https://doi.org/10.35662/unine-thesis-3256

Faculté des Sciences
Institut de mathématiques
Rue Emile-Argand 11, 2000 Neuchâtel

On the symplectic topology of rational homology balls

Thèse

présentée à la Faculté des Sciences
pour l'obtention du grade de docteur ès Sciences en mathématiques

par

Johannes Hauber

Acceptée sur proposition du jury :

Prof. Felix Schlenk, Université de Neuchâtel, directeur de thèse
Prof. Peter Feller, Université de Neuchâtel, membre interne
Prof. Paul Biran, ETH Zürich, expert externe
Dr. Nikolas Adaloglou, Sorbonne Université, expert externe

Soutenue le 24 juin 2026

Université de Neuchâtel
Faculté des sciences

Rue Emile-Argand 11
CH-2000 Neuchâtel
doctorat.sciences@unine.ch

IMPRIMATUR POUR THESE DE DOCTORAT

La Faculté des sciences de l'Université de Neuchâtel autorise l'impression de la présente thèse soutenue par

Monsieur Johannes HAUBER

Titre :

"On the symplectic topology of rational homology balls"

sur le rapport des membres du jury composé comme suit :

- **Prof. Felix Schlenk,** directeur de thèse, Université de Neuchâtel, Suisse
- **Prof. Peter Feller,** Université de Neuchâtel, Suisse
- **Prof. Paul Biran,** ETHZ, Suisse
- **Dr Nikolas Adaloglou,** Sorbonne Université, France

Neuchâtel, le 4 août 2026

La Doyenne, Prof. P. Junier

Lebe mit deinem Jahrhundert, aber sei nicht sein Geschöpf; leiste deinen Zeitgenossen, aber, was sie bedürfen, nicht, was sie loben.

— Friedrich Schiller, *Ueber die ästhetische Erziehung des Menschen*

Abstract

Rational homology balls $B_{p,q}$, for coprime integers $0 < q < p$, and their quantitative refinements, the pin-ellipsoids $E_{p,q}(\alpha,\beta)$ and pin-balls $B_{p,q}(\alpha) := E_{p,q}(\alpha,\alpha)$, form a natural class of Liouville domains whose skeleta are Lagrangian (p,q)-pinwheels. These are immersed Lagrangian disks with singular locus along their boundary. From the perspective of almost toric fibrations, pin-ellipsoids generalise ordinary symplectic ellipsoids; from the perspective of singularity theory, the rational homology balls $B_{p,q}$ arise as the Milnor fibres of certain cyclic quotient singularities. This thesis studies both quantitative and qualitative aspects of their symplectic topology.

On the quantitative side, we investigate two symplectic embedding problems: embeddings of pin-balls into pin-cylinders and embeddings of pin-ellipsoids into $\mathbb{C}P^2$. For the first problem, we prove that Gromov's non-squeezing theorem continues to hold in the pin-setting. For the second problem, we show that, for each pair (p,q) consisting of a Markov number p and one of its companion numbers q, the set of all $(\alpha,\beta) \in \mathbb{R}^2_{>0}$ for which there exists a symplectic embedding $E_{p,q}(\alpha,\beta) \hookrightarrow \mathbb{C}P^2$ has the structure of an infinite staircase, generalising the classical Fibonacci staircase discovered by McDuff and Schlenk. A key ingredient is an isotopy theorem asserting that any two such embeddings are ambiently Hamiltonian isotopic. We also prove sharp constraints for pairs of disjoint pin-ball embeddings.

On the qualitative side, we establish Arnold's nearby Lagrangian conjecture for pinwheels: any two embeddings of Lagrangian (p,q)-pinwheels in $B_{p,q}$ are related by a compactly supported Hamiltonian isotopy. The proof combines neck-stretching, the symplectic rational blow-up, and the computation of the compactly supported symplectic mapping class group of $B_{p,q}$, which is generated by a natural twist about the pinwheel, the pintwist $\tau_{p,q}$, analogous to the usual Dehn–Seidel twist.

As applications of the methods developed in this thesis, we give a new proof of the local Lagrangian unknotting theorem of Eliashberg and Polterovich, and show that the only Lagrangian pinwheels in $B_{p,q}$ are of (p,q)-type.

We hope that, taken together, these results establish the rational homology balls $B_{p,q}$, their Lagrangian skeleta, and their quantitative refinements as natural objects of study in symplectic topology by showing that they provide robust rational-homology analogues of balls, ellipsoids, and cotangent bundles, extending several classical theorems in symplectic geometry to a singular and arithmetically rich setting.

Keywords Symplectic topology, singular Lagrangians, Lagrangian pinwheels, Markov numbers, rational homology balls, rational blow-up/down, nearby Lagrangian conjecture.

Résumé

Les boules d'homologie rationnelle $B_{p,q}$, pour des entiers premiers entre eux $0 < q < p$, ainsi que leurs raffinements quantitatifs, les pin-ellipsoïdes $E_{p,q}(\alpha,\beta)$ et les pin-boules $B_{p,q}(\alpha) := E_{p,q}(\alpha,\alpha)$, forment une classe naturelle de domaines de Liouville dont les squelettes sont des pinwheels lagrangiens de type (p,q). Ce sont des disques lagrangiens immergés dont le lieu singulier est situé le long du bord. Du point de vue des fibrations presque toriques, les pin-ellipsoïdes généralisent les ellipsoïdes symplectiques ordinaires ; du point de vue de la théorie des singularités, les boules d'homologie rationnelle $B_{p,q}$ apparaissent comme les fibres de Milnor de certaines singularités quotients cycliques. Cette thèse étudie à la fois les aspects quantitatifs et qualitatifs de leur topologie symplectique.

Du côté quantitatif, nous étudions deux problèmes de plongement symplectique: les plongements de pin-boules dans des pin-cylindres et les plongements de pin-ellipsoïdes dans $\mathbb{C}P^2$. Pour le premier problème, nous démontrons que le théorème de non-entassement de Gromov reste valable dans le cadre des pinwheels. Pour le second problème, nous montrons que, pour chaque paire (p,q) formée d'un nombre de Markov p et de l'un de ses nombres compagnons q, l'ensemble des $(\alpha,\beta) \in \mathbb{R}^2_{>0}$ pour lesquels il existe un plongement symplectique $E_{p,q}(\alpha,\beta) \hookrightarrow \mathbb{C}P^2$ possède la structure d'un escalier infini, généralisant l'escalier de Fibonacci classique découvert par McDuff et Schlenk. Un ingrédient essentiel est un théorème d'isotopie affirmant que deux tels plongements quelconques sont isotopes par une isotopie hamiltonienne ambiante. Nous démontrons également des contraintes optimales pour les plongements de paires de pin-boules disjointes.

Du côté qualitatif, nous établissons la conjecture des lagrangiens voisins d'Arnold pour les pinwheels: deux plongements quelconques de pinwheels lagrangiens de type (p,q) dans $B_{p,q}$ sont reliés par une isotopie hamiltonienne à support compact. La preuve combine l'étirement du cou, l'éclatement rationnel symplectique et le calcul du groupe des classes d'applications symplectiques à support compact de $B_{p,q}$, qui est engendré par un twist naturel autour du pinwheel, le pin-twist $\tau_{p,q}$, analogue au twist de Dehn–Seidel usuel.

Comme applications des méthodes développées dans cette thèse, nous donnons une nouvelle preuve du théorème local de dénouement lagrangien d'Eliashberg et Polterovich, et nous montrons que les seuls pinwheels lagrangiens dans $B_{p,q}$ sont de type (p,q).

Nous espérons que, pris ensemble, ces résultats établissent les boules d'homologie rationnelle $B_{p,q}$, leurs squelettes lagrangiens et leurs raffinements quantitatifs comme des objets naturels d'étude en topologie symplectique, en montrant qu'ils fournissent des analogues robustes, en homologie rationnelle, des boules, des ellipsoïdes et des fibrés cotangents, et qu'ils étendent plusieurs théorèmes classiques de géométrie symplectique à un cadre singulier et arithmétiquement riche.

Mots clés topologie symplectique, lagrangiens singuliers, pinwheels lagrangiens, nombres de Markov, boules d'homologie rationnelle, éclatement/contraction rationnel(le), conjecture des lagrangiens voisins.

Acknowledgments

I am very happy, first of all, to express my gratitude to my advisor, Felix Schlenk. Thank you, Felix, for giving me so much space during my time as your PhD student. This time could have ended really badly, since sociology was much more interesting than mathematics for the first two years, but it didn't! I'm very happy to have had you as an advisor, personal mentor, collaborator and ultimately, as a friend!

There are two other people that I'm particularly grateful to have had around during my time as a graduate student: Nikolas Adaloglou and Jonny Evans. A lot of what I've learned in the last four years I learned from conversations with Nik and Jonny and I don't think they realise how impactful their thinking about mathematics is. Thank you for sharing your knowledge and your intuitions!

On that note I'd also like to thank my other collaborators: Joel Schmitz, with whom I gladly shared an office and who taught me how to do mathematics without ever writing anything, Gerard Bargalló i Gómez and Joé Brendel. Moreover, I'd like to thank Matthew Buck, who I was fortunate enough to meet at many conferences and discuss mathematics with extensively.

Of course, developing the mathematical ideas contained in my thesis during the last years would not have been possible without all the wonderful colleagues I had during my time at the University of Neuchâtel. I thank every single one of them!

Special thanks to Paul Biran and Peter Feller, who agreed to serve on my committee and who, maybe unknowingly on their part, influenced my mathematical thinking.

There is a circle of friends that I was very happy to have around during my time as a PhD student: Alissa, Clara, Jakob, Léo, Nik (again), Nicolas, Noah, Philip, Pit and Ziad.

Lastly, I want to thank Lena and my family for being so warm to me during this whole time.

Contents

1 Introduction

1.1 Classical Questions and Theorems

Symplectic geometry studies manifolds X equipped with a closed, non-degenerate two-form ω. The resulting structure is flexible enough to be locally standard, by Darboux's theorem, and yet rigid enough to support subtle global phenomena. Darboux's theorem implies that sufficiently small regions of any symplectic manifold are symplectomorphic to subsets of $(\mathbb{C}^n, \omega_{\mathrm{st}})$ and therefore it is always possible to symplectically embed small balls in symplectic manifolds. Nevertheless, the problem of understanding which standard domains can be embedded into a fixed symplectic manifold is highly non-trivial. The foundational example and one of the points of departure for symplectic topology is Gromov's famous non-squeezing theorem. This result initiated the systematic use of symplectic embeddings as a way of measuring symplectic size. For $\boldsymbol{\alpha} \in \mathbb{R}^n_{>0}$ denote the standard symplectic ellipsoid by

$$E(\boldsymbol{\alpha}) := \left\{ (z_1, \dots, z_n) \in \mathbb{C}^n \;\middle|\; \sum_{j=1}^{n} \frac{\pi \left|z_j\right|^2}{\alpha_j} \leq 1 \right\} \subseteq (\mathbb{C}^n, \omega_{\mathrm{st}}).$$

For $\lambda \in \mathbb{R}_{>0}$ define the ball by $B^{2n}(\lambda) := E(\lambda, \dots, \lambda)$ and, by slight abuse of notation, the cylinder by $Z^{2n}(\lambda) := E(\lambda, \infty, \dots, \infty)$.

Theorem 1.1.1 (Gromov's non-squeezing theorem [Gro85])**.** *There exists a symplectic embedding $\iota : B^{2n}(\lambda) \hookrightarrow Z^{2n}(1)$ if and only if $\lambda \leq 1$.*

Remark 1.1.2. A remarkable aspect of Gromov's non-squeezing theorem is that it holds in all dimensions. By contrast, many rigidity results in symplectic geometry, including the results of this thesis, rely essentially on four-dimensional techniques.

Starting from Gromov's non-squeezing theorem, embedding problems have been pivotal for symplectic topology ever since. See [Sch18] for an excellent introduction to these problems and their history. Let us give a biased selection of cornerstone results on symplectic embedding problems that will be relevant in the following: McDuff–Polterovich's work on the connection between ball embeddings and algebraic geometry [MP94]; connectedness of spaces of ball embeddings [McD91; McD93; Bir96]; ball packing results [Gro85; MP94; Bir97] and McDuff and Schlenk's computation in [MS12] of the set

$$\mathcal{A} := \left\{ (\alpha, \beta) \in \mathbb{R}^2_{>0} \;\middle|\; E(\alpha, \beta) \overset{s}{\hookrightarrow} B(1) \right\},$$

which led to the discovery of the first symplectic staircase.

Another fundamental aspect of symplectic geometry is the study of Lagrangian submanifolds and their rigidity properties. One of the central open problems in this direction is Arnold's nearby Lagrangian conjecture [Arn86], which asserts that every closed exact Lagrangian submanifold of a cotangent bundle T^*Q, where Q is a closed manifold, should be Hamiltonian isotopic to the zero-section. In dimension four, the nearby Lagrangian conjecture has been established only in a few cases, such as T^*S^2 [Hin04; LW12], T^*T^2 [DGI16], $T^*\mathbb{R}P^2$ [HPW16;

Ada25] and in the non-compact case $T^*\mathbb{R}^2$ [EP96]. A closely related theme is the computation of symplectomorphism groups and symplectic mapping class groups. For instance, the Dehn–Seidel twist [Sei; Sei99] about a Lagrangian sphere is a central example of a compactly supported symplectomorphism, and such twists play a central role in the study of the topology of symplectic mapping class groups. From this point of view, cotangent bundles provide a natural local model in which questions about Lagrangian uniqueness, Hamiltonian isotopy, and symplectic mapping classes interact.

This thesis studies the rational homology balls $B_{p,q}$, together with their quantitative refinements $E_{p,q}(\alpha,\beta)$, $B_{p,q}(\alpha)$, $E_{p,q}(\alpha,\infty)$, and $E_{p,q}(\infty,\alpha)$, called pin-ellipsoids, pin-balls and pin-cylinders. These domains are natural rational-homology analogues of balls and ellipsoids, and their Lagrangian skeleta are the (p,q)-pinwheels $L_{p,q}$. From the Lagrangian point of view, $B_{p,q}$ can be viewed as a singular analogue of a cotangent bundle: the pinwheel $L^{\mathrm{st}}_{p,q}$ plays the role of the zero-section, and $B_{p,q}$ plays the role of its standard neighbourhood. This thesis studies both the quantitative embedding theory of these domains and the qualitative Lagrangian rigidity of pinwheels inside $B_{p,q}$. The guiding principle is that many classical theorems, in particular those mentioned above, about balls, ellipsoids, cotangent bundles, and Lagrangian uniqueness admit natural rational-homology analogues in this setting.

1.2 Rational Homology Balls

The symplectic manifolds $B_{p,q}$ can be viewed from several different perspectives.

First, they arise in algebraic geometry. Consider the cyclic quotient singularity $\frac{1}{p^2}(1, pq-1)$, that is, the quotient of $\mathbb{C}^2$ by the action of the group of p^2th roots of unity with weights $(1, pq-1)$. This singularity admits a smoothing whose Milnor fibre is the rational homology ball $B_{p,q}$. Equivalently, one may view $B_{p,q}$ as a quotient of the Milnor fibre of the A_{p-1}-singularity. Namely, define

$$A_{p-1} := \{(z_1, z_2, z_3) \in \mathbb{C}^3 \mid z_1 z_2 = z_3^p + 1\},$$

with the restriction of the standard symplectic form. The group of pth roots of unity acts freely on A_{p-1} with weights $(1, -1, q)$ and the quotient of A_{p-1} by this action equivalently defines $B_{p,q}$. Thus, $B_{p,q}$ is the Milnor fibre of the cyclic quotient singularity.

Second, $B_{p,q}$ is also a natural object to study from the point of view of low-dimensional topology and contact geometry. The link of the cyclic quotient singularity $\frac{1}{p^2}(1, pq-1)$ is the lens space $L(p^2, pq-1)$, and $B_{p,q}$ is a minimal symplectic filling of this lens space. Rational homology balls appear prominently in the classification of fillings of lens spaces and in the rational blow-down/up construction. While not every symplectic rational homology ball arises as a lens space filling, this class is the best studied and most useful one for applications. Moreover, this class also interacts most directly with algebraic surface singularities.

Third, $B_{p,q}$ has a distinguished Lagrangian skeleton, which is the vanishing cycle of the smoothing. This vanishing cycle is a Lagrangian (p,q)-pinwheel, which we denote by $L_{p,q}$ in the following. Topologically, a p-pinwheel is just a disk attached to a circle along the usual degree p map. In the $p=1$ case we obtain a disk, the $p=2$ case defines an $\mathbb{R}P^2$ and for $p \geq 3$ the resulting space is singular. A Lagrangian (p,q)-pinwheel is an immersed Lagrangian disk which is embedded away from the boundary circle and has a prescribed local model along the boundary depending on q. See Figure 1.1 (b) for an illustration of how the immersion looks along the boundary of the disk. The rational homology ball $B_{p,q}$ retracts onto this pinwheel skeleton. In particular, its reduced rational homology vanishes, so it has the rational homology of a ball even though its skeleton is generally not contractible. Taking this "Lagrangian skeleton"-perspective also explains why $B_{p,q}$ should be viewed as the "cotangent bundle" of the Lagrangian pinwheel $L^{\mathrm{st}}_{p,q}$. Indeed, Khodorovskiy's neighbourhood theorem [Kho13b] for Lagrangian pinwheels says that a (p,q)-pinwheel in a symplectic manifold (X,ω) admits a neighbourhood that is symplectomorphic to a star-shaped neighbourhood of the Lagrangian skeleton of $B_{p,q}$.

1.3 Almost toric geometry and pin-ellipsoids

The rational homology ball $B_{p,q}$ is also natural from the viewpoint of almost toric fibrations (ATFs). ATFs, introduced by Symington [Sym98; Sym01; Sym03], are Lagrangian torus fibrations on symplectic four-manifolds in which elliptic and focus-focus singularities are allowed. Their base diagrams are two-dimensional integral-affine objects from which much of the symplectic geometry of the symplectic four-manifold can be read off.

The rational homology ball $B_{p,q}$ admits an almost toric fibration whose base contains a single node. Using this diagram one can define the quantitative analogues of the standard ellipsoid, ball, and cylinder. For $\alpha,\beta > 0$, the pin-ellipsoid $E_{p,q}(\alpha,\beta)$ is the compact symplectic manifold corresponding to the almost toric base diagram as shown in Figure 1.1 (a).

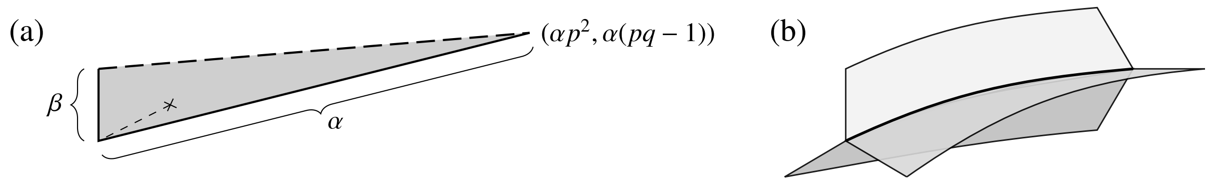

Figure 1.1: (a) The almost toric base diagram $\mathfrak{A}_{p,q}(\alpha,\beta)$ of $E_{p,q}(\alpha,\beta)$. (b) An illustration of how a (p,q)-pinwheel looks around its core circle.

This definition recovers familiar objects, while also introducing new ones. In the trivial case $(p,q) = (1,1)$ we recover the ordinary ellipsoid $E_{1,1}(\alpha,\beta) = E(\alpha,\beta)$. For $(p,q) = (2,1)$, the pinwheel $L_{2,1}$ is just a Lagrangian $\mathbb{R}P^2$ and $E_{2,1}(\alpha,\beta)$ defines a neighbourhood of the zero-section in $T^*\mathbb{R}P^2$. In particular, $B_{2,1}(\alpha)$ is the codisk bundle $D^*_{\leq\alpha}\mathbb{R}P^2$ with respect to the round metric on $\mathbb{R}P^2$, while $E_{2,1}(\alpha,\beta)$ in general is a codisk bundle with respect to a Finsler metric on $\mathbb{R}P^2$.

1.4 Outline of the thesis

This thesis covers half of the work the author did during his PhD at the Université de Neuchâtel under the guidance of Professor Felix Schlenk. It does not include the two papers that the author wrote at the beginning of his PhD, the first of which [BHS24] studies Lagrangian tori in smoothings of Wahl singularities and was written in collaboration with Brendel and Schmitz, and the second [AH25], which was written together with Adaloglou, studies symplectic embeddings of rational homology balls into rational and ruled surfaces, while also answering a question by Kronheimer about Lagrangian $\mathbb{R}P^2$s. Hence, the thesis contains two papers, the first one [ABH26], written in collaboration with Adaloglou and Bargalló i Gómez, and the second one [Ada+25], a joint work with Adaloglou, Brendel, Evans and Schlenk.

[Ada+25] appeared before [ABH26] and is certainly the intellectual inspiration for [ABH26]. The ordering is justified by the fact that [ABH26] studies the "intrinsic" symplectic topology of $B_{p,q}$, while [Ada+25] studies the symplectic phenomena, in particular quantitative ones, of symplectic embeddings $B_{p,q} \hookrightarrow \mathbb{C}P^2$.

Chapter 2 ([ABH26]). This chapter focuses on the symplectic topology of $B_{p,q}$ itself. The two main sections, Section 2.3.3 and Section 2.4, establish that the Lagrangian (p,q)-pinwheels in $B_{p,q}$ are unique up to compactly supported symplectomorphism and that the compactly supported mapping class group of $B_{p,q}$, namely $\pi_0(\mathrm{Symp}_c(B_{p,q}))$, is generated by a natural twist about the Lagrangian skeleton $L^{\mathrm{st}}_{p,q} \subseteq B_{p,q}$. Taken together these sections establish the nearby Lagrangian conjecture for Lagrangian pinwheels. To illustrate the methods developed within this chapter, Section 2.5 provides three applications: first the non-squeezing theorem for rational homology balls, second an alternative proof of the famous Eliashberg–Polterovich theorem about unknottedness of Lagrangian planes, and third a result determining which Lagrangian pinwheels embed into $B_{p,q}$.

Chapter 3 ([Ada+25]). The are three main results. First, any two symplectic embeddings of a fixed pin-ellipsoid $E_{p,q}(\alpha,\beta)$ into $\mathbb{C}P^2$ are ambiently Hamiltonian isotopic. Second, for fixed (p,q), either no (p,q)-pinwheel embeds into $\mathbb{C}P^2$, or the set

$$\mathcal{A}_{p,q} := \left\{(\alpha,\beta) \in \mathbb{R}^2_{>0} \;\middle|\; E_{p,q}(\alpha,\beta) \overset{s}{\hookrightarrow} \mathbb{C}P^2(1)\right\}$$

has the structure of an infinite staircase in a certain range. Third, we prove sharp constraints for pairs of disjoint pin-ball embeddings into $\mathbb{C}P^2$, giving a rational-homology analogue of Gromov's two-ball theorem.

Taken together, Chapter 2 and Chapter 3 prove the nearby Lagrangian conjecture for pinwheels, compute the symplectic mapping class group of $B_{p,q}$, and establish rational-homology analogues of Gromov's non-squeezing theorem, Gromov's two-ball theorem, and the Fibonacci staircase of McDuff–Schlenk.

2 The Nearby Lagrangian Conjecture for pinwheels

2.1 Introduction

2.1.1 Short summary

The goal of this paper is to establish the nearby Lagrangian conjecture (NLC), for a family of singular Lagrangians called *Lagrangian (p, q)-pinwheels*, for any coprime integers $0 < q < p$. Our point of departure is that a Lagrangian (p, q)-pinwheel naturally appears as the Lagrangian skeleton, denoted by $L^{\mathrm{st}}_{p,q}$, of the Liouville manifold $B_{p,q}$. By a fundamental result of Khodorovskiy, who initiated the study of Lagrangian pinwheels in [Kho13b], every Lagrangian (p, q)-pinwheel in a symplectic 4-manifold has a Weinstein-type neighbourhood which is symplectomorphic to a neighbourhood of the skeleton of $B_{p,q}$. Therefore, one can view $B_{p,q}$ as the "cotangent bundle" of $L^{\mathrm{st}}_{p,q}$. With this perspective, we will show that:

Main Theorem. *Any Lagrangian (p, q)-pinwheel in $B_{p,q}$ is Hamiltonian isotopic to $L^{\mathrm{st}}_{p,q}$.*

The technical core of the paper consists of two largely independent parts:

1. We determine how the *symplectic rational blow-up* along a (p, q)-pinwheel $L_{p,q} \subset B_{p,q}$ affects the symplectic topology of $B_{p,q}$. In particular, we show that there is a compactly supported symplectomorphism $\psi : B_{p,q} \to B_{p,q}$ which maps $L_{p,q}$ to $L^{\mathrm{st}}_{p,q}$. A key point is using the almost toric fibration on $B_{p,q}$ to construct a convenient compactification $X_{p,q}$ which makes the analysis of pseudoholomorphic curves manageable.

2. We compute the weak homotopy type of $\mathrm{Symp}_c(B_{p,q})$. By understanding how $\mathrm{Symp}_c(B_{p,q})$ acts on a distinguished symplectic cylinder $K_{p,q}$, we are able to show that $\mathrm{Symp}_c(B_{p,q})$ is weakly homotopy equivalent (w.h.e.) to $\mathbb{Z}$, generated by a very natural compactly supported symplectomorphism $\tau_{p,q}$, which we call the *pintwist*.

In Section 2.2, we collect the relevant background material concerning the symplectic and algebraic geometry of $B_{p,q}$ and pinwheels. Section 2.3 concerns the first technical part of the paper and Section 2.4 contains the second technical part. In Section 2.5, we provide three applications of the results of the previous sections to related quantitative and Lagrangian embedding problems of the rational homology balls $B_{p,q}$. Finally, in the Appendix we clarify an obtuse point in the definition of pinwheels and show that all pinwheels are locally visible, in the sense that the Khodorovskiy neighbourhood admits an almost toric fibration compatible with the pinwheel.

