents homologous to L_0 , we include all sufficiently smooth Lagrangians (e.g., immersed, T^2 -cones singularities, etc.) which admit unobstructed brane structures isomorphic to L_0 in $D^b \text{Fuk}(X)$.
- Then take the closure under the varifold/current topology.

Remark 5.15. Joyce’s LMCF is expected to preserve the exactness, the quantitative almost calibrated condition, and the unobstructedness of the brane structure, so sufficiently smooth objects in \mathcal{L} should remain in \mathcal{L} under Joyce’s LMCF. It is interesting to ask when the flow also preserves the positivity condition on the bordism current.

While at present several ingredients are missing, if this program can indeed be carried through, it would prove the existence of special Lagrangians under the assumption of Thomas–Yau semistability (see Definition 3.32).

L^p -Smoothing property and Joyce’s LMCF. Allard compactness requires an a priori bound $\int_L |\vec{H}| \leq C$, which cannot be implied by the quantitative almost calibrated condition, since the mean curvature involves one more derivative than the Lagrangian angle. However, for the purpose of our variational strategy, it is enough to ensure any minimization sequence of \mathcal{S} can be replaced by a sequence with $\int_{L_i} |\vec{H}| \leq C$.

Conjecture 5.20 (L^p -smoothing property). *There exists $p \geq 1$ and a uniform constant C , such that for any $L \in \mathcal{L}$, we can find $L' \in \mathcal{L}$ with $\mathcal{S}(L') \leq \mathcal{S}(L)$, and $\int_{L'} |\vec{H}|^p \leq C$.*

Remark 5.16. This is called a ‘smoothing property’ because L' quantitatively improves the regularity of L . It does not suggest L' is smooth, and indeed we expect the special

Lagrangians which minimize \mathcal{S} may have codimension two singularities. Since the volume is a priori bounded, the Hölder inequality shows the L^p -smoothing property is stronger for bigger p , and in particular L^2 -smoothing implies L^1 -smoothing.

We think the smoothing property may be quite deep, and our limited attempt here is to explain how it relates to Joyce’s LMCF program, which suggests the smoothing property may hold with $p = 2$. Recall the defining feature of the Solomon functional is its variation property under exact isotopies among unobstructed objects:

$$\delta \mathcal{S} = \int_L h \operatorname{Im}(e^{-i\hat{\theta}} \Omega),$$

which holds under sufficient smoothness assumptions. Under a sufficiently smooth LMCF (L_t) in a Calabi–Yau manifold, the Lagrangians evolve by the local Hamiltonian function $-\theta_t$ up to an inconsequential additive constant (see Section 4.1), so \mathcal{S} evolves by

$$\begin{aligned} \partial_t \mathcal{S} &= - \int_{L_t} (\theta_t - \hat{\theta}) \operatorname{Im}(e^{-i\hat{\theta}} \Omega) \\ &= - \int_{L_t} (\theta_t - \hat{\theta}) \operatorname{Im}(e^{i(\theta - \hat{\theta})}) d\operatorname{vol}_{L_t} \\ &= - \int_{L_t} (\theta_t - \hat{\theta}) \sin(\theta_t - \hat{\theta}) d\operatorname{vol}_{L_t}. \end{aligned} \tag{5.4}$$

If L_t is almost calibrated, then $-\pi < \theta - \hat{\theta} < \pi$, so $\partial_t \mathcal{S} \leq 0$. We conclude that *the Solomon functional decreases in time along Joyce’s LMCF under the almost calibrated assumption*, at least for the time between the surgeries. It is plausible \mathcal{S} is either continuous or jumps downwards at the surgeries in Joyce’s LMCF,⁶¹ which would then imply the Solomon functional is monotone decreasing for all time.

Now recall that the heat equation on the Lagrangian angle implies an integral bound on the mean curvature (4.2). In particular, if the LMCF can be run for a *definite amount of time* T , then there exists some $t \leq T$, with

$$\int_{L_t} |\vec{H}|^2 d\operatorname{vol}_{L_t} \leq T^{-1} \int_{L_{t=0}} \theta^2 d\operatorname{vol}_{L_{t=0}} \leq C T^{-1},$$

where crucially the a priori constant C does not depend on any quantitative smoothness assumption on the initial Lagrangian, provided it is quantitatively almost calibrated. Such L_t would be a good candidate for L' , subject to the hypothesis that Joyce’s LMCF remains within the class of Lagrangians \mathcal{L} .

⁶¹A somewhat analogous phenomenon in the Brakke flow is that the total volume mass is either continuous or can only jump downwards in time. The mass loss is typically related to the disappearance of a component of the evolving varifold, which is conceptually similar to ‘collapsing zero objects’ in Joyce’s LMCF. This is ruled out by the almost calibrated condition, so optimistically one can even hope for the continuity of the Solomon functional in the almost calibrated setting.

Morally the class \mathcal{L} arises as varifold/current limits of those Lagrangians admissible in Joyce's program. Under the plausible assumption that Joyce's LMCF can be passed to the varifold/current limit, then the L^2 -smoothing property can be well explained. The $\mathcal{S}(L') \leq \mathcal{S}(L)$ condition comes from the decrease of the Solomon functional along the flow, and the $\int_{L'} |\tilde{H}|^2 d\text{vol}_{L'} \leq C$ condition would follow if Joyce's LMCF can be run for a uniform amount of time T . If T can be taken arbitrarily large, then we can demand further that the L^2 mean curvature is arbitrarily small.

Remark 5.17. In minimal surface theory, the ability to approximate an unknown object by objects with quantitative derivative controls, is frequently the key of the regularity theory. Notable examples include the Lipschitz and harmonic approximations that lie at the core of De Giorgi's ε -regularity theorem, and the centre manifolds at the core of Almgren's big regularity theorem. An excellent survey is [24]. While there are plenty of techniques for constructing area competitors in geometric measure theory, we lack useful ways to construct competitors within the Lagrangian world. Developing such techniques is essential to the L^p -smoothing property, and possibly also to the Floer theoretic aspects of the variational program.

5.4. Floer theory under weak regularity

The variational program must incorporate singular Lagrangians as objects of $D^b \text{Fuk}(X)$, which naturally raises many Floer theoretic questions, such as the following:

- Suppose a sequence of (exact, quantitatively almost calibrated, smooth) Lagrangians converge in the varifold/current topology to some singular Lagrangian, then what Floer theoretic information can be passed to the limit?
- What does it mean for two Lagrangian currents to lie in the same derived Fukaya category class?
- Does it still make sense to talk about Floer theoretic obstructions in the weak regularity setting?

In this section we will offer some general remarks and speculations about the nature of these problems, but will not solve them in any definitive way.

Remark 5.18. There is a field called C^0 -symplectic topology [66], which studies properties stable with respect to convergence of Lagrangians under C^0 -Hamiltonian isotopies, especially spectral-type invariants. This is morally related to our concerns here, but as far as the author understands, Floer theory for Lagrangian varifolds/currents is not yet explicitly treated in this field.

Floer theoretic difficulties. If one wishes to build Floer theory for Lagrangian currents by mimicking the smooth case constructions, then one immediately runs into a large number of severe difficulties.

- (i) For exact embedded Lagrangians, the *self Floer cohomology* of a Lagrangian is isomorphic to the singular cohomology: $HF^*(L, L) \simeq H^*(L)$. Now in the light of Almgren's

big regularity theorem, our best hope is that in the variational argument we only encounter codimension two singularities in the Lagrangian. We have no right to assume the topology of the Lagrangian is fixed in the variational framework. The homology groups $H_{n-m}(L)$ for $m \geq 1$ are highly unstable under varifold/current convergence if codimension two singularities can form, so for $m \geq 1$ we do not expect a direct geometric definition of $HF^m(L, L)$ for Lagrangian currents, that possesses any reasonable continuity property under convergence.

(ii) The standard way to set up Floer theory between two Lagrangians is to consider the *transverse intersection* points as the generators of the Floer complex, and counts of holomorphic strips as differentials between generators. This viewpoint depends heavily on the differential topology of the Lagrangians, which runs into troubles for Lagrangian currents, where tangent spaces only need to exist almost everywhere in a measure theoretic sense.

(iii) Once Lagrangian intersections are not well behaved, we cannot define the *bounding cochains* supported at intersection points in the usual way.

(iv) Parallel transport along *local systems* may break down.

(v) It is unclear how to define (relative) *spin structures* on Lagrangian currents.

(vi) Standard Floer theory depends heavily on *transversality arguments* based on differential topology, which is lost on Lagrangian currents.

In short, a direct geometric construction of the A_∞ structure is unlikely for Lagrangian currents.

Formal limit perspective. One natural idea is that we only develop Floer theory for sufficiently smooth Lagrangians (e.g., immersed Lagrangians, isolated T^2 -cones, etc), and formally treat Lagrangian currents using *approximation by smooth objects*. Suppose L_i are sufficiently smooth Lagrangian branes in the same $D^b \text{Fuk}(X)$ class, and $L_i \rightarrow L$ in the varifold/current topology, and assume the brane structures provide a Cauchy sequence in some appropriate sense, then one *formally* declare the Lagrangian current L as carrying an object in the same $D^b \text{Fuk}(X)$ class. A weak Lagrangian brane would then tautologically be an equivalence class of Cauchy sequences. The same Lagrangian current may in principle support many different formal brane structures, not necessarily all in the same derived category class.

In this perspective, weak Lagrangian branes are indirect constructions, whose properties amount to quantitative properties of sufficiently smooth Lagrangians that can be bounded in terms of a priori quantities such as the distance on the branes, the flat norm on the currents, the Hausdorff distance between the Lagrangians, etc.

Question 11. Is there a notion of distance between two Lagrangian branes L, L' in the same $D^b \text{Fuk}(X)$ class, that has *precompactness property modulo gauge* under varifold/current topology, in the setting of exact, quantitative almost calibrated Lagrangians with bounded Lagrangian potential?

One concrete notion of distance is as follows (see [34, Definition 2.2] and [10, Section 5]). We can look for the α, β representing generators in $HF^0(L, L')$ and $HF^0(L', L)$ with cohomological compositions equal to the identity; in the almost calibrated case $CF^{-1}(L, L') = 0$, so α, β are unique up to scaling. Since all bounding cochains and A_∞ products have non-negative Novikov exponents, and the sum of Novikov exponents add up to zero, we must have some negative Novikov exponent for α or β . In our context, the Novikov exponent amounts to $(f_L - f_{L'})(p)$ at $p \in CF^0(L, L')$ and $(f_{L'} - f_L)(q)$ at $q \in CF^0(L', L)$. The quantity

$$-\min\{\text{Novikov exponents among all intersection points of } \alpha, \beta\}$$

provides a candidate notion of distance $d(L, L')$ between Lagrangian branes. Notice this distance bounds the energy of the holomorphic discs with boundary on L, L' . Given three objects L, L', L'' , by considering the composition of the generators, it is easy to deduce

$$d(L, L'') \leq d(L, L') + d(L', L'').$$

Does this notion of distance have any precompactness property? Namely, given a sequence of sufficiently smooth Lagrangian objects $L_i \in \mathcal{L}$, (e.g., a minimizing sequence for the Solomon functional), and assuming the Lagrangian potentials are uniformly bounded, then up to making gauge equivalent choices of local systems and bounding cochains, when can we extract a Cauchy subsequence?

Remark 5.19. As an illustration of the subtlety, consider immersed Lagrangians L built as the cone of $L_2 \xrightarrow{\gamma} L_1[1]$. Replacing γ by $c\gamma$ for $c > 0$ results in new bounding cochain structures on $L_1 \cup L_2$, but the distance between these brane structures is zero. The limit $c \rightarrow 0$ however belongs to a different $D^b \text{Fuk}(X)$ class. This suggests our formulation of weak Lagrangian branes is probably not sufficient to distinguish between several derived category classes.

One may also ask if the weak Lagrangian branes agree with ordinary Lagrangian branes in the case of smooth immersed Lagrangians.

Question 12. Suppose L_i is a sequence of immersed Lagrangian branes, all in the same $D^b \text{Fuk}(X)$ class, and is a Cauchy sequence with respect to the distance on the branes. Suppose L is an immersed Lagrangian, and $L_i \rightarrow L$ in the varifold/current topology. Then does there exist a suitable brane structure on L so that $L_i \rightarrow L$ with respect to the distance on the branes?

Geometric perspective: Bordism currents and triangulated categories. It is interesting to ask if *any* Floer theoretic geometric construction may be performed on Lagrangian currents at all. While the A_∞ -category structure on the Fukaya category may not necessarily be robust under varifold/current convergence of Lagrangians, only a subset of the structures are essential to the Thomas–Yau conjecture:

- The notion of derived Fukaya category classes.

- The notion of distinguished triangles, within the class of Lagrangians \mathcal{L} . This is the categorical shadow of the phenomenon that Lagrangians can be broken into several components under weak limits.
- The central charge function.

The central charge is of numerical nature, and is continuous under convergence in the current topology. A key feature of L, L' lying in the *same derived category class* is that there is a bordism current \mathcal{C} constructed from holomorphic curves, such that $\partial\mathcal{C} = L - L'$. Likewise for *distinguished triangles* $L_1 \rightarrow L \rightarrow L_2 \rightarrow L_1[1]$ in the weak regularity setting, a key expected property is that there should be a bordism current between L and $L_1 + L_2$, constructed from families of holomorphic curves.

Question 13. Let L_i, L'_i be unobstructed (sufficiently smooth) exact Lagrangians, all in the same derived category class. Assume convergence $L_i \rightarrow L$ and $L'_i \rightarrow L'$ in the varifold/current topology. Can we assign an $(n + 1)$ -bordism current \mathcal{C} between L and L' , constructed from the moduli space of holomorphic curves with boundary on L and L' ?

The basic idea is to take the bordism current \mathcal{C}_i with $\partial\mathcal{C}_i = L_i - L'_i$, constructed from the universal family of holomorphic curves, and attempt to extract the limit as currents. This could be morally viewed as a version of *Gromov compactness* for families. As rather strong evidence, in the quantitatively almost calibrated setting we derived uniform energy bound for holomorphic curves contributing to \mathcal{C}_i , by proving the potential clustering property (see Section 5.2.1, and Proposition 3.41). If we work with *Fukaya category over the integers*, the bordism currents \mathcal{C}_i would be *integral currents*, and we can hope to extract limit by some compactness argument. The problem is that we do not know \mathcal{C}_i have uniform mass upper bounds. Moreover, it is an interesting question how to formulate the parametrized family structure of the bordism current in the geometric measure theory language.