In the rest of the introduction we will first provide more context around the nearby Lagrangian conjecture and Lagrangian pinwheels. Then we will give a thorough discussion of our main results and the techniques we will develop to obtain them.

2.1.2 Context and motivation

The nearby Lagrangian conjecture

The *nearby Lagrangian conjecture*, attributed to V. I. Arnol'd, asserts that *given a smooth closed manifold Q, any exact Lagrangian in T^*Q is Hamiltonian isotopic to the zero section*. It is one of the foundational open problems in symplectic geometry, as it implies, among other things, that the smooth topology of Q and the symplectic topology of T^*Q are strongly related.

The NLC can be broken into two different subconjectures:

NLC$_1$ Any exact Lagrangian embedding $Q \hookrightarrow T^*Q$ is Hamiltonian isotopic (possibly after reparametrization) to the zero section.

NLC$_2$ Any exact Lagrangian $L \subset T^*Q$ is diffeomorphic to Q.

The two subconjectures have a wildly different character depending on whether T^*Q has dimension 4 or ≥ 6. When $\dim T^*Q = 4$, NLC$_2$ becomes manageable by the classification of compact surfaces and therefore research has focused on NLC$_1$, where the effectiveness of the theory of pseudoholomorphic curves has borne results. Using positivity of intersections of pseudoholomorphic curves, which is specific to dimension 4, the full NLC has been established for some simple smooth 2-dimensional manifolds, namely T^*S^2 [Hin04; LW12], T^*T^2 [DGI16], $T^*\mathbb{R}P^2$ [HPW16; Ada25] and in the non-compact case $T^*\mathbb{R}^2$ [EP96]. In addition, a version of the NLC has also been established in [Dim19] for the Lagrangian *Whitney sphere*, i.e. a Lagrangian sphere with a single positive self-intersection.

Considerable effort was needed to establish each of these results, and there does not seem to be a unified approach. However, all of these ultimately leverage the fact that there exists some compactification of the respective cotangent bundle to a simple rational symplectic 4-manifold, namely $\mathbb{C}P^2$ or $S^2 \times S^2$, where rational pseudoholomorphic curves are very well understood. The reader is pointed to the comprehensive survey [PS24] for more information and anecdotes around these low-dimensional results and their history.

The situation is starkly different in higher dimensions. There, already NLC$_2$ becomes extremely intricate; for example, it is not at all clear to what extent the smooth structure of Q dictates the symplectic structure of T^*Q. The most general step towards proving NLC$_2$ culminated in [AK18], building upon [FSS08; Abo12; Kra13] amongst others, where the authors show that any exact Lagrangian is simple homotopy equivalent to the zero section. See also [Abo+25b; Abo+25a] for more recent results. However, in high dimensions there has not been any progress towards NLC$_1$, because of the lack of control over the intersections of pseudoholomorphic curves.

Pinwheels

A p-pinwheel is the topological space obtained by attaching a 2-cell to a 1-cell via the p-to-1 covering map. A Lagrangian (p, q)-pinwheel is a p-pinwheel whose 2-cell is Lagrangian. In addition, we require that the 2-cell meets the 1-cell in a specified way, which depends on q. Since the definitions are somewhat technical, we refer to the Background Section 2.2.2 for the details. We note that $(1, 1)$-pinwheels are just Lagrangian discs and $(2, 1)$-pinwheels are Lagrangian projective planes.

Pinwheels were first defined by Khodorovskiy in [Kho13b], in order to define the *symplectic rational blow-up*, which is an operation that replaces a standard neighbourhood of a pinwheel with a chain of symplectic spheres, called the *Wahl chain* $C_{p,q}$, where the self-intersection numbers of the exceptional spheres in the Wahl chain depend on the Hirzebruch–Jung continued fraction associated to $\frac{p^2}{pq-1}$. The symplectic rational blow-up is the opposite operation of the *symplectic rational blow-down*, defined by Symington in [Sym98].

Lagrangian pinwheels lie at the interface between symplectic geometry and algebraic geometry, specifically surface singularities. In particular, one can view the (p,q)-pinwheel as the *vanishing cycle* of the $\frac{1}{p^2}(1, pq-1)$ cyclic quotient singularity. In this setting, the rational blow-up operation can be viewed as exchanging the smoothing of the singularity with the chain of exceptional spheres $C_{p,q}$ that appear in the minimal resolution of the $\frac{1}{p^2}(1, pq-1)$ singularity.[1] This point of view was first adopted by Evans and Smith in their influential papers [ES18; ES20] where they show that Lagrangian pinwheels are the right objects whose symplectic rigidity reflects the algebraic geometry of quotient singularities and their smoothings.

The rational homology ball $B_{p,q}$ is the smoothing, or Milnor fibre, of the $\frac{1}{p^2}(1, pq-1)$ cyclic quotient singularity. As such, it admits a universal p-fold covering by the Milnor fibre of the A_{p-1}-singularity, from which one can infer most basic symplectic properties of $B_{p,q}$; see Section 2.2.1. By Khodorovskiy's neighbourhood theorem [Kho13b, Lemma 3.3/3.4], every (p,q)-pinwheel has a neighbourhood symplectomorphic to a neighbourhood of $L^{\mathrm{st}}_{p,q}$. Hence, $L^{\mathrm{st}}_{p,q}$ can be viewed as the zero section of $B_{p,q}$. Our main theorem is:

Theorem A (NLC$_1$ for pinwheels; Theorems 2.3.24 and 2.4.11)**.** *Let L be a (p,q)-pinwheel in $B_{p,q}$. Then, there exists a compactly supported Hamiltonian isotopy ψ such that $\psi(L) = L^{\mathrm{st}}_{p,q}$.*

By now, there is an extensive list of results studying qualitative and quantitative embeddings of the rational homology balls $B_{p,q}$ and pinwheels, see [Kho13a; ES18; ES20; SS20; BS24; Ada26; AH25; Ada+25; Buc25]. These results show that Lagrangian pinwheels and their neighbourhoods naturally extend several classical theorems of symplectic geometry, e.g. Gromov's non-squeezing [AH25], the staircases of McDuff–Schlenk [Ada+25], and Biran's Lagrangian barrier phenomena [BS24].

2.1.3 Main results and outline of the paper

In Section 2.2 we set up notation and collect the relevant facts we will use in the rest of the paper. To keep the paper at a reasonable length we point the reader to precise references for the various facts we use. In particular, we will assume working familiarity with notions of almost toric fibrations and toric diagrams, at the level of Evans' invaluable book [Eva23].

Theorem A is the combination of two independent results which form the technical body of the paper. To show the nearby Lagrangian conjecture for a pinwheel $L \subset B_{p,q}$ we first show it to be true *up to symplectomorphism*:

Theorem B. *Let L be a Lagrangian (p,q)-pinwheel in $B_{p,q}$. Then, there exists $\psi \in \mathrm{Symp}_c(B_{p,q})$ such that $\psi(L) = L^{\mathrm{st}}_{p,q}$.*

Section 2.3 concerns the proof of Theorem B. First, we set up a convenient compactification $X_{p,q}$ of $B_{p,q}$, which allows us to pass to the more standard setting of closed pseudoholomorphic curves. Most of the geometry of $X_{p,q}$ is encoded in the compactifying divisor $\mathcal{D}_{p,q}$. Then, we show that rationally blowing up a (p,q)-pinwheel produces a symplectic manifold $\widetilde{X}_{p,q}$ which is

[1] In addition, the symplectomorphism $\tau_{p,q}$ can be viewed as the symplectic monodromy associated to the smoothing of the singularity.

regulated, in the sense of [Ada+25]. Roughly, this means that $\widetilde{X}_{p,q}$ admits a foliation by rational curves, all but finitely many of which are smooth. The main technical part of this section is to show that given *any* Lagrangian (p,q)-pinwheel, the resulting blown-up manifolds are symplectomorphic via a symplectomorphism that preserves certain aspects of the corresponding regulations. In particular, Theorem 2.3.21 shows that the regulation consists of exactly one broken fibre, and its shape (i.e. number of spheres and their self-intersections) depends only on p and q. This implies that the complements of the Khodorovskiy neighbourhoods of any two Lagrangian pinwheels are symplectomorphic. The last part of this section then shows that such a symplectomorphism can be extended to the interior of the neighbourhoods, thus providing the desired symplectomorphism.

To upgrade the symplectomorphism of Theorem B to a Hamiltonian diffeomorphism, we compute the weak homotopy type of $\mathrm{Symp}_c(B_{p,q})$ to obtain control over the symplectomorphism group. We show that a certain natural symplectomorphism $\tau_{p,q} \in \mathrm{Symp}_c(B_{p,q})$, first studied in Buck's PhD thesis [Buc25], generates the symplectic mapping class group. Buck showed, using notions from Lagrangian Floer theory, that $\tau_{p,q}$ has infinite order in $\mathrm{Symp}_c(B_{p,q})$. Using completely different techniques, we show that, in fact, it generates the symplectic mapping class group. The pintwist $\tau_{p,q}$ is very reminiscent of the usual Dehn–Seidel twist about a Lagrangian sphere or $\mathbb{R}P^2$.

Theorem C. *The group $\mathrm{Symp}_c(B_{p,q})$ is weakly homotopy equivalent to $\mathbb{Z}$, with a generator being the pintwist $\tau_{p,q}$. Therefore, for any $\psi \in \mathrm{Symp}_c(B_{p,q})$ there exists $\phi \in \mathrm{Ham}_c(B_{p,q})$ such that $\psi = \tau^n_{p,q} \circ \phi$.*

Section 2.4 contains the proof of Theorem C. The main point is to understand $\mathrm{Symp}_c(B_{p,q})$ via its action on the space of symplectic cylinders $\mathcal{S}(K_{p,q})$ which agree outside a compact set with the anticanonical divisor $K_{p,q}$, which we call *the central cylinder* of $B_{p,q}$; this is detailed in Section 2.2.1. The study of this action is carried out through the following diagram, which we call a *Comb diagram*[2]:

$$\begin{tikzcd}
\mathrm{Symp}_c(B_{p,q}\backslash K_{p,q}) \arrow[r] & \mathrm{Fix}(K_{p,q}) \arrow[r] \arrow[d, "D"] & \mathrm{Pr}(K_{p,q}) \arrow[r] \arrow[d, "\rho"] & \mathrm{Symp}_c(B_{p,q}) \arrow[d, dashed, "\text{transitive action}"] \\
 & \mathrm{Aut}(\nu K_{p,q}) \simeq \mathbb{Z} & \mathrm{Symp}_c(K_{p,q}) \simeq \mathbb{Z} & \mathcal{S}(K_{p,q}) \simeq pt
\end{tikzcd}$$

The spaces appearing in this diagram are the following: $\mathrm{Pr}(K_{p,q})$ (resp. $\mathrm{Fix}(K_{p,q})$) is the space of compactly supported symplectomorphisms that preserve (resp. fix) the central cylinder $K_{p,q}$ and $\mathrm{Aut}(\nu K_{p,q})$ is the space of symplectic automorphisms of the normal bundle of $K_{p,q}$. The diagram is built from the three fibrations obtained from the natural action of $\mathrm{Symp}_c(B_{p,q})$ on $\mathcal{S}(K_{p,q})$; the restriction map ρ; and the normal derivative map D. The top row that holds the "teeth" of the comb is formed by inclusions.

The essential feature of the comb diagram is that it relates the homotopy groups of $\mathrm{Symp}_c(B_{p,q})$ to those of $\mathrm{Symp}_c(B_{p,q}\backslash K_{p,q})$. The reason to consider the central cylinder as the divisor in the comb diagram is that the Liouville manifolds $B_{p,q}\backslash K_{p,q}$ can be identified with star-shaped domains in a single complete Liouville manifold T^*Wh. We then show, again via the appropriate comb diagram, that the symplectomorphism group of T^*Wh is weakly contractible and therefore $\mathrm{Symp}_c(B_{p,q}\backslash K_{p,q})$ is also weakly contractible.

[2] Looking at the diagram from afar, it resembles a horizontal *hair-comb*.

The symplectic geometry of T^*Wh, the cotangent bundle of the Whitney sphere, was first studied by Dimitroglou-Rizell in [Dim19], and we make good use of his analysis of the holomorphic curves therein.

The comb diagram has been the main tool in computing new symplectomorphism groups from old ones. The idea is already implicit in Gromov's original paper [Gro85, Section 2.4] and was used in most subsequent computations, for example [Abr98; AM00], as well as [Eva11, Section 6] and very recently [AMW25, Equation (2.3)]. Given a symplectic manifold (X, ω) and some symplectic divisor S, the comb diagram allows one to relate $\mathrm{Symp}_c(X)$, $\mathrm{Symp}_c(X\backslash S)$ and the orbit space of S under the action of $\mathrm{Symp}(X)$, which is commonly shown to be a moduli space $\mathcal{S}(S)$ of symplectic surfaces. Usually, this allows one to pass from a *compact* manifold X to the *non-compact* $X\backslash S$. One of the important novelties of our work is to show that this technique can also work when X is not closed. Indeed, to compute $\mathrm{Symp}_c(B_{p,q})$, we will *not* compactify $B_{p,q}$ but, rather, we will be viewing $B_{p,q}$ as a partial compactification of a space D^*Wh, the unit cotangent bundle of the Whitney sphere. The reader is invited to compare our computation of $\mathrm{Symp}_c(B_{2,1})$ to that of Evans, in [Eva11, Theorem 1.5].[3,4]

The key hard input from holomorphic curves in this section is the content of Proposition 2.4.8, analysing the properties of rational pseudoholomorphic curves in the compactification $X_{p,q}$. We show that the compactified curves in the class of the regulation of $X_{p,q}$ cannot break, thereby establishing that $\mathcal{S}(K_{p,q})$ is contractible. $\mathrm{Symp}_c(B_{p,q})$ retracts to $\mathrm{Pr}(K_{p,q})$ and this is shown to be generated by $\tau_{p,q}$.

In the rest of the paper, we provide various applications of the techniques we develop to prove Theorems B and C. The first application concerns a pin-version of Gromov's non-squeezing theorem. Using the almost toric base diagrams of $B_{p,q}$, we can define various domains that generalize the usual toric ones. In particular, we define the *pin-ellipsoids* $E_{p,q}(\alpha,\beta)$ for $0 < \alpha,\beta \leq \infty$. In analogy with the standard situation, we let $B_{p,q}(r) = B_{p,q}(r,r)$ be the *pin-ball*. We show that they satisfy a pin-version of Gromov's non-squeezing theorem:

Application I (Pin-ball non-squeezing)**.** *The pin-ball $B_{p,q}(1)$ symplectically embeds into the pin-cylinders $E_{p,q}(\alpha,\infty)$ and $E_{p,q}(\infty,\beta)$ if and only if $1 \leq \alpha,\beta$.*

For the proof, we assume that a symplectic embedding $\iota : B_{p,q}(1) \hookrightarrow E_{p,q}(\infty,\beta)$ exists and then rationally blow up along ι. The unique broken fibre of the regulation obtained on the blown-up manifold contains a single exceptional curve, whose area provides the desired obstruction.

The second application concerns an alternative proof of a classical theorem of Eliashberg and Polterovich [EP96] concerning Lagrangian planes. This theorem is of fundamental importance in 4-dimensional symplectic geometry as it shows that one cannot produce Lagrangian knotting by considering some local operation on a Lagrangian surface. (However, Lagrangian knotting *does* occur on the global level, as was first exhibited by Seidel [Sei99].)

Application II (Eliashberg–Polterovich local unknottedness)**.** *Let Δ_{st} be a standard Lagrangian disc with boundary on $S^3 \subset \mathbb{R}^4$ and let Δ be a Lagrangian which agrees with Δ_{st} near the*

[3] It is very tempting to try to mimic the Seidel–Evans approach, i.e. to first compute the $\mathbb{Z}_p$-equivariant symplectomorphism group and then to relate it to $\mathrm{Symp}_c(B_{p,q})$. However, we were not able to make this approach work; in A_{p-1} there are p holomorphic foliations, and when $p > 2$, they might intersect with each other in uncontrolled ways.

[4] The recent preprint of Keating–Smith–Wemyss [KSW26] contains a thorough discussion of computing certain algebraic aspects of symplectomorphism groups through mirror symmetry and the action of Symp on the Fukaya category.

boundary. Then, there exists a compactly supported Hamiltonian diffeomorphism $\psi \in Ham_c(B^4)$ *such that* $\psi(\Delta) = \Delta_{\mathrm{st}}$.

The idea of our proof follows a relative version of Theorem B, for the case $(p, q) = (2, 1)$. Since the Lagrangian discs $\Delta, \Delta_{\mathrm{st}}$ are standard near the boundary, we can use the symplectic cut operation to compactify the ball B^4 into $\mathbb{C}P^2$ and the Lagrangian discs into Lagrangian real projective planes. Application II is equivalent to the fact that there exists a symplectomorphism of $\mathbb{C}P^2$ which is the identity near the compactifying line $\mathbb{C}P^1_\infty$ and maps the two Lagrangian projective planes onto each other. We show this by carefully analysing the rational blow-up in this setting.

The final application concerns the *nearby Lagrangian subconjecture 2*, adapted for pinwheels, namely answering the natural question: *Which Lagrangian pinwheels exist in* $B_{p,q}$*?* For example, the paper of Lekili–Maydanskiy [LM14], which originally studied symplectic aspects[5] of $B_{p,q}$, shows that, for $p > 2$, there are no smooth exact Lagrangian submanifolds in $B_{p,q}$.[6] We extend these results to show that the only Lagrangian pinwheels in $B_{p,q}$ are those of (p, q)-type, i.e. of the same type as the Lagrangian skeleton of $B_{p,q}$.

Application III. *If there is a symplectic embedding of a star-shaped subdomain of* $B_{n,m}$ *into* $B_{p,q}$*, then* $(n, m) = (p, q)$*. In particular, a Lagrangian pinwheel in* $B_{p,q}$ *is of* (p, q)*-type.*

The proof is reminiscent of the one in [Ada+25, Section 5.4], which itself is an elaboration of [ES18]. In particular, we assume the existence of a symplectic embedding $\iota : B_{n,m} \hookrightarrow B_{p,q}$ and then stretch the neck around the contact boundary of $\iota(B_{n,m})$ to find specific rational curves in the rational blow-up $\widetilde{X}_{p,q}$. The numerics for these curves, encoded in the adjunction formula, imply that $(n, m) = (p, q)$. The corresponding statement for pinwheels follows by considering the associated Khodorovskiy neighbourhood. This result fits into the growing literature on embeddings of rational homology balls; see, for instance, Evans and Smith [ES18] and the more topologically flavoured works [Etn+23; LP22; LP24; Owe18; Owe20; GO25].

2.1.4 Relation with ABEHS

Many of the ideas contained in this paper, especially Section 2.3.3, grew out of the collaboration of the first and third authors with Brendel, Evans, and Schlenk that culminated in [Ada+25]. In a certain sense, this paper can be seen as a prequel to [Ada+25], as the theorems here can be seen as local versions of the isotopy and embedding theorems of [Ada+25], concerning embeddings of $B_{p,q}$ into $\mathbb{C}P^2$.

Theorem B, which constructs the symplectomorphism relating the pinwheels, should be compared with [Ada+25, Corollary 5.2.2]. In fact, our situation is much more straightforward; $B_{p,q}$ is the natural manifold containing $L_{p,q}$; therefore, the arithmetic appearing in the adjunction formula of the neck-stretched curves in Section 2.3 is somewhat more manageable.

On the other hand, upgrading the symplectomorphism to a Hamiltonian one is automatic in the setting of [Ada+25], because $\mathrm{Symp}(\mathbb{C}P^2)$ is connected by Gromov [Gro85]. However, apart from the calculation of $\mathrm{Symp}_c(B_{p,1})$, which is due to Gromov [Gro85] for $p = 1$ and to Evans

[5]Interestingly, in [LM14] it is also shown that the symplectic homology of $B_{p,q}$ does not vanish in spite of having no exact Lagrangians for $p > 2$. Another notable symplectic feature of $B_{p,q}$, for $q = 1$, is that by [Kar18] the Kontsevich cosheaf conjecture is known to hold for them.

[6]Since $B_{2,1}$ is symplectomorphic to $T^*\mathbb{R}P^2$, it contains many Lagrangian $\mathbb{R}P^2$s, all Hamiltonian isotopic by [HPW16; Ada25] and by the main theorem of this paper.

[Eva11] for $p = 2$, the symplectomorphism group of $\mathrm{Symp}_c(B_{p,q})$ for other $p > 2$ was entirely unknown before.[7] Partial progress was made by Buck [Buc25], who showed that $\tau_{p,q}$ has infinite order.

Finally, the quantitative embedding obstructions of Application I are obtained similarly to *Markov staircases*, in the sense that in both cases the obstructions are extracted from the area of exceptional curves in the broken fibres after blowing up.

2.2 Background

In this section we will collect basic definitions and properties of the rational homology ball $B_{p,q}$ and Lagrangian pinwheels. We have tried to keep the exposition tight and self-contained. However, in order to keep the paper at a reasonable length, for certain technical aspects (which will not be relevant for the rest of the paper), the reader will be referred to the literature. In particular, we assume working familiarity with almost toric fibrations (ATFs), at the level of [Eva23; Sym03].

Unless stated otherwise, we always assume that $0 < q < p$ are two coprime integers, and we will denote the group of the pth roots of unity by $\boldsymbol{\mu}_p := \{\zeta \in \mathbb{C} \mid \zeta^p = 1\}$.

2.2.1 The rational homology ball $B_{p,q}$ and its cylinder $K_{p,q}$

Here we review some important facts about the symplectic geometry of the rational homology balls $B_{p,q}$. Following Lekili–Maydanskiy [LM14], we view $B_{p,q}$ as a p-fold quotient. Consider the complex hypersurface $A_{p-1} = \{z_1 z_2 = z_3^p + 1\} \subset \mathbb{C}^3$ with the restriction of the standard symplectic form ω_{st} on $\mathbb{C}^3$.[8] There is a free $\mathbb{Z}_p$-Kähler action on A_{p-1} given by

$$\zeta \cdot (z_1, z_2, z_3) = (\zeta z_1, \zeta^{-1} z_2, \zeta^q z_3) \quad \text{where } \zeta \in \boldsymbol{\mu}_p. \tag{2.2.1}$$

Definition 2.2.1. The symplectic manifold $(B_{p,q}, \omega_{p,q})$ is defined to be the quotient of $(A_{p-1}, \omega_{\mathrm{st}})$ by this $\mathbb{Z}_p$-action.

There are two different fibrations on $B_{p,q}$ that we will use to understand its symplectic geometry. One is a *Lefschetz (orbi-)fibration*[9] and the other is an *almost toric fibration*.

The Lefschetz fibration

Consider the usual Lefschetz fibration $\tilde{\pi} : A_{p-1} \to \mathbb{C}$ given by projecting onto the third coordinate z_3. The only nodal fibres sit over the pth roots of unity $\boldsymbol{\mu}_p$, and the $\mathbb{Z}_p$-action defined in (2.2.1) respects the fibration (in the sense that it maps fibres to fibres). This is shown in Figure 2.1. Notably, the action preserves the central fibre $\tilde{\pi}^{-1}(0)$.

Definition 2.2.2 (Central cylinder)**.** The *central cylinder* of $(A_{p-1}, \omega_{\mathrm{st}})$ is the fibre $K_{p-1} := \tilde{\pi}^{-1}(0)$. The *central cylinder* of $(B_{p,q}, \omega_{p,q})$ is the quotient of K_{p-1} under the corresponding action. We will denote this cylinder by $K_{p,q}$.

[7] Recall that $B_{1,1}$ is just the standard ball B^4 and $B_{2,1}$ is the unit disc bundle of $\mathbb{R}P^2$.

[8] In fact, this is the Milnor fibre of the A_{p-1}-singularity, but for the sake of brevity, we choose this more compact notation.

[9] We will not focus on the orbifold aspects of this fibration, since we will mostly work on the relevant honest Lefschetz fibration on the universal cover, which is A_{p-1}.

Remark 2.2.3. The hyperplane section $\{z_3 = 0\} \cap \{z_1 z_2 = z_3^p\} =: \widetilde{K}_{p-1}$ is just the usual nodal conic in the plane $\{z_3 = 0\}$. The smoothing $\{z_1 z_2 = z_3^p + 1\}$ of $\{z_1 z_2 = z_3^p\}$ induces the usual smoothing of $\widetilde{K}_{p-1}$ to a smooth affine conic, which is precisely K_{p-1}. In other words, the two nested singularities $\widetilde{K}_{p-1} \subset \widetilde{A}_{p-1} = \{z_1 z_2 = z_3^p\}$ get simultaneously smoothed to $K_{p-1} \subset A_{p-1}$. We will use this later to show that the symplectic monodromy of A_{p-1} restricts to the symplectic monodromy of K_{p-1}.

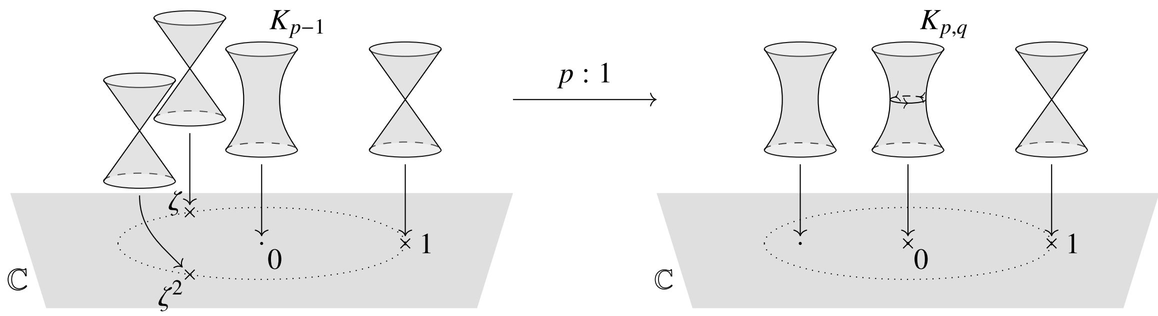

Figure 2.1: The Lefschetz fibrations defined on A_{p-1} and $B_{p,q}$. Note that the central cylinder $K_{p,q}$ is p-fold covered by a generic fibre of the Lefschetz fibration defined on $B_{p,q}$.

The almost toric fibration

The symplectic manifold $B_{p,q}$ also admits an almost toric fibration, as is thoroughly explained in [Eva23; Eva24]. Its almost toric base diagram $\mathfrak{A}_{p,q}$ is as shown in Figure 2.2. Representing $B_{p,q}$ via its almost toric base diagram has many advantages. For example, it is easy to talk about geometrically interesting objects in $B_{p,q}$ and to come up with potential ways to prove theorems about them. Domains in $B_{p,q}$ of interest to us are so-called pin-ellipsoids and pin-balls, which were introduced in [AH25] and explored further in [Ada+25].

Definition 2.2.4 (Pin-ellipsoid and pin-ball)**.** For $\alpha, \beta \in \mathbb{R}_{>0}$, the (p,q)-*pin-ellipsoid* $E_{p,q}(\alpha,\beta) \subseteq B_{p,q}$ is defined by its almost toric base diagram $\mathfrak{A}_{p,q}(\alpha,\beta)$, which is shown in Figure 2.2. The *pin-ball* is defined to be $B_{p,q}(\alpha) = E_{p,q}(\alpha,\alpha)$.

Remark 2.2.5. As can be seen in the almost toric base diagrams in Figure 2.2, the contact boundary of $E_{p,q}(\alpha,\beta)$ is the lens space $L(p^2, pq-1)$. Equivalently, one can argue that the boundary of $B_{p,q}$ is the same (as contact manifolds) as the boundary of a small neighbourhood of the quotient singularity $\frac{1}{p^2}(1, pq-1)$ which is the lens space $L(p^2, pq-1)$ by definition; see Remark 2.2.23 for more details.

While the fibrations are in a sense dual to each other, meaning that the Lefschetz fibration is a symplectic fibration and the ATF is a Lagrangian fibration, certain information is visible in both. For example, the central cylinder $K_{p,q} \subseteq B_{p,q}$ is defined to be the fibre over zero, but can also be seen over the toric boundary of $\mathfrak{A}_{p,q}$, i.e. as a visible surface in the sense of Symington [Sym03, Section 7], as explained in [Eva23, Section 7].

Lemma 2.2.6. *The visible surface over the toric boundary on the usual ATF on $B_{p,q}$ (resp. A_{p-1}) is the central cylinder $K_{p,q}$ (resp. K_{p-1}). In particular, the central cylinders are Poincaré dual to the first Chern class.*

Proof. The first part is proven in [Eva23, Example 7.6]. The statement about the Chern class follows from [Sym03, Proposition 8.2].[10] □

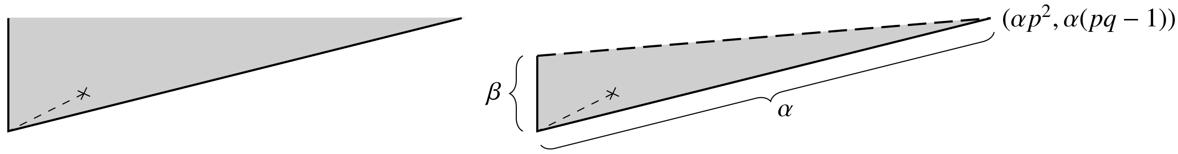

Figure 2.2: On the left the almost toric base diagram $\mathfrak{A}_{p,q}$ of $B_{p,q}$. The branch cut points into the (p,q)-direction. Note that over the toric boundary lies the embedded symplectic cylinder $K_{p,q}$ and the Lagrangian pinwheel's core circle is an essential loop on the central cylinder $K_{p,q}$. On the right, the almost toric base diagram $\mathfrak{A}_{p,q}(\alpha,\beta)$ of the pin-ellipsoid $E_{p,q}(\alpha,\beta)$.

2.2.2 Pinwheels

Definition 2.2.7. A *topological p-pinwheel* P_p is the topological space $D^2/\sim$, where the equivalence relation on the unit disc $D^2 \subseteq \mathbb{C}$ is given by $z \sim z'$ if and only if $z, z' \in \partial D^2$ and $z = e^{2\pi i k/p} z'$ for some $k \in \mathbb{Z}$.

Definition 2.2.8 (Pinwheel)**.** A *Lagrangian* (p,q)*-pinwheel* in a symplectic manifold (X,ω), usually denoted by $L_{p,q}$, is a Lagrangian immersion $f : D^2 \looparrowright X$ such that

1. f factors through a continuous embedding $P_p \hookrightarrow X$, where P_p is a p-pinwheel;
2. the restriction $f|_{D^2\setminus\partial D^2}$ is a Lagrangian embedding;
3. there is a Darboux chart of a neighbourhood U of $f(\partial D^2)$ in (X,ω) to an open neighbourhood $V \subseteq T^*S^1 \times \mathbb{C}$ that maps the *core circle* $f(\partial D^2)$ to the zero section of T^*S^1 times $0 \in \mathbb{C}$, such that in these coordinates the map f takes the form
$$f(t,s) = \left(pt, \frac{q}{2p}s^2, se^{iqt}\right),$$
where $(t,s) \in \partial D^2 \times [0,\varepsilon) \subseteq D^2$ are collar coordinates near the boundary of the parametrizing disc. See Figure 2.3.

Remark 2.2.9. A Lagrangian $(2,1)$-pinwheel is just a Lagrangian $\mathbb{R}P^2$, which can be described as a disc capping off a Möbius band. Similarly, for a general $L_{p,q}$, a neighbourhood of the core circle consists of p flanges meeting at the core and winding around it q times, as can be seen in the coordinates in Definition 2.2.8 (3). The boundary circle of this neighbourhood projects to a (p,q)-torus knot[11]; this circle is then capped off by a disc. See Figure 2.3.

[10] Symington only states the result for closed symplectic manifolds. However, the proof generalizes immediately. To obtain a precise statement one then has to either work with Borel–Moore homology or with compact subsets, e.g. $E_{p,q}(\alpha,\beta) \subseteq B_{p,q}$.
[11] Here the projection is from $S^1 \times D^3$ onto $S^1\times$ the last two components of D^3.

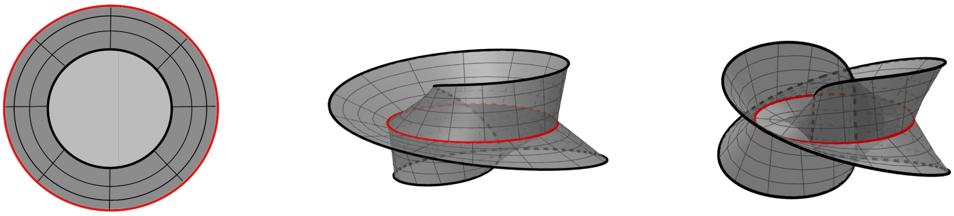

Figure 2.3: On the left the domain of parametrization of a Lagrangian pinwheel with a collar of its boundary (that becomes the core, in red) highlighted. On the right, the image of the collar neighbourhood for the (3, 1)- and (3, 2)-cases.

Remark 2.2.10. Our definition of a Lagrangian pinwheel is stronger than the original one given by Khodorovskiy [Kho13b, Definition 3.1]. There, the way the flanges meet the core circle is only prescribed up to homotopy, whereas we require that they meet in a specific geometric way. The reason for our stronger definition is very natural: for purely linear algebra reasons, two (p, q)-pinwheels cannot be related by a smooth ambient Hamiltonian isotopy if their flanges do not form at each point of the core circle p-tuples of Lagrangian planes that are related by a linear symplectic map.

In contrast to the problem of Hamiltonian isotopies between pinwheels, for problems of existence of pinwheels or, more generally, of embeddings of rational homology balls $B_{p,q}$, the difference between these two definitions is immaterial. Indeed, it is proved in [Kho13b, Lemma 3.3] that any pinwheel can be C^0-perturbed, through Lagrangian immersions, so that its flanges meet the core in the standard way.

Example 2.2.11. (Visible Pinwheel) As explained in [Eva23, Chapter 5] and [Sym03, Section 7], over the branch cut in $\mathfrak{A}_{p,q}$ one can explicitly construct Lagrangian (p, q)-pinwheels. If the intersection of the constructed Lagrangian pinwheel with any regular fibre over the branch cut consists of a single circle, then we will call it the *visible Lagrangian* (p, q)*-pinwheel* and denote it by $L^{\text{vis}}_{p,q}$.

Remark 2.2.12. A priori calling such a pinwheel "the" visible Lagrangian (p, q)-pinwheel is not justified. However, one can prove that any two choices of visible pinwheels are related by a Hamiltonian isotopy localised around the branch cut. Moreover, having chosen a visible Lagrangian (p, q)-pinwheel, the nodal slide can be realised in such a way that it preserves the pinwheel. See [GV23, Subsection 7.4] for more details on this. These two observations then justify talking about "the" visible Lagrangian (p, q)-pinwheel in $B_{p,q}$.

Definition 2.2.13 (Khodorovskiy neighbourhood)**.** A neighbourhood with contact-type boundary $U \subseteq (X, \omega)$ of a Lagrangian pinwheel $L_{p,q} \subseteq (X, \omega)$ is called a *Khodorovskiy neighbourhood* if its completion $\overline{U}$ along its contact boundary yields a pair $(\overline{U}, L_{p,q})$ that is symplectomorphic to $(B_{p,q}, L^{\text{vis}}_{p,q})$.

Extending Khodorovskiy's neighbourhood construction in [Kho13b, Section 3.4], we prove that any Lagrangian pinwheel admits a Khodorovskiy neighbourhood.

Theorem 2.2.14 (Theorem 2.6.6)**.** *Suppose that* $L_{p,q} \subseteq (X, \omega)$ *is a Lagrangian pinwheel in a symplectic manifold. Then for some* $\varepsilon > 0$ *there exists a symplectic embedding* $\psi : B_{p,q}(\varepsilon) \hookrightarrow (X, \omega)$ *such that* $\psi(L^{\text{vis}}_{p,q}) = L_{p,q}$.

From the dual point of view, one can also construct pinwheels using the Lefschetz-type fibration on $B_{p,q}$.

Example 2.2.15. Consider the Lefschetz fibration $\tilde{\pi} : A_{p-1} \to \mathbb{C}$, as introduced in Section 2.2.1. The fibration has p critical points, so there are also p vanishing thimbles. These are just p Lagrangian discs with common boundary on the central cylinder K_{p-1}. As shown in [Kar18, Proposition 5.1.9], one can take these thimbles to consist of the Lagrangian skeleton of the standard Liouville vector field. Passing to the quotient, these Lagrangian discs are identified, and since the central cylinder is preserved, one obtains a Lagrangian (p, q)-pinwheel in $B_{p,q}$, with its core circle on $K_{p,q}$. This is shown in Figure 2.4.

Since the vanishing thimbles in A_{p-1} are visible in the almost toric fibration, the (p, q)-pinwheel constructed via the quotient above is also visible for the standard almost toric fibration on $B_{p,q}$. The visibility statements are explained in detail in [Eva23, Remark 7.10].

Definition 2.2.16 (Standard pinwheel)**.** The Lagrangian (p, q)-pinwheel resulting after taking the quotient of the Lagrangian skeleton of A_{p-1} will be denoted by $L^{\mathrm{st}}_{p,q} \subseteq B_{p,q}$.

Remark 2.2.17. The standard pinwheel $L^{\mathrm{st}}_{p,q} \subseteq B_{p,q}$ and the visible pinwheel $L^{\mathrm{vis}}_{p,q} \subseteq B_{p,q}$ coincide. They are the same Lagrangian pinwheel from two different viewpoints. See for example [Eva24].

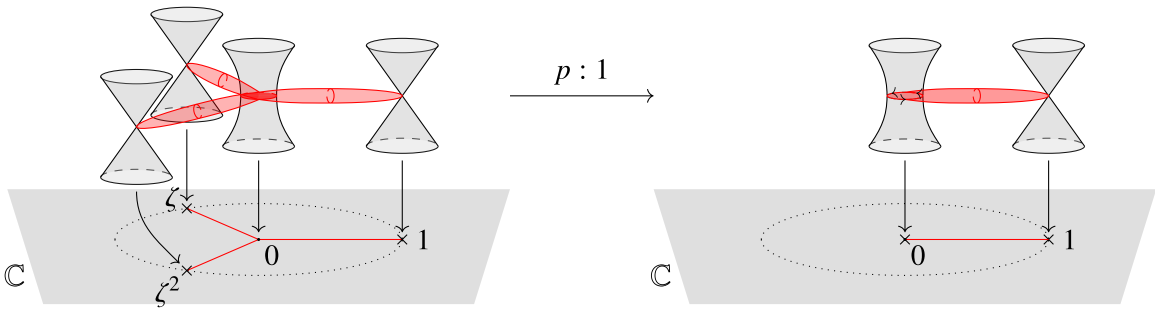

Figure 2.4: The Lefschetz fibrations on A_{p-1} and on $B_{p,q}$. Also shown are the vanishing thimbles and vanishing paths.

2.2.3 Pintwists

We now introduce the map $\tau_{p,q} \in \mathrm{Symp}_c(B_{p,q})$ which will act as the analogue of the Dehn–Seidel twist about a Lagrangian sphere. It was originally defined by Buck [Buc25, Proposition 5.4], adapting arguments from [Sei00], and we will follow his approach. Since $\tau_{p,q}$ has only recently been introduced, we repeat the general strategy of the definition for the benefit of the reader, and refer to [Buc25] for details. First, we define a twist $\tau_{p-1} \in \mathrm{Symp}_c(A_{p-1})$ on A_{p-1} and then push it down to $\mathrm{Symp}_c(B_{p,q})$ via the universal p-fold covering. Therefore, most properties of $\tau_{p,q}$ will be derived from those of τ_{p-1}.

For a thorough discussion of how one defines the symplectic monodromy of smoothings of singularities, we refer to [Buc25], as well as [MS16, Section 6.3] for a more introductory discussion of symplectic connections.

Definition 2.2.18. The map $\tau_{p-1} \in \mathrm{Symp}_c(A_{p-1})$ is defined to be the symplectic monodromy map associated to the A_{p-1} singularity.

Remark 2.2.19. Recall that a priori the symplectic monodromy of the fibration $(z_1, z_2, z_3) \to z_1 z_2 - z_3^p$ will have non-compact support. Seidel's construction in [Sei00] uses a cut-off function to produce the compactly supported map τ_{p-1}. Also, notice that τ_{p-1} is very far from generating $\mathrm{Symp}_c(A_{p-1})$. Indeed, as was shown in [Eva11] and [Wu14], $\mathrm{Symp}_c(A_{p-1})$ is weakly homotopy equivalent to the braid group with p strands.

Since the $\frac{1}{p}(1, -1, q)$-action, as defined in (2.2.1), is unitary, it preserves the horizontal connection, and therefore τ_{p-1} is equivariant with respect to this action and it descends to $B_{p,q}$.

Definition 2.2.20 (Pintwist)**.** The *pintwist* $\tau_{p,q} \in \mathrm{Symp}_c(B_{p,q})$ is the map induced by τ_{p-1} after taking the quotient of A_{p-1}.

Remark 2.2.21. The pintwist can also be understood intrinsically: from the perspective of Remark 2.2.23, the pintwist is the monodromy of the smoothing of the $\frac{1}{p^2}(1, pq-1)$ cyclic quotient singularity.

We now collect from [Buc25, Section 5.2] the main properties of the pintwist that we will need later. Points (1) and (3) follow by direct adaptations of arguments in [Buc25, Lemma 5.6].

Lemma 2.2.22 (Basic properties of pintwists)**.** *Let $\tau_{p-1} : A_{p-1} \to A_{p-1}$ be the symplectic monodromy map and let $\tau_{p,q} : B_{p,q} \to B_{p,q}$ be the pintwist.*

1. *The map τ_{p-1} preserves the central cylinder K_{p-1} and $\tau_{p-1}|_{K_{p-1}} = \tau$, where $\tau \in \mathrm{Symp}_c(K_{p-1})$ is the usual Dehn twist.*

2. *The pintwist preserves the central cylinder $K_{p,q}$. In particular, $\tau_{p,q}|_{K_{p,q}}$ is isotopic to τ^p, where $\tau \in \mathrm{Symp}_c(K_{p,q})$ is the usual Dehn twist.*

3. *The standard pinwheel $L^{\mathrm{st}}_{p,q}$ is Lagrangian, and thus Hamiltonian, isotopic to $\tau_{p,q}(L^{\mathrm{st}}_{p,q})$.*

Proof. The first point follows from the fact that the natural symplectic connection (coming from the symplectic form on $\mathbb{C}^3$) of the fibration $(z_1, z_2, z_3) \to z_1 z_2 - z_3^p$ restricts to a symplectic connection on the sub-fibration given by $K^w_{p-1} = \{z_1 z_2 - z_3^p = w\} \cap \{z_3 = 0\}$, where K^1_{p-1} is the central cylinder. Therefore, the symplectic monodromy τ_{p-1} induces a symplectic monodromy of cylinders, which can be seen to be the usual monodromy of the smoothing of the double line to the smooth conic.

The second point follows by inspection of the commutative diagram:

$$\begin{tikzcd}
(A_{p-1}, K_{p-1}) \arrow[r, "\tau_{p-1}"] \arrow[d, "\frac{1}{p}(1,-1,q)"] & (A_{p-1}, K_{p-1}) \arrow[d, "\frac{1}{p}(1,-1,q)"] \\
(B_{p,q}, K_{p,q}) \arrow[r, "\tau_{p,q}"] & (B_{p,q}, K_{p,q})
\end{tikzcd}$$

Since τ_{p-1} preserves K_{p-1}, after taking the quotient we see that $\tau_{p,q}$ preserves $K_{p,q}$. The map $\tau_{p-1}|_{K_{p-1}}$ twists a latitude of K_{p-1} once. Since $K_{p-1} \to K_{p,q}$ is a p-fold cover, $\tau_{p,q}|_{K_{p,q}}$ induces p twists of the latitude of K_{p-1} and is therefore a pth power of the usual Dehn twist.

The last point follows verbatim from the proof of [Buc25, Lemma 5.6], by observing that the arguments work for a vanishing thimble in $B_{p,q}$ over any path γ in the base, and in particular, the pinwheel $L^{\mathrm{st}}_{p,q}$. Recall that $L^{\mathrm{st}}_{p,q}$ is the vanishing thimble of the Lefschetz fibration on $B_{p,q}$, which lives over a path γ connecting the values 0 and 1. The *central twist* τ defined in [Buc25, Section 5.1], which is a mapping class of the disc with a marking at 1, fixes γ; therefore $\tau_{p,q}(L^{\mathrm{st}}_{p,q})$ is Lagrangian isotopic to $L^{\mathrm{st}}_{p,q}$. See also Figure 2.4.

The Lagrangian isotopy may be upgraded to a Hamiltonian one by appealing to Banyaga's Symplectic isotopy extension theorem ([Ban78], see also [MS16, Theorem 3.3.2]), since all positive-degree rational cohomology groups of $B_{p,q}$ and of its (p,q)-pinwheels vanish.[12] □

2.2.4 Resolutions, rational blow-ups and neck-stretching

This subsection closely follows the setup of [Ada+25, Sections 2 and 3].

Let $0 < a < n$ be two coprime integers and consider the action of the group of nth roots of unity $\boldsymbol{\mu}_n$ on $\mathbb{C}^2$ with weights $(1,a)$, i.e. $\zeta \cdot (z_1, z_2) = (\zeta z_1, \zeta^a z_2)$ for $\zeta \in \boldsymbol{\mu}_n$. Denote the quotient of $\mathbb{C}^2$ by this action by $\mathbb{C}^2/_{1/n(1,a)}$. This quotient is a toric orbifold, whose singular point is modelled on the cyclic quotient singularity $\frac{1}{n}(1,a)$. Since the action is Kähler, the quotient is an orbifold which admits a Kähler form. Depending on the context, we will focus more on the complex structure or the symplectic form. *Wahl singularities* are just the family of cyclic quotient singularities of the form $\frac{1}{p^2}(1, pq-1)$.

Remark 2.2.23. By a deep result [KS88], Wahl singularities are the cyclic quotient singularities which admit $\mathbb{Q}$-Gorenstein smoothings whose Milnor fibre is a rational homology ball. Such a Milnor fibre is the rational homology ball $B_{p,q}$ introduced above, and the (p,q)-pinwheel is the vanishing cycle of the smoothing.

Denote $\mathcal{O}_{p,q} := \mathbb{C}^2/_{1/p^2(1,pq-1)}$ and recall that its moment image $\Delta_{p,q}$ is the wedge bounded by the vectors $(0,1)$ and $(p^2, pq-1)$, i.e. the same as $\mathfrak{A}_{p,q}$ without the branch cut and the node. This is discussed, for example, in [Eva23, Chapter 3]. By analogy with the almost toric base diagram $\mathfrak{A}_{p,q}(\alpha,\beta)$, define the compact toric orbifold $\mathcal{O}_{p,q}(\alpha,\beta)$ via its moment image, which is shown in Figure 2.5 (a), for positive real numbers α and β. Since we will be interested mainly in uniqueness questions of Lagrangian pinwheels, we will fix for $\lambda > 0$ the "shape" $\Delta_{p,q}(\lambda,\lambda)$ in the moment image $\Delta_{p,q}$ and denote the corresponding compact orbifold by $\mathcal{B}_{p,q}(\lambda)$, whose moment image is shown in Figure 2.5 (b).

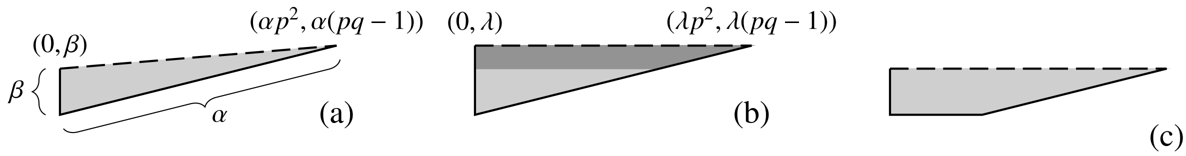

Figure 2.5: (a) The moment image $\Delta_{p,q}(\alpha,\beta)$ of the compact toric orbifold $\mathcal{O}_{p,q}(\alpha,\beta)$. (b) The moment image of $\mathcal{B}_{p,q}(\lambda)$ with a region used to define the neck-stretching setup shaded in darker gray. (c) The moment image for a choice of Kähler structure on $\widetilde{U}_{\mathrm{st}}$. All of the figures are illustrations of the situation in the $(p,q) = (2,1)$ case.

Cyclic quotient surface singularities admit a unique *minimal resolution*, meaning that any other resolution is obtained from the minimal one via a sequence of complex blow-ups. To describe the minimal resolution of $\mathbb{C}^2/_{1/n(1,a)}$, it is enough to prescribe the self-intersections of the spheres in the chain of exceptional spheres C introduced by the resolution. A minimal

[12] One could also show that the usual Flux arguments for Lagrangian isotopies carry through, thus only using that $H^1(L_{p,q}, \mathbb{Q}) = 0$.

resolution is given by a chain of spheres $C = (C_1, \ldots, C_m)$ with self-intersections $C_i^2 = -b_i$, where $[b_1, \ldots, b_m]$ is the *Hirzebruch–Jung continued fraction expansion* (HJ continued fraction in the following) of n/a, i.e.

$$[b_1, \ldots, b_m] := b_1 - \cfrac{1}{b_2 - \cfrac{1}{\ddots - \frac{1}{b_m}}} = \frac{n}{a}.$$

In the case of Wahl singularities the HJ continued fraction of $p^2/(pq-1) = [b_1, \ldots, b_m]$ is called a *Wahl chain*. If the context is clear, we will also refer to the chain of rational curves appearing in the minimal resolution of a Wahl singularity as a *Wahl chain*. Before we move on to discuss the rational blow-up and the neck-stretching setup that we are going to use later, let us record some useful algebraic properties of HJ continued fractions. Recall that the *dual* HJ continued fraction of $\frac{p}{q} = [x_1, \ldots, x_r]$ is given by $\frac{p}{p-q} = [y_1, \ldots, y_s]$.

Lemma 2.2.24 ([UZ25, Proposition-Definition 2.3]). *Assume that* $\frac{p}{q} = [x_1, \ldots, x_r]$ *and* $\frac{p}{p-q} = [y_1, \ldots, y_s]$. *Denote by* $[q]^{-1}$ *the unique integer* $0 < [q]^{-1} < p$ *that satisfies* $[q]^{-1}q \equiv 1 \mod p$.

i) We have $[x_1, \ldots, x_r, 1, y_s, \ldots, y_1] = 0$ *and* $\frac{p}{[q]^{-1}} = [x_r, \ldots, x_1]$.

ii) The Wahl chain, as well as its dual, can be computed from the HJ continued fractions of $\frac{p}{q}$ *and* $\frac{p}{p-q}$ *via*

$$\frac{p^2}{pq-1} = [x_1, \ldots, x_{r-1}, x_r + y_s, y_{s-1}, \ldots, y_1] \tag{2.2.2}$$

and

$$\frac{p^2}{p(p-q)+1} = [y_1, \ldots, y_s, 2, x_r, \ldots, x_1]. \tag{2.2.3}$$

iii) We have

$$\begin{pmatrix} p & -[q]^{-1} \\ q & \frac{1-q[q]^{-1}}{p} \end{pmatrix} = \begin{pmatrix} x_1 & -1 \\ 1 & 0 \end{pmatrix} \cdots \begin{pmatrix} x_r & -1 \\ 1 & 0 \end{pmatrix}. \tag{2.2.4}$$

Since we will mainly be interested in resolutions of $\mathcal{B}_{p,q}(\lambda)$, let us define a resolution in this case rigorously.

Definition 2.2.25. A *chain-shaped resolution* of $\mathcal{B}_{p,q}(\lambda)$ is given by a tuple $((\widetilde{U}, \widetilde{J}), \pi)$, where $(\widetilde{U}, \widetilde{J})$ is a smooth complex manifold and $\pi \colon \widetilde{U} \to \mathcal{B}_{p,q}(\lambda)$ is a holomorphic map that satisfies the following property: $\widetilde{U}$ contains a linear chain of embedded rational complex curves $C = \{C_1, \ldots, C_m\}$ such that $\pi(C) = x$, where $x \in \mathcal{B}_{p,q}(\lambda)$ is the unique orbifold point, and the restriction

$$\widetilde{U} \setminus C \to \mathcal{B}_{p,q}(\lambda) \setminus \{x\}$$

is a biholomorphism.