Remark 5.20. We note that Question 13 is formulated without any smoothness assumption on L and L' . In the light of the conjectural L^p -smoothing property (see Section 5.3), one may be able to assume some a priori L^2 -bound on the mean curvature.

While Lagrangian intersections, bounding cochains, spin structures, local systems etc. do not make sense directly on Lagrangian currents, the bordism current has a chance to make sense, and encodes substantial information. For instance, the orientations of the moduli spaces reflect the spin structures, and the weighting factors for the moduli spaces encode the combined effect of bounding cochain elements and the parallel transport along the local system.

Remark 5.21. As mentioned in Section 3.5, the mere requirement for the Floer theoretic obstruction criterion (i.e., the stability condition) to make sense for Lagrangian currents is already very constraining. Most statements are simply impossible to make without concepts that need at least C^1 -regularity, and the bordism currents between integration cycles

are among the rare exceptions. This was one of the heuristic arguments in Section 3.5 that obstructions must come from bordism currents.

Question 14. How much of the triangulated category structure works for weak regularity exact Lagrangians? How much of Floer theory can be developed upon the notion of bordism currents? Is it possible to encode weak Lagrangian branes à la the formal limit perspective, in terms of bordism currents?

We mentioned in Remark 3.6 that when more than two Lagrangians are present, Floer theory would also produce $(n + 2)$ -dimensional currents whose boundary exhibit homological relations between the $(n + 1)$ -dimensional bordism currents. Such ‘bordisms between bordisms’ may encode further information about the triangulated category.

Privileged role of HF^0 . We consider quantitative almost calibrated Lagrangians. We mentioned above that $HF^m(L, L)$ for $m \geq 1$ is problematic, by analogy with singular cohomology. On the other hand, $H^0(L) \simeq H_n(L)$ is much more *robust* compared to higher cohomologies, in the sense that the fundamental cycle of L can deform in a *continuous* way, under topological changes such as the shrinking of a codimension two cycle. Continuing with the analogy, we expect the geometric information in HF^0 behaves more continuously under current/varifold limits than the higher degree Floer groups. This is compatible with the fact that the bordism current \mathcal{C} between L, L' encodes the compositions $\alpha \circ \beta = 1_{L'}$ and $\beta \circ \alpha = 1_L$, with $\alpha \in HF^0(L, L')$ and $\beta \in HF^0(L', L)$, and we expect bordism currents have some continuity properties under varifold/current convergence.

Remark 5.22. This privileged role of HF^0 is reflected in the usual Thomas–Yau argument (see Section 2.2), which only makes use of HF^0 , not the higher Floer cohomologies, nor full set of higher A_∞ products.

Remark 5.23. In the passage from the Fukaya category to the derived category, the morphism space only retains HF^0 , not the full HF^* . The usual way the derived category remembers higher Floer cohomology is via the shift operator $[m]$. However, in the Thomas–Yau–Joyce picture, working with the almost calibrated setting means conjecturally that we are picking out an abelian subcategory, which breaks the shift symmetry of the derived category. This gives a categorical explanation why HF^0 may behave very differently from the higher Floer groups.

Multiplicity issues. The same underlying geometric Lagrangian can conceivably support many different objects in the Fukaya category. A possible source of this problem is a sequence of immersed Lagrangians L_i converging to a *multiple* of a Lagrangian current L . The underlying Lagrangian current contains only the support information and the multiplicity, which can be imagined as the number of sheets in L_i . Much geometric information, however, is not captured this way

Take two Lagrangians L_1, L_2 which are both C^∞ close to a given immersed Lagrangian L' , but whose Lagrangian potentials differ by approximately a constant. In the limit

$L_1 \cup L_2 \rightarrow 2L'$ as currents, but the *potential information* is lost. On the other hand, the potential clustering property can restore this information.

Immersed Lagrangians may be non-trivial (branched) covers over other immersed Lagrangians. When this happens, the *monodromy information* is not remembered by the underlying current. On the other hand, it is conceivable that some (generalized) local system data can restore this information.

Let Q be a closed smooth manifold. Abouzaid [2] showed that the wrapped Fukaya category of the cotangent bundle T^*Q is generated by any cotangent fibre T_q^*Q , and the wrapped Floer cochain complex of T_q^*Q is A_∞ -equivalent to $C_{-*}(\Omega_q Q)$ for the based loop space $\Omega_q Q$. In particular, for any (compact, embedded, exact) Lagrangian $L \subset T^*Q$, the Floer cohomologies $HW^*(T_q^*Q, L)$ and $HF^*(L, L)$ are representations of $H_{-*}(\Omega_q Q)$. This cotangent bundle case can be viewed as the local model of Lagrangians contained in a small neighbourhood of a given embedded Lagrangian.

It is interesting to ask how much of such information can still make sense for Lagrangian currents.

Remark 5.24. Multiple covers of Lagrangians may be related to the following problem of the Fukaya category. Given a class in the Grothendieck group of $D^b \text{Fuk}(X)$ represented by a Lagrangian, one may ask if the primitive of this class is also represented by a Lagrangian. Such questions are related to the idempotent closure problem of $D^b \text{Fuk}(X)$ in Joyce’s program, which seems very delicate.

Remark 5.25. Construction of special Lagrangian branched multiple covers over given special Lagrangians is currently studied by Donaldson [30] and He [37] among others.

Remark 5.26. A holomorphic vector bundle analogue for multiply covered Lagrangians is the (multiple) extension of the bundle by itself, such as the E' fitting into a short exact sequence $0 \rightarrow E \rightarrow E' \rightarrow E \rightarrow 0$. In the HYM setting these are prototypical sources of semistable but not stable bundles, and it would not be surprising if similar phenomenon happens in the Thomas–Yau program.

5.5. Asymptotes of the Solomon functional

We have emphasized that the Solomon functional depends not only on L , but also the potential f_L , and that the freedom of additive constants causes the space of (L, f_L) to be non-compact, even though the space of Lagrangians \mathcal{L} is more or less compact under the varifold/current topology. We now wish to explain why the asymptotic behaviour of the Solomon functional should be controlled by Thomas–Yau semistability. The key tool is an a priori bound on the difference between the Solomon functional and the elementary functional, for which we gave sufficient conditions in Sections 3.8.3 and 5.2.2.

In the setup of (N, A) -potential clustering (see Corollary 5.9, Section 5.2.1), we will rewrite the elementary functional $\bar{\mathcal{S}}$ (see (3.23)). Recall we have a Lagrangian L built from L_1, \dots, L_N ; in the unobstructed immersed Lagrangian context, this structure comes from

a twisted complex (see Section 3.8.3). We introduce the new Lagrangian currents

$$\mathcal{E}_k = L_1 + L_2 + \cdots + L_k, \quad k = 0, 1, 2, \dots, N,$$

which in the immersed context corresponds to the twisted complex (3.2). In particular, $\mathcal{E}_N = L$, which is homologous to L_0 . Thus

$$\begin{aligned} \bar{s} &= \operatorname{Im} \left(\sum_1^N \left(\sup_{L_i} f_{L_i} \right) e^{-i\hat{\theta}} \left(\int_{\mathcal{E}_i} \Omega - \int_{\mathcal{E}_{i-1}} \Omega \right) \right) - \left(\sup_{L_0} f_{L_0} \right) \operatorname{Im} \left(e^{-i\hat{\theta}} \int_{L_0} \Omega \right) \\ &= \operatorname{Im} \left(\sum_1^{N-1} \left(\sup_{L_i} f_{L_i} - \sup_{L_{i+1}} f_{L_{i+1}} \right) e^{-i\hat{\theta}} \int_{\mathcal{E}_i} \Omega \right) \\ &\quad + \left(\sup_{L_N} f_{L_N} - \sup_{L_0} f_{L_0} \right) \operatorname{Im} \left(e^{-i\hat{\theta}} \int_{L_0} \Omega \right). \end{aligned}$$

But we chose in the beginning $\hat{\theta} = \arg \int_{L_0} \Omega$. Thus $\operatorname{Im}(e^{-i\hat{\theta}} \int_{L_0} \Omega) = 0$ and

$$\bar{s} = \sum_1^{N-1} \left(\sup_{L_i} f_{L_i} - \sup_{L_{i+1}} f_{L_{i+1}} \right) \operatorname{Im} \left(e^{-i\hat{\theta}} \int_{\mathcal{E}_i} \Omega \right). \quad (5.5)$$

As part of the potential clustering property, we have

$$\sup_{L_1} f_{L_1} \leq \sup_{L_2} f_{L_2} \leq \cdots \leq \sup_{L_N} f_{L_N}.$$

We arrive at the following key *dichotomy*:

In the *unstable* case, there exists some $1 \leq k \leq N-1$, such that

$$\operatorname{Im} \left(e^{-i\hat{\theta}} \int_{\mathcal{E}_k} \Omega \right) > 0,$$

or equivalently

$$\arg \int_{\mathcal{E}_k} \Omega > \hat{\theta}.$$

Notice L fits into a distinguished triangle

$$\mathcal{E}_k \rightarrow L \rightarrow \cup_{i \geq k+1} L_i \rightarrow \mathcal{E}_k[1],$$

We explained in Theorem 3.21 under the extra hypotheses of automatic transversality and the positivity condition, that this leads to a *Floer theoretic obstruction*. In Conjecture 3.31 we heuristically argued that even without these extra hypotheses, the Floer theoretic obstruction should follow from the Thomas–Yau–Joyce program.

From a different perspective, we can add an arbitrarily large *positive* number $a > 0$ to the Lagrangian potential on L_{k+1}, \dots, L_N . This is compatible with the Novikov positivity condition, so L stays unobstructed, but \bar{S} changes by an unbounded amount

$$-a \operatorname{Im} \left(e^{-i\hat{\theta}} \int_{\mathcal{E}_i} \Omega \right) < 0.$$

We conclude that *in the unstable case, the elementary functional is unbounded from below*.

In the *semistable* case, for any L in the class \mathcal{L} that can be written in the twisted complex form as above, we always have

$$\operatorname{Im}\left(e^{-i\hat{\theta}} \int_{\mathcal{E}_k} \Omega\right) \leq 0 \quad \forall k = 1, 2, \dots, N-1. \quad (5.6)$$

Then the elementary functional (5.5) is *non-negative*.

In Section 5.2 we argued that since the homology class of L is prescribed a priori, subject to the quantitative almost calibrated assumption, only finitely many possibilities of homology classes can arise for L_i in any decomposition. Thus the stronger condition

$$\operatorname{Im}\left(e^{-i\hat{\theta}} \int_{\mathcal{E}_k} \Omega\right) < 0 \quad \forall k = 1, 2, \dots, N-1$$

would be equivalent to a uniform bound: for some small $c > 0$,

$$\operatorname{Im}\left(e^{-i\hat{\theta}} \int_{\mathcal{E}_k} \Omega\right) \leq -c < 0 \quad \forall k = 1, 2, \dots, N-1.$$

This holds when the class \mathcal{L} is *strictly stable* (see Definition 3.32). Together with potential clustering, it implies

$$\bar{\mathcal{S}}(L) \geq c(\sup_L f_L - \inf_L f_L + A).$$

Therefore, if the Lagrangian potential oscillation becomes unbounded, then the elementary functional goes to positive infinity. The geometric intuition is the properness of the Solomon functional modulo a global additive constant for f_L .

Since the Solomon functional and the elementary functional only differ by a bounded amount, the above conclusions transfer to the Solomon functional. Thus *the Solomon functional is bounded below in the semistable case, and unbounded from below in the unstable case*. A key slogan here is that *the asymptotic behaviour of the Solomon functional is governed by Floer theory*. This is analogous to the partially conjectural picture in the variational approach to the HYM equation, where the asymptotic behaviour of the Donaldson functional is governed by algebraic geometry (see Section 2.5).

Remark 5.27. In Definition 3.32, the Thomas–Yau semistability makes use of distinguished triangles for all almost calibrated Lagrangian objects, not just those with $|\theta| \leq \frac{\pi}{2} - \varepsilon$. This makes the Thomas–Yau semistability a priori stronger than the semistable situation of the above dichotomy. We expect from the Thomas–Yau–Joyce picture that both stability notions are actually equivalent under our initial assumption that there is a representative L_0 with $|\theta| < \frac{\pi}{2} - \varepsilon$. But for our main purpose, that Thomas–Yau semistability implies the existence of special Lagrangians, we do not mind Thomas–Yau semistability being stronger than necessary.

5.5.1. Thomas–Yau conjecture. The following is our interpretation of the *Thomas–Yau existence conjecture*.

Conjecture 5.21. *Suppose L_0 is an exact, quantitatively almost calibrated, unobstructed Lagrangian object in \mathcal{L} . Assuming Thomas–Yau semistability for L_0 , then the following (equivalent) statements hold:*

- (1) *There is a special Lagrangian representative in \mathcal{L} .*
- (2) *There is no distinguished triangle in \mathcal{L} satisfying the destabilizing condition.*
- (3) *The Solomon functional is bounded from below on \mathcal{L} .*
- (4) *The Solomon functional has a minimizer in \mathcal{L} .*

Here is a glossary of the evidence presented previously.

(2) is tautological from Thomas–Yau semistability (see Remark 5.27).

(1) implies Thomas–Yau semistability; see the Floer theoretic obstructions in Theorem 3.21, Theorem 3.26, Conjecture 3.31, where we justified this for immersed Lagrangians under the automatic transversality and the positivity condition, or alternatively by assuming Joyce’s program.

(2) \iff (3). This is a consequence of (N, A) -potential clustering (see Corollary 5.9, Section 5.2.1), the uniform bound for $\mathcal{S} - \bar{\mathcal{S}}$ (see Sections 5.2.2 and 3.8.3), and formula (5.5) for the elementary functional.

(1) \implies (4). See Proposition 3.40, where we justified this under automatic transversality and the positivity condition.

(4) \implies (3). Obvious.

The rest of this section concerns (2) \implies (4), and the next concerns (4) \implies (1). The arguments will rely on several unproven statements, which we consider plausible, but may involve rather significant difficulties or substantial foundational work. Nevertheless, we think it is instructive to see heuristically how everything fits together.