Remark 2.2.26. Since all the resolutions that we will encounter are chain-shaped, a resolution is always understood to be chain-shaped in the following.

Example 2.2.27. The minimal resolution of $\mathcal{B}_{p,q}(\lambda)$ is just given by the minimal subdivision of the fan spanned by the inward normals to the two edges of $\Delta_{p,q}$. This minimal subdivision can be computed by the HJ continued fraction of $p^2/(pq-1) = [b_1, \dots, b_m]$ and determines the resolution as a normal toric variety, see [Ada+25, Definition 2.1.4]. We will denote this toric variety by $\widetilde{U}_{\mathrm{st}}$.

Choosing appropriate additional data consisting of areas for the exceptional spheres in the minimal resolution and a potential function, we can induce a Kähler structure on $\widetilde{U}_{\mathrm{st}}$. We will denote the corresponding Kähler form by $\widetilde{\omega}$. Moreover, the usual moment image of a minimal resolution of $\Delta_{p,q}(\lambda, \lambda)$ obtained by symplectic cuts, as shown in Figure 2.5 (c), is the moment image of the symplectic manifold $(\widetilde{U}_{\mathrm{st}}, \widetilde{\omega})$.[13]

We will now define the *rational blow-up* in a way that naturally incorporates the neck-stretching setup needed later. Let N be a neighbourhood of $\partial B_{p,q}(\lambda)$ corresponding to one in $\mathcal{B}_{p,q}(\lambda)$ as shown in Figure 2.5. Using such an N, we can define a complex structure J_N on N by pulling back the natural complex structure on $\mathcal{B}_{p,q}(\lambda) \subseteq \mathcal{O}_{p,q}$.

Definition 2.2.28. Assume that $\iota : B_{p,q}(\lambda) \hookrightarrow (X, \omega)$ is a symplectic embedding. Then define

$$\mathcal{J}_N := \{J \in \mathcal{J}_\tau(X, \omega) \mid J|_{\iota(N)} = \iota_* J_N\}$$

to be the space of all ω-tame almost complex structures on X that agree with $\iota_* J_N$ on $\iota(N)$.

In particular, the almost complex structures $J \in \mathcal{J}_N$ are adjusted to the neck N and therefore yield neck-stretching sequences of tame almost complex structures: starting with $J =: J_0 \in \mathcal{J}_N$ we obtain a neck-stretching sequence $\{J_t\}_t \subseteq \mathcal{J}_\tau(X, \omega)$.[14]

Assume that $\iota : B_{p,q}(\lambda) \hookrightarrow (X, \omega)$ is a symplectic embedding and denote $U := \iota(B_{p,q}(\lambda))$ and $V := X \setminus U$. Moreover, let $\overline{U}$ and $\overline{V}$, respectively, be the symplectic completions of U and V. By considering the neck N and its completion along the concave end we obtain a complex manifold $(\overline{N}_-, \overline{J}_N)$ that is isomorphic to a punctured neighbourhood of the orbifold $\mathcal{O}_{p,q}$. Denote the orbifold obtained by adding the point at negative infinity to $\overline{V}$ by $(\widehat{X}, \widehat{J})$. Note that $(\widehat{N}, \widehat{J})$, which we define to be the orbifold obtained by adding the point at negative infinity to $(\overline{N}_-, \overline{J}_N)$, is a copy of $\mathcal{B}_{p,q}(\lambda)$ inside $(\widehat{X}, \widehat{J})$.

Definition 2.2.29 (Rational blow-up)**.** Given a symplectic embedding $\iota : B_{p,q}(\lambda) \hookrightarrow (X, \omega)$, a resolution $\widetilde{X}$ of $\widehat{X}$ is given by exchanging $\widehat{N} \subsetneq \widehat{X}$ with the resolution $\widetilde{U}_{\mathrm{st}}$ of $\mathcal{B}_{p,q}(\lambda)$. The symplectic manifold thereby obtained, $\widetilde{X} = V \cup \widetilde{U}_{\mathrm{st}}$, is called the *rational blow-up* of X along ι.

Remark 2.2.30. The symplectic form $\widetilde{\omega}$ on the rational blow-up $\widetilde{X}$ coincides with $\omega|_V$ on V and with $\widetilde{\omega}$ in $\widetilde{U}_{\mathrm{st}}$. Moreover, this procedure also respects the almost complex geometry in the sense that $\widetilde{X}$ carries an almost complex structure $\widetilde{J}$ that coincides with $J|_V$ on V and agrees with the integrable complex structure on $\widetilde{U}_{\mathrm{st}}$.

The setup we are now in is summarized in the following figure, taken from [Ada+25].

The upshot of this setup is the following lemma for pseudoholomorphic curves in the three manifolds $(\overline{V}, \overline{J})$, $(\widehat{X}, \widehat{J})$ and $(\widetilde{X}, \widetilde{J})$.

Lemma 2.2.31 ([Ada+25, Lemma 3.1.7/3.1.8])**.** *There are bijections between the following collections of objects: (a) irreducible finite-energy punctured $\overline{J}$-holomorphic curves C in $\overline{V}$; (b) irreducible orbifold $\widehat{J}$-holomorphic curves $\widehat{C}$ in $\widehat{X}$; and (c) irreducible $\widetilde{J}$-holomorphic curves $\widetilde{C}$ in $\widetilde{X}$ other than $C_1, \dots, C_m$. Moreover:*

[13]Of course the moment image will depend on the choices of data on $\widetilde{U}_{\mathrm{st}}$.

[14]For details on this, see [Bou+03; CM05; EGH00]. Neck-stretching along the contact boundary of normal neighbourhoods of Lagrangian pinwheels is discussed in [Fer16; Ada+25; ES18].

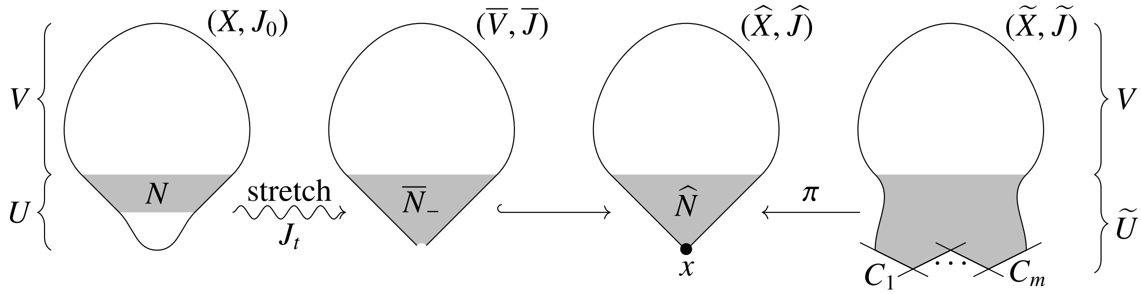

Figure 2.6: The relations between the spaces defined in Section 2.2.4.

i) Any irreducible $\widetilde{J}$-holomorphic curve in $\widetilde{X}$ which is not one of the C_i must enter the interior of V.

ii) By choosing J generically on V, we can ensure that all curves that enter V are regular so that the only non-regular $\widetilde{J}$-holomorphic curves are amongst the $C_1, \ldots, C_m$.

2.2.5 The topology of the rational blow-up

To understand the topology of $\widetilde{X}$, we want to leverage the fact that the boundary of $B_{p,q}(\lambda)$ is given by a rational homology sphere, since it is diffeomorphic to the lens space $L(p^2, pq-1)$. Therefore, the rational homology of $\widetilde{X}$ splits into two parts: the rational homology of $\widetilde{U}_{\mathrm{st}}$ and the rational homology of V. We will mostly do this in order to compute intersection numbers of curves in $\widetilde{X}$ by computing them locally in $\widetilde{U}_{\mathrm{st}}$ and V and to express the Chern class of $\widetilde{X}$ in terms of the Chern classes of $\widetilde{U}_{\mathrm{st}}$ and V. Given a class $A \in H_2(\widetilde{X}; \mathbb{Z})$, let $A_{\mathbb{Q}}$ be the corresponding class in $H_2(\widetilde{X}; \mathbb{Q})$.

Lemma 2.2.32 ([Ada+25, Lemma 3.2.2])**.** *Let $\widetilde{X}$ be the rational blow-up along a symplectic embedding $\iota : B_{p,q}(\lambda) \hookrightarrow (X, \omega)$. Given a class $A \in H_2(\widetilde{X}; \mathbb{Z})$ there exist rational numbers $a_1, \ldots, a_m$ and a class $A_X \in H_2(X; \mathbb{Z})$ such that*

$$A_{\mathbb{Q}} = \frac{1}{p} A_X + \sum_{j=1}^{m} a_j [C_j]. \tag{2.2.5}$$

Moreover, the integral class pA_X can be represented by a cycle contained in V.

The point of the lemma is that we can now determine the numbers a_j just by knowing how A intersects the classes C_i. Indeed, let $\tau_j = A \cdot C_j$ and consider the tuples $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_m)$ and $\boldsymbol{a} = (a_1, \ldots, a_m)$. Define the intersection matrix of the Wahl chain by $M_{i,j} = C_i \cdot C_j$. By construction, we have that $\boldsymbol{\tau} = M\boldsymbol{a}$. Therefore to determine the tuple $\boldsymbol{a}$, it is enough to compute M^{-1}, if we know the intersection profile of A with the Wahl chain.

In the sequel it will be crucial not only to know the inverse of the intersection matrix of a Wahl chain, but also the inverse of a chain of exceptional curves introduced during the minimal resolution of a general cyclic quotient singularity.

Lemma 2.2.33 ([Ada+25, Lemma 3.2.4])**.** *Suppose that $0 < a < n$ are coprime integers and that the minimal resolution of the cyclic quotient singularity $\frac{1}{n}(1, a)$ is governed by the HJ continued fraction $\frac{n}{a} = [b_1, \ldots, b_m]$. Consider the intersection matrix $M_{i,j} = C_i \cdot C_j$ of the* *associated chain-shaped resolution, where $C_i^2 = -b_i$. Then M is invertible and the entries of its inverse M^{-1} are given by*

$$M_{ij}^{-1} = \begin{cases} -(e_i f_j)/n & \text{if } i \leq j \\ -(e_j f_i)/n & \text{if } i > j \end{cases},$$

where the integers e_i and f_i are defined via the negative continued fractions $e_i/e_{i-1} = [b_{i-1}, \ldots, b_1]$ and $f_i/f_{i+1} = [b_{i+1}, \ldots, b_m]$. We will call the numbers e_i and f_i left and right "accompanying numbers" in the following.[15]

Remark 2.2.34. [Ada+25, Lemma 3.2.4] is stated only for Wahl chains, but the proof does not depend on the HJ continued fraction being Wahl. Recall from the proof of [Ada+25, Lemma 3.2.4] that these accompanying numbers satisfy the recursive formulas:

$$e_{i+1} = b_i e_i - e_{i-1} \qquad \text{and} \qquad f_{i+1} = b_i f_i - f_{i-1} \tag{2.2.6}$$

and the identities:

$$e_{i+1} f_i - e_i f_{i+1} = n \qquad \text{and} \qquad e_i f_i - e_{i-1} f_{i+1} = a. \tag{2.2.7}$$

The initial conditions for the recursions are

$$e_0 = 0, \qquad e_1 = 1, \qquad f_0 = n, \qquad f_1 = a. \tag{2.2.8}$$

It follows from the definition of the accompanying numbers and the recursive formula that the sequence $(e_i)_i$ is increasing while the sequence $(f_i)_i$ is decreasing, and that all consecutive accompanying numbers are coprime. Moreover, the extremal values of the accompanying numbers are given by

$$e_m = [a]^{-1}, \qquad e_{m+1} = n, \qquad f_m = 1, \qquad f_{m+1} = 0, \tag{2.2.9}$$

where $0 < [a]^{-1} < n$ is the unique integer satisfying $[a]^{-1} a \equiv 1 \mod n$.

The accompanying numbers are also very useful in determining the Chern class of a small neighbourhood of the chain of spheres C in the resolution of a singularity. Let $\widetilde{U}$ be a neighbourhood of C. By Alexander–Lefschetz duality we obtain $H^2(\widetilde{U}; \mathbb{Q}) \cong H_2(\widetilde{U}; \mathbb{Q})$, which implies that we can write the canonical class $K_{\widetilde{U}}$ of $\widetilde{U}$ as

$$K_{\widetilde{U}} = \sum_{j=1}^{m} k_j C_j,$$

where the C_j are classes Poincaré dual to the corresponding exceptional curves.

Definition 2.2.35 (Discrepancies)**.** The coefficients k_j are called the *discrepancies* of the singularity.

Lemma 2.2.36 ([HTU17, Section 2.1.])**.** *Suppose that $0 < a < n$ are two coprime integers. The discrepancies of a cyclic quotient singularity $\frac{1}{n}(1, a)$ can be computed from its accompanying numbers via*

$$k_j = -1 + \frac{e_j + f_j}{n}. \tag{2.2.10}$$

[15] Here "left" and "right" should indicate that the e-sequence is defined from the prefixes, whereas the f-sequence is defined from the suffixes of $[b_1, \ldots, b_m]$.

Moreover, the discrepancies of Wahl singularities, i.e. cyclic quotient singularities of the form $\frac{1}{p^2}(1, pq-1)$*, satisfy* $k_j \in (-1, 0]$.

Remark 2.2.37 ([ES20, Section 5.2]). Given a symplectic embedding $\iota : B_{p,q}(\lambda) \hookrightarrow X$, denote the rational blow-up along ι by $\widetilde{X} = V \cup \widetilde{U}_{\mathrm{st}}$, where $V = X \setminus \iota(B_{p,q}(\lambda))$, as before. Mayer–Vietoris yields that $H^2(\widetilde{X}, \mathbb{Q}) \cong H^2(V, \mathbb{Q}) \oplus H^2(\widetilde{U}_{\mathrm{st}}, \mathbb{Q})$, since $\widetilde{U}_{\mathrm{st}} \cap V$ is just the boundary lens space $\Sigma_{p,q}$ of $B_{p,q}(\lambda)$, which is a rational homology sphere. By naturality of the Chern class, we have that the Chern class of $\widetilde{X}$ splits as the sum of the Chern classes of $\widetilde{U}_{\mathrm{st}}$ and V. Calculating the discrepancies k_j amounts precisely to calculating the Chern class of $\widetilde{U}_{\mathrm{st}}$.

2.3 Part 1: Uniqueness up to symplectomorphism

In Section 2.3.1 we introduce our compactification $X_{p,q}$ of $B_{p,q}$ and explain how the central cylinder $K_{p,q}$ becomes a fibre class under the compactification procedure, as well as how $X_{p,q}$ is birationally derived from a Hirzebruch surface. In Section 2.3.2 we prove the existence of a regulation of $X_{p,q}$ (Figure 2.8). We then show how this regulation "persists" under rationally blowing up $X_{p,q}$ along an embedding $\iota : B_{p,q}(\lambda) \hookrightarrow X$ (Theorem 2.3.21). Finally, in Section 2.3.3 we conclude the proof of Theorem B.

2.3.1 Compactifying $B_{p,q}$

We start by explaining a compactification of the rational homology ball $B_{p,q}$. There is a lot of freedom in doing so, and we present the compactification most suited for our purposes.[16]

Example 2.3.1. It is well known that $B_{2,1} \cong T^*\mathbb{R}P^2$ admits a compactification to $\mathbb{C}P^2$, as discussed in [Ada25; Par+18]; that[17] $B_{2;1,1} \cong T^*S^2$ compactifies to $S^2 \times S^2$, see [Hin04; LW12]; and that the family $B_{n,1}$ admits compactifications to rational and ruled surfaces, as explained in [AH25; Sym01].

Given the almost toric base diagram $\mathfrak{A}_{p,q}$ of $B_{p,q}$, as shown in Figure 2.2, for a positive real number $\alpha > 0$ let $X^{\mathrm{orb}}_{p,q}(\alpha)$ be the possibly singular symplectic manifold obtained by performing a horizontal symplectic cut at height α, capping off the almost toric base diagram to a closed triangle. Notice that the left upper corner is Delzant while the upper right corner is non-Delzant if $(p, q) \neq (2, 1)$. Let x_{orb} be the orbifold point defined by the upper right corner.

Lemma 2.3.2. *The orbifold point* x_{orb} *is equivalent to the quotient singularity* $\frac{1}{pq-1}(1, p^2)$.

Proof. After a reflection along the $x = y$ axis, the non-Delzant corner is equivalent to the quotient singularity defined by the vectors $(0, 1)$ and $(pq - 1, p^2)$, which is the claimed quotient singularity. □

Now, we would like to minimally resolve the singularity x_{orb} in order to obtain a smooth compactification of $B_{p,q}$. However, since $pq - 1 < p^2$ we cannot directly apply the usual algorithm that computes the Hirzebruch–Jung continued fraction in order to determine a minimal

[16] For example, different ones are considered in [BO12; Par+18; Buc25]. However, it is not hard to understand how to pass from one to another via blow-ups and blow-downs of appropriate divisors. The one in [Eva23, Section 9.3] is intimately related to ours and we will comment on this relation later.

[17] The space $B_{2;1,1}$ fits into a larger family $B_{d;p,q}$, which specializes to $B_{p,q}$ for $d = 1$. In line with Remark 2.2.23, these can be characterized as the Milnor fibres of those cyclic quotient singularities that admit $\mathbb{Q}$-Gorenstein smoothings. See [Eva23, Section 7] for details on this.

resolution, as described for example in [Eva23, Chapter 9]. So we need to further modify the singularity to be in normal form.

Lemma 2.3.3. *The quotient singularity x_{orb} is equivalent to $\frac{1}{pq-1}(1, p^2 - d_0(pq-1))$, where d_0 is the unique positive integer such that $(pq-1)d_0 < p^2 \leq (pq-1)(d_0+1)$. In particular, the thereby introduced compactifying divisor has intersection profile according to the mixed continued fraction*

$$\frac{p^2}{pq-1} = [-d_0; d_1, \dots, d_n] = d_0 + \frac{1}{d_1 - \frac{1}{\dots - \frac{1}{d_n}}},$$

i.e. the spheres $D_1, \dots, D_n$ introduced in the process have self-intersection $-d_1, \dots, -d_n$, while the sphere D_0 has self-intersection $+d_0$.

Proof. Let d_0 be as in the statement of the lemma and define $0 < l = p^2 - d_0(pq-1)$. Then, the pair of vectors $\{(0,1), (pq-1, p^2)\}$ is related to the pair of vectors $\{(0,1), (pq-1, l)\}$ by a shear along $(0,1)$. Therefore the quotient singularity x_{orb} is equivalent to the quotient singularity $\frac{1}{pq-1}(1, l)$. Since by construction $l \leq pq-1$ the usual algorithm to compute a minimal resolution can be applied to $\frac{pq-1}{l}$. This is exactly the result of

$$\frac{p^2}{pq-1} = d_0 + \frac{1}{\frac{pq-1}{l}}$$

after taking the negative continued fraction expansion of $\frac{pq-1}{l}$. □

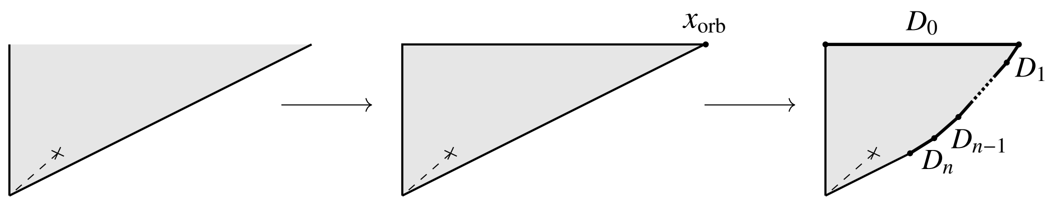

Figure 2.7: The steps taken to compactify the rational homology ball $B_{p,q}$ to $X_{p,q}$. This introduces the compactifying divisor $\mathcal{D}_{p,q} = (D_0, \dots, D_n)$.

Remark 2.3.4. Note that in the computation of the type of the quotient singularity in Lemma 2.3.3 we went through a transformation with determinant equal to -1, i.e. the reflection along the $x = y$ axis. One can avoid this and just apply a shear along the branch cut given by

$$\begin{pmatrix} 1 - pq & p^2 \\ -q^2 & 1 + pq \end{pmatrix}$$

in the middle figure of Figure 2.7. Then the horizontal edge will point in the $(pq-1, q^2)$-direction after applying the shear and therefore we see that x_{orb} is modelled on the quotient singularity $\frac{1}{pq-1}(1, q^2)$. Computing the Hirzebruch–Jung continued fraction, and therefore the minimal resolution of the singularity, gives the Hirzebruch–Jung continued fraction associated to the singularity $\frac{1}{pq-1}(1, p^2 - d_0(pq-1))$ in reverse order. This is because $0 < p^2 - d_0(pq-1) \leq pq - 1$ is the unique integer such that $q^2(p^2 - d_0(pq-1)) \equiv 1 \mod pq-1$ and hence the claim follows from Lemma 2.2.24.

We will be switching between these different viewpoints freely, as they highlight different aspects of the compactification. For example, viewing the orbifold point as one that is modelled on the $\frac{1}{pq-1}(1,q^2)$ singularity allows for quick and easy computations, while the mixed continued fraction introduced in Lemma 2.3.3 captures the structure of the compactification more transparently.

Definition 2.3.5. Define the compactification of the rational homology ball $B_{p,q}$ to be the closed symplectic manifold $\big(X_{p,q}(\alpha;\boldsymbol{\rho}),\omega(\alpha;\boldsymbol{\rho})\big)$ obtained by capping the almost toric base diagram $\mathfrak{A}_{p,q}$ at height $\alpha > 0$ and resolving the quotient singularity x_{orb}, giving the chain of spheres $D_0,\dots,D_n$ areas according to $\boldsymbol{\rho}\in\mathbb{R}^{n+1}_{>0}$, i.e. $\int_{D_i}\omega(\alpha;\boldsymbol{\rho})=\rho_i$.

Remark 2.3.6. From now on "compactifying $B_{p,q}$" will mean compactifying $B_{p,q}$ according to Definition 2.3.5 unless specified otherwise, and we will often drop all the decorations in the definition above, whenever it is understood that we just consider some, meaning a suitable, compactification. We will write $(X_{p,q},\omega)$ for a compactification as above and write $\mathcal{D}_{p,q} = \{D_0,\dots,D_n\}$ for the compactifying divisor.

Moreover, even though the horizontal cut on $B_{2,1}$ does not produce a non-Delzant corner, we adopt the convention to introduce a compactifying divisor of the form $\mathcal{D}_{2,1} = (D_0, D_1)$, where $D_0^2 = +3$ and $D_1^2 = -1$, such that the compactification $X_{2,1}$ fits into the family $X_{n,1}$.

Example 2.3.7. We will use two examples to illustrate the theorems and lemmata pictorially throughout the paper. We give the relevant associated integers here. If $(p,q) = (5,2)$ then the relevant Hirzebruch–Jung continued fraction and the corresponding d_0 are:

$$\frac{p^2}{pq-1} = \frac{25}{9} = [3,5,2], \qquad \frac{pq-1}{q^2} = \frac{9}{4} = [3,2,2,2] \quad \text{and} \qquad d_0 = c_1 - 1 = 2,$$

where $c_1 = 3$ is the first component of the Wahl chain. This means that the Wahl chain $C_{5,2}$ has intersection profile $(-3,-5,-2)$ and the compactifying divisor $\mathcal{D}_{5,2}$ has intersection profile $(+2,-2,-2,-2,-3)$.

For the family $(p,q) = (n,1)$ we have:

$$\frac{n^2}{n-1} = [n+2,\underbrace{2,\dots,2}_{(n-2)\text{-times}}], \qquad \frac{n-1}{1} = [n-1] \quad \text{and} \qquad d_0 = c_1 - 1 = n+1.$$

We collect some small facts about the compactification we constructed. The most crucial object is the embedded symplectic sphere that is obtained from the central cylinder $K_{p,q}$ under the compactification procedure, as we will show that this sphere is part of a well-behaved holomorphic foliation of $X_{p,q}$.

Lemma 2.3.8. *The central cylinder $K_{p,q}$, living over the toric boundary of $\mathfrak{A}_{p,q}$, corresponds in the compactification to an embedded symplectic sphere of square zero.*

Proof. After applying a shear along the branch cut to the middle figure of Figure 2.7, the $\frac{1}{pq-1}(1,q^2)$ singularity will sit on the bottom left edge of the ATF triangle. Then the first symplectic cut of the minimal resolution of the $\frac{1}{pq-1}(1,q^2)$ singularity is horizontal, meaning that the central cylinder yields a square zero embedded symplectic sphere after the compactification. □

Remark 2.3.9. In the following we will therefore refer to this square zero embedded symplectic sphere F as a "fibre" and call $F \in H_2(X_{p,q}; \mathbb{Z})$ the "fibre class". See Figure 2.8 for another representation of the ATF on the compactification $X_{p,q}$.

Moreover, we have that $F \cdot L_{p,q} \equiv q \mod p$, which is a direct consequence of the fact that the Poincaré dual of $c_1(B_{p,q})$ is represented by $K_{p,q}$ and that $c_1(B_{p,q}) = q \in H^2(B_{p,q}; \mathbb{Z}) \cong \mathbb{Z}_p$ as shown in [ES18, Section 2]. This can also be deduced from the way that F and $L_{p,q}$ intersect, as shown in Figure 2.8.