Conjecture 5.22. *In the semistable case, the Solomon functional has a minimizer.*

Proof (heuristic). First, we claim that for a minimizing sequence $L^{(k)}$ of the Solomon functional, without loss of generality the Lagrangian potential $f_L^{(k)}$ is a priori bounded:

$$\sup_k \|f_L^{(k)}\|_{L^\infty} \leq C. \quad (5.7)$$

Consider the potential clustering setup. We can adjust the Lagrangian potentials on L_i by constants separately, and as long as $\sup_{L_j} f_{L_j} \leq \inf_{L_i} f_{L_i}$ for $j < i$, this process will not affect the Novikov positivity requirement, so the Lagrangian branes should remain in \mathcal{L} . We view $\sup_{L_1} f_{L_1}, \sup_{L_2} f_{L_2} - \sup_{L_1} f_{L_1}, \dots, \sup_{L_N} f_{L_N} - \sup_{L_{N-1}} f_{L_{N-1}}$ as independent constants. Adjusting all potentials by a common constant does not affect the Solomon functional, but allows us to set $\sup_{L_1} f_{L_1} = 0$. Decreasing $\sup_{L_i} f_{L_i} - \sup_{L_{i-1}} f_{L_{i-1}}$ subject to the Novikov positivity requirement will *decrease* the elementary functional (5.5), crucially because of the *semistability condition* (5.6). The part $\mathcal{S} - \bar{\mathcal{S}}$ is unchanged. Thus after this adjustment, the sequence is still minimizing for the Solomon functional. We can

thus achieve

$$\sup_{L_{i-1}} f_{L_{i-1}} = \inf_{L_i} f_{L_i}$$

for all i . By the potential clustering property, we then have (5.7).

Next we need the *compactness* from geometric measure theory. As discussed in Sections 5.1 and 5.2, under quantitative almost calibratedness there is an a priori volume bound, and the Lagrangians all remain in a fixed bounded subset of X , so Federer–Fleming compactness (see Theorem 5.2) holds automatically. The uniform potential bound (5.7) would then justify that the weak limit is an almost calibrated Lagrangian current L with bounded potential f_L (see Lemma 5.7). The continuity of the Solomon functional (see Lemma 5.8) then shows $\mathcal{S}(L) = \inf_{\mathcal{L}} \mathcal{S}$.

In Section 5.3 we presented the evidence for the conjectural L^2 -smoothing property, which would allow us to assume a uniform a priori bound on the minimizing sequence

$$\int_L |\vec{H}| \leq C.$$

so we can use Allard compactness theorem 5.3. In effect, we can assume the minimizing sequence converges subsequentially both as currents and as varifolds. By assumption the class \mathcal{L} is closed under the varifold/current topology of the Lagrangian, so the limit L lies in \mathcal{L} , whence provides a *minimizer* in \mathcal{L} . ■

Remark 5.28. If we demand \mathcal{L} is closed under the flat topology of currents, without requiring varifold convergence, then we would not need the difficult L^2 -smoothing property in the argument. However, this would allow the pathological behaviour in Example 5.4, which would increase the difficulty of Floer theory for weak regularity Lagrangians.

Remark 5.29. For the geometric measure theoretic purpose of finding special Lagrangians, the existence of a minimizer as a *Lagrangian current* L is probably sufficient. However, for applications to the Fukaya category, it is highly desirable to know that L carries a *formal brane structure* (see Section 5.4), which likely requires resolving Question 11. Some analogies suggest the question may be subtle:

In geometric invariant theory (GIT), there are niceties concerning semistable, polystable and stable objects. If we take a sequence of semistable objects in a fixed reductive group orbit, the limit may jump outside the orbit, so that the orbit does not admit a polystable representative. Several semistable orbits may be ‘ S -equivalent’, and each S -equivalence class contains a unique polystable orbit.

In the gauge theory of holomorphic bundles, likewise a sequence of connections in the same complexified gauge orbit may jump outside the orbit in the limit; algebro-geometrically, this jumping of bundle structure is usually related to bundle extensions.

One motivation for the Thomas–Yau program is to form the moduli space of (semi)-stable Lagrangian branes. The Hausdorff property of the moduli space is a delicate question.

For these reasons, as well as Remark 5.19, we are not certain if the Lagrangian minimizer should be interpreted as a representative in the chosen $D^b \text{Fuk}(X)$ class, or if several semistable $D^b \text{Fuk}(X)$ classes should be identified under some suitable S -equivalence relation. We think this question requires further developments in Floer theory. The question is also reflected in the delicacy of the infinite time limit in Joyce's Bridgeland stability proposal.

5.6. Minimizers and special Lagrangians

Conjecture 5.23. *A minimizer L of the Solomon functional inside \mathcal{L} is a special Lagrangian of phase $\hat{\theta}$.*

We will give several heuristic reasons. The essential issue is that there should be enough Lagrangian competitors within the class \mathcal{L} .

LMCF viewpoint. In Section 5.3 we discussed that the Solomon functional should be non-increasing under Joyce's LMCF. Suppose the flow extends weakly to Lagrangians in \mathcal{L} . The flow starting from a minimizer L must have constant $\mathcal{S}(L_t)$, but the evolution (5.4) would then force $\theta = \hat{\theta}$, namely L is a special Lagrangian, and the flow is in fact constant.

Hamiltonian variations. If the Lagrangian angle of the minimizer satisfies

$$-\pi/2 + \varepsilon < \inf_L \theta \leq \sup_L \theta < \pi/2 - \varepsilon,$$

then we have a more elliptic argument. Given any compactly supported global C^∞ Hamiltonian function H on X , we can associate a 1-parameter family of symplectomorphisms ϕ_t by exponentiating the Hamiltonian vector field. Since $d\phi_t$ only moves the tangent planes by $O(|t|)$ for small $|t| \ll 1$, the Lagrangian angle of $\phi_t(L)$ is still within $(-\pi/2 + \varepsilon, \pi/2 - \varepsilon)$, namely the quantitatively almost calibrated condition is preserved.

Under global Hamiltonian deformations, the first variation of the Solomon functional is

$$\delta \mathcal{S}(H) = \frac{d}{dt} \mathcal{S}(\phi_t(L))|_{t=0} = \int_L H \operatorname{Im}(e^{-i\hat{\theta}} \Omega).$$

We need another ingredient which is expected to hold once the Floer theory is sufficiently developed in the weak regularity setting:

- The class of *unobstructed* exact Lagrangian objects is preserved by Hamiltonian isotopies. As such $\phi_t(L)$ should remain inside the class \mathcal{L} .

These would imply that the minimizer L satisfies

$$\int_L H \operatorname{Im}(e^{-i\hat{\theta}} \Omega) = 0.$$

for any compactly supported C^∞ function on X . This means $\operatorname{Im}(e^{-i\hat{\theta}} \Omega) = 0$ as currents, which is equivalent to $\theta = \hat{\theta}$ under the almost calibrated setting.

Remark 5.30. The assumption that $-\pi/2 + \varepsilon < \inf_L \theta \leq \sup_L \theta < \pi/2 - \varepsilon$ for the minimizer is not innocent, but represents a principal gap in our program to find special Lagrangian currents. The problem is that if on the minimizer $\sup_L \theta_L = \frac{\pi}{2} - \varepsilon$, and a priori L has no regularity assumption (e.g., the Lagrangian angle may a priori be highly oscillatory), then we lack techniques to construct Lagrangian competitors which remain quantitatively almost calibrated.

5.7. Thomas–Yau uniqueness revisited

The Thomas–Yau uniqueness argument has a conceptually rather mysterious aspect: from local computations of Floer degrees, one arrives at the global conclusion that the two special Lagrangians share the same support. We shall now present a different argument, which is not completely rigorous, but unlike the standard arguments, it could potentially work on Lagrangians with mild singularities.

Conjecture 5.24 (Thomas–Yau uniqueness in the weak setting). *Suppose L, L' are two special Lagrangian integral currents L, L' with the same phase angle $\hat{\theta}$, equipped with suitable unobstructed brane structures, such that $L \simeq L'$ in $D^b \text{Fuk}(X)$. Then $L = L'$ as currents.*

Proof (heuristic). In general, we expect there is an $(n + 1)$ -dimensional rectifiable current \mathcal{C} with $\partial \mathcal{C} = L - L'$ constructed from universal families of holomorphic curves with boundary on L and L' . The holomorphic curves $u: \Sigma \rightarrow X$ can appear in three types:

(i) Automatically transverse holomorphic curves: there exist first order deformations v_1, \dots, v_{n-1} such that $dF = \Omega(\cdot, v_1, \dots, v_{n-1})$ does not vanish identically as a 1-form on Σ (see Section 3.3).

(ii) Non-constant holomorphic curves, which are not automatically transverse. We expect their boundary evaluation to be contained in a Hausdorff dimension $\leq n - 1$ subset of $\text{supp}(L) \cup \text{supp}(L')$ (see Section 3.3).

(iii) Constant holomorphic maps $u: \Sigma \rightarrow \text{supp}(L) \cap \text{supp}(L')$. These would only arise if L and L' have some overlapping support, so did not appear in our previous discussions. For dimensional reasons, these cannot contribute to the $(n + 1)$ -dimensional current \mathcal{C} .

The key difference from the second case is that at interior points of $\text{supp}(L) \cap \text{supp}(L')$, there are n linearly independent first order deformations, such that v_1, \dots, v_n span TL upon boundary evaluation. This behaviour can only be compatible with $dF = 0$ for constant curves.

We now impose the special Lagrangian condition, and consider the automatically transverse case. Along $\partial \Sigma$, the counterclockwise directional derivative of F has argument equal to the constant Lagrangian angle $\hat{\theta}$ modulo $\pi \mathbb{Z}$. As such we expect $F(\partial \Sigma)$ to be contained in a line segment with incline angle $\hat{\theta}$. By the maximum principle on the holomorphic function F , the entire $F(\Sigma) \subset \mathbb{C}$ is contained in a line segment. However,

the open mapping theorem in complex analysis then implies F is constant, which *rules out the automatically transverse curves*.

Now the only contributions to \mathcal{C} would come from the non-constant, not automatically transverse curves. This forces $\text{supp}(\partial\mathcal{C}) \cap (\text{supp}(L) \cup \text{supp}(L'))$ to be contained in a Hausdorff $(n-1)$ -dimensional subset. However $\partial\mathcal{C} = L - L'$ as integral currents, so the n -dimensional current $L - L'$ has support dimension $\leq n-1$, which forces it to vanish. This shows $L = L'$. ■

Question 15. When can we say furthermore that the formal brane structures on $L = L'$ are related by some gauge equivalence?

Remark 5.31. In the subsequent work [53], the author proved a *weak-strong uniqueness theorem*, which extends the Thomas–Yau uniqueness to certain special Lagrangians which arise as limits of smooth and unobstructed Lagrangian branes in the varifold topology.

5.7.1. Special Lagrangians are minimizers. We now revisit Proposition 3.40. Our goal is to suggest that the automatic transversality, positivity condition, and even smoothness assumptions can be removed in Proposition 3.40, at the cost of assuming the entire force of the Thomas–Yau conjecture, under the setting of this section.

Conjecture 5.25. *If there exists a special Lagrangian L in the class \mathcal{L} , then it is a minimizer of the Solomon functional.*

Proof (heuristic). We begin by noting that the existence of a special Lagrangian representative should imply Thomas–Yau semistability (see Conjecture 3.31). By the Thomas–Yau existence Conjecture 5.21 this implies the Solomon functional has a minimizer L' , which must be a special Lagrangian. Then the Thomas–Yau uniqueness Conjecture 5.24 implies $L = L'$ as currents. ■

5.8. Comparison with Joyce’s LMCF program

We have already made extensive comparisons between the variational approach and the LMCF program of Joyce, but it may help to summarize a few highlights.

(i) Joyce’s program is much more ambitious in that it tackles the entire derived Fukaya category, not just the almost calibrated Lagrangians. We feel the quantitative almost calibratedness is so pervasively used in the variational approach that it cannot be removed. Dropping the almost calibratedness will give rise to significantly more difficulties in Joyce’s program: the collapsing of zero objects can then happen, and the Solomon functional no longer needs to decrease. Neves’s example of finite time singularity [62] is a concrete manifestation of the difficulty. The almost calibrated condition is also natural from the viewpoint of the continuity method (see Section 4.2), which deals with special Lagrangians inside varying ambient almost Calabi–Yau structures.

(ii) Joyce does not specify the Bridgeland stability in a priori Floer theoretic terms. An a priori guess on the nature of the stability condition is central to the variational method.

Even though our picture is largely conjectural, it seems to be the most precise description hitherto of how stability condition comes into the existence questions of special Lagrangians.

(iii) Joyce primarily focuses on compact Calabi–Yaus, and mentions the exact case only as an easier analogue. We have focused on the exact case, although we feel some parts of our picture may extend to compact Calabi–Yaus, if one is prepared to overcome (even more) significant Floer theoretic technical hurdles. However, we do not know what would replace the a priori estimates on the Lagrangian potentials, and notably the potential clustering condition.

(iv) Joyce’s LMCF involves objects with a priori higher regularity, although its infinite time convergence behaviour may well require understanding weak regularity Lagrangians. The variational method requires working with varifold/current like objects throughout.

(v) Joyce’s LMCF needs to make essential use of genericity conditions. This in particular requires extremely precise classification of all possible generic singularities in order to perform surgeries, a task that becomes overwhelmingly difficult for complex dimension ≥ 3 . Our variational program is less sensitive to such arguments. On the other hand, we still potentially need to understand some generic singularities, so that the class \mathcal{L} contains enough competitors, to enable the proof of the L^p -smoothing property for some $p \geq 1$, and Conjecture 5.23.

(vi) Although time and again we appealed to Joyce’s LMCF to heuristically justify certain claims, it is only because we lack other ways of constructing Lagrangian competitors with sufficient control, and the basic logical framework of the variational approach is independent of the LMCF. It seems desirable (on account of the extraordinary difficulty of Joyce’s program) to keep this logical independence manifest in the program to rigourize our variational proposal.

(vii) Joyce’s program has a number of highly non-trivial categorical predictions discussed in Section 3.6, such as the idempotent closedness of $D^b \text{Fuk}(X)$. Even if these predictions turn out to be false, it would not affect the validity of the variational method.

A. The Fukaya category

This appendix is a brief reminder about Floer theory in the exact setting. There exist both excellent surveys on the Fukaya category of embedded Lagrangians, such as Auroux [9] and Smith [75], and many in depth treatments such as Seidel [71], Akaho–Joyce [4] and Fukaya–Oh–Ohta–Ono [35]. Our very limited goal is to recall some key notions prevalent in the main text, and explain some basic intuitions, but we will not get into the more technical aspects, such as the details of perturbation schemes, which are treated carefully in these standard references.

For the Thomas–Yau–Joyce program, one needs to incorporate immersed Lagrangians. The canonical reference is Akaho–Joyce [4] for a treatment using virtual techniques,

and Woodward et al. [67, 68] which avoids virtual counting by using stabilizing divisors. The exact assumption affords some technical simplifications, for which a sketchy account is found in [45, Section 4.1]. Another technical treatment in the exact setting, not allowing certain teardrop curves, is in Alston–Bao [7].

A.1. Fukaya category for embedded exact Lagrangians

Floer cohomology and A_∞ -structure with mod 2 coefficients. Let (X, ω, J) be a Stein manifold, namely a Kähler manifold with $\omega = \sqrt{-1}\partial\bar{\partial}\phi$ for a plurisubharmonic exhaustion function ϕ . In particular, X is an exact symplectic manifold, meaning $\omega = d\lambda$, where λ is the Liouville 1-form. All almost complex structure perturbations are assumed to agree with the fixed complex structure outside some compact set.

Given two transversely intersecting exact embedded⁶² Lagrangians L, L' with potential $f_L, f_{L'}$, namely $df_L = \lambda|_L$ and $df_{L'} = \lambda|_{L'}$, and some extra brane data, one can associate an algebraic invariant called the *Floer cohomology*. A general feature of Floer theory, is that the constructions depend on many auxiliary choices, but the invariants depend on only a small number of data, and should always be invariant under global Hamiltonian isotopies.

We assume $c_1(TX) = 0$ and let Ω be a complex volume form on TX . We shall always assume the Lagrangians to be *graded*, namely the phase function $\theta = \arg \Omega|_L: L \rightarrow S^1$ lifts to a real valued function. The grading is part of the brane data. Working first with \mathbb{Z}_2 coefficients, the Floer cohomology can be defined as the cohomology of a complex $(CF^*(L, L'), d)$. Here $CF^*(L, L')$ is generated by the transverse intersection points $p \in L \cap L'$, whose *degrees* are given by

$$\mu_{L, L'}(p) = \frac{1}{\pi} \left(\sum \phi_i + \theta_L(p) - \theta_{L'}(p) \right), \quad (\text{A.1})$$

where we put the tangent planes of L, L' inside $T_p X \simeq \mathbb{C}^n$ into the standard form

$$T_p L = \mathbb{R}^n \subset \mathbb{C}^n, \quad T_p L' = (e^{i\phi_1}, \dots, e^{i\phi_n})\mathbb{R}^n \subset \mathbb{C}^n, \quad 0 < \phi_i < \pi.$$

Notice if we reverse the roles of L, L' , then we can regard $p \in CF^*(L', L)$, but this affects the degree by $\mu_{L, L'}(p) = n - \mu_{L', L}(p)$. For alternative formulations of the degree in terms of Lagrangian Grassmannians, see [70].

Remark A.1. The degree convention $\mu_{L, L'}$ here follows Joyce [45], which corresponds to $\mu_{L', L}$ in [9, 71, 75]. The advantage of this convention is its compatibility with the central charge formula $Z(L) = \int_L \Omega$ in the Bridgeland stability. If instead one uses the convention of [9, 71, 75], then adding π to the Lagrangian phase would correspond to the shift $[-1]$ in $D^b \text{Fuk}(X)$, so the central charge would be $Z(L) = \int_L \Omega$.

⁶²In our terminology, embedded Lagrangians are always connected, while immersed Lagrangians can have disconnected domains.

Remark A.2. For almost calibrated Lagrangians $-\pi < \theta_L - \theta_{L'} < \pi$, whence $-1 < \mu_{L,L'}(p) < n + 1$. Since the degrees are always integers, we must have $0 \leq \mu_{L,L'}(p) \leq n$.

We consider the moduli space $\mathcal{M}(p, q; J, [u])$ of finite energy holomorphic strips with ends at p, q and boundary on L, L' , in the homotopy class $[u] \in \pi_2(X, L \cup L')$:

$$u: \Sigma = \mathbb{R} \times [0, 1] \rightarrow X, \quad u(s, 0) \in L, \quad u(s, 1) \in L', \quad \lim_{s \rightarrow -\infty} u = p, \quad \lim_{s \rightarrow +\infty} u = q,$$

$$\partial_s u + J(t, u) \partial_t u = 0, \quad E(u) = \int u^* \omega = \iint |\partial u / \partial s|^2 ds dt < \infty.$$

Index theory of the Cauchy–Riemann operator with Lagrangian boundary conditions implies this moduli space has virtual dimension $\deg q - \deg p$. The holomorphic strip equation is invariant under domain translation in the \mathbb{R} direction. Using generic domain dependent almost complex structures which are fixed outside a large compact set, one can achieve suitable *transversality* on the moduli spaces, and in particular $\mathcal{M}(p, q; J, [u])/\mathbb{R}$ are isolated points for $\deg q - \deg p = 1$. A key advantage of the exact setting is that the energy can be computed a priori by the topological formula

$$\int_{\Sigma} u^* \omega = \int_{\partial \Sigma} \lambda = \int_{-\infty}^{\infty} df_L - \int_{-\infty}^{\infty} df_{L'} = (f_L - f_{L'})(q) - (f_L - f_{L'})(p). \quad (\text{A.2})$$

By *Gromov compactness*, the number of isolated points is finite, and only finitely many homotopy classes $[u]$ admit holomorphic strips. To save some notations, we sometimes write

$$\mathcal{M}(p, q) = \bigcup_{[u] \in \pi_2(M, L \cup L')} \mathcal{M}(p, q; J, [u]).$$

Remark A.3. The role of convexity assumptions at the infinity of X (such as the existence of a plurisubharmonic exhaustion function) is to ensure that for a finite given collection of Lagrangians, all holomorphic curves remain inside a fixed bounded region. This is needed to apply Gromov compactness.

Remark A.4. More generally, one can add a Hamiltonian term in the Cauchy–Riemann equation, and replace transverse intersection points by Hamiltonian chords. The Cauchy–Riemann equation then gets modified to the *Floer equation*

$$\partial_s u + J(t, u)(\partial_t u - X_H) = 0, \quad (\text{A.3})$$

where X_H is a Hamiltonian vector field. This perturbation is not needed for Floer theoretic transversality statements if L and L' are already transverse, but is an essential ingredient in showing the Hamiltonian invariance of Floer cohomology.

Remark A.5. We generally distinguish between the $s \rightarrow -\infty$ end, and the $s \rightarrow +\infty$ end. The main difference is the ordering of the Lagrangians at the intersection point. Here we are following the Joyce convention [45], which is opposite to Auroux [9]. This is dictated by compatibility with the degree formula (A.1). Similarly, later the A_{∞} product also

requires the Lagrangian boundaries to be ordered clockwise, as opposed to the counter-clockwise convention in Auroux [9].

The Floer differential $d: CF^k(L, L') \rightarrow CF^{k+1}(L, L')$ is $dp = \sum n_{p,q}q$, where $n_{p,q}$ is the mod 2 count of $\mathcal{M}(p, q)/\mathbb{R}$. The key fact of Floer theory is that $d^2 = 0$. For this, one considers the holomorphic strips between p, r with $\deg r - \deg p = 2$, modulo the translation invariance \mathbb{R} direction. This moduli space $\mathcal{M}(p, r)/\mathbb{R}$ is 1-dimensional. Generally in Floer theory, the boundary of the compactified moduli spaces comes from *disc breaking* and *disc and sphere bubbling*. The latter is ruled out for energy reasons by the exactness assumption, and the former gives

$$\partial(\overline{\mathcal{M}(p, r)/\mathbb{R}}) = \bigcup_q \mathcal{M}(p, q)/\mathbb{R} \times \mathcal{M}(q, r)/\mathbb{R}.$$

In terms of mod 2 counts, $\sum_q n_{p,q}n_{q,r} = 0$, namely $d^2 = 0$. This fact allows one to take the cohomology, which is $HF^*(L, L')$. Although suppressed in this notation, the homotopy classes of discs are additive under disc breaking. This fact allows one to introduce some extra weighting factors involving energy and holonomy of local systems.

The general strategy to show the Floer cohomology is independent of the choices of almost complex structures and Hamiltonian perturbations, is to consider *continuity equations*, whose counts define chain maps at the level of CF^* , so descend to comparison maps between Floer cohomologies defined by different auxiliary data (see Auroux [8, Section 1.5]).

Floer cohomology admits rich algebraic structures, but the deeper structure is better set up at the chain level $CF^*(L, L')$. We temporarily avoid the issue of signs and self Floer cohomology. The Fukaya category can be seen as the generalization of Floer cohomology in two directions:

(i) We allow the interplay of many (transverse) Lagrangians. Each Lagrangian is labelled by an object in the Fukaya category. This labelling is the main difference between an algebra and a category.

(ii) The holomorphic strips are replaced by *holomorphic polygons*, with boundary segments mapped to a clockwise ordered sequence of at least three Lagrangians L_0, L_1, \dots, L_k , and clockwise ordered boundary marked points x_0, x_1, \dots, x_k mapped to the Lagrangian intersection points $q \in L_0 \cap L_k, p_1 \in L_0 \cap L_1, \dots, p_k \in L_{k-1} \cap L_k$. We distinguish x_0 as the output, and regard x_1, \dots, x_k as inputs. The marked points x_0, x_1, x_2 on the boundary of the domain disc are fixed, while the other $k - 2$ marked points are allowed to move freely preserving their cyclic ordering.

The moduli spaces of such polygons (with suitably domain dependent perturbations) are denoted as $\mathcal{M}(p_1, \dots, p_k, q)$. Moduli spaces with at least three marked points do not have the domain translation invariance, so there is no need to divide by \mathbb{R} .

Remark A.6. From the viewpoint of gluing theory, it is convenient to regard the boundary marked points of the holomorphic polygons as punctures, where the Riemann surface

structure is locally modelled on strip like ends. The moduli of abstract holomorphic polygons with $k + 1$ marked points has a compactification known as the Stasheff associahedron $\bar{\mathcal{R}}_{k+1}$. On account of the geometric picture of polygons in \mathbb{R}^2 , we often refer to the strip like ends as corners.

The energy formula (A.2) generalizes to the holomorphic polygon case:

$$\int_{\Sigma} u^* \omega = (f_{L_0} - f_{L_k})(q) - \sum_1^k (f_{L_{i-1}} - f_{L_i})(p_i). \quad (\text{A.4})$$

The virtual dimension formula is

$$\text{vdim } \mathcal{M}(p_1, \dots, p_k, q) = \deg q - \sum_1^k \deg p_i + k - 2.$$

Here $\deg q - \sum_1^k \deg p_i$ comes from the index theory of the Cauchy–Riemann operator, and $k - 2$ comes from the freedom to move the marked points on the boundary. Under suitable domain dependent perturbation schemes, in this exact setting one can ensure transversality, so that the moduli space is smooth. For setting up the Fukaya category, the zero-dimensional moduli spaces are particularly important, since counting points give rise to operations, and 1-dimensional moduli spaces are important for producing A_∞ -relations. In the main text, we have also given considerable attention to $(n - 1)$ -dimensional moduli spaces, since these are relevant for producing bordism currents.

Within the exact setting, disc and sphere bubbling is impossible. After compactification, the moduli space of holomorphic polygons can have two kinds of boundaries, due to two kinds of disc breaking:

Disc breaking at the corners. The disc may break at $L_i \cap L_{i+1}$. The polygons near the breaking limit are obtained from gluing polygons with corners mapped to $p_1, \dots, p_{l-1}, r, p_{l+1}, \dots, q$, and strips with boundary on L_{l-1}, L_l and two ends mapped to p_l, r . (Of course, disc breaking can also happen at the outgoing corner q .)

Disc splitting at the edges. When there are at least 4 Lagrangians, the domain disc can split into two discs with $l + 1 \geq 3$ and $k + 1 - l \geq 3$ marked points. The edges of one disc map to $L_0, \dots, L_j, L_{j+l}, \dots, L_k$, with cyclically marked points mapping to $p_1, \dots, p_j, r \in L_j \cap L_{l+j}, p_{l+j+1}, \dots, p_k$ and q . The edges of the other disc map to L_j, \dots, L_{l+j} , with marked points mapping to p_{j+1}, \dots, p_{l+j} and $r \in L_j \cap L_{l+j}$.

When disc breaking and disc splitting are taken into account, the moduli spaces can be compactified into $\bar{\mathcal{M}}(p_1, \dots, p_k, q)$. We then have

$$\begin{aligned} \partial \bar{\mathcal{M}}(p_1, \dots, p_k, q) &= \bigcup \bar{\mathcal{M}}(p_l, r)/\mathbb{R} \times \bar{\mathcal{M}}(p_1, \dots, p_{l-1}, r, p_{l+1}, \dots, p_k, q) \\ &\cup \bigcup \bar{\mathcal{M}}(p_1, \dots, p_j, r, p_{l+j+1}, \dots, p_k, q) \times \bar{\mathcal{M}}(p_{j+1}, \dots, p_{l+j}, r). \end{aligned} \quad (\text{A.5})$$

In particular, the (virtual) dimensions of both sides are equal, which constrains $\deg r$.

Remark A.7. More generally, disc breaking and disc splitting can happen in a bubble tree fashion. Such multiple splitting/breaking do not concern us, because under sufficient transversality conditions, they occur only in codimension at least two in the moduli space. To set up Fukaya categories in the exact setting, only zero and 1-dimensional moduli spaces are needed, so the multiple bubble trees do not occur. When we make use of higher-dimensional moduli spaces in the main text, the bubble trees do occur, but the codimension two condition means the deeper boundary strata do not contribute to the boundary of the bordism current, in the sense of currents.

The A_∞ -structure is the algebraization of the disc breaking/splitting phenomenon. It consists of multilinear maps

$$m_k: CF^*(L_{k-1}, L_k) \otimes \cdots \otimes CF^*(L_1, L_2) \otimes CF^*(L_0, L_1) \rightarrow CF^*(L_0, L_k)[2-k]$$

satisfying the A_∞ -relation

$$\sum_{l=1}^k \sum_{j=0}^{k-l} \pm m_{k+1-l}(p_k, \dots, p_{j+l+1}, m_l(p_{j+l}, \dots, p_{j+1}), p_j, \dots, p_1) = 0.$$

Here m_1 is the degree one Floer differential d , and for $k \geq 2$ the operation m_k is defined by counting holomorphic polygons in moduli spaces of virtual dimension zero:

$$m_k(p_k, \dots, p_1) = \sum_q \# \mathcal{M}(p_1, \dots, p_k, q)q.$$

Virtual dimension zero requires $\deg q = \sum_1^k \deg p_i + 2 - k$, which explains the degree of m_k . The A_∞ -relation is the direct translation of (A.5), with disc breaking at corners contributing the m_1 terms, and disc splitting contributing the other terms.