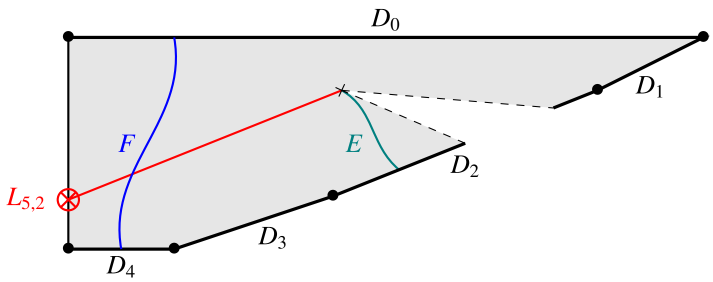

Figure 2.8: The compactification $X_{5,2}$. This almost toric base diagram is related to the one partially shown on the right in Figure 2.7 via a rotation of the branch cut. All the relevant objects of the compactification are included in the figure: (1) the compactifying divisor $\mathcal{D}_{5,2} = \{D_0, \ldots, D_4\}$; (2) the Lagrangian pinwheel $L_{5,2}$; (3) a fibre class F and (4) the distinguished exceptional divisor E. In Figure 2.9 it is shown how $X_{5,2}$ is related to the Hirzebruch surface $\mathbb{F}_2$.

Lemma 2.3.10. *If $q \neq 1$ there is a visible exceptional divisor E in the ATF defined on $X_{p,q}$ which connects the node to one of the edges appearing in the toric resolution of the $\frac{1}{pq-1}(1, q^2)$ singularity. We refer to this exceptional divisor as the distinguished exceptional class $E \in H_2(X_{p,q}; \mathbb{Z})$.*

Remark 2.3.11. The situation of the lemma is shown in the example $(p, q) = (5, 2)$ in Figure 2.8. Recall from Example 2.3.7 that in the case $(p, q) = (n, 1)$ the compactification is already a rational and ruled surface, because the compactification only introduces two compactifying divisors.

Proof of Lemma 2.3.10. Assume that $\frac{p}{q} = [x_1, \ldots, x_r]$ and $\frac{p}{p-q} = [y_1, \ldots, y_s]$. It is easy to see that a resolution of the form $[x_1, \ldots, x_r, 2, y_s, \ldots, y_1, 1]$ is a non-minimal resolution of the $\frac{1}{pq-1}(1, q^2)$ singularity.[18] This resolution covers the minimal one by consecutively blowing down (-1)-spheres, and this blow-down procedure will contract the whole string $[y_1, \ldots, y_s]$ if and only if it is a string of the form $[2, \ldots, 2]$, meaning that $(p, q) = (n, 1)$. Note that in this case $\frac{n}{1} = [n]$, and since the connecting (-2)-sphere is also blown down in the process we see that this yields exactly the compactification $X_{n,1}$, which is a rational and ruled surface. If not, the string

[18] This resolution is obtained by introducing a vertical symplectic cut in the middle step of Figure 2.7. Then the thereby introduced singularity is exactly the dual singularity to $\frac{p^2}{pq-1}$, which has a minimal resolution as discussed in equation (2.2.3). The compactification constructed in this manner is exactly the one used in [Buc25] and [Eva23, Figure 9.11]. The string is in the reverse order to the one given in (2.2.3) because of a reversal of orientation.

$[y_1, \ldots, y_s]$ contains a first entry y_j that is not equal to 2 and therefore

$$\frac{1}{pq-1}(1, q^2) = [x_1, \ldots, x_r, 2, y_s, \ldots, y_j - 1]. \tag{2.3.11}$$

This concludes the proof, because $\frac{p}{q} = [x_1, \ldots, x_r]$ implies that the edge associated to the connecting 2 in the HJ continued fraction of $\frac{1}{pq-1}(1, q^2)$ is pointing in the (p, q)-direction. □

Corollary 2.3.12. *The compactification $X_{p,q}$ is obtained from the standard toric model of $\mathbb{F}_{d_0}$ via a sequence of toric blow-ups and a non-toric blow-up of a distinguished edge, meaning that it is birationally derived from the Hirzebruch surface $\mathbb{F}_{d_0}$.*

Remark 2.3.13. This sequence of blow-ups/downs is shown in Figure 2.9 for the example $(p, q) = (5, 2)$.

Proof of Corollary 2.3.12. In the case $(n, 1)$ there is nothing to prove, since the compactification $X_{n,1}$ is already $\mathbb{F}_{d_0}$, as $d_0 = n + 1$ in this case according to Example 2.3.7. Hence, we assume that $(p, q) \neq (n, 1)$. In this case the minimal resolution of the singularity $\frac{1}{pq-1}(1, q^2)$ is given by

$$\frac{1}{pq-1}(1, q^2) = [x_1, \ldots, x_r, 2, y_s, \ldots, y_j - 1],$$

as deduced in (2.3.11). It follows from Lemma 2.3.10, see also Figure 2.8, that

$$[x_1, \ldots, x_r, 1, y_s, \ldots, y_j - 1] = 0,$$

i.e. that the HJ continued fraction is a zero continued fraction, and therefore there is a unique sequence of interior blow-downs to $[1, 1]$. Now contracting the (-1)-sphere that intersects the negative section will leave the positive section unchanged, and therefore the "minimal model" of $X_{p,q}$ is $\mathbb{F}_{d_0}$. □

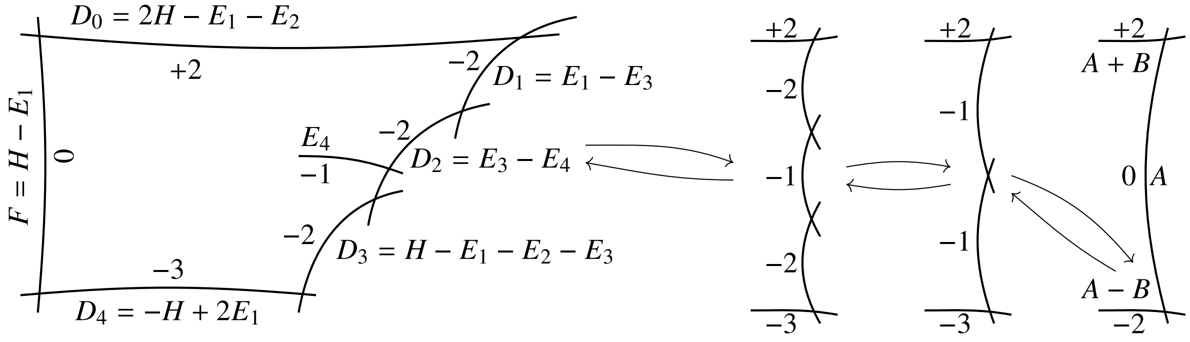

Figure 2.9: This figure shows how the compactification $X_{5,2}$ is birationally derived from the Hirzebruch surface $\mathbb{F}_2$ via a sequence of blow-ups. On the right we only display the blow-down sequence of the broken fibre. The homology classes of the divisors and the fibre on the left are easily calculated from the sequence of blow-ups, using the usual identification $S^2 \times S^2 \# \overline{\mathbb{C}P^2} \cong X_2$. In the language of Lemma 2.3.10 this means that E_4 is the distinguished exceptional class.

As a last point on the compactification $X_{p,q}$ we discuss some homological considerations that we will need later. Assume that $\iota : B_{p,q}(\varepsilon) \hookrightarrow B_{p,q}$ is a symplectic embedding. After choosing suitable compactification data, meaning that we compactify away from the embedding, we view ι as an embedding into the compactification X of $B_{p,q}$.

Rationally blowing up along ι yields a closed symplectic manifold $\widetilde{X}$ that contains the Wahl chain $C_{p,q}$ and the compactifying divisor $\mathcal{D}_{p,q}$. These two chains form a basis of $H_2(\widetilde{X};\mathbb{Q})$. Given any curve $\widetilde{C}$ in $\widetilde{X}$ we can express this curve in rational homology as:

$$\widetilde{C} = \sum_{i=0}^{n} a_i D_i + \sum_{j=1}^{m} c_j C_j \in H_2(\widetilde{X};\mathbb{Q}).$$

Define $(M_{\mathcal{D}})_{ij} = D_i \cdot D_j$, the intersection matrix of the compactifying divisor, and write $\xi_i := \widetilde{C} \cdot D_i$. This results in $\boldsymbol{\xi} = M_{\mathcal{D}}\boldsymbol{a}$. Recall from Definition 2.3.5 that the subchain $\mathcal{D}' := (D_1, \ldots, D_n)$ of the compactifying divisor corresponds to taking a minimal resolution of the cyclic quotient singularity $\frac{1}{pq-1}(1, [q^2]^{-1})$, where $0 < [q^2]^{-1} < pq - 1$ is the unique integer that satisfies $q^2[q^2]^{-1} \equiv 1 \mod pq - 1$. In particular, if we assume that $(pq-1)/[q^2]^{-1} = [d_1, \ldots, d_n]$ and consider the accompanying numbers g_i, h_j as defined in Lemma 2.2.33, where they are denoted by e_i, f_j, we obtain

$$(M_{\mathcal{D}'})^{-1}_{ij} = \begin{cases} -(g_i h_j)/(pq-1) \text{ if } i \le j \\ -(g_j h_i)/(pq-1) \text{ if } i > j \end{cases}. \tag{2.3.12}$$

Defining $\boldsymbol{e}_1 := (1, 0, \ldots, 0)$, the first canonical basis vector of length n, and $\boldsymbol{v}_{\mathcal{D}'} := (pq-1)(M_{\mathcal{D}'})^{-1}\boldsymbol{e}_1$, the first column of $(pq-1)(M_{\mathcal{D}'})^{-1}$, we can give a compact formula for $(M_{\mathcal{D}})^{-1}$ using Schur complements:

$$M_{\mathcal{D}} = \begin{pmatrix} d_0 & \boldsymbol{e}_1^T \\ \boldsymbol{e}_1 & M_{\mathcal{D}'} \end{pmatrix} \quad \text{and} \quad (M_{\mathcal{D}})^{-1} = \frac{1}{p^2}\begin{pmatrix} pq-1 & \boldsymbol{v}_{\mathcal{D}'}^T \\ \boldsymbol{v}_{\mathcal{D}'} & p^2(M_{\mathcal{D}'})^{-1} + \frac{1}{(pq-1)}\boldsymbol{v}_{\mathcal{D}'} \cdot \boldsymbol{v}_{\mathcal{D}'}^T \end{pmatrix}, \tag{2.3.13}$$

where d_0 was defined in Lemma 2.3.3.

2.3.2 Persistence of regulations

Recall the structure of the compactification $X_{p,q}$ of $B_{p,q}$ that we have constructed so far: it is birationally derived from a rational and ruled surface; there is a distinguished fibre class $F \in H_2(X;\mathbb{Z})$; the compactifying divisor $\mathcal{D}_{p,q}$ contains a positive section and a negative section; and if $\dim(H_2(X;\mathbb{Q})) > 2$, which means that $q \neq 1$, there is a distinguished exceptional divisor E, which intersects the compactifying divisor in a single point. This was proven in Corollary 2.3.12, Lemma 2.3.8, Lemma 2.3.10 and the setup is shown in Figure 2.8.

We now move on to showing that F is part of a well-behaved foliation of X, which is constructed by holomorphic-curves techniques. Recall the language of "regulations" and "rulings" from [Ada+25, Section 4.1]: an almost complex manifold (Y, J) is said to admit a J-holomorphic regulation in the class $A \in H_2(Y;\mathbb{Z})$ if A is a fibre class, i.e. has self-intersection equal to zero, and pairs with the first Chern class of Y according to the adjunction formula, and the evaluation map $\mathrm{ev} : \overline{\mathcal{M}}_{0,1}(A, J) \to Y$ has degree one. Moreover, curves in $\overline{\mathcal{M}}_{0,0}(A, J)$ are referred to as rulings and they are called "broken" if they belong to $\partial\overline{\mathcal{M}}_{0,1}(A, J)$ and "smooth" otherwise. Define the space of almost complex structures

$$\mathcal{J}_\tau(\mathcal{D}) := \{J \in \mathcal{J}_\tau(X, \omega) \mid \text{all irreducible components of } \mathcal{D} \text{ are } J\text{-holomorphic}\},$$

where $\mathcal{J}_\tau(X, \omega)$ denotes the space of tame almost complex structures on (X, ω). This space is non-empty since $\mathcal{D}$ is a normal crossing divisor, and it is weakly contractible [Eva11, Appendix]. The following two lemmas were first proven by Buck [Buc25, Section 2.4] and our proofs are inspired by Buck's proofs. In our case we can give simplified proofs of the lemmas, since our setup is significantly simpler.

Lemma 2.3.14. *If* $\dim(H_2(X; \mathbb{Q})) > 2$*, i.e.* $q \neq 1$*, then the distinguished exceptional class* E *is* J*-holomorphically represented by a unique smooth rational curve, which we denote by* E_J*, for any* $J \in \mathcal{J}_\tau(\mathcal{D})$*.*

Proof. The proof proceeds in two steps: we first show that any curve in the moduli space $\mathcal{M}_{0,0}(E, J)$ is smooth, i.e. $\overline{\mathcal{M}}_{0,0}(E, J) = \mathcal{M}_{0,0}(E, J)$, and then conclude that $\mathcal{M}_{0,0}(E, J)$ is non-empty for all $J \in \mathcal{J}_\tau(\mathcal{D})$.

For the first step suppose that $J \in \mathcal{J}_\tau(\mathcal{D})$ and that $S \in \overline{\mathcal{M}}_{0,0}(E, J)$ is a stable curve with irreducible components $S_1, \dots, S_N$. Since $X \setminus \mathcal{D}$ is a Liouville domain, no component of S can be completely contained in $X \setminus \mathcal{D}$ and, furthermore, since E intersects $\mathcal{D}$ exactly once by positivity of intersections, we must have that $[S_j] = E$ for one j. This immediately implies $S = S_j$.

Now denote by $\mathcal{J}_\tau(\mathcal{D}, E) \subseteq \mathcal{J}_\tau(\mathcal{D})$ the subset of all almost complex structures such that E is J-holomorphically represented. We know by Lemma 2.3.10 that this subset is non-empty, and applying automatic transversality, see [HLS97] or [Wen18, Chapter 2], we obtain that each $\mathcal{M}_{0,0}(E, J)$ is a compact smooth manifold of dimension 0 for each $J \in \mathcal{J}_\tau(\mathcal{D}, E)$ and that $\mathcal{J}_\tau(\mathcal{D}, E)$ is open. The automatic transversality criterion and dimension statement follow from a direct computation:

$$\operatorname{ind}(u) = -2 + 2c_1(E) = 0 > -2,$$

for $u \in \mathcal{M}_{0,0}(E, J)$. Moreover, the first argument shows that $\mathcal{J}_\tau(\mathcal{D}, E)$ is also closed. Indeed, assume that $(J_\nu)_\nu \subseteq \mathcal{J}_\tau(\mathcal{D}, E)$ is a sequence converging to $J \in \mathcal{J}_\tau(\mathcal{D})$. Then for each ν there is a unique embedded curve $u_\nu \in \mathcal{M}_{0,0}(E, J_\nu)$. Gromov compactness implies that there is a subsequence of $(u_\nu)_\nu$ converging to a stable curve $u_\infty \in \overline{\mathcal{M}}_{0,0}(E, J)$. Consequently, $J \in \mathcal{J}_\tau(\mathcal{D}, E)$ and connectedness of $\mathcal{J}_\tau(\mathcal{D})$ then implies the lemma, i.e. $\mathcal{J}_\tau(\mathcal{D}) = \mathcal{J}_\tau(\mathcal{D}, E)$. □

Lemma 2.3.15. *For any tame almost complex structure* $J \in \mathcal{J}_\tau(\mathcal{D})$*, the compactification* X *admits a non-degenerate* J*-holomorphic regulation in the class* F *with at most one broken ruling. There exists a broken ruling if and only if* $\dim(H_2(X; \mathbb{Q})) > 2$*. Moreover, if it exists it is given by* $\mathcal{T}_J = (\mathcal{D} \setminus \{D_0, D_n\}) \cup E_J$*, i.e. the compactifying divisor without the positive and negative section and the exceptional divisor representing the distinguished class* E*.*

Proof. The proof is very similar to that of the previous lemma and is easy if there is no distinguished exceptional class; therefore, we omit this case. Assume that a distinguished exceptional class exists. Fix a point $x \in D_0 \setminus D_1$, i.e. a point on the positive section away from the intersection of D_0 with the chain of negative spheres $\mathcal{D}' = \{D_1, \dots, D_n\}$. Define

$$\mathcal{J}_\tau(\mathcal{D}, F, x) := \{J \in \mathcal{J}_\tau(\mathcal{D}) \mid \exists \text{ an embedded } J\text{-holomorphic sphere in the class } F \text{ through } x\}.$$

By Lemma 2.3.8, see also Figure 2.8, this space is non-empty and it is open by automatic transversality, see again [HLS97] or [Wen18, Chapter 2], because for $u \in \mathcal{M}_{0,0}(F, J)$ we have

$$\operatorname{ind}(u) = -2 + 2c_1(F) = 2 > 0.$$

We now want to show that $\mathcal{J}_\tau(\mathcal{D}, F, x)$ is closed. To do so, assume that $(J_\nu)_\nu \subseteq \mathcal{J}_\tau(\mathcal{D}, F, x)$ is a sequence of almost complex structures converging to $J \in \mathcal{J}_\tau(\mathcal{D})$. Then consider for all ν the unique (by positivity of intersection) element $u_\nu \in \mathcal{M}_{0,0}(F, J_\nu)$ passing through x. By Gromov compactness, there is a subsequence of $(u_\nu)_\nu$ that Gromov-converges to a stable J-holomorphic curve $u_\infty \in \overline{\mathcal{M}}_{0,0}(F, J)$. Now write S for the union of all irreducible components of u_∞ that are not contained in $\mathcal{D}' \cup E_J$. Then we can write

$$[F] = [S] + a[E_J] + \sum_{i=1}^{n} b_i[D_i], \tag{2.3.14}$$

where a and the b_i are the covering multiplicities. Intersecting (2.3.14) with D_0 yields $1 = D_0 \cdot S + b_1$, but we know that x is on some component of S, which implies $D_0 \cdot S \geq 1$ and so we obtain $b_1 = 0$ and $D_0 \cdot S = 1$. Now intersect (2.3.14) with $\mathcal{D}' = \{D_1, \dots, D_n\}$. Writing this as a system of equations we have

$$\boldsymbol{e}_n = \boldsymbol{w} + a\boldsymbol{e}_r + M_{\mathcal{D}'}\boldsymbol{b}, \tag{2.3.15}$$

where $\boldsymbol{e}_j$ is the jth canonical basis vector, $\boldsymbol{w} = S \cdot \mathcal{D}'$, r is the unique index associated to the intersection between the distinguished exceptional class E and $\mathcal{D}'$; and $M_{\mathcal{D}'}$ is the intersection matrix of $\mathcal{D}'$. Then (2.3.15) is equivalent to

$$\boldsymbol{b} = (M_{\mathcal{D}'})^{-1}(\boldsymbol{e}_n - \boldsymbol{w} - a\boldsymbol{e}_r). \tag{2.3.16}$$

Using the formula for the inverse of $M_{\mathcal{D}'}$ given in (2.3.12) and that $b_1 = 0$, the first row of this equation is given by

$$\sum_{i=1}^{n} h_i w_i + a h_r = 1, \tag{2.3.17}$$

where the h_i are the right accompanying numbers of the $\frac{1}{pq-1}(1, [q^2]^{-1})$ singularity, meaning in particular that $h_n = 1$ and $h_i > 1$ for $i < n$. Because $r < n - 1$ by Lemma 2.3.10, which proves that the distinguished exceptional class does not intersect $\mathcal{D}'$ at one of its ends, we see that (2.3.17) forces $\boldsymbol{w} = \boldsymbol{e}_n$ as well as $a = 0$. This implies that (2.3.15) reads

$$M_{\mathcal{D}'}\boldsymbol{b} = \boldsymbol{0},$$

which means $\boldsymbol{b} = \boldsymbol{0}$, i.e. $[F] = [S]$.

What is left to prove is that S contains exactly one component. As before, every component in S has to intersect $\mathcal{D}$, as otherwise it would be contained in $B_{p,q}$, which is Liouville. Therefore, by positivity of intersection, there are at most two components of S because $F \cdot D_0 = 1$ and $F \cdot D_n = 1$. So assume that there are two components and pick the component A that is intersecting D_0. Then we can write $A = \sum_{i=1}^{n} a_i[D_i]$ for some rational coefficients a_i, and because the intersection profile of A with $\mathcal{D}$ is given by $(1, 0, \ldots, 0)$ we obtain $\boldsymbol{e}_1 = M_{\mathcal{D}}\boldsymbol{a}$.[19] The self-intersection of the rational curve A is therefore given by

$$A \cdot A = \boldsymbol{a}^T M_{\mathcal{D}} \boldsymbol{a} = a_1.$$

We can compute a_1 via $\boldsymbol{a} = M_{\mathcal{D}}^{-1}\boldsymbol{e}_1$, which yields $a_1 = \frac{pq-1}{p^2}$ by (2.3.13), an obvious contradiction. Hence the limit curve u_∞ is an embedded rational curve passing through x and $\mathcal{J}_\tau(\mathcal{D}, F, x)$ is closed. The connectedness of $\mathcal{J}_\tau(\mathcal{D})$ therefore implies $\mathcal{J}_\tau(\mathcal{D}, F, x) = \mathcal{J}_\tau(\mathcal{D})$ as desired.

By [Ada+25, Proposition 4.1.4] this means that for all $J \in \mathcal{J}_\tau(\mathcal{D})$ the compactification X admits a non-degenerate J-holomorphic regulation in the class F. Also, since the point $x \in D_0 \setminus D_1$ was arbitrary, this means that through every point $x \in D_0 \setminus D_1$ and for every $J \in \mathcal{J}_\tau(\mathcal{D})$ there is a smooth J-holomorphic ruling passing through x.

Moreover, by [Ada+25, Proposition 4.1.4] we know that for $J \in \mathcal{J}_\tau(\mathcal{D})$ there is a ruling passing through $x \in D_0 \cap D_1$. But then inductively we can show that this ruling has to contain all of $(\mathcal{D}' \setminus D_n) \cup E_J$: it has to contain D_1, because $D_1 \cdot F = 0$ and then this reasoning propagates. This means that this ruling $\mathcal{T}_J$ is broken and is given by $(\mathcal{D}' \setminus D_n) \cup E_J$. □

With these two lemmas in place we now proceed to study the behaviour of the regulation after rationally blowing up along an embedding $\iota : B_{p,q}(\varepsilon) \hookrightarrow X$. The precise description of how the blown-up fibre class interacts with the compactifying divisor and Wahl chain is stated in Theorem 2.3.21. As in Section 2.2.4, denote $U = \iota(B_{p,q}(\varepsilon)) \subseteq X$ and $V = X \setminus U$, and consider a neck-stretching sequence of almost complex structures $\{J_t\}_t$ such that $J_0 \in \mathcal{J}_N(\mathcal{D})$, where $\mathcal{J}_N(\mathcal{D}) \subseteq \mathcal{J}_\tau(\mathcal{D})$ denotes the subspace of almost complex structures that are standard on a neck of the embedded ball $B_{p,q}(\varepsilon)$, as defined in Definition 2.2.28.

Lemma 2.3.16 ([Ada+25, Lemma 5.4.1])**.** *Fix a point x on the positive section $D_0 \subseteq X$ and for every t consider the unique J_t-holomorphic curve F_t in the class $F \in H_2(X;\mathbb{Z})$ passing through x. For a subsequence $t_j \to \infty$ the sequence F_{t_j} converges to a limit building with components in both $\overline{U}$ and $\overline{V}$.*

Proof. The existence of the subsequence that yields a holomorphic limit building follows from the SFT compactness theorem [Bou+03; CM05]. That the completions contain components follows from the fact that each F_t has non-trivial intersection with the divisor D_0 and the Lagrangian pinwheel $\iota(L_{p,q})$, as discussed in Remark 2.3.9. □

Now consider the component C of the holomorphic limit building that is contained in $\overline{V}$ and intersects the divisor D_0. We will continue investigating this component in the following lemmas and it should always be understood that C is exactly this component. Rationally blowing up along the embedding ι yields a symplectic manifold $\widetilde{X}$. We denote the curve that corresponds (cf. Lemma 2.2.31) to C under this procedure by $\widetilde{C}$. The Wahl chain $C_{p,q}$ and the compactifying divisor $\mathcal{D}_{p,q}$ form a basis of $H_2(\widetilde{X};\mathbb{Q})$. We express $\widetilde{C}$ as well as the canonical class $K = -\mathrm{PD}(c_1(\widetilde{X}))$ of $\widetilde{X}$ in this basis:

$$K = -F - \sum_{i=0}^{n} D_i + \sum_{j=1}^{m} k_j C_j, \qquad \widetilde{C} = \sum_{i=0}^{n} a_i D_i + \sum_{j=1}^{m} c_j C_j, \tag{2.3.18}$$

[19] Note that if S had one component the intersection profile would be $(1, 0, \ldots, 0, 1)$.

where $\boldsymbol{k} = (k_1, \dots, k_m)$ are the discrepancies.[20]

Recall that $(M_C)_{ij} = C_i \cdot C_j$ denotes the intersection matrix of the Wahl chain and that defining $\chi_j := \widetilde{C} \cdot C_j$ yields the equation $\boldsymbol{\chi} = M_C \boldsymbol{c}$. Analogously we write $(M_{\mathcal{D}})_{ij} = D_i \cdot D_j$ for the intersection matrix of the compactifying divisor, and writing $\xi_i := \widetilde{C} \cdot D_i$ means $\boldsymbol{\xi} = M_{\mathcal{D}} \boldsymbol{a}$. The two pairings that will play a role in the following are the self-intersection of $\widetilde{C}$

$$\widetilde{C} \cdot \widetilde{C} = \boldsymbol{\xi}^T M_{\mathcal{D}}^{-1} \boldsymbol{\xi} + \boldsymbol{\chi}^T M_C^{-1} \boldsymbol{\chi} \tag{2.3.19}$$

and the pairing of $\widetilde{C}$ with the canonical class K

$$K \cdot \widetilde{C} = -a_0 - a_n - \mathbf{1}^T \boldsymbol{\xi} + \boldsymbol{k}^T \boldsymbol{\chi}. \tag{2.3.20}$$

These two equations immediately follow from the fact that $C_{p,q}$ and $\mathcal{D}_{p,q}$ are disjoint and that F intersects $\mathcal{D}_{p,q}$ once positively at each end.