The first few A_∞ -relations have clear geometric meanings:

- The Floer differential squares to zero.
- The Floer product m_2 satisfies the Leibniz rule. As such m_2 descends to a product structure on the mod 2 coefficient Floer cohomology

$$HF^*(L_1, L_2) \otimes HF^*(L_0, L_1) \rightarrow HF^*(L_0, L_2).$$

- The m_2 is associative up to a homotopy given by the m_3 terms. In particular, the Floer product is associative on cohomology.

The higher structures naturally lead to the Fukaya category of embedded exact Lagrangians. This requires some discussion on signs, brane structures, and self Floer cohomologies.

Self Floer cohomology. It is desirable to take Floer cohomology of L with itself. One major feature of $HF^*(L, L)$ is that it contains *units*, at least at cohomological level.

One challenge to implement *self Floer cohomology* is that L is not transverse to itself, so the Cauchy–Riemann equation needs perturbation. There are many frameworks

to address this problem, and one idea dating back to Floer is to use the Hamiltonian invariance of Floer cohomology, to think of self Floer cohomology via

$$HF^*(L, L) \simeq HF^*(L, \phi_{\varepsilon H}(L)),$$

where $\phi_{\varepsilon H}$ is the time one flow of the small generic Hamiltonian εH [9, Section 1.6]. For $\varepsilon \ll 1$, the Lagrangian $\phi_{\varepsilon H}(L)$ can be identified as a graph over L inside T^*L , the transverse intersection $L \cap \phi_{\varepsilon H}(L)$ are the critical points of $H|_L$, and a suitable setup of the Floer trajectories (A.3) can be identified as Morse flowlines of H . Thus $HF^*(L, L)$ is isomorphic to the Morse cohomology of L , so $HF^*(L, L) \simeq H^*(L)$. In the exact case, the ring structure on $HF^*(L, L)$ defined from perturbed holomorphic triangles agrees with the cup product ring structure on $H^*(L)$. The unit can be represented by the Morse generator of $H^0(L)$, or more non-perturbatively via the Pionikhin–Salamon–Schwarz map.

It takes some effort to promote the self Floer cohomology to the Fukaya category framework, and ensure the consistency in the perturbation schemes (see Auroux [9, Section 2.1] for a sketch and Seidel [71] for details). In applications it is often more convenient to avoid Hamiltonian perturbations as much as possible.

Example A.1 (Floer products involving the identity). We wish to heuristically explain a special case relevant to Joyce–Imagi–Santos (see Section 2.3), concerning the geometric interpretation of the Floer product mod 2

$$HF^0(L', L) \otimes HF^0(L, L') \rightarrow HF^0(L, L).$$

Here L, L' are assumed to be transverse. Hamiltonian invariance means we can alternatively think of

$$HF^0(L', \phi_{\varepsilon H}(L)) \otimes HF^0(L, L') \rightarrow HF^0(L, \phi_{\varepsilon H}(L)).$$

This is defined by the count of holomorphic triangles with input corners at $CF^0(L, L')$, $CF^0(L', \phi_{\varepsilon H}(L))$, and an output corner at $CF^0(L, \phi_{\varepsilon H}(L))$. We may assume that the Morse function $H|_L$ has only one maximum point r on L , which represents the unit of $HF^*(L, L)$. When $\varepsilon \rightarrow 0$, then L and $\phi_{\varepsilon H}(L)$ coincide, and the holomorphic triangles become holomorphic strips with ends at $CF^0(L, L')$, $CF^0(L', L)$ (alternatively seen as a degree n output) and passing through the point $r \in L$. This last incidence condition is independent of the position of r on L , since we can choose H to have its maximum at any generic prescribed point. Notice in this strip interpretation, there is no longer any Hamiltonian perturbation. This interpretation featured in Lemma 2.7.

Sign issues and brane structures. To go beyond mod 2 coefficients, we need to orient moduli spaces. Good references can be found in Seidel’s book [71] and Abouzaid [2, Appendix]. All Lagrangians are assumed to be graded, with second Stiefel–Whitney class equal to the restriction of a fixed class in $H^2(X, \mathbb{Z}/2)$, and we equip the Lagrangians with *relative spin structures*. At any transverse Lagrangian intersection point $p \in L_+ \cap L_-$, there is a unique up to homotopy path Λ_p of Lagrangian planes in $T_p X \simeq \mathbb{C}^n$ with graded

lift interpolating TL_+ and TL_- . We fix a relative spin structure on Λ_p , compatible with the relative spin structure on L_\pm . We can associate a vector space o_p as the determinant line of the Cauchy–Riemann operator D_p on the upper half plane with boundary data Λ_p . The dual of o_p is denoted o_p^\vee , namely $o_p \otimes o_p^\vee \simeq \mathbb{R}$ canonically. The orientation line $|o_y|$ is the free abelian group generated by the two possible orientations of o_y with the relation that their sum vanishes. Furthermore, we equip the Lagrangians L with (rank one) *local systems* E , and write the Floer cochain complex as the graded vector space

$$CF^*(L, L') = \bigoplus_{p \in L \cap L'} \text{Hom}(E|_p, E'|_p) \otimes |o_p|.$$

Remark A.8. There are some variants on the coefficient ring/field of the local system. The simplest case is the trivial local system, in which case we simply delete the Hom factor. Other popular choices have parallel transport in \mathbb{Q}^* , \mathbb{R}^* , \mathbb{C}^* , or the units in the Novikov ring.⁶³ Different choices could in principle lead to slightly different versions of the derived Fukaya category. The smaller the coefficient ring/field, the more stringent is the notion of derived isomorphism of objects. For the purpose of extending the Solomon functional (see Section 3.4) to be real valued, we require all coefficients to be at least contained in \mathbb{R} , so we will usually work simultaneously with \mathbb{R} , \mathbb{Q} and \mathbb{Z} local systems. On the other hand, it is claimed in [67, Remark 4.5] that in the exact setting the immersed Fukaya algebras can be defined *over the integers*. The specific advantage of working over integers, as discussed in the main text, is primarily that the bordism current \mathcal{C} between Lagrangians is then an *integral current*, rather than \mathbb{R} -linear combinations of integral currents.

Given a holomorphic polygon $u: \Sigma \rightarrow X$, with inputs x_1, \dots, x_k and output x_0 mapping to p_1, \dots, p_k and q , the det line of the linearized Cauchy–Riemann operator D_u can be computed from gluing kernel and cokernels:

$$\det(D_u \# D_{p_k} \dots \# D_{p_1}) \simeq \det(D_u) \otimes o_{p_k} \otimes \dots \otimes o_{p_1}.$$

The role of the *relative spin structure*, is to specify a homotopically unique choice of isotopy between the glued operator $D_u \# D_{x_k} \dots \# D_{x_1}$ and D_{x_0} (i.e., an *isotopy between Lagrangian boundary conditions*), hence a preferred isomorphism

$$\det D_u \simeq o_q \otimes o_{p_1}^\vee \otimes \dots \otimes o_{p_k}^\vee.$$

The tangent space of the moduli space of holomorphic polygons $\mathcal{M}(p_1, \dots, p_k, q)$ involves not only the linearized Cauchy–Riemann operator, but also the variation of the complex structure of the domain of the polygon, controlled by the Stasheff associahedron \mathcal{R}_{k+1} . Denote $\lambda^{\text{top}}(V)$ as the top wedge product of a vector space V . Then there are preferred isomorphisms depending on the relative spin structure choice

$$\lambda^{\text{top}}(T\mathcal{M}(p_1, \dots, p_k, q)) \simeq \lambda^{\text{top}}(\mathcal{R}_{k+1}) \otimes o_q \otimes o_{p_1}^\vee \otimes \dots \otimes o_{p_k}^\vee. \quad (\text{A.6})$$

⁶³The $U(1)$ -local systems are popular in the physics literature, but appear rarely in Floer theory.

Remark A.9. Fix an orientation on L , then $CF^0(L, L')$ is naturally dual to $CF^n(L', L)$. The local system factor $\text{Hom}(E, E')$ is naturally dual to $\text{Hom}(E', E)$. Given a Lagrangian path Λ_p associated to a Lagrangian intersection p , the reverse path is also associated with a determinant line bundle, which can be identified with $o_p^\vee \otimes \lambda(TL)$, since the two half planes with Lagrangian boundaries can be glued to a disc, such that the det line of the Cauchy–Riemann operator is canonically isomorphic to $\lambda(TL)$.

For $k \geq 2$, when the moduli spaces are zero-dimensional, so carry canonical orientations, then a universal orientation choice for $\lambda^{\text{top}}(\mathcal{R}_{k+1})$ determines an operator

$$|c_u|: |o_{p_k}| \otimes \cdots \otimes |o_{p_1}| \rightarrow |o_q|.$$

In our degree conventions the corners $x_0, x_1, x_2, \dots, x_k$ on the domain disc boundary are ordered clockwise, so a natural orientation on the Stasheff associahedron can be obtained by fixing x_0, x_1, x_2 and allowing the other corner points to move in the clockwise orientation. The parallel transports along the local systems contribute another factor

$$\text{Hom}(E_{k-1}, E_k)|_{p_k} \otimes \cdots \otimes \text{Hom}(E_1, E_2)|_{p_2} \otimes \text{Hom}(E_0, E_1)|_{p_1} \rightarrow \text{Hom}(E_0, E_k)|_q.$$

Each pseudoholomorphic polygon contributes to the operation

$$m_k: CF^*(L_{k-1}, L_k) \otimes \cdots \otimes CF^*(L_0, L_1) \rightarrow CF^*(L_0, L_k)[2-k]$$

via the tensor product of the orientation factor $|c_u|$ and the local system factor, multiplied by another sign factor depending only on the degrees (see [71, equation (12.24)])

$$(-1)^{\deg p_1 + 2 \deg p_2 + \cdots + k \deg p_k}.$$

In the case of holomorphic strips, we have a natural isomorphism

$$T\mathcal{M}(p, q) = \mathbb{R}(-\partial_s) \oplus T(\mathcal{M}(p, q)/\mathbb{R}), \quad (\text{A.7})$$

where $-\partial_s$ denotes the translation vector field pointing towards the input point. When $\mathcal{M}(p, q)/\mathbb{R}$ consists of isolated points, it carries canonical orientations, whence by (A.6) we obtain

$$|c_u|: |o_p| \rightarrow |o_q|.$$

The local system parallel transport produces another factor

$$\text{Hom}(E_0, E_1)|_p \rightarrow \text{Hom}(E_0, E_1)|_q.$$

Each pseudoholomorphic strip contributes to the Floer differential

$$d: CF^*(L_0, L_1) \rightarrow CF^{*+1}(L_0, L_1)$$

by the product of these two factors. We write

$$m_1: CF^*(L_0, L_1) \rightarrow CF^{*+1}(L_0, L_1), \quad m_1 = (-1)^{\deg p_1} d.$$

When the signs and local system weighting factors are taken into account, the A_∞ -relation reads

$$\sum_{l=1}^k \sum_{j=0}^{k-l} (-1)^\dagger m_{k+1-l}(p_k, \dots, p_{j+l+1}, m_l(p_{j+l}, \dots, p_{j+1}), p_j, \dots, p_1) = 0,$$

where $\dagger = j + \deg p_1 + \dots + \deg p_j$. The *Fukaya category* for the compact embedded Lagrangians comprises the following data:

- The objects are embedded Lagrangians (with additional brane data, such as grading, Lagrangian potential, orientation, relative spin structure, and local system).
- The morphisms $\text{Hom}^*(L, L')$ are the vector spaces $CF^*(L, L')$ (where L can coincide with L').
- The A_∞ -composition maps are the multilinear maps m_k satisfying the A_∞ relations.

The Fukaya category is an example of an A_∞ -category.

In particular, the Floer differential squares to zero, so we can define the Floer cohomology groups $HF^*(L, L')$ for embedded exact Lagrangian branes. The Floer product on cohomology is given by

$$[p_2] \circ [p_1] = (-1)^{\deg p_1} m_2(p_2, p_1),$$

which is associative.

Example A.2. (Floer products involving the identity, continued) In the context of Example A.1, the orientation isomorphism (A.6) for the holomorphic triangle is determined by whether the isotopy of the Lagrangian boundary conditions respects the relative spin structure. Since the relative spin structure on $\phi_{\varepsilon H}(L)$ is induced from L , this problem is equivalent to the corresponding isotopy problem for the limiting holomorphic strip. The holonomy factor of the local systems for the holomorphic triangle, is also reduced to that of the limiting strip.

In the simplest case when we are given closed elements $\alpha \in CF^0(L, L')$ and $\beta \in CF^0(L', L)$ each involving only one intersection point, the local systems are trivial, and only one holomorphic curve contributes to the Floer product, then

$$\beta \circ \alpha = 1_L \in HF^0(L, L)$$

means that for the holomorphic strip from α to β passing through a generically chosen point $r \in L$, the Lagrangian boundary condition on the disc obtained by gluing TL, TL' and the two Lagrangian paths at the two strip like ends, can be contracted to constant, respecting the prescribed relative spin structures on L, L' and the two ends. More generally, many intersections points and holomorphic strips may contribute to the Floer product, and $\beta \circ \alpha = 1_L \in HF^0(L, L)$ means a weighted signed count of holomorphic strips is equal to one.

Under sufficient transversality assumptions, we can form the $(n - 1)$ -dimensional moduli space of holomorphic strips from α to β , and thereby produce an $(n + 1)$ -dimensional universal family \mathcal{C} , as in the main text in Section 3.1. Using the relative spin structures on L, L' and the Lagrangian paths associated with the ends, we use (A.6), (A.7) and Remark A.9 to induce a canonical orientation on the moduli space from

$$\beta \otimes \alpha \in CF^0(L', L) \otimes CF^0(L, L').$$

Using the complex orientation on the holomorphic curve Σ , and inserting an extra minus sign, we obtain an orientation on \mathcal{C} . This tricky minus sign accounts for the difference between the counterclockwise orientation of $\partial\Sigma$ compatible with the complex orientation, and the clockwise orientation of $\partial\Sigma$ compatible on the L -boundary with the translation vector field $-\partial_s$. Putting everything together, $\beta \circ \alpha = 1_L \in HF^0(L, L)$ means in the sense of weighted counts, that $\partial\mathcal{C}$ passes once through a generic point $r \in L$ in the same orientation as $\lambda(TL)$. In other words, the L -boundary evaluation of $\partial\mathcal{C}$ *sweeps out the oriented cycle L* .

The same argument says that if

$$\alpha \circ \beta = 1_{L'} \in HF^0(L', L'),$$

then the moduli space of holomorphic strips from β to α produces a universal family \mathcal{C}' , whose L' -boundary evaluation map sweeps out the oriented cycle L' . The subtle point is that due to the reversal of the \mathbb{R} -translation vector fields, \mathcal{C}' has the reverse orientation as \mathcal{C} . Therefore, the L' -boundary evaluation of $\partial\mathcal{C}$ sweeps out the oriented cycle $-L'$ instead of L' . Here ends the example.

Twisted complexes, distinguished triangles, derived category. A fundamental problem of the embedded Fukaya category is that it lacks enough geometric objects. Morally, Fukaya category is a construction that inputs the symplectic geometry of Lagrangian branes, and outputs the representation theory of an A_∞ -category. Now the general feature of A_∞ -module categories is that one can take cones and idempotent summands, two properties which are useful for classifying such categories, and desirable for mirror symmetry. The problem is that cones and idempotent summands are not obviously represented by embedded Lagrangian objects under the Yoneda embedding. The common solution is to sideline this issue by the formal algebraic construction of twisted complexes and idempotent completions. This is not quite adequate for the Thomas–Yau conjecture. However, we will discuss how the introduction of immersed Lagrangian objects geometrizes the twisted complexes (see Section A.2). The geometric meaning of idempotents is an open problem.

The formal algebraic constructions are well explained in [9, Section 3] and [75, Section 4], to which we refer the reader for more details. Given objects L_1, \dots, L_N of the Fukaya category \mathcal{A} , a *twisted complex* (L, b_L) consists of

- the formal shifted direct sum $L = \bigoplus_1^N L_i[k_i]$ with $k_i \in \mathbb{Z}$ formally keeping track of degrees (the geometric meaning of the shift [1] is to add a constant π to the Lagrangian

phase, which reverses the orientation of the Lagrangian, with a corresponding twist to the spin structure);

- a strictly triangular differential $b_L \in \text{End}(L)$, i.e., a collection of maps

$$b_{ij} \in \text{Hom}^{k_j - k_i + 1}(L_i, L_j)$$

for $i > j$,⁶⁴

satisfying the equation

$$\sum_{k \geq 1} m_k(b_L, \dots, b_L) = 0,$$

that is,

$$\sum_{k \geq 1} \sum_{i=i_0 > \dots > i_k=j} m_k(b_{i_{k-1}i_k}, \dots, b_{i_0i_1}) = 0.$$

Notice the strict triangularity implies the sum is finite. One can define morphisms between these twisted complexes, and assign A_∞ -structures to make twisted complexes into an A_∞ -category $Tw\mathcal{A}$, into which \mathcal{A} naturally embeds fully faithfully. Using the A_∞ -structure, it makes sense to talk about closed morphisms and cohomologies, similar to the construction of Floer cohomology.

Given twisted complexes $A = (L, b_L)$, $B = (L', b_{L'}) \in Tw\mathcal{A}$, and a closed morphism $f \in \text{Hom}^0((L, b_L), (L', b_{L'}))$, the abstract *mapping cone* of f is the twisted complex

$$\text{Cone}(f) = \left(L[1] \oplus L', \begin{pmatrix} b_L & 0 \\ f & b_{L'} \end{pmatrix} \right).$$

Generally, a mapping cone of f is an object of $Tw\mathcal{A}$ quasi-isomorphic to $\text{Cone}(f)$. This gives rise to a *distinguished triangle*

$$A \rightarrow B \rightarrow \text{Cone}(f) \xrightarrow{[1]} A.$$

This illustrates the advantage of introducing twisted complexes: $Tw\mathcal{A}$ is a *triangulated category*.

The cohomological category of $Tw\mathcal{A}$ is commonly denoted $D^b \text{Fuk}(X)$. This has the same objects as $Tw\mathcal{A}$, but the Floer cochain spaces are replaced by their H^0 , namely we remember the Floer cohomology.

Under the Yoneda embedding, $Tw\mathcal{A}$ embeds into its module category. The idempotent closure $Tw^\pi\mathcal{A}$ is obtained by formally adding the direct summands of the Yoneda image of twisted complexes in $Tw\mathcal{A}$. The cohomological category of $Tw^\pi\mathcal{A}$ is commonly denoted $D^\pi \text{Fuk}(X)$. In the variant setting of compact X , it is usually $D^\pi \text{Fuk}(X)$ instead of $D^b \text{Fuk}(X)$ that shows up in mirror symmetry, since the derived category of coherent sheaves is automatically idempotent closed.

⁶⁴In most symplectic references such as [9] the morphisms b_{ij} go in the opposite direction $i < j$. This just amounts to reversing the ordering of L_1, \dots, L_N . We find our reversed convention a little more convenient for the Harder–Narasimhan decomposition.

Remark A.10. Once immersed Lagrangians are admitted as objects of Fukaya categories, the twisted complexes are largely redundant. Joyce claims that by including immersed and singular Lagrangians with rank one local systems, then $D^b \text{Fuk}(X)$ is automatically idempotent closed, so there is no difference between $D^b \text{Fuk}(X)$ and $D^\pi \text{Fuk}(X)$ (see [45, Conjecture 3.6]). However, it is highly non-obvious why direct summands are Yoneda represented by geometric Lagrangian objects,⁶⁵ so this claim is regarded by many experts as a weakness of Joyce’s proposal. For this reason, in our more restrictive proposal we stick with the more geometric $D^b \text{Fuk}(X)$ (including immersed and singular objects, but not formal idempotent summands) in favour of $D^\pi \text{Fuk}(X)$, and the idempotent closure problem does not falsify our program.

A.2. Immersed exact Lagrangians

According to Joyce’s LMCF program, immersed Lagrangians are a necessary part of any Fukaya category adequate for the Thomas–Yau conjecture. As far as the author is aware, only immersed Floer cohomology [4], rather than the full categorical framework, has been written down in the literature, although in the exact setting this is commonly believed to be a relatively routine matter, as sketched in [45, Section 4.1]. Our limited goal is to highlight the main difference with the embedded case, namely the issues of *obstructions and bounding cochains*. Once these two issues are taken into account, what works in the embedded case will also work in the immersed case.

Teardrop curves and obstructions. The assumptions on (X, ω) are as in the previous section. Immersed Lagrangians are immersions $\iota: L \rightarrow X$ with $\omega|_L = 0$, and all self intersections are transverse. The domain of L is allowed to be disconnected, so the union of finitely many transversely intersecting embedded Lagrangians are examples of immersed Lagrangians. Each *self intersection* point of two local sheets L_+, L_- corresponds to two different points x_\pm on the domain of L . It is important to distinguish x_+ and x_- , because for the boundary of the holomorphic curve to pass through x_+ in the clockwise direction means crossing from L_+ to L_- , and x_- signifies the opposite crossing.

We say L is exact, if there is a function f_L on the domain of L , such that df_L agrees with the Liouville 1-form restricted to L . For energy reasons, this forbids non-trivial holomorphic discs with boundary on L which never change local sheets at any boundary point. The caveat is that the relative homology class $[\omega] \in H_2(X, L)$ may still be non-zero. The brane structures on L are as in the embedded case. The construction of $CF^*(L, L)$ depends on the approach, but a common feature is that it includes

$$CF_{\text{self}}^*(L, L) = \bigoplus_{\text{self intersection } p} CF^*(L_+, L_-) \oplus CF^*(L_-, L_+)$$

⁶⁵There exist some wild speculations, such as incorporating coisotropic branes into the Fukaya category in order to have more geometric objects.

generated by the local system factor $\text{Hom}(E_+, E_-)|_p$ (resp. $\text{Hom}(E_-, E_+)|_p$) tensored with the orientation line.

The Gromov compactness discussion is largely similar to the embedded case. A new phenomenon is the *teardrop curves*, namely the holomorphic curves with boundary on L and a single output corner at a self intersection point $r \in CF^*(L_+, L_-)$. Of particular importance is the case with

$$\mu_{L_+, L_-}(r) = 2.$$

The number 2 is intuitively explained by the two degrees of freedom of the domain Möbius transforms fixing the corner point $\text{Aut}(D^2, 1)$, modulo which such teardrop curves occur in dimension zero moduli spaces.

Now if we attempt to run the usual argument for $d^2 = 0$ in Floer cohomology, we would consider the moduli space of holomorphic strips between p, q with

$$\deg q - \deg p = 2,$$

modulo the translation \mathbb{R} . However, in addition to the usual strip breaking, the holomorphic strips can also break into a holomorphic triangle with inputs p, r and output q , and a teardrop curve with corner at r . In summary, teardrop curves with corner at a degree 2 intersection point *obstruct* Floer cohomology.

The automorphism group $\text{Aut}(D^2, 1)$ forbids the naïve domain dependent perturbation schemes, which in turn causes transversality problems. In the literature there are two approaches to solve this problem: Joyce and Akaho [4] use virtual perturbation techniques for bordered Riemann surfaces, while Woodward et al. [67, 68] circumvent the virtual perturbations by utilizing stabilizing divisors. Both approaches assign curved A_∞ algebra structures (m_0, m_1, \dots) to the Floer cochain spaces $CF^*(L, L)$ of immersed Lagrangians. In the exact setting, the $m_0 \in CF_{\text{self}}^2(L, L)$ term amounts to a count of teardrop curves with corner at degree 2 self intersection points, with weighting factors coming from the holonomy of the local system. Since in the main text the emphasis is on the automatic transversality assumption, we shall not dwell on the details of perturbation schemes, but only identify a few simplifications in the exact setting.

Remark A.11. The rough idea of Woodward et al. is to introduce interior marked points, constrained to lie on a Donaldson divisor D disjoint from the Lagrangians. The virtual dimension is not affected by these divisor constraints, since each interior marked point increases it by 2, while each divisor constraint decreases it by 2. One needs to arrange D to be of sufficiently high degree, so that each non-trivial pseudoholomorphic disc with boundary on the Lagrangians has at least one intersection with D . On a teardrop curve, imposing the divisor constraint at interior marked points kills the domain automorphisms $\text{Aut}(D^2, 1)$, so one can then introduce domain dependent perturbation of almost complex structures compatible with D to achieve sufficient transversality to make sense of counts. The appealing feature of this approach, is that adding marked points does not alter the geometric interpretation of the holomorphic curves, so stays closer to geometry than the virtual approach.

The framework of Woodward et al. [67, 68] is not restricted to exact settings, and works also for compact symplectic manifolds with rational $[\omega] \in H^2(X)$. Producing the Donaldson divisor with the intersection properties is easier if $[\omega] \in H_2(X, L)$ is a rational class, although the methods in [14, Section 3.1] allows one to largely relax this assumption.

In exact manifolds, as mentioned in [67, Remark 4.5], one can avoid the spherical components of the treed discs. In the exact Lagrangian setting, the only bubbling happens at the self intersection points. These afford significant simplifications to the construction, and allows one to think of the treed discs in [14, 15, 67, 68] in terms of a tree of holomorphic polygons connected at the self intersection points. By avoiding the troublesome sphere bubbles, one can also relax the restriction of moduli spaces of dimension at most one.

Remark A.12. A very technical aspect of Akaho–Joyce [4] is that the A_∞ structure is not constructed directly, but through a sequence of approximations involving energy cutoff scales. In the exact setting, the topological energy formula implies a priori energy bounds, so this complication would not arise.

Cancellation of obstructions. To make sense of Floer cohomology one needs to cancel the obstructions by introducing *bounding cochains* $b \in CF_{\text{self}}^1(L, L)$, which represents a formal sum of $b_p \in \text{Hom}(E_+, E_-)|_p \otimes |o_p|$ associated to degree one intersection points $p \in CF^1(L_+, L_-)$. We require

- the *Novikov positivity* condition $f_{L_+}(p) \geq f_{L_-}(p)$ for each of the intersection points appearing in b ;
- the *Maurer–Cartan equation*

$$m_0^b = m_0 + m_1(b) + m_2(b, b) + \cdots = 0 \in CF_{\text{self}}^2(L, L).$$

Geometrically, the coefficients of $q \in CF_{\text{self}}^2(L, L)$ in the $m_k(b, \dots, b)$ term represent the zero-dimensional counts of holomorphic polygons with the inputs at the summands b_p of b , and the output at q , weighted by the holonomy and orientation factors. Using the Novikov positivity requirement of the bounding cochain, the topological energy formula (A.4) for the polygon then implies

$$\int_{\Sigma} \omega \leq f_{L_-}(q) - f_{L_+}(q),$$

where the boundary of Σ passes from L_+ to L_- at q in the clockwise direction. By Gromov compactness, this uniform energy bound implies there are only *finitely many* terms involved in the Maurer–Cartan equation. When such a bounding cochain b exists, we say (L, b) defines an unobstructed Lagrangian brane. In this case, both the Akaho–Joyce and the Woodward–Palmer approaches assign *self Floer cohomology groups*

$$HF^*((L, b), (L, b)),$$

defined as the cohomology of a degree one operator

$$m_1^b: CF^*(L, L) \rightarrow CF^{*+1}(L, L), \quad m_1^b(x) = \sum_{k, k' \geq 0} m_{k+k'+1}(\underbrace{b, \dots, b}_{k'}, x, \underbrace{b, \dots, b}_k).$$

This cohomology is invariant under global Hamiltonian deformations. Two bounding cochains b, b' on L are said to be *gauge equivalent*, if there is $h \in CF^0(L, L)$ satisfying the Novikov positivity condition, such that

$$b - b' = \sum_{k, k' \geq 0} m_{k+k'+1}(b', \dots, b', h, b, \dots, b).$$

Gauge equivalent bounding cochains give rise to isomorphic Floer cohomology.

Remark A.13. In the embedded case, there are no self intersections, so the Maurer–Cartan equation is vacuous, and the Lagrangian is automatically unobstructed, with zero bounding cochain. The unobstructed condition is not automatic in general for immersed Lagrangians, and a significant aspect of the Joyce program in [45] is that unobstructed Lagrangians ought to be better behaved in the LMCF.

Now suppose (L, b) and (L', b') are two unobstructed Lagrangian branes, intersecting transversally avoiding the self intersections of L and L' . Then we can define the Floer cohomology $HF^*((L, b), (L', b'))$. The Floer cochain space $CF^*(L, L')$ is the same as in the embedded case, generated by the local system factor tensored with the orientation factor, associated to the transverse intersection points. The Floer differential is

$$m_1^{b, b'}(p) = \sum_{k, k' \geq 0} m_{k+k'+1}(b', \dots, b', p, b, \dots, b),$$

where the sum has k' insertions of b' , and k insertions of b . The coefficient of $q \in CF^{*+1}(L, L')$ are morally defined by the weighted count of holomorphic polygons with boundary marked points mapping to the summands of $b, \dots, p, b', \dots, q$, arranged in clockwise order. A similar a priori energy bound argument shows the sum is finite.