Lemma 2.3.17 ([Ada+25, Proposition 5.4.3])**.** *The intersection profile of $\widetilde{C}$ with $\mathcal{D}$ is given by $\boldsymbol{\xi} = (1, 0, \dots, 0)$. Moreover, $\widetilde{C}$ is an embedded curve of self-intersection $\widetilde{C}^2 \leq 0$.*

Proof. Assume for the sake of a contradiction that the intersection profile with the compactifying divisor $\mathcal{D}$ is given by $\boldsymbol{\xi} = (1, 0 \dots, 0, 1)$.[21] Then (2.3.19) reads $\widetilde{C} \cdot \widetilde{C} = a_0 + a_n + \boldsymbol{\chi}^T M_C^{-1} \boldsymbol{\chi}$ and (2.3.20) reads $K \cdot \widetilde{C} = -a_0 - a_n - 2 + \boldsymbol{k}^T \boldsymbol{\chi}$. Plugging these two equations into the adjunction formula we obtain

$$\sum_{x \in \mathrm{Sing}(\widetilde{C})} \delta_x = \frac{1}{2}(\widetilde{C} \cdot \widetilde{C} + K \cdot \widetilde{C}) + 1 = \frac{1}{2}(\boldsymbol{k}^T \boldsymbol{\chi} + \boldsymbol{\chi}^T M_C^{-1} \boldsymbol{\chi}).$$

Since all the δ_x contribute positively and the terms in the bracket on the right-hand side are non-positive[22], this means that $\boldsymbol{\chi} = \mathbf{0}$, which is a contradiction since this would imply that $\widetilde{C}$ is an embedded holomorphic sphere that corresponds to a holomorphic representative of F in X that does not intersect the class of the pinwheel defined via the embedding ι.

Since we picked $\widetilde{C}$ to be the component in the limit building that intersects D_0 we therefore have $\boldsymbol{\xi} = (1, 0, \dots, 0)$. This implies that (2.3.19) becomes $\widetilde{C} \cdot \widetilde{C} = a_0 + \boldsymbol{\chi}^T M_C^{-1} \boldsymbol{\chi}$ and (2.3.20) becomes $K \cdot \widetilde{C} = -a_0 - a_n - 1 + \boldsymbol{k}^T \boldsymbol{\chi}$. Consider the adjunction formula for $\widetilde{C}$:

$$\sum_{x \in \mathrm{Sing}(\widetilde{C})} \delta_x = \frac{1}{2}(\widetilde{C} \cdot \widetilde{C} + K \cdot \widetilde{C}) + 1 = \frac{1}{2}(-a_n - 1 + \boldsymbol{k}^T \boldsymbol{\chi} + \boldsymbol{\chi}^T M_C^{-1} \boldsymbol{\chi}) + 1.$$

From (2.3.13) it immediately follows that $a_n = 1/p^2$ and therefore $-a_n - 1 = -(p^2 + 1)/p^2$. As before, the term in the bracket is negative, which implies that $\mathrm{Sing}(\widetilde{C})$ is empty, because the δ_x contribute positively. In particular, $\widetilde{C}$ is a smoothly embedded sphere. The claim about

[20] The formula for the canonical class K of $\widetilde{X}$ follows from the fact that the canonical class K_X of X is given by $K_X = -F - \sum_{i=0}^n D_i$, since X is toric and the toric boundary is Poincaré dual to the first Chern class. Moreover, the splitting of the first Chern class is discussed in Remark 2.2.37.

[21] Recall that $\widetilde{C}$ is the component of the limit building that intersects D_0, therefore this intersection is forced and hence the two possible intersection profiles are $(1, 0 \dots, 0)$ and $(1, 0 \dots, 0, 1)$.

[22] Recall that the discrepancies can be computed via $k_j = -1 + \frac{e_j + f_j}{p^2}$, see Lemma 2.2.36, which means that $k_j \in (-1, 0]$; that all the entries of the matrix M_C^{-1} are negative, which was shown in Lemma 2.2.33; and that, by positivity of intersection, the components of $\boldsymbol{\chi}$ are non-negative integers.

the self-intersection number of $\widetilde{C}$ is then a consequence of the fact that $a_0 = \frac{pq-1}{p^2} < 1$, which follows from (2.3.13) and shows that

$$\widetilde{C} \cdot \widetilde{C} = a_0 + \chi^T M_C^{-1} \chi \leq a_0 < 1. \qquad \square$$

As a result of this lemma, choosing the almost complex structure generically on the complement of the embedding ι, we can ensure that $\widetilde{C}$ is a (0)-sphere or a (-1)-sphere. However, since there is a $\mathbb{C}$-family of rational curves representing the fibre class F that we can stretch, we must find a (0)-sphere.[23] In the following we will therefore assume that $\widetilde{C}$ is a curve of square zero.

Lemma 2.3.18 ([Ada+25, Lemma 5.4.6])**.** *The curve $\widetilde{C}$ intersects the Wahl chain precisely once. Moreover, the intersection is with the first sphere in the Wahl chain.*

Proof. By Lemma 2.3.17 we know that the adjunction formula for $\widetilde{C}$ reads

$$1 = \frac{p^2+1}{2p^2} - \frac{1}{2}(\boldsymbol{k}^T \chi + \chi^T M_C^{-1} \chi), \tag{2.3.21}$$

where the first term is calculated using $a_n = 1/p^2$, which follows from the equation $\boldsymbol{a} = (M_{\mathcal{D}})^{-1}\boldsymbol{\xi}$ and (2.3.13). Observe that all the entries in χ are non-negative integers by positivity of intersection and that all the entries of M_C^{-1} are negative. Therefore, the second summand on the right-hand side contributes at least

$$\sum_{j=1}^{m} \left(\frac{1}{2p^2} \chi_j^2 e_j f_j + \frac{1}{2} \chi_j \left(1 - \frac{e_j + f_j}{p^2} \right) \right), \tag{2.3.22}$$

which is the diagonal term, to a sum that is equal to 1. Assume that there is a $\chi_j \geq 2$. Then the sum (2.3.22) is at least

$$\frac{1}{2p^2} \chi_j^2 e_j f_j + \frac{1}{2} \chi_j \left(1 - \frac{e_j + f_j}{p^2} \right) \geq 1 + \frac{1}{p^2} \left(2e_j f_j - e_j - f_j \right) \geq 1,$$

since $e_j, f_j \geq 1$, in contradiction to (2.3.21).

If $\chi_j = 1$ for some index j we have

$$\frac{1}{2p^2} \chi_j^2 e_j f_j + \frac{1}{2} \chi_j \left(1 - \frac{e_j + f_j}{p^2} \right) = \frac{1}{2} \left(1 + \frac{1}{p^2} \left(e_j f_j - e_j - f_j \right) \right) \geq \frac{1}{2} \left(1 - \frac{1}{p^2} \right).$$

This means that if there are at least two χ_j that are equal to 1 the right-hand side of (2.3.21) is at least

$$\frac{p^2+1}{2p^2} + 1 - \frac{1}{p^2} > 1,$$

which is a contradiction because $p \geq 2$. In conclusion we see that $\chi_j = 1$ for exactly one j and we are left to determine this index. By (2.3.21) we have

$$0 = \frac{p^2+1}{2p^2} + \frac{1}{2p^2}(p^2 + e_j f_j - (e_j + f_j)) - 1 = \frac{1}{2p^2}(e_j f_j - e_j - f_j + 1) = \frac{1}{2p^2}(e_j - 1)(f_j - 1)$$

[23] That there is a $\mathbb{C}$-family of rational curves was proven in Lemma 2.3.15.

which is the case if and only if $e_j = 1$ or $f_j = 1$. By definition of the accompanying numbers this is only the case if either $j = 1$ or $j = m$, i.e. $\widetilde{C}$ must intersect the Wahl chain at one of its ends. Since $\widetilde{C}$ is a (0)-curve we also know that

$$0 = a_0 + \chi^T M_C^{-1} \chi = \frac{pq-1}{p^2} - \frac{1}{p^2} e_j f_j.$$

This implies that $j = 1$, because $(e_1, f_1) = (1, pq-1)$ and $(e_m, f_m) = ([pq-1]^{-1}, 1)$, where $0 < [pq-1]^{-1} < p^2$ is the integer such that $(pq-1)[pq-1]^{-1} \equiv 1 \mod p^2$ as discussed in (2.2.8) and (2.2.9), i.e. $[pq-1]^{-1} = p^2 - (pq+1)$.[24] □

To summarise, we have found a square zero curve $\widetilde{C}$ in $\widetilde{X}$ that intersects both D_0 and C_1 once positively and does not intersect any of the other spheres in $\mathcal{D}_{p,q}$ and $C_{p,q}$, which implies that $\widetilde{C}$ is a smooth ruling of a regulation of $\widetilde{X}$ in the class $\widetilde{C}$. What these considerations amount to is that the regulation of X in the class F "persists" under the rational blow-up along ι. In the case that $\iota = \iota_{\text{vis}}$ this is obvious from the almost toric base diagram, as illustrated in Figure 2.10.

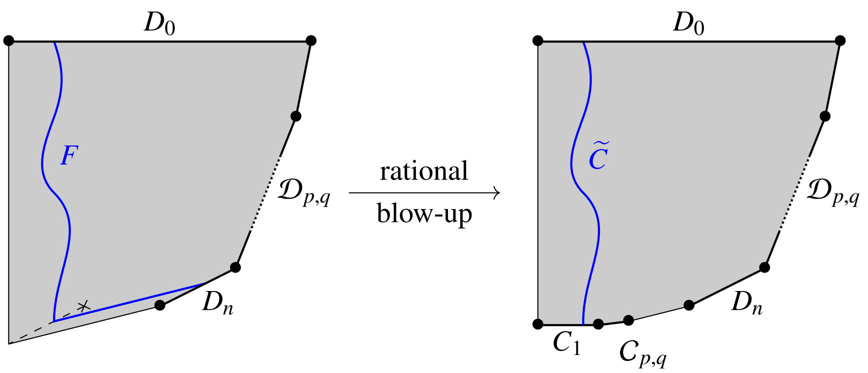

Figure 2.10: The transformation of the regulation in the visible case.

Remark 2.3.19. Choosing $\widetilde{J}$ generically we can ensure that the irreducible components of $C_{p,q}$ and $\mathcal{D}_{p,q}$ are the only $\widetilde{J}$-holomorphic spheres with self-intersection strictly less than -1, according to Lemma 2.2.31. In particular, this means that all the other $\widetilde{J}$-holomorphic embedded rational curves in $\widetilde{X}$ are either exceptional spheres or have non-negative square and that all the curves in $C_{p,q} \setminus C_1$ and $\mathcal{D}_{p,q} \setminus D_0$ have to appear as irreducible components of a broken ruling. Since every broken ruling contains an (-1)-sphere, successively contracting these in the broken ruling will yield a "minimal model" Y of $\widetilde{X}$, which contains no broken rulings. The broken ruling before the last step will consist of two transversally intersecting (-1)-spheres.[25]

Lemma 2.3.20. *The number of contractions that lead from $\widetilde{X}$ to the minimal model Y is equal to $m + n - 1$. In particular, the second Betti number of $\widetilde{X}$ is equal to $m + n + 1$.*

Proof. We know that a basis of $H_2(\widetilde{X}, \mathbb{Q})$ is given by $C_{p,q}$ and $\mathcal{D}_{p,q}$, and therefore we have

$$\dim(H_2(\widetilde{X}, \mathbb{Q})) = n + (m+1).$$

[24] Of course there is one case in which $[pq-1]^{-1} = pq - 1$, namely $(p, q) = (2, 1)$. However, this is the case in which the Wahl chain is just a single (-4)-sphere.

[25] See [Ada+25, Corollary 4.2.8] and the surrounding discussion.

Denoting the number of contractions that lead from $\widetilde{X}$ to Y by k, the equation $k+2=n+m+1$ follows, since the second Betti number of Y is equal to 2. $\square$

The discussion above then implies the following theorem.

Theorem 2.3.21. *Assume that the almost complex structure $\widetilde{J}$ is chosen generically on the subset $V \subseteq \widetilde{X}$. Then there is an embedded $\widetilde{J}$-holomorphic sphere $\widetilde{C} \subset \widetilde{X}$ of square zero such that*

$$\widetilde{C} \cdot C_j = \begin{cases} 0 & \text{if } j \neq 1, \\ 1 & \text{if } j = 1 \end{cases} \qquad \text{and} \qquad \widetilde{C} \cdot D_i = \begin{cases} 0 & \text{if } i \neq 0, \\ 1 & \text{if } i = 0 \end{cases} .$$

Moreover, the regulation defined by the $\widetilde{J}$-holomorphic sphere $\widetilde{C}$ has exactly one broken ruling that is given by

$$\mathcal{T} = (C_2 \cup \ldots \cup C_m) \cup (D_1 \cup \ldots \cup D_n) \cup E_{\widetilde{J}} = (\mathcal{C}_{p,q} \setminus C_1) \cup (\mathcal{D}_{p,q} \setminus D_0) \cup E_{\widetilde{J}},$$

where $E_{\widetilde{J}}$ is an exceptional sphere that intersects the Wahl chain and the compactifying divisor at their terminal components.

The structure of the regulation in the class $\widetilde{C}$ in Theorem 2.3.21 is shown in Figure 2.11. Before proving this theorem, let us formulate a crucial corollary of Theorem 2.3.21.

Corollary 2.3.22. *In the setup of Theorem 2.3.21 the complement of $\mathcal{T}$, i.e. $\widetilde{X} \setminus \mathcal{T}$, is minimal.*

Proof. By Theorem 2.3.21 the symplectic manifold $\widetilde{X}$ is derived from a Hirzebruch surface Y by consecutive blow-ups of a single ruling. This means that $\widetilde{X} \setminus \mathcal{T}$ is contained in $Y \setminus F$, where F is a fibre of the Hirzebruch surface.[26] Therefore, $\widetilde{X} \setminus \mathcal{T}$ is minimal. $\square$

Proof of Theorem 2.3.21. The first part of the theorem is a summary of the discussion before and the second part is obvious if $(p,q)=(2,1)$, since in that case $\mathcal{C}_{2,1}$ consists of a single sphere and $\mathcal{D}_{2,1}$ consists of a positive sphere and an exceptional divisor. Therefore we move on to prove the second part and assume that $(p,q) \neq (2,1)$.

By Remark 2.3.19 we know that there has to be at least one broken ruling since the tails of the chains $\mathcal{C}_{p,q}$ and $\mathcal{D}_{p,q}$ have to be part of a broken ruling. The goal is to show that there is precisely one broken ruling given by the tails of $\mathcal{C}_{p,q}$ and $\mathcal{D}_{p,q}$, denoted by $\mathcal{C}' := \mathcal{C}_{p,q} \setminus C_1$ and $\mathcal{D}' := \mathcal{D}_{p,q} \setminus D_0$ in the following, connected by a single exceptional sphere that intersects the tails at their ends.

Assume that both tails are contained in two distinct broken rulings,[27] which we will call $\mathcal{T}_C$ and $\mathcal{T}_{\mathcal{D}}$. We know that the broken rulings have to contain at least one exceptional sphere and hence we have

$$\dim(H_2(\widetilde{X};\mathbb{Q})) \geq ((n-1)+1)+(m+1).$$

However, we know by Lemma 2.3.20 that this is an equality and therefore the regulation of $\widetilde{X}$ has exactly two broken rulings, namely $\mathcal{T}_C$ and $\mathcal{T}_{\mathcal{D}}$, each containing a single (-1)-sphere, E_C

[26] Note that denoting this fibre by F makes sense, since a smooth ruling of $\widetilde{X}$ is a fibre of Y under the sequence of blow-downs.

[27] Recall that broken rulings are either disjoint, or they coincide.

and $E_{\mathcal{D}}$. We will now show that such a configuration cannot exist. Consider the expression

$$\mathcal{T}_C = \sum_{j=2}^{m} \alpha_j C_j + \epsilon E_C,$$

where the α_j and ϵ are the covering multiplicities. Because $\mathcal{T}_C$ is a broken ruling of the regulation in the class $[\widetilde{C}]$ we know that

$$\mathcal{T}_C \cdot C_1 = \widetilde{C} \cdot C_1 = 1 \qquad \text{and} \qquad \mathcal{T}_C \cdot D_0 = \widetilde{C} \cdot D_0 = 1,$$

which implies $\alpha_2 = 1$ and $\epsilon = 1$. Now assume that E_C intersects C' in C_r. Then $\mathcal{T}_C \cdot C_j = \widetilde{C} \cdot C_j = 0$ implies

$$M_{C'}\boldsymbol{\alpha} + \boldsymbol{e}_r = 0, \quad \text{i.e. } \boldsymbol{\alpha} = -(M_{C'})^{-1}\boldsymbol{e}_r,$$

where $M_{C'}$ is the intersection matrix of C' and $\boldsymbol{e}_r$ is the rth standard basis vector. In particular, we must have $1 = \alpha_2 = -(M_{C'})^{-1}_{1r}$. In the notation of Lemma 2.2.33 this implies $1 = f_r/(pq-1)$, because the HJ continued fraction for C' is $(pq-1)/(b_1(pq-1)-p^2)$. However, this is a contradiction, because $f_r < f_0 = pq-1$.

Therefore, the regulation of $\widetilde{X}$ has a single broken ruling $\mathcal{T}$ which contains the tails C' and $\mathcal{D}'$ and a single exceptional sphere E.[28] What is left to show is the intersection pattern of E and the tails. As before we write

$$\mathcal{T} = \sum_{j=2}^{m} \alpha_j C_j + \sum_{i=1}^{n} \beta_i D_i + \epsilon E,$$

where again the α_j, β_i and ϵ are the multiplicities. Because $[\mathcal{T}] = [\widetilde{C}]$ we have

$$\mathcal{T} \cdot C_1 = \widetilde{C} \cdot C_1 = 1 \qquad \text{and} \qquad \mathcal{T} \cdot D_0 = \widetilde{C} \cdot D_0 = 1,$$

which implies $\alpha_2 = \beta_1 = 1$. Assume that E meets C' in C_r and $\mathcal{D}'$ in D_s. This means that we have to show that $r = m$ and $s = n$. The intersection identities $\mathcal{T} \cdot C_j = \widetilde{C} \cdot C_j$ and $\mathcal{T} \cdot D_i = \widetilde{C} \cdot D_i$ imply

$$M_{C'}\boldsymbol{\alpha} + \epsilon\boldsymbol{e}_r = 0, \quad \text{i.e. } \boldsymbol{\alpha} = -\epsilon(M_{C'})^{-1}\boldsymbol{e}_r, \qquad \text{and} \qquad M_{\mathcal{D}'}\boldsymbol{\beta} + \epsilon\boldsymbol{e}_s = 0, \quad \text{i.e. } \boldsymbol{\beta} = -\epsilon(M_{\mathcal{D}'})^{-1}\boldsymbol{e}_s.$$

Considering the first component of each of these equations we get

$$\epsilon = -\frac{1}{(M_{C'})^{-1}_{1r}} = \frac{pq-1}{f_r} \qquad \text{and} \qquad \epsilon = -\frac{1}{(M_{\mathcal{D}'})^{-1}_{1s}} = \frac{pq-1}{h_s},$$

which implies that $f_r = h_s$. Since $(pq-1)/[q^2]^{-1}$ is the dual of $(pq-1)/(c_1(pq-1)-p^2)$, they only share one right accompanying number, namely 1. This implies that $r = m$ and $s = n$, which concludes the proof.[29] □

[28] That there cannot be more than one exceptional sphere again follows from the fact that we can compute the dimension of $H_2(\widetilde{X}, \mathbb{Q})$ by counting the number of contractions and alternatively by adding the lengths of the chains C' and $\mathcal{D}'$. This also shows that there is a single broken ruling.

[29] Recall that $[q^2]^{-1} = p^2 - d_0(pq-1)$ and that $d_0 = b_1 - 1$. The fact that these dual HJ continued fractions only share a single right accompanying number is easy to see. Write $N = (pq-1)$ and $a = [q^2]^{-1}$. Then we have

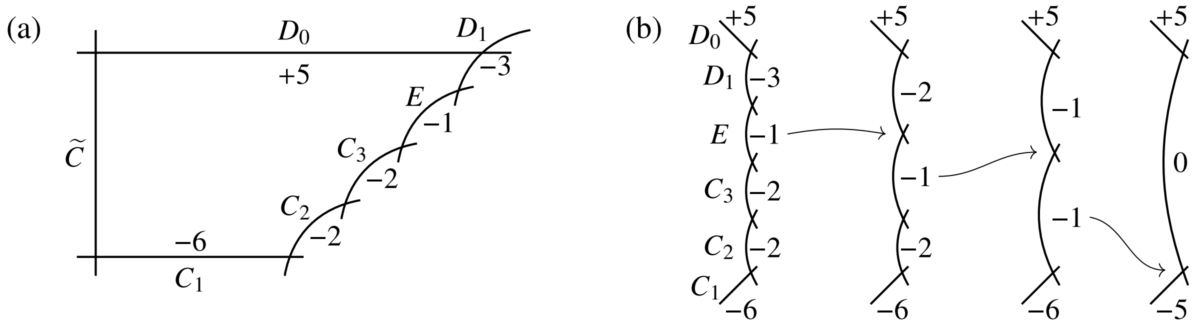

Figure 2.11: (a) An example that illustrates the structure of the regulation in the class $\widetilde{C}$ proved in Theorem 2.3.21 in the case $(p, q) = (4, 1)$. (b) The contraction process of the broken ruling for this example. In this example $d_0 = 5$ and the figure shows how $\widetilde{X}_{4,1}$ is derived from $\mathbb{F}_5$.

Remark 2.3.23. Theorem 2.3.21 also implies that the minimal model of Y can be arranged to be the Hirzebruch surface $\mathbb{F}_{d_0}$. This follows from the structure of $\mathcal{T}$ proven in Theorem 2.3.21. The last contraction is a contraction of a broken fibre that consists of two (-1)-spheres. Contracting the curve that intersects the negative section yields the Hirzebruch surface $\mathbb{F}_{d_0}$.

2.3.3 Constructing the symplectomorphism

In this subsection we want to show that Lagrangian pinwheels in their normal neighbourhoods are unique up to symplectomorphism. More precisely, we want to prove the following theorem.

Theorem 2.3.24. *Suppose that $\iota : B_{p,q}(\varepsilon) \hookrightarrow B_{p,q}$ is a symplectic embedding. Then there exists a compactly supported symplectomorphism $\Phi \in \mathrm{Symp}_c(B_{p,q})$ such that $\Phi \circ \iota = \iota_{\text{vis}}$.*

Remark 2.3.25. Note that the uniqueness of Lagrangian (p, q)-pinwheels in $B_{p,q}$ up to symplectomorphism follows from this theorem, since we showed that every such pinwheel is "locally visible" in Theorem 2.2.14.

This theorem is very similar to [Ada+25, Theorem 1.4.1] and its proof will indeed also follow along similar lines. However, we will base the proof of Theorem 2.3.24 on a "different" method.[30] We shall need the following analogue, for star-shaped sets in $T^*S^1 \times \mathbb{R}^2$, of the Gromov–McDuff Theorem [MS04, Theorem 9.4.2], which is formulated for star-shaped sets in $\mathbb{R}^4$.

Definition 2.3.26. A star-shaped subset W of a Liouville manifold $(X, d\lambda)$ is a subset that contains the Liouville skeleton of λ and $\phi_t^\lambda(x) \in W$ for all $t \leq 0$ and $x \in W$.

$f_r \equiv -e_r a \mod N$ and $h_s \equiv -g_r a \mod N$, because of the recursive formula (2.2.6), which yields $(e_r + g_s)a \equiv 0 \mod N$. Because $\gcd(a, N) = 1$ this implies $e_r + g_s \equiv 0 \mod N$. However, the left accompanying numbers are increasing and they terminate at $e_{m-1} = N - [a]^{-1}$ and $g_n = [a]^{-1}$. This implies $0 < e_r + g_s \leq N$ and therefore the claim.

[30] It will be clear from the method of proof and from the references cited therein that the root of this method is Gromov's [Gro85] $S^2 \times S^2$-foliation result, which is also the underlying result for the proof of [Ada+25, Theorem 1.4.1].

Let p be the coordinate in the cotangent direction in T^*S^1 and r be the radius function on $\mathbb{R}^2$. Let $\mathcal{X} = p\partial_p + \frac{1}{2}r\partial_r$ be the standard Liouville vector field on $T^*S^1 \times \mathbb{R}^2$.

Theorem 2.3.27. *Let (X, ω) be a connected minimal symplectic 4-manifold and $\mathcal{K} \subset X$ be a compact set such that there exists a symplectomorphism $\psi\colon \left(T^*S^1 \times \mathbb{R}^2\right) \setminus \mathcal{V} \to X \setminus \mathcal{K}$, where $\mathcal{V} \subset T^*S^1 \times \mathbb{R}^2$ is a star-shaped compact set. Then (X, ω) is symplectomorphic to $(T^*S^1 \times \mathbb{R}^2, \omega_0)$. Moreover, for every open neighbourhood $\mathcal{U} \subset X$ of $\mathcal{K}$, the symplectomorphism can be chosen to agree with ψ^{-1} on $X \setminus \mathcal{U}$.*

For the proof of Theorem 2.3.27 we need the following result.

Theorem 2.3.28 ([CMM25, Lemma 3.1.3])**.** *Let $\omega_{a,b}$ be the usual split symplectic form on $S^2 \times S^2$ that gives the factors areas a and b, respectively. Denote by p_N and p_S the North and South poles of S^2, and consider the three spheres $S_1 = S^2 \times \{p_S\}$ and $S_2^N = \{p_N\} \times S^2$, $S_2^S = \{p_S\} \times S^2$ in $S^2 \times S^2$. Let ω be a symplectic form on $S^2 \times S^2$ that agrees with $\omega_{a,b}$ on a neighbourhood of the configuration $\mathcal{S} := S_1 \cup S_2^N \cup S_2^S$. Then there exists a symplectomorphism $\psi\colon (S^2 \times S^2, \omega_{a,b}) \to (S^2 \times S^2, \omega)$ that is the identity on a neighbourhood of $\mathcal{S}$.*

Theorem 2.3.27 follows from Theorem 2.3.28 exactly as the Gromov–McDuff Theorem [MS04, Theorem 9.4.2] follows from [MS04, Theorem 9.4.7]. Note that [MS04, Lemma 9.4.10] is also readily adapted; one just needs to replace the scaling maps $x \mapsto tx$ on $\mathbb{R}^4$ by the dilations on $T^*S^1 \times \mathbb{R}^2$ induced by the Liouville flow of $\mathcal{X}$.