It is instructive to see why $(m_1^{b, b'})^2 = 0$. We consider the breaking of 1-dimensional moduli spaces, associated with $p, r \in CF^*(L, L')$ with $\deg r - \deg p = 2$. There are several mechanisms for disc bubbling and disc splittings:

- The polygon breaks into two parts, connected at a nodal point mapping to some $q \in CF^*(L, L')$ with $\deg q - \deg p = 1$. The sum of all such contributions give rise to $\langle m_1^{b, b'}(p), q \rangle \langle m_1^{b, b'}(q), r \rangle$, and summing over q, r produces $(m_1^{b, b'})^2(p)$.
- The polygon bubbles off a teardrop curve at a self intersection point q of degree 2 on either L or L' .
- The polygon splits into two parts, connected at a node mapping to a degree 2 self intersection point q on either L or L' .

The combined effect of the last two contributions, is a sum of the weighted counts of polygons with boundary mapping to $b, b, \dots, p, b', \dots, q, b', \dots, r$ multiplied by the coefficient of q in

$$m_0^{b'} = m_0 + m_1(b') + \dots + \in CF_{\text{self}}^2(L', L')$$

in the case of $q \in CF_{\text{self}}^2(L', L')$ (the case with $q \in CF_{\text{self}}^2(L, L)$ gives an entirely similar contribution related to $m_0^b \in CF_{\text{self}}^2(L, L)$). By the unobstructed assumption,

$$m_0^b = 0 \quad \text{and} \quad m_0^{b'} = 0,$$

so these contributions vanish. But the grand sum of all contributions from all boundaries of the moduli spaces should be zero, which implies $(m_1^{b,b'})^2 = 0$.

The generalization to many Lagrangians is a matter of bookkeeping. We have the A_∞ compositions

$$m_k^{b_0, \dots, b_k}: CF^*(L_{k-1}, L_k) \otimes \cdots \otimes CF^*(L_0, L_1) \rightarrow CF^*(L_0, L_k)[2-k]$$

and

$$\begin{aligned} m_k^{b_0, \dots, b_k}(p_k, \dots, p_1) \\ = \sum m_l(b_k, \dots, b_k, p_k, b_{k-1}, \dots, p_{k-1}, \dots, p_1, b_0, \dots, b_0). \end{aligned} \quad (\text{A.8})$$

In particular, this induces a product structure on Floer cohomology

$$HF^*(L_1, L_2) \otimes HF^*(L_0, L_1) \rightarrow HF^*(L_0, L_2)$$

(with bounding cochains suppressed in the notation),

$$[\beta] \circ [\alpha] = (-1)^{\deg \alpha} m_2^{b_0, b_1, b_2}(\beta, \alpha).$$

We say two unobstructed Lagrangian branes L, L' are *isomorphic* in $D^b \text{Fuk}(X)$, if there exist $[\alpha] \in HF^0(L, L')$ and $[\beta] \in HF^0(L', L)$, such that their compositions are the cohomological units: $[\beta] \circ [\alpha] = 1_L \in HF^0(L, L)$ and $[\alpha] \circ [\beta] = 1_{L'} \in HF^0(L', L')$.

The union of several components. In our convention an immersed Lagrangian can have several components. Of particular interest is the case where L is the union of transverse immersed Lagrangians L_1, \dots, L_N with bounding cochains b_1, \dots, b_N respectively, and we have morphisms $b_{ij} \in CF^1((L_i, b_i), (L_j, b_j))$ for $i > j$. The key assumption here is that the morphisms only go in one direction from L_i to L_j , not vice versa. We assume that

$$b = \sum b_i + \sum_{i>j} b_{ij}$$

is a bounding cochain for the immersed Lagrangian L , and in particular all intersection points in b_{ij} satisfy the Novikov positivity condition $f_{L_i} \geq f_{L_j}$. We can write out the Maurer–Cartan equation

$$m_0 + m_1(b) + m_2(b, b) + \cdots = 0$$

in component form: for any $i > j$,

$$\sum_l \sum_{k \leq l} \sum_{i=i_0 > \cdots > i_k=j} m_l(b_{i_k}, \dots, b_{i_k}, b_{i_{k-1}i_k}, \dots, b_{i_1}, \dots, b_{i_1}, b_{i_0i_1}, b_{i_0}, \dots, b_{i_0}) = 0.$$

The key observation is that this is precisely how one would define *twisted complexes* built on L_1, \dots, L_N , in the presence of the bounding cochains b_1, \dots, b_N and the data b_{ij} , when no further degree shifts are involved (see the exact setting in Section A.1). In this sense, we say that ‘immersed Lagrangians geometrize twisted complexes’. In other words, if the unobstructed immersed Lagrangians are admitted into the Fukaya category, then there is no need to formally add twisted complexes.

Lemma A.3 (Blocking together connected components based on potential clustering). *Assume L is the finite union of transversely intersecting immersed Lagrangians, with a bounding cochain b . Then L can be decomposed as a twisted complex built from some L_1, \dots, L_N , such that $\inf_{L_i} f_{L_i} > \sup_{L_j} f_{L_j}$ whenever $i > j$, and the Lagrangian potential f_{L_i} has connected range for each L_i .*

Proof. The decomposition can continue as long as there exists a real number c , such that the Lagrangian components can be partitioned into two types, with Lagrangian potential strictly smaller than c (resp. greater than c). As long as $\inf_{L_i} f_{L_i} > \sup_{L_j} f_{L_j}$ whenever $i > j$, the Novikov positivity condition on the Lagrangian intersection points would imply that the entries $b_{ij} \in CF^1(L_i, L_j)$ of b can only go in the direction $i > j$ and not vice versa, so the immersed Lagrangian L is necessarily of the twisted complex form. This algorithm stops in finitely many steps since there are only finitely many components involved. ■

Orientation signs on bordism currents. In Sections 3.1 and 3.1.2 we encountered the $(n-1)$ -dimensional moduli spaces such as $\mathcal{M}(b, \dots, b, \alpha, b', \dots, \beta)$ and $\mathcal{M}(b, \dots, b, \gamma)$. The special case of holomorphic strips was already mentioned in Example A.2.

We now consider the moduli $\mathcal{M}(p_1, \dots, p_k)$ of polygons with at least 3 corners p_1, \dots, p_k , all regarded as inputs, arranged in clockwise order on $\partial\Sigma$, each carrying the local system factors $\text{Hom}(E_+, E_-)|_{p_i}$ and the orientation factors $|o_{p_i}|$. The clockwise composition of the local system hom factors and the parallel transport along $\partial\Sigma$, produces a holonomy factor around $\partial\Sigma$, which is a number in $\mathbb{Q}, \mathbb{R}, \mathbb{Z}$ depending on the coefficient ring choice. Using (A.6) and Remark A.9, as well as the clockwise orientation convention on the Stasheff associahedron, we acquire a (naïve) orientation on $T\mathcal{M}(p_1, \dots, p_k)$. To assign *orientation* and *weighting factors* to $\mathcal{M}(p_1, \dots, p_k)$, we take the product of the holonomy factor, the naïve orientation on $T\mathcal{M}(p_1, \dots, p_k)$, and another universal sign factor

$$(-1)^{\deg p_1 + 2 \deg p_2 + \dots + k \deg p_k} (-1)^{\deg p_k}.$$

The appearance of this universal sign adjustment is a familiar convention in the open-closed map, see [1, equation (5.24)]. The notation \mathcal{M} is a shorthand for the weighted sum of all the $(n-1)$ -dimensional moduli spaces involved in the construction of the bordism current.

We equip the domain Σ with the complex orientation, and together with an *extra minus sign*, the orientation on \mathcal{M} induces the orientation on \mathcal{C} . This minus sign arises

for the same reason as in Example A.2, namely the discrepancy between our clockwise convention on $\partial\Sigma$, with the standard complex orientation on Σ .

Acknowledgements. The author is currently an MIT CLE Moore Instructor and a Clay Research Fellow. This paper owes much intellectual debt to the proposal of Thomas–Yau and Joyce. The author thanks MIT for providing a stimulating research environment, and especially P. Seidel, S. Rezhikov and T. Collins for useful discussions. He also thanks the referees for helpful suggestions.

References

- [1] M. Abouzaid, [A geometric criterion for generating the Fukaya category](#). *Publ. Math. Inst. Hautes Études Sci.* (2010), no. 112, 191–240 Zbl [1215.53078](#) MR [2737980](#)
- [2] M. Abouzaid, [On the wrapped Fukaya category and based loops](#). *J. Symplectic Geom.* **10** (2012), no. 1, 27–79 Zbl [1298.53092](#) MR [2904032](#)
- [3] M. Abouzaid and I. Smith, [Exact Lagrangians in plumbings](#). *Geom. Funct. Anal.* **22** (2012), no. 4, 785–831 Zbl [1266.53073](#) MR [2984118](#)
- [4] M. Akaho and D. Joyce, [Immersed Lagrangian Floer theory](#). *J. Differential Geom.* **86** (2010), no. 3, 381–500 Zbl [1226.53085](#) MR [2785840](#)
- [5] W. K. Allard, [On the first variation of a varifold](#). *Ann. of Math. (2)* **95** (1972), 417–491 Zbl [0252.49028](#) MR [0307015](#)
- [6] F. J. Almgren, Jr., *Almgren’s big regularity paper*. World Sci. Monogr. Ser. Math. 1, World Scientific, River Edge, NJ, 2000 Zbl [0985.49001](#) MR [1777737](#)
- [7] G. Alston and E. Bao, [Exact, graded, immersed Lagrangians and Floer theory](#). *J. Symplectic Geom.* **16** (2018), no. 2, 357–438 Zbl [1398.53091](#) MR [3834425](#)
- [8] D. Auroux, Mirror symmetry and T -duality in the complement of an anticanonical divisor. *J. Gökova Geom. Topol. GGT* **1** (2007), 51–91 Zbl [1181.53076](#) MR [2386535](#)
- [9] D. Auroux, [A beginner’s introduction to Fukaya categories](#). In *Contact and symplectic topology*, pp. 85–136, Bolyai Soc. Math. Stud. 26, János Bolyai Mathematical Society, Budapest, 2014 Zbl [1325.53001](#) MR [3220941](#)
- [10] P. Biran, O. Cornea, and E. Shelukhin, [Lagrangian shadows and triangulated categories](#). Astérisque 426, Société Mathématique de France, Paris, 2021 Zbl [1491.53091](#) MR [4340068](#)
- [11] S. Boucksom, [Variational and non-archimedean aspects of the Yau–Tian–Donaldson conjecture](#). In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. II. Invited lectures*, pp. 591–617, World Scientific, Hackensack, NJ, 2018 Zbl [1447.32030](#) MR [3966781](#)
- [12] K. A. Brakke, *The motion of a surface by its mean curvature*. Math. Notes 20, Princeton University Press, Princeton, NJ, 1978 Zbl [0386.53047](#) MR [0485012](#)
- [13] T. Bridgeland, [Stability conditions on triangulated categories](#). *Ann. of Math. (2)* **166** (2007), no. 2, 317–345 Zbl [1137.18008](#) MR [2373143](#)
- [14] F. Charest and C. Woodward, [Floer trajectories and stabilizing divisors](#). *J. Fixed Point Theory Appl.* **19** (2017), no. 2, 1165–1236 Zbl [1376.53103](#) MR [3659006](#)
- [15] F. Charest and C. T. Woodward, Floer theory and flips. [v1] 2015, [v11] 2022, arXiv:[1508.01573v11](#)

- [16] G. Chen, [The J-equation and the supercritical deformed Hermitian–Yang–Mills equation](#). *Invent. Math.* **225** (2021), no. 2, 529–602 Zbl [1481.53093](#) MR [4285141](#)
- [17] X. Chen and J. Cheng, [On the constant scalar curvature Kähler metrics \(II\)—Existence results](#). *J. Amer. Math. Soc.* **34** (2021), no. 4, 937–1009 Zbl [1477.14067](#) MR [4301558](#)
- [18] T. C. Collins, A. Jacob, and S.-T. Yau, [\(1, 1\) forms with specified Lagrangian phase: a priori estimates and algebraic obstructions](#). *Camb. J. Math.* **8** (2020), no. 2, 407–452 Zbl [1442.14124](#) MR [4091029](#)
- [19] T. C. Collins and Y. Shi, [Stability and the deformed Hermitian–Yang–Mills equation](#). In *Surveys in differential geometry 2019. Differential geometry, Calabi–Yau theory, and general relativity. Part 2*, pp. 1–38, Surv. Differ. Geom. 24, International Press, Boston, MA, 2022 Zbl [1539.53097](#) MR [4479717](#)
- [20] T. C. Collins, D. Xie, and S.-T. Yau, [The deformed Hermitian–Yang–Mills equation in geometry and physics](#). In *Geometry and physics. Vol. I*, pp. 69–90, Oxford University Press, Oxford, 2018 Zbl [1421.35300](#) MR [3932257](#)
- [21] T. C. Collins and S.-T. Yau, [Moment maps, nonlinear PDE, and stability in mirror symmetry](#). 2018, arXiv:[1811.04824](#)
- [22] T. C. Collins and S.-T. Yau, [Moment maps, nonlinear PDE and stability in mirror symmetry, I: geodesics](#). *Ann. PDE* **7** (2021), no. 1, article no. 11 Zbl [1469.58007](#) MR [4242135](#)
- [23] C. De Lellis, [Almgren’s \$Q\$ -valued functions revisited](#). In *Proceedings of the International Congress of Mathematicians. Volume III*, pp. 1910–1933, Hindustan Book Agency, New Delhi, 2010 Zbl [1226.49041](#) MR [2827872](#)
- [24] C. De Lellis, [Almgren’s center manifold in a simple setting](#). In *Harmonic analysis and applications*, pp. 257–288, IAS/Park City Math. Ser. 27, American Mathematical Society, Providence, RI, 2020 Zbl [1468.49050](#) MR [4249630](#)
- [25] R. Dervan, J. B. McCarthy, and L. M. Sektnan, [\$Z\$ -critical connections and Bridgeland stability conditions](#). *Camb. J. Math.* **12** (2024), no. 2, 253–355 Zbl [07945662](#) MR [4779674](#)
- [26] S. K. Donaldson, [Anti self-dual Yang–Mills connections over complex algebraic surfaces and stable vector bundles](#). *Proc. London Math. Soc. (3)* **50** (1985), no. 1, 1–26 Zbl [0529.53018](#) MR [0765366](#)
- [27] S. K. Donaldson, [Infinite determinants, stable bundles and curvature](#). *Duke Math. J.* **54** (1987), no. 1, 231–247 Zbl [0627.53052](#) MR [0885784](#)
- [28] S. K. Donaldson, [Moment maps and diffeomorphisms](#). In *Surveys in differential geometry*, pp. 107–127, Surv. Differ. Geom. 7, International Press, Somerville, MA, 2000 Zbl [1106.53037](#) MR [1919424](#)
- [29] S. K. Donaldson, [Lower bounds on the Calabi functional](#). *J. Differential Geom.* **70** (2005), no. 3, 453–472 Zbl [1149.53042](#) MR [2192937](#)
- [30] S. K. Donaldson, [Deformations of multivalued harmonic functions](#). *Q. J. Math.* **72** (2021), no. 1-2, 199–235 Zbl [1492.53084](#) MR [4271385](#)
- [31] M. R. Douglas, B. Fiol, and C. Römersberger, [Stability and BPS branes](#). *J. High Energy Phys.* (2005), no. 9, article no. 006 MR [2171358](#)
- [32] H. Federer and W. H. Fleming, [Normal and integral currents](#). *Ann. of Math. (2)* **72** (1960), 458–520 Zbl [0187.31301](#) MR [0123260](#)
- [33] K. Fukaya, [Counting pseudo-holomorphic discs in Calabi–Yau 3-folds](#). *Tohoku Math. J. (2)* **63** (2011), no. 4, 697–727 Zbl [1244.57046](#) MR [2872962](#)
- [34] K. Fukaya, [Gromov–Hausdorff distance between filtered \$A_\infty\$ categories 1: Lagrangian Floer theory](#). 2021, arXiv:[2106.06378v1](#)

- [35] K. Fukaya, Y.-G. Oh, H. Ohta, and K. Ono, *Lagrangian intersection Floer theory: anomaly and obstruction. Part I*. AMS/IP Stud. Adv. Math. 46, American Mathematical Society, Providence, RI; International Press, Somerville, MA, 2009 Zbl 1181.53002 MR 2553465
- [36] M. Haskins, *The geometric complexity of special Lagrangian T^2 -cones*. *Invent. Math.* **157** (2004), no. 1, 11–70 Zbl 1064.53032 MR 2135184
- [37] S. He, *The branched deformations of the special Lagrangian submanifolds*. *Geom. Funct. Anal.* **33** (2023), no. 5, 1266–1321 Zbl 1527.53066 MR 4646409
- [38] H.-J. Hein and S. Sun, *Calabi–Yau manifolds with isolated conical singularities*. *Publ. Math. Inst. Hautes Études Sci.* **126** (2017), 73–130 Zbl 1397.32009 MR 3735865
- [39] K. Hori, S. Katz, A. Klemm, R. Pandharipande, R. Thomas, C. Vafa, R. Vakil, and E. Zaslow, *Mirror symmetry*. Clay Math. Monogr. 1, American Mathematical Society, Providence, RI; Clay Mathematics Institute, Cambridge, MA, 2003 Zbl 1044.14018 MR 2003030
- [40] Y. Imagi, *Generalized Thomas–Yau uniqueness theorems*. In *Birational geometry, Kähler–Einstein metrics and degenerations*, pp. 323–342, Springer Proc. Math. Stat. 409, Springer, Cham, 2023 Zbl 1546.53082 MR 4606643
- [41] Y. Imagi, D. Joyce, and J. Oliveira dos Santos, *Uniqueness results for special Lagrangians and Lagrangian mean curvature flow expanders in \mathbb{C}^m* . *Duke Math. J.* **165** (2016), no. 5, 847–933 Zbl 1341.53115 MR 3482334
- [42] D. Joyce, *On counting special Lagrangian homology 3-spheres*. In *Topology and geometry: commemorating SISTAG*, pp. 125–151, Contemp. Math. 314, American Mathematical Society, Providence, RI, 2002 Zbl 1060.53059 MR 1941627
- [43] D. Joyce, *Special Lagrangian submanifolds with isolated conical singularities. V. Survey and applications*. *J. Differential Geom.* **63** (2003), no. 2, 279–347 Zbl 1076.53064 MR 2015549
- [44] D. Joyce, *Riemannian holonomy groups and calibrated geometry*. Oxf. Grad. Texts Math. 12, Oxford University Press, Oxford, 2007 Zbl 1200.53003 MR 2292510
- [45] D. Joyce, *Conjectures on Bridgeland stability for Fukaya categories of Calabi–Yau manifolds, special Lagrangians, and Lagrangian mean curvature flow*. *EMS Surv. Math. Sci.* **2** (2015), no. 1, 1–62 Zbl 1347.53052 MR 3354954
- [46] D. Joyce, Y.-I. Lee, and M.-P. Tsui, *Self-similar solutions and translating solitons for Lagrangian mean curvature flow*. *J. Differential Geom.* **84** (2010), no. 1, 127–161 Zbl 1206.53071 MR 2629511
- [47] M. Kontsevich and Y. Soibelman, *Stability structures, motivic Donaldson–Thomas invariants and cluster transformations*. 2008, arXiv:0811.2435v1
- [48] M. Kontsevich and Y. Soibelman, *Motivic Donaldson–Thomas invariants: Summary of results*. In *Mirror symmetry and tropical geometry*, pp. 55–89, Contemp. Math. 527, American Mathematical Society, Providence, RI, 2010 Zbl 1214.14014 MR 2681792
- [49] M. Kontsevich and Y. Soibelman, *Cohomological Hall algebra, exponential Hodge structures and motivic Donaldson–Thomas invariants*. *Commun. Number Theory Phys.* **5** (2011), no. 2, 231–352 Zbl 1248.14060 MR 2851153
- [50] M. Kontsevich and Y. Soibelman, *Wall-crossing structures in Donaldson–Thomas invariants, integrable systems and mirror symmetry*. In *Homological mirror symmetry and tropical geometry*, pp. 197–308, Lect. Notes Unione Mat. Ital. 15, Springer, Cham, 2014 Zbl 1326.14042 MR 3330788
- [51] D. Kwon and Y.-G. Oh, *Structure of the image of (pseudo)-holomorphic discs with totally real boundary condition*. *Comm. Anal. Geom.* **8** (2000), no. 1, 31–82 Zbl 0951.32025 MR 1730896

- [52] G. Lawlor, [The angle criterion](#). *Invent. Math.* **95** (1989), no. 2, 437–446 Zbl [0662.49018](#) MR [0974911](#)
- [53] Y. Li, Quantitative Thomas–Yau uniqueness. 2022, [arXiv:2207.08047v1](#)
- [54] Y. Li, [Metric SYZ conjecture and non-Archimedean geometry](#). *Duke Math. J.* **172** (2023), no. 17, 3227–3255 Zbl [1532.32019](#) MR [4688155](#)
- [55] Y.-S. Lin, [Correspondence theorem between holomorphic discs and tropical discs on K3 surfaces](#). *J. Differential Geom.* **117** (2021), no. 1, 41–92 Zbl [1468.53072](#) MR [4195752](#)
- [56] J. D. Lotay and A. Neves, [Uniqueness of Lagrangian self-expanders](#). *Geom. Topol.* **17** (2013), no. 5, 2689–2729 Zbl [1302.53088](#) MR [3190297](#)
- [57] J. D. Lotay and T. Pacini, [Complexified diffeomorphism groups, totally real submanifolds and Kähler–Einstein geometry](#). *Trans. Amer. Math. Soc.* **371** (2019), no. 4, 2665–2701 Zbl [1420.53038](#) MR [3896093](#)
- [58] J. D. Lotay and T. Pacini, [From Lagrangian to totally real geometry: Coupled flows and calibrations](#). *Comm. Anal. Geom.* **28** (2020), no. 3, 607–675 Zbl [1509.53108](#) MR [4124139](#)
- [59] D. Matessi, [Lagrangian pairs of pants](#). *Int. Math. Res. Not. IMRN* (2021), no. 15, 11306–11356 Zbl [1496.53086](#) MR [4294119](#)
- [60] R. C. McLean, [Deformations of calibrated submanifolds](#). *Comm. Anal. Geom.* **6** (1998), no. 4, 705–747 Zbl [0929.53027](#) MR [1664890](#)
- [61] F. Morgan, [Geometric measure theory](#). Fifth edn., Elsevier/Academic Press, Amsterdam, 2016 Zbl [1338.49089](#) MR [3497381](#)
- [62] A. Neves, [Singularities of Lagrangian mean curvature flow: Zero-Maslov class case](#). *Invent. Math.* **168** (2007), no. 3, 449–484 Zbl [1119.53052](#) MR [2299559](#)
- [63] A. Neves, Recent progress on singularities of Lagrangian mean curvature flow. In *Surveys in geometric analysis and relativity*, pp. 413–438, Adv. Lect. Math. (ALM) 20, International Press, Somerville, MA, 2011 Zbl [1268.53073](#) MR [2906935](#)
- [64] A. Neves, [Finite time singularities for Lagrangian mean curvature flow](#). *Ann. of Math. (2)* **177** (2013), no. 3, 1029–1076 Zbl [1315.53074](#) MR [3034293](#)
- [65] Y.-G. Oh, [Fredholm theory of holomorphic discs under the perturbation of boundary conditions](#). *Math. Z.* **222** (1996), no. 3, 505–520 Zbl [0863.53024](#) MR [1400206](#)
- [66] Y.-G. Oh and S. Müller, [The group of Hamiltonian homeomorphisms and \$C^0\$ -symplectic topology](#). *J. Symplectic Geom.* **5** (2007), no. 2, 167–219 Zbl [1144.37033](#) MR [2377251](#)
- [67] J. Palmer and C. Woodward, [Invariance of immersed Floer cohomology under Maslov flows](#). *Algebr. Geom. Topol.* **21** (2021), no. 5, 2313–2410 Zbl [1511.53081](#) MR [4334514](#)
- [68] J. Palmer and C. Woodward, Invariance of immersed Floer cohomology under Lagrangian surgery. [v1] 2019, [v10] 2023, [arXiv:1903.01943v10](#)
- [69] R. Schoen and J. Wolfson, [Minimizing area among Lagrangian surfaces: The mapping problem](#). *J. Differential Geom.* **58** (2001), no. 1, 1–86 Zbl [1052.53056](#) MR [1895348](#)
- [70] P. Seidel, [Graded Lagrangian submanifolds](#). *Bull. Soc. Math. France* **128** (2000), no. 1, 103–149 Zbl [0992.53059](#) MR [1765826](#)
- [71] P. Seidel, [Fukaya categories and Picard–Lefschetz theory](#). Zur. Lect. Adv. Math., European Mathematical Society, Zürich, 2008 Zbl [1159.53001](#) MR [2441780](#)
- [72] P. Seidel, Fukaya A_∞ -structures associated to Lefschetz fibrations. V. [v1] 2020, [v2] 2021, [arXiv:2010.13692v2](#)
- [73] L. Simon, [Lectures on geometric measure theory](#). Proc. Cent. Math. Anal. Aust. Natl. Univ. 3, Australian National University, Centre for Mathematical Analysis, Canberra, 1983 Zbl [0546.49019](#) MR [0756417](#)

- [74] C. T. Simpson, [Constructing variations of Hodge structure using Yang–Mills theory and applications to uniformization](#). *J. Amer. Math. Soc.* **1** (1988), no. 4, 867–918 Zbl [0669.58008](#) MR [0944577](#)
- [75] I. Smith, [A symplectic prolegomenon](#). *Bull. Amer. Math. Soc. (N.S.)* **52** (2015), no. 3, 415–464 Zbl [1317.53003](#) MR [3348443](#)
- [76] K. Smoczyk, [Mean curvature flow in higher codimension: Introduction and survey](#). In *Global differential geometry*, pp. 231–274, Springer Proc. Math. 17, Springer, Heidelberg, 2012 Zbl [1247.53004](#) MR [3289845](#)
- [77] J. P. Solomon, [The Calabi homomorphism, Lagrangian paths and special Lagrangians](#). *Math. Ann.* **357** (2013), no. 4, 1389–1424 Zbl [1282.53067](#) MR [3124936](#)
- [78] J. P. Solomon, [Curvature of the space of positive Lagrangians](#). *Geom. Funct. Anal.* **24** (2014), no. 2, 670–689 Zbl [1296.53157](#) MR [3192038](#)
- [79] J. P. Solomon and A. M. Yuval, [Geodesics of positive Lagrangians from special Lagrangians with boundary](#). 2020, arXiv:2006.06058v2
- [80] L. Spolaor, [Almgren’s type regularity for semicalibrated currents](#). *Adv. Math.* **350** (2019), 747–815 Zbl [1440.49048](#) MR [3948685](#)
- [81] A. Strominger, S.-T. Yau, and E. Zaslow, [Mirror symmetry is \$T\$ -duality](#). *Nuclear Phys. B* **479** (1996), no. 1-2, 243–259 Zbl [0896.14024](#) MR [1429831](#)
- [82] R. P. Thomas, [Moment maps, monodromy and mirror manifolds](#). In *Symplectic geometry and mirror symmetry (Seoul, 2000)*, pp. 467–498, World Sci. Publ., River Edge, NJ, 2001 Zbl [1076.14525](#) MR [1882337](#)
- [83] R. P. Thomas and S.-T. Yau, [Special Lagrangians, stable bundles and mean curvature flow](#). *Comm. Anal. Geom.* **10** (2002), no. 5, 1075–1113 Zbl [1115.53054](#) MR [1957663](#)
- [84] K. Uhlenbeck and S.-T. Yau, [On the existence of Hermitian–Yang–Mills connections in stable vector bundles](#). *Frontiers of the Mathematical Sciences: 1985 (New York, 1985)*. *Comm. Pure Appl. Math.* **39** (1986), suppl., 257–293 MR [0861491](#)
- [85] E. Witten, [Supersymmetry and Morse theory](#). *J. Differential Geometry* **17** (1982), no. 4, 661–692 Zbl [0499.53056](#) MR [0683171](#)
- [86] H. Yuan, [Family Floer program and non-Archimedean Syz mirror construction](#). PhD thesis, State University of New York at Stony Brook, 2021

Received 8 November 2022; revised 21 June 2024.

Yang Li

Department of Mathematics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA; current address: Department of Pure Mathematics and Mathematical Statistics, Cambridge University, Wilberforce Road, Cambridge CB3 0WB, UK; yl454@cam.ac.uk