With these preparations in place we are ready to prove Theorem 2.3.24.

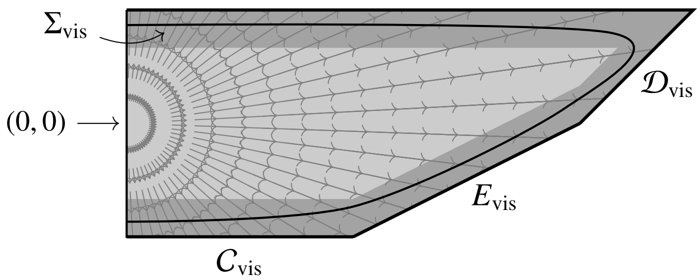

Figure 2.12: The visible situation in the case $(p, q) = (2, 1)$. The linear chain of spheres $\mathcal{T}_{\mathrm{vis}}$ is formed by the Wahl chain C_{vis}, consisting of a single (-4)-sphere in this case, and the compactifying divisor $\mathcal{D}_{\mathrm{vis}}$, formed by a $(+3)$-sphere and a (-1)-sphere, connected by an exceptional divisor E_{vis}. The boundary Σ_{vis} of a normal neighbourhood of the configuration $\mathcal{T}_{\mathrm{vis}} = C_{\mathrm{vis}} \cup E_{\mathrm{vis}} \cup \mathcal{D}_{\mathrm{vis}}$ is a contact hypersurface. Σ_{vis} is diffeomorphic to $S^1 \times S^2$ and is star-shaped with respect to the Liouville field shown in gray.

Proof of Theorem 2.3.24. Pick compactification data $(\alpha, \boldsymbol{\rho})$, as in Definition 2.3.5, such that the embedding ι also defines an embedding of $B_{p,q}(\varepsilon)$ into $X_{p,q}(\alpha, \boldsymbol{\rho})$. We obtain two symplectic manifolds by rationally blowing up X along ι and ι_{vis}, which we denote by $\widetilde{X}$ and $\widetilde{X}_{\mathrm{vis}}$. Recall that both of these symplectic manifolds contain a linear chain of symplectic spheres $\mathcal{T}$ and $\mathcal{T}_{\mathrm{vis}}$ that consists of the Wahl chain and the compactifying divisor, connected by an exceptional sphere, as shown in Theorem 2.3.21. Figure 2.12 shows the visible situation.

The symplectic neighbourhood theorem gives symplectic embeddings $\psi_{\mathrm{vis}}\colon \nu \hookrightarrow \widetilde{X}_{\mathrm{vis}}$ and $\psi\colon \nu \hookrightarrow \widetilde{X}$ of a plumbing of normal bundles, whose boundary Σ_{vis} and Σ, respectively, are contact-type hypersurfaces contactomorphic to $S^1 \times S^2$. The hypersurface Σ_{vis} can be assumed to be a visible contact hypersurface in a toric domain whose completion is $T^*S^1 \times \mathbb{R}^2$ and is star-shaped with respect to a specific Liouville field, see Figure 2.12. Let $\mathcal{K}_{\mathrm{vis}}$ be the closure in $\widetilde{X}_{\mathrm{vis}}$ of the complement of $\psi_{\mathrm{vis}}(\nu)$; its boundary is Σ_{vis}. Similarly, let $\mathcal{K}$ be the closure in $\widetilde{X}$ of the complement of $\psi(\nu)$. The boundary of $\mathcal{K}$ is Σ, and $\mathcal{K}$ is minimal by Corollary 2.3.22. Our goal is to extend the symplectomorphism $\psi \circ \psi_{\mathrm{vis}}^{-1} \colon \psi_{\mathrm{vis}}(\nu) \to \psi(\nu)$ to a symplectomorphism

$$\Psi\colon \widetilde{X}_{\mathrm{vis}} = \psi_{\mathrm{vis}}(\nu) \cup \mathcal{K}_{\mathrm{vis}} \to \psi(\nu) \cup \mathcal{K} = \widetilde{X}$$

that takes $\mathcal{T}_{\mathrm{vis}}$ to $\mathcal{T}$. Choose coordinates on the almost toric base diagram in Figure 2.12 so that the source of the Liouville vector field is at $(0,0)$, and take $\delta > 0$ so small that the "ball" of radius δ is contained in this almost toric base diagram. Write $S_{\mathrm{vis}}(\delta)$ for the set in $\widetilde{X}_{\mathrm{vis}}$ represented by this ball. It is canonically symplectomorphic to the set $S(\delta) := \{p^2+r^2 \le \delta^2\} \subseteq T^*S^1 \times \mathbb{R}^2$, where p is the cotangent direction in T^*S^1 and r is the radius function on $\mathbb{R}^2$. Let $\mathcal{X}_{\mathrm{vis}} := \sum_{i=1}^2 p_i \partial_{p_i}$ be the Liouville vector field centred on $(0,0)$ defined by the action-angle coordinates, and let $\mathcal{X} := p\partial_p + \frac{1}{2} r \partial_r$ be the standard Liouville vector field on $T^*S^1 \times \mathbb{R}^2$.

Using the flows of $\mathcal{X}_{\mathrm{vis}}$ and $\mathcal{X}$ we construct a symplectic embedding $\mathcal{K}_{\mathrm{vis}} \hookrightarrow T^*S^1 \times \mathbb{R}^2$ onto a closed star-shaped region with smooth boundary. We call the image $\mathcal{K}_{\mathrm{star}}$ and its boundary Σ_{star}. Explicitly,

$$\mathcal{K}_{\mathrm{star}} = S^1 \times \{(0,0)\} \cup \left\{ \phi^t_{\mathcal{X}}(x) \mid x \in \partial S(\delta),\ -\infty < t \le t^*(x) \right\}$$

where $t^*(x)$ is defined by $\phi^{t^*(x)}_{\mathcal{X}}(x) \in \Sigma_{\mathrm{vis}}$. Since $\mathcal{X}_{\mathrm{vis}}$ is transverse to Σ_{vis}, we find $\varepsilon > 0$ such that the trajectories $\phi^t_{\mathcal{X}_{\mathrm{vis}}}(x)$ exist for $x \in \Sigma_{\mathrm{vis}}$ and $0 \le t \le \varepsilon$. Set

$$\mathcal{U}_{\mathrm{vis}} := \left\{ \phi^t_{\mathcal{X}_{\mathrm{vis}}}(x) \mid x \in \Sigma_{\mathrm{vis}},\ 0 < t < \varepsilon \right\} \subset \widetilde{X}_{\mathrm{vis}} \setminus \mathcal{T}_{\mathrm{vis}} \quad \text{and} \quad \mathcal{U}_{\mathrm{star}} := \left\{ \phi^t_{\mathcal{X}}(x) \mid x \in \Sigma_{\mathrm{star}},\ 0 < t < \varepsilon \right\}.$$

Using again the flows of $\mathcal{X}_{\mathrm{vis}}$ and $\mathcal{X}$ we construct the symplectomorphism $\xi_{\mathrm{vis}} \colon \mathcal{U}_{\mathrm{star}} \to \mathcal{U}_{\mathrm{vis}}$ that maps flow segments to flow segments. Also set $\mathcal{U} := (\psi \circ \psi_{\mathrm{vis}}^{-1})(\mathcal{U}_{\mathrm{vis}}) \subset \widetilde{X} \setminus \mathcal{T}$. Then we also have the symplectomorphism

$$\xi := \psi \circ \psi_{\mathrm{vis}}^{-1} \circ \xi_{\mathrm{vis}} \colon \mathcal{U}_{\mathrm{star}} \to \mathcal{U}.$$

Define two open symplectic manifolds X_{vis} and X by

$$X_{\mathrm{vis}} = \mathcal{U}_{\mathrm{vis}} \bigcup_{\xi_{\mathrm{vis}}} (T^*S^1 \times \mathbb{R}^2 \setminus \mathcal{K}_{\mathrm{star}}) \quad \text{and} \quad X = \mathcal{U} \bigcup_{\xi} (T^*S^1 \times \mathbb{R}^2 \setminus \mathcal{K}_{\mathrm{star}}).$$

Then the maps

$$\begin{aligned} \Psi_{\mathrm{vis}} \colon T^*S^1 \times \mathbb{R}^2 \setminus \mathcal{K}_{\mathrm{star}} &\to X_{\mathrm{vis}}, & \Psi_{\mathrm{vis}}|_{\mathcal{U}_{\mathrm{star}}} &= \xi_{\mathrm{vis}}, & \Psi_{\mathrm{vis}}|_{T^*S^1\times\mathbb{R}^2\setminus\mathcal{K}_{\mathrm{star}}} &= \mathrm{id}, \\ \Psi \colon T^*S^1 \times \mathbb{R}^2 \setminus \mathcal{K}_{\mathrm{star}} &\to X, & \Psi|_{\mathcal{U}_{\mathrm{star}}} &= \xi, & \Psi|_{T^*S^1\times\mathbb{R}^2\setminus\mathcal{K}_{\mathrm{star}}} &= \mathrm{id}, \end{aligned}$$

are symplectomorphisms. They both extend over $\mathcal{K}_{\mathrm{star}}$ and hence to $T^*S^1 \times \mathbb{R}^2$. More precisely, take $\mathcal{V}_{\mathrm{star}} := \bigcup_{\frac{\varepsilon}{2} < t < \varepsilon} \phi^t_{\mathcal{X}'}(\Sigma_{\mathrm{star}})$. Since $\mathcal{K}$ is minimal, Theorem 2.3.27 guarantees the existence of symplectomorphisms

$$\widehat{\Psi}_{\mathrm{vis}} \colon T^*S^1 \times \mathbb{R}^2 \to X_{\mathrm{vis}} \cup \mathcal{K}_{\mathrm{vis}} \quad \text{and} \quad \widehat{\Psi} \colon T^*S^1 \times \mathbb{R}^2 \to X \cup \mathcal{K}$$

such that

$$\widehat{\Psi}_{\mathrm{vis}}|_{\mathcal{V}_{\mathrm{star}}} = \xi_{\mathrm{vis}} \quad \text{and} \quad \widehat{\Psi}|_{\mathcal{V}_{\mathrm{star}}} = \xi.$$

Set $\widetilde{\mathcal{K}}_{\mathrm{vis}} := \mathcal{K}_{\mathrm{vis}} \cup \mathcal{U}_{\mathrm{vis}}$. Then the map $\widehat{\chi}\colon \widetilde{X}_{\mathrm{vis}} \setminus \mathcal{T}_{\mathrm{vis}} \to \widetilde{X} \setminus \mathcal{T}$ defined by

$$\widehat{\chi}|_{\widetilde{\mathcal{K}}_{\mathrm{vis}}} = \widehat{\Psi} \circ \widehat{\Psi}_{\mathrm{vis}}^{-1} \quad \text{and} \quad \widehat{\chi}|_{\widetilde{X}_{\mathrm{vis}} \setminus (\mathcal{T}_{\mathrm{vis}} \cup \widetilde{\mathcal{K}}_{\mathrm{vis}})} = \psi \circ \psi_{\mathrm{vis}}^{-1}$$

is a symplectomorphism that agrees near $\mathcal{T}_{\mathrm{vis}}$ with $\psi \circ \psi_{\mathrm{vis}}^{-1}$. Therefore, $\widehat{\chi}$ descends to a symplectomorphism $\chi\colon \widetilde{X}_{\mathrm{vis}} \to \widetilde{X}$ taking $\mathcal{T}_{\mathrm{vis}}$ to $\mathcal{T}$. The last step is then to glue this symplectomorphism appropriately with the symplectomorphism $\iota \circ \iota_{\mathrm{vis}}^{-1}$ on the interface that is given by the linear chain of spheres to obtain a symplectomorphism $\Phi \in \mathrm{Symp}_c(B_{p,q})$ that intertwines the embeddings ι_{vis} and ι. The proof of this step is verbatim the proof of [Ada+25, Corollary 5.2.5] and we therefore omit it here. □

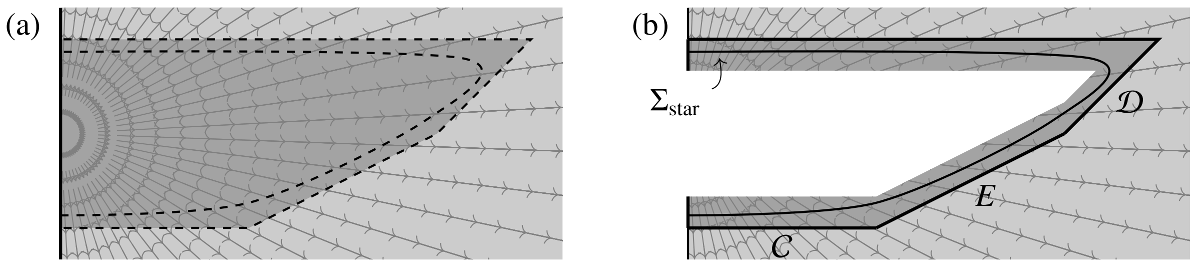

Figure 2.13: The "visible" part of the constructions in Section 2.3.3. (a) shows how the Delzant polytope of $\widetilde{X}_{\mathrm{vis}} \setminus \mathcal{T}_{\mathrm{vis}}$ embeds into the standard Delzant polytope of $T^*S^1 \times \mathbb{R}^2$, respecting the Liouville vector fields. How Σ_{vis} embeds is also shown. (b) shows how $\psi(\nu)$ can be completed to match the situation in (a).

2.4 Part 2: From symp to Ham

Our goal in this section is to compute the homotopy groups of $\mathrm{Symp}_c(B_{p,q})$, in order to show that any compactly supported symplectomorphism is, up to Hamiltonian isotopy, a power of a pintwist (see Definition 2.2.18). First, we show how the space $W_{p,q} = B_{p,q} \backslash K_{p,q}$ can be seen as a star-shaped domain in T^*Wh, the cotangent bundle of the Whitney sphere, and then we show how this inclusion induces a weak homotopy equivalence (w.h.e.) $\mathrm{Symp}_c(W_{p,q}) \hookrightarrow \mathrm{Symp}_c(T^*Wh)$. Then we explain how to compute $\mathrm{Symp}_c(B_{p,q})$ by relating it to $\mathrm{Symp}_c(W_{p,q})$ and $\mathcal{S}(K_{p,q})$, the space of symplectic cylinders which agree with $K_{p,q}$ outside a compact subset.

Remark 2.4.1. We remind the reader that, since $H^1(B_{p,q}; \mathbb{R}) = 0$, any symplectic isotopy is also Hamiltonian. We therefore will often blur the distinction between the two notions as we see fit.

Remark 2.4.2. Since the compactly supported symplectomorphism groups we are going to consider have the homotopy type of countable CW-complexes, the notions of homotopy equivalence and weak homotopy equivalence coincide for them. However, working with homotopy equivalences is enough for our purposes.

2.4.1 T^*Wh as a Liouville completion of $B_{p,q} \setminus K_{p,q}$

The observation that will allow us to compute the weak homotopy type of $\mathrm{Symp}_c(B_{p,q})$ is the fact that the manifold $W_{p,q} = B_{p,q} \backslash K_{p,q}$, equipped with the restriction of the standard symplectic form

$\omega_{p,q}$, has contractible symplectomorphism group. Here, we show that there exists a (complete) Liouville manifold $(T^*Wh, d\lambda_{\text{can}})$ such that $(W_{p,q}, \omega_{p,q})$ symplectically embeds in it as a star-shaped open subset. Recall from Definition 2.3.26 that a star-shaped subset of a Liouville manifold is a subset that contains the Liouville skeleton and is negatively invariant. The main property of star-shaped subsets we will use is the following.

Lemma 2.4.3. *Let W be a star-shaped subset of a Liouville manifold $(X, d\lambda)$. Then the inclusion $i : W \hookrightarrow X$ induces a weak homotopy equivalence $i_* : \mathrm{Symp}_c(W, d\lambda) \to \mathrm{Symp}_c(X, d\lambda)$.*

Our proof follows that of [Eva11, Proposition 2.1] with minor changes. We include it for the convenience of the reader.

Proof. To prove surjectivity of i_* at the level of homotopy groups, assume that g_s is a continuous family of compactly supported symplectomorphisms of X, with $s \in S^n$. Since S^n is compact, there is a compact subset $K \subset X$ such that the support of all elements in the image of g is contained in K. Since W contains the λ-skeleton, there exists a time $T > 0$ such that $\phi_t^\lambda(K) \subset W$, for all $t \leq -T$. Therefore, for any element g_s the endpoint $g_{s,-T}$ of the symplectic isotopy $g_{s,t} = \phi_{-t}^\lambda \circ g_s \circ \phi_t^\lambda$ has support in W. Hence conjugation with ϕ_t^λ homotopes the family $g \in \pi_n(\mathrm{Symp}_c(X))$ to a family in $\pi_n(\mathrm{Symp}_c(W))$, and so i_* is surjective at the level of homotopy groups.

Similarly, for injectivity, consider two families $g, f : S^n \to \mathrm{Symp}_c(W)$ and suppose they are homotopic through maps in $\mathrm{Symp}_c(X)$. Conjugating by $\phi_{\tau(t)}^\lambda$, where $\tau : [0,1] \to (-\infty, 0]$ is a sufficiently U-shaped smooth function satisfying $\tau(0) = \tau(1) = 0$, we obtain a homotopy from g to f supported in W. □

We will now show that $W_{p,q}$ can be viewed as a star-shaped subset of the Liouville manifold T^*Wh, which we will introduce in a moment. In combination with the lemma above, this implies that $\mathrm{Symp}_c(T^*Wh)$ and $\mathrm{Symp}_c(W_{p,q})$ are weakly homotopy equivalent.

The symplectic manifold T^*Wh was introduced by Dimitroglou-Rizell in [Dim19], with an eye towards classification of Lagrangian submanifolds therein. Let us swiftly recall the points that are relevant for us, referring to [Dim19, Section 3] for details. Let Wh be a Lagrangian *Whitney sphere*, i.e. an immersed sphere with a single transverse positive self-intersection. The space $(D^*Wh, d\lambda)$ is the model neighbourhood of a Lagrangian Whitney sphere and it is constructed by self-plumbing the cotangent disc bundle of a sphere. By carefully completing this Liouville domain, one obtains a Liouville manifold $(T^*Wh, d\lambda_{\text{can}})$ whose Liouville skeleton is the Lagrangian Whitney sphere, and in this sense $(T^*Wh, d\lambda_{\text{can}})$ can be viewed as the cotangent bundle of the Whitney sphere. For a detailed exposition of this construction, see [Dim19, Subsections 3.1 and 3.2].

The space $(T^*Wh, d\lambda_{\text{can}})$ also carries an explicit almost toric fibration $T^*Wh \xrightarrow{\pi} \mathbb{R}^2$, where the skeletal Whitney sphere of T^*Wh is the unique singular torus fibre over the point $(1, 1)$. Furthermore, the Liouville flow of λ_{can} respects π, in the sense that it takes fibres to fibres, and the differential $D\pi$ pushes forward the Liouville vector field to the radial vector field on $\mathbb{R}^2$ emanating from the unique node at $(1, 1)$, outside a small disc around the node. This is proven in [Dim19, Propositions 3.7 and 3.8]. The discussion is summarized in Figure 2.14.

Having collected these facts, we can now use the general theory of almost toric fibrations, as explained for example in [Sym03; Eva23] to construct the wanted symplectic embeddings $W_{p,q} \hookrightarrow T^*Wh$. First, we do this on the level of the corresponding base diagrams.

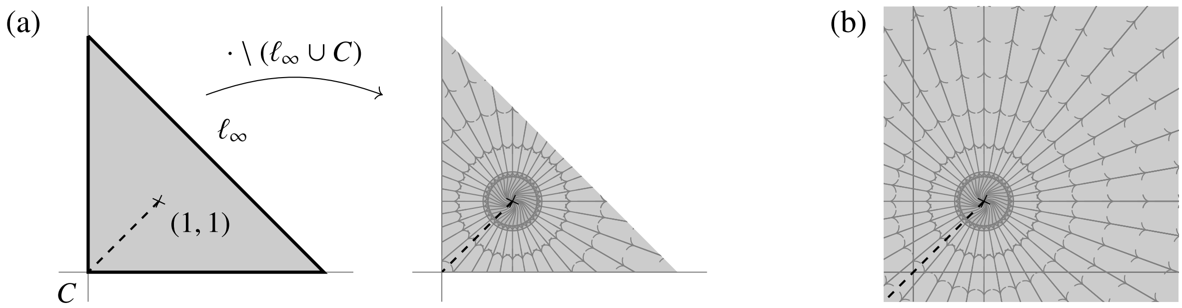

Figure 2.14: (a) an almost toric base diagram of $\mathbb{C}P^2$ and the same almost toric base diagram with the line ℓ_∞ and the conic C removed. Moreover, the Liouville vector field emanating from the Whitney sphere is indicated. The Liouville vector field is not radial in a neighbourhood of the node, which is shown schematically in the figure. (b) The almost toric base diagram of $\mathbb{C}P^2 \setminus (\ell_\infty \cup C)$ extends to an almost toric base diagram of T^*Wh by completing with respect to the Liouville flow.

Proposition 2.4.4. *The almost toric base diagram of $W_{p,q}$ admits an affine embedding into the almost toric base diagram of T^*Wh. Therefore, there exists a symplectic embedding $W_{p,q} \hookrightarrow T^*Wh$ which is fibred in the complement of an arbitrarily small neighbourhood of the singular fibre. In particular, we can view $W_{p,q}$ as a star-shaped subspace of $(T^*Wh, d\lambda_{can})$.*

Proof. Since $0 < q < p$ are coprime, there exist integers $a, b \geq 0$ such that $bq - ap = 1$. For such a, b, the affine transformation

$$M = \begin{pmatrix} q-a & b-p \\ -a & b \end{pmatrix} \in \mathrm{SL}_2(\mathbb{Z})$$

maps the standard wedge $\mathfrak{A}_{p,q}$ to one in which the branch cut points in the $(1,1)$-direction. Since we have removed the boundary from $\mathfrak{A}_{p,q}$, it is clear that the sheared base diagram is an affine subset of the affine base diagram of T^*Wh. This is shown in Figure 2.15.

Now [Sym03, Proposition 4.8] guarantees that the affine embedding of the base diagrams lifts to the desired symplectic embedding. Finally, the fact that the embedding thereby obtained $\iota : W_{p,q} \hookrightarrow T^*Wh$ is star-shaped follows because the base diagram contains the node. Furthermore, since ι respects the almost toric fibrations outside of a small disc around the node, the star-shaped property for $\iota(W_{p,q})$ follows directly from the fact that the base diagram of $W_{p,q}$ is a star-shaped subset of $\mathbb{R}^2$. □

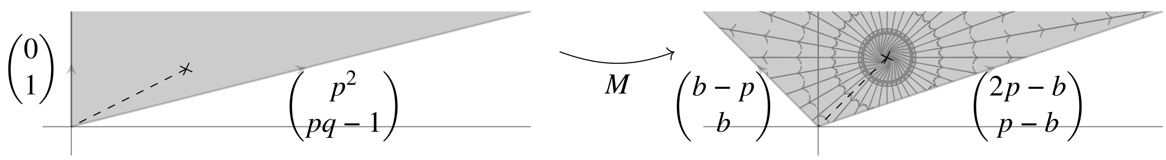

Figure 2.15: The almost toric base diagram of $W_{p,q}$, as a subset of the toric base diagram of T^*Wh.

Remark 2.4.5. For the special case $(p, q) = (1, 1)$, the above proof shows that D^*Wh, the unit cotangent disc bundle of the Whitney sphere, can be symplectically identified with $B_{1,1}(1, 1)\backslash C_{1,1}$, which is exactly Rizell's point of view in [Dim19].

Combining Lemma 2.4.3 and Proposition 2.4.4 we immediately get:

Corollary 2.4.6. *The symplectic embedding $W_{p,q} \hookrightarrow T^*Wh$ induces a weak homotopy equivalence at the level of compactly supported symplectomorphism groups.*

2.4.2 Comb diagrams

Let C be a symplectic submanifold of (M, ω) and let $\mathcal{S}(C)$ be the space of (unparametrized) symplectic submanifolds of M symplectomorphic to C which agree with C outside a (not fixed) compact set. If $\mathcal{S}(C)$ is path-connected, the natural action of $\mathrm{Symp}_c(M)$ on $\mathcal{S}(C)$ is transitive since any path of symplectic submanifolds induces an ambient symplectic isotopy. See [Aur97, Proposition 4]. In that case, we have the homotopy-action sequence

$$\begin{tikzcd} \mathrm{Pr}(C) \arrow[r] & \mathrm{Symp}_c(M) \arrow[d] \\ & \mathcal{S}(C) \end{tikzcd} \tag{2.4.23}$$

where $\mathrm{Pr}(C)$[31] is the stabilizer of C under the action of Symp_c. To relate $\mathrm{Symp}_c(M \setminus C)$ to $\mathrm{Pr}(C)$ we consider the restriction map ρ:

$$\mathrm{Fix}(C) \to \mathrm{Pr}(C) \xrightarrow{\rho} \mathrm{Symp}_c(C) \tag{2.4.24}$$

and the normal derivative map D:

$$\ker(D) \hookrightarrow \mathrm{Fix}(C) \xrightarrow{D} \mathrm{Aut}(\nu C), \tag{2.4.25}$$

where $\mathrm{Aut}(\nu C)$ is the group of symplectic automorphisms of the normal bundle of C, which are the identity outside of a compact set. By a standard Moser argument, $\ker(D)$ is weakly homotopy equivalent to $\mathrm{Symp}_c(M \setminus C)$. Therefore all these groups fit together in the following diagram which we will refer to as a *Comb diagram*.

$$\begin{tikzcd} \mathrm{Symp}_c(M \setminus C) \arrow[r] & \mathrm{Fix}(C) \arrow[r] \arrow[d, "D"] & \mathrm{Pr}(C) \arrow[r] \arrow[d, "\rho"] & \mathrm{Symp}_c(M) \arrow[d] \\ & \mathrm{Aut}(\nu C) & \mathrm{Symp}_c(C) & \mathcal{S}(C) \end{tikzcd} \tag{2.4.26}$$

The key structural property of the Comb diagram that we will use multiple times is that it allows us to relate the (compactly supported) symplectomorphisms of M and those of $M \setminus C$, based on the behaviour of C under the action of the symplectomorphism group. Thus comb diagrams help us to leverage known results about symplectomorphism groups to compute new ones.

[31] The notation Pr refers to the fact that we are considering the subgroup of symplectomorphisms that set-wise preserve the submanifold C, in contrast with the subgroup of symplectomorphisms Fix of those symplectomorphisms that point-wise fix C.

We now demonstrate how to use this diagram to compute $\mathrm{Symp}_c(T^*Wh)$, which will be needed for the computation of $\mathrm{Symp}_c(B_{p,q})$. Again we will rely on [Dim19] for certain technical aspects of the symplectic geometry of T^*Wh. The key input is that the relevant moduli space of symplectic cylinders has the homotopy type of a punctured plane.

Theorem 2.4.7. *The space* $\mathrm{Symp}_c(D^*Wh)$ *is weakly contractible.*

Proof. For $M = B^4$ and C a smooth conic tangent to the coordinate planes at the boundary, we have that $M \setminus C = D^*Wh$, as explained in detail in Proposition 2.4.4; compare also with the original argument in [Dim19, Section 2 and 3]. Gromov shows that $\mathrm{Symp}_c(B^4)$ is contractible, so our plan is to use the Comb diagram to deduce that $\mathrm{Symp}_c(D^*Wh)$ is also weakly homotopy equivalent to a point.

We start by showing that $\mathcal{S}(C) \simeq \mathbb{C}^*$. Let C_s be a family of smooth conics tangent to the coordinate planes at the boundary of the ball B^4, where s is a spherical parameter of S^k for some k. Pick a family of tame almost complex structures J_s on B^4 such that C_s is J_s holomorphic. Since the space of almost complex structures is contractible, we may extend J_s to a bigger family $J_{t,s}$ parametrized by $(t, s) \in I \times S^k$ such that $J_{0,s}$ is constantly the standard integrable complex structure and $J_{1,s} = J_s$. By Theorem 4.3 in [Dim19], the family $J_{t,s}$ gives a corresponding family of symplectic fibrations $f_{t,s} : \mathbb{C}P^2 \backslash \ell_\infty \to \mathbb{C}$ with a unique singular value, say $w_{t,s}$. Because this singular value is unique, we can find smoothly varying complex numbers $z_{t,s} \neq w_{t,s}$. The $I \times S^k$-family of smooth symplectic conics

$$C_{t,s} = f_{t,s}^{-1}(z_{t,s})$$

interpolates between the original family $C_s = C_{1,s}$ and the family $C_{0,s} = f_{0,s}^{-1}(z_{0,s})$. Therefore any symplectic family of smooth conics can be isotoped to a family of smooth fibres of the standard Lefschetz fibration on B^4. Since the space of smooth fibres is just the smooth locus of the standard Lefschetz fibration, it is readily identified with $\mathbb{C}^*$ and thus our desired result follows.

Let us move to investigating the restriction map $\rho : \mathrm{Pr}(C) \to \mathrm{Symp}_c(C)$ via the fibration (2.4.24). Since $\mathrm{Symp}_c(B^4)$ is contractible, we have that $\mathrm{Pr}(C)$ is w.h.e. to $\mathbb{Z}$ via the fibration (2.4.23), where a generator of $\mathrm{Pr}(C)$ is given by the monodromy over a simple loop of smooth symplectic conics around the singular one. This is the standard monodromy of the Milnor fibration of the A_1-curve singularity, which is known to be the standard Dehn twist of a cylinder, which also generates the mapping class group of C. Therefore ρ sends a generator of $\mathrm{Pr}(C)$ to a generator of $\mathrm{Symp}_c(C)$ and so induces a weak homotopy equivalence. From this we also deduce that $\mathrm{Fix}(C) = \ker \rho$ is weakly contractible.

For the last step of the proof, we need to understand the derivative map $D : \mathrm{Fix}(C) \to \mathrm{Aut}(\nu C)$ by investigating the fibration (2.4.25). Since $\mathrm{Fix}(C)$ is weakly contractible and $\mathrm{Aut}(\nu C) \simeq \mathbb{Z}$, the map D necessarily induces a homotopy equivalence to its image and therefore its kernel, namely $\mathrm{Symp}_c(M \setminus C) \simeq \mathrm{Symp}_c(D^*Wh)$ must have trivial homotopy groups.[32]

As explained in Section 2.4.1, T^*Wh is the Liouville completion of D^*Wh and therefore $\mathrm{Symp}_c(D^*Wh)$ is w.h.e. to $\mathrm{Symp}_c(T^*Wh)$, which finishes the proof. □

[32] Our group $\mathrm{Aut}(\nu C)$ is equivalent to the group $\mathcal{G}_2$ in [Eva11], i.e. the group of automorphisms of a symplectic rank 2 bundle over a sphere, where over the north and south pole the automorphism equals the identity.

2.4.3 The extrinsic geometry of the central cylinders $K_{p,q}$

The Comb diagram for $M = B_{p,q}$ and $C = K_{p,q}$ takes the following form:

$$\begin{tikzcd}
\mathrm{Symp}_c(D^*Wh) \arrow[r] & \mathrm{Fix}(K_{p,q}) \arrow[r] \arrow[d, "D"] & \mathrm{Pr}(K_{p,q}) \arrow[r] \arrow[d, "\rho"] & \mathrm{Symp}_c(B_{p,q}) \arrow[d, dashed, "\text{transitive action}"] \\
& \mathrm{Aut}(\nu K_{p,q}) \simeq \mathbb{Z} & \mathrm{Symp}_c(K_{p,q}) \simeq \mathbb{Z} & \mathcal{S}(K_{p,q}) \simeq pt
\end{tikzcd} \tag{2.4.27}$$

Since we will view $B_{p,q}$ as a partial compactification of D^*Wh by adding the cylinder $K_{p,q}$ we need to understand the symplectic moduli space $\mathcal{S}(K_{p,q})$, as well as the maps D and ρ appearing in the Comb diagram, which will then be used to deduce the weak homotopy equivalences claimed in (2.4.27).

Proposition 2.4.8. *The symplectic moduli space $\mathcal{S}(K_{p,q})$ associated to the central cylinder $K_{p,q}$ in $B_{p,q}$ for $p \geq 2$ is weakly contractible.*

Proof. The proof follows the same strategy as that of Theorem 2.4.7. We momentarily drop p, q from the notation for the sake of readability. Take an S^k-family of cylinders K_s in $\mathcal{S}(K)$. Since the indexing set is compact there is a large compact set outside which all of the K_s agree with the central cylinder K. Compactify the rational homology ball to X such that the compactification does not interact with the compact set and write F for the symplectic sphere corresponding to K and F_s for those corresponding to K_s.[33] By construction we have that $x_0 = F \cap D_0 = F_s \cap D_0$ and $x_n = F \cap D_n = F_s \cap D_n$ and that $[F_s] = [F]$ for all s. Now pick an almost complex structure $J \in \mathcal{J}_\tau(\mathcal{D})$ that makes F holomorphic. Recall that by Lemma 2.3.15 this implies that F is the unique smooth ruling passing through x_0. Choose a family $J_s \in \mathcal{J}_\tau(\mathcal{D})$ of almost complex structures that agree near $\mathcal{D}$ with J and that make the corresponding F_s holomorphic. We can now extend this family to an $I \times S^k$-family $J_{t,s}$ such that $J_{0,s} = J$ and $J_{1,s} = J_s$ since the space $\mathcal{J}_\tau(\mathcal{D})$ is weakly contractible.[34]

By Lemma 2.3.15, there is a unique family of embedded symplectic spheres $F_{t,s}$ through x_0 and x_n that are $J_{t,s}$-holomorphic. This provides a deformation of the family F_s supported away from the compactifying divisor $\mathcal{D}$ through embedded symplectic spheres to F, and hence it provides the desired compactly supported homotopy of K_s to K. □

Remark 2.4.9. The essential difference between the geometry of the ball and the geometry of a $B_{p,q}$ when $p \geq 2$ is exactly the fact that the cylinders $K_{p,q}$ are *rigid*, i.e. there exists a *unique* one for each J when $p \geq 2$. In contrast, the central cylinder $K_{1,1}$ comes in a $\mathbb{C}^*$-family.

We now move on to investigate the restriction map ρ and the derivative map D of the Comb diagram. These maps are usually easier to understand when the ambient manifold is compact because the symplectic curve K is then a closed sphere and so $\mathrm{Symp}_c(K)$ and $\mathrm{Aut}(\nu K)$ are actually connected. Therefore it is easy to show that the maps ρ and D are surjective, by building local Hamiltonian isotopies. However, when K is a cylinder, both $\mathrm{Symp}_c(K)$ and $\mathrm{Aut}(\nu K)$ have infinitely many (contractible) connected components, so understanding which of these components are hit by the maps becomes more delicate.

[33] Recall that the compactification procedure was explained in Section 2.3.1 and the geometry of the sphere F was discussed in Lemma 2.3.8.

[34] That $\mathcal{J}_\tau(\mathcal{D})$ is weakly contractible is, for example, shown in [Eva11, Appendix].

Proposition 2.4.10 (Infinitesimal symplectic variations of $K_{p,q}$)**.** *The image of the restriction map* $\rho : \mathrm{Pr}(K_{p,q}) \to \mathrm{Symp}_c(K_{p,q})$ *is exactly the connected components containing the multiples of* τ^p*, the* p*-th power of the standard Dehn twist along a cylinder. The image of the derivative map* $D : \mathrm{Fix}(K_{p,q}) \to \mathrm{Aut}(\nu K_{p,q})$ *is exactly the connected component of* $\mathrm{Aut}(\nu K_{p,q})$ *containing the identity.*

Proof. For both statements, we will consider the universal p-fold covering $A_{p-1} \xrightarrow{\pi} B_{p,q}$ introduced in Section 2.2.1.

The first part follows from topological considerations. First, we notice that any $\psi \in \mathrm{Pr}(K_{p,q})$ lifts to an element $\tilde{\psi} \in \mathrm{Symp}(A_{p-1})$ that preserves the central cylinder of A_{p-1}. Since π is a p-fold covering, if the restriction of $\tilde{\psi}$ on the central cylinder is τ^k, then the restriction of ψ is τ^{pk}, as explained in Lemma 2.2.22. Also by this lemma, for the pintwist $\tau_{p,q}$ we have that the restriction to the central cylinder is τ^p and so $\rho(\tau_{p,q})$ generates the biggest subgroup actually possible.

Moving to the derivative D, let us consider the corresponding, lifted, map on the universal cover $\widetilde{D} : \mathrm{Fix}(K_{p-1}) \to \mathrm{Aut}(\nu K_{p-1})$. Since K_{p-1} is exactly a p-fold cover of $K_{p,q}$, if the image of D contained a component of $\mathrm{Aut}(\nu K_{p,q})$ other than the identity, so would the image of $\widetilde{D}$. Therefore, it is enough to show that $\widetilde{D}$ is homotopically constant.

To do this we argue as follows: it is well known that the Milnor fibre of the A_{p-1} singularity is symplectomorphic to the linear plumbing of the cotangent bundles of $p-1$ Lagrangian spheres. Therefore there exists a Liouville form, say λ, whose flow has as skeleton a chain of Lagrangian spheres, which can also be taken as the matching spheres of the usual Lefschetz fibration on A_{p-1}. Evidently, the matching spheres are disjoint from the central cylinder, so by conjugating with the flow of λ, we see that for any $\psi \in \mathrm{Fix}(K_{p-1}) \subseteq \mathrm{Symp}_c(A_{p-1})$ there exists some Hamiltonian map ϕ such that $\psi \circ \phi$ is supported away from K_{p-1}. Since ψ and $\psi \circ \phi$ are in $\mathrm{Fix}(K_{p-1})$, so is ϕ. Therefore our claim will follow if we can show that $[\widetilde{D}(\psi)] = [\widetilde{D}(\psi \circ \phi)]$ for such a Hamiltonian symplectomorphism $\phi \in \mathrm{Ham}_c(A_{p-1})$. Equivalently, we will show that $[\widetilde{D}(\phi)] = 0 \in \mathbb{Z} \simeq \mathrm{Aut}(\nu K)$. Since Proposition 2.4.8 can also be proven for the central K_{p-1} in A_{p-1}, we have that $\pi_0(\mathcal{S}(K_{p-1})) \cong 1$.[35] Thus, the fact $[\widetilde{D}(\phi)] = 0$ can be deduced by examining the long exact sequence on homotopy groups:

$$\pi_1(\mathrm{Symp}_c(A_{p-1})) \longrightarrow \pi_1(\mathcal{S}(K_{p-1})) \longrightarrow \pi_0(\mathrm{Pr}(K_{p-1})) \longrightarrow \pi_0(\mathrm{Symp}_c(A_{p-1})) \longrightarrow \pi_0(\mathcal{S}(K_{p-1})).$$

Since ϕ is Hamiltonian, it is trivial in $\pi_0(\mathrm{Symp}_c(A_{p-1}))$ and therefore the exact sequence shows that it is the image of the symplectic monodromy around some loop in $\mathcal{S}(K_{p-1})$. Any such monodromy, by definition, preserves the horizontal vector field defined by the velocity of the loop. Therefore $\widetilde{D}(\phi)$ has to be nullhomotopic, concluding the proof. □

2.4.4 Computing $\mathrm{Symp}_c(B_{p,q})$ and finishing the proof

We now have all the ingredients to perform our calculation.

Theorem 2.4.11. *The group* $\mathrm{Symp}_c(B_{p,q})$ *is weakly homotopy equivalent to* $\mathbb{Z}$*, where the generator can be taken to be the pintwist* $\tau_{p,q}$.

[35] This can be proven in exactly the same way as Proposition 2.4.8. However, one then needs the result analogous to Lemma 2.3.15, which is proven in [Buc25, Section 2].

Proof. We know all the necessary parts of the Comb diagram for a $B_{p,q}$:

$$\begin{tikzcd} \mathrm{Symp}_c(D^*Wh) \arrow[r, hook] & \mathrm{Fix}(K_{p,q}) \arrow[r, hook] \arrow[d, "D"] & \mathrm{Pr}(K_{p,q}) \arrow[r, hook] \arrow[d, "\rho"] & \mathrm{Symp}_c(B_{p,q}) \arrow[d] \\ & \mathrm{Aut}(\nu K_{p,q}) & \mathrm{Symp}_c(K_{p,q}) & \mathcal{S}(K_{p,q}) \end{tikzcd} .$$

By Proposition 2.4.10 we have that D is homotopic to the constant map and therefore the inclusion map $\mathrm{Symp}_c(D^*Wh) \hookrightarrow \mathrm{Fix}(K_{p,q})$ is a w.h.e., so Theorem 2.4.7 implies that $\mathrm{Fix}(K_{p,q})$ is weakly contractible. Therefore $\mathrm{Pr}(K_{p,q})$ is w.h.e. to its image under ρ, which by Proposition 2.4.10 and the fact that the components of $\mathrm{Symp}_c(K_{p,q})$ are contractible is w.h.e. to $\mathbb{Z}\langle\tau_{p,q}\rangle$.

Finally, since by Proposition 2.4.8 $\mathcal{S}(K_{p,q})$ is weakly contractible, the inclusion of $\mathbb{Z}\langle\tau_{p,q}\rangle$ to $\mathrm{Symp}_c(B_{p,q})$ is a weak homotopy equivalence. □

Corollary 2.4.12 (Extracting the Hamiltonian)**.** *Suppose that $Y \subseteq B_{p,q}$ is star-shaped and that $L \subseteq Y$ is a Lagrangian (p,q)-pinwheel. Then, there exists a compactly supported Hamiltonian diffeomorphism $\phi^H \in \mathrm{Ham}_c(Y)$ such that $\phi^H(L) = L^{\mathrm{st}}_{p,q}$.*

Proof. Pick some compactification $X_{p,q}$ as in Definition 2.3.5. The subset $Z = X_{p,q} \backslash \mathcal{D}_{p,q}$ is a star-shaped domain of $B_{p,q}$. Therefore, there exists $T \leq 0$ such that $\phi^\lambda_T(Z) := Z' \subset Y$. Since Y is star-shaped, $\phi^\lambda_T(L) := L' \subset Z'$. Recall that ϕ^λ_t acts on the symplectic form only by rescaling, so all the arguments of Section 2.3, in particular Theorem 2.4.11, can be applied to Z' in order to construct a symplectomorphism $\psi \in \mathrm{Symp}_c(Z')$ such that $\psi(L') = L^{\mathrm{st}}_{p,q}$.

Similarly, the arguments of Section 2.4 apply: Z' is a star-shaped subset of $B_{p,q}$ and therefore the inclusion $Z' \hookrightarrow B_{p,q}$ induces a weak homotopy equivalence on the groups of compactly supported symplectomorphisms. Combined with the fact that the pintwist $\tau_{p,q}$ can be arranged to have support arbitrarily close to $L^{\mathrm{st}}_{p,q}$, we get that $\tau_{p,q}$ generates the symplectic mapping class group of Z' and thus ψ, by Theorem 2.4.11, factors as $\psi = \tau^k_{p,q} \circ \phi^H$, for $k \in \mathbb{Z}$ and H a time-dependent Hamiltonian.

The Hamiltonian diffeomorphism $\phi^H = \tau^{-k}_{p,q} \circ \psi$ maps L' to $\tau^{-k}_{p,q}(L^{\mathrm{st}})$ which is Lagrangian isotopic to L^{st}, by Lemma 2.2.22. Therefore, L is Lagrangian isotopic to L^{st} and ultimately Hamiltonian isotopic to it, since $H^1(L^{\mathrm{st}}, \mathbb{R}) = 0$, as in Lemma 2.2.22 (3). □

2.5 Applications

2.5.1 Non-squeezing theorems

This section sheds light on the non-squeezing theorem proven by the first and third authors in [AH25]. There, a non-squeezing theorem for $B_{n,1}$ was proven for all $n \geq 1$. We reprove this theorem here and extend the theorem to all $B_{p,q}$. Moreover, we prove the non-squeezing theorem for two symplectically different pin-cylinders.

Recall that the pin-ellipsoids and pin-balls were defined via the almost toric base diagram $\mathfrak{A}_{p,q}$ of $B_{p,q}$ in analogy to how the usual ball and ellipsoids in $\mathbb{C}^2$ can be defined by their moment image as shown in Figure 2.2, and that $B_{p,q}(\lambda) := E_{p,q}(\lambda, \lambda)$. By a slight abuse of notation, we define for $\lambda > 0$ the two pin-cylinders $E_{p,q}(\lambda, \infty)$ and $E_{p,q}(\infty, \lambda)$ via the almost toric base diagram $\mathfrak{A}_{p,q}(\lambda, \infty)$ and $\mathfrak{A}_{p,q}(\infty, \lambda)$, as shown in Figure 2.16.

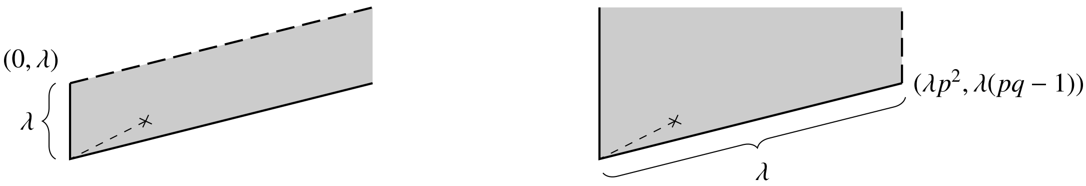

Figure 2.16: On the left, the almost toric base diagram $\mathfrak{A}_{p,q}(\infty, \lambda)$ defining $E_{p,q}(\infty, \lambda)$; on the right, the corresponding almost toric base diagram for $E_{p,q}(\lambda, \infty)$.

Remark 2.5.1. It is important to observe that if $p \geq 3$, the two pin-cylinders $E_{p,q}(\lambda, \infty)$ and $E_{p,q}(\infty, \lambda)$ are not symplectomorphic. That the usual standard cylinder $E(\lambda, \infty)$ is symplectomorphic to $E(\infty, \lambda)$ is obvious by exchanging the coordinate planes and that $E_{2,1}(\lambda, \infty)$ and $E_{2,1}(\infty, \lambda)$ are symplectomorphic follows from the fact that $\mathfrak{A}_{2,1}(\lambda, \infty)$ and $\mathfrak{A}_{2,1}(\infty, \lambda)$ are related by an integral affine transformation. That this symmetry no longer holds true for $p \geq 3$ will be clear from the proofs of Theorem 2.5.2 and Theorem 2.5.5 and is carried out explicitly in Remark 2.5.6.

Theorem 2.5.2. *If there exists a symplectic embedding $B_{p,q}(\alpha) \overset{s}{\hookrightarrow} E_{p,q}(\lambda, \infty)$, then $\alpha \leq \lambda$.*

Remark 2.5.3. It should be mentioned that the proof here follows, in spirit, the strategy of [Ada+25, Theorem 1.5.2] rather than the one in [AH25, Theorem D]. Namely, it uses Theorem 2.3.24—the fact that Lagrangian pinwheels are unique up to symplectomorphism in their neighbourhoods—and then employs visible curves to obstruct the symplectic embedding, using techniques akin to those used in [MS25b; MS25a]. The main difference is that in [AH25], the obstruction arises in a more ad hoc way, by finding a well-chosen exceptional curve in the compactification $\widetilde{X}_{p,q}$. The main observation of [Ada+25] is that the exceptional curve that provides the obstruction naturally appears as a component of a broken fibre of a regulation of $\widetilde{X}_{p,q}$. Thus we avoid the complicated combinatorics of [AH25].

Proof of Theorem 2.5.2. Assume that there is a symplectic embedding $\iota : B_{p,q}(\alpha) \hookrightarrow E_{p,q}(\lambda, \infty)$. By Corollary 2.4.12 we can assume that ι and the visible embedding ι_{vis} coincide on a small neighbourhood of $L^{\text{st}}_{p,q}$, i.e. on $B_{p,q}(\varepsilon)$ for some small $0 < \varepsilon$. Now, view ι as an embedding into a "non-minimal" compactification as shown in Figure 2.17 (a).[36] Rationally blowing up the balls $B_{p,q}(t\alpha)$ for $t \in (0, 1]$ along the embedding ι yields a family of symplectic manifolds $\left(\widetilde{X}_{p,q}(t), \widetilde{\omega}(t)\right)$. We choose the rational blow-up data such that the last sphere in the Wahl chain C_m has maximal symplectic area. For the sake of clarity of the proof, we assume that all the other spheres have vanishing symplectic area. It is not hard to see that this setup results in the following:

(i) The symplectic area of C_m is

$$\int_{C_m} \widetilde{\omega}(t) = \frac{t\alpha p^2}{p^2 - (pq+1)},$$

as shown in Figure 2.17 (c).[37]

[36]This compactification is exactly the one that was discussed in Footnote 18.
[37]A computation in toric geometry relying on (2.2.4) shows that the direction associated to C_m in the minimal resolution is given by $(p^2 - (pq+1), pq - (q^2+1))$. Then the length of the longest segment that can fit into $\mathfrak{A}_{p,q}(t\alpha, t\alpha)$ is given by the term in the equation. See also [AH25, Appendix B] for the $(p, q) = (n, 1)$ case.

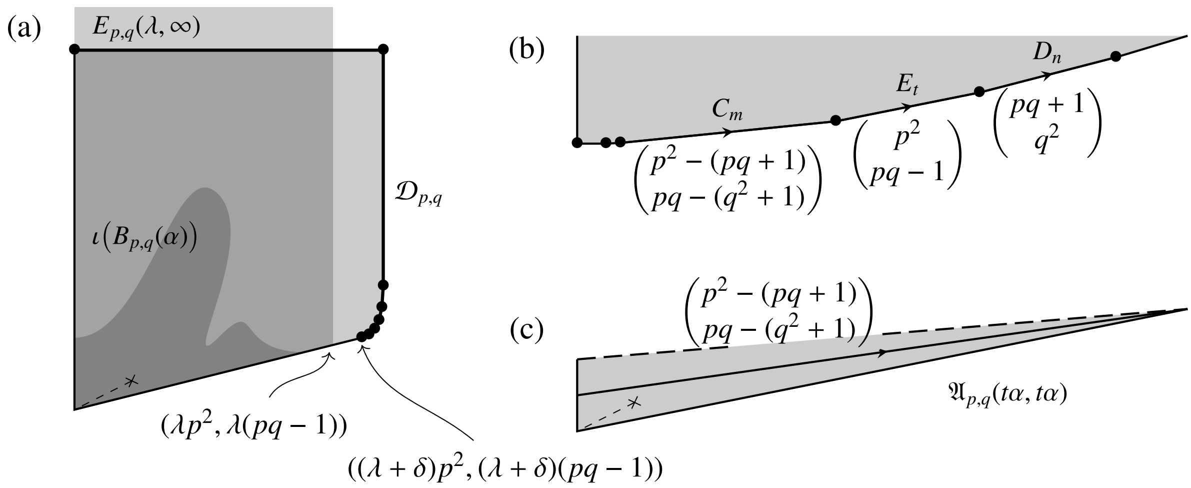

Figure 2.17: (a) The almost toric base diagram of the compactification used in the proof of Theorem 2.5.2. The darker shaded region indicates how the embedded pin-ball $\iota\big(B_{p,q}(\alpha)\big)$ sits inside the compactification after the initial adjustment. (b) After rationally blowing up the ball $B_{p,q}(t\alpha)$ along the embedding ι for $t\alpha \leq \varepsilon$, which means precisely along ι_{vis}, the Wahl chain and the compactifying divisor are connected by a visible exceptional divisor on the toric boundary. The vectors can easily be computed from (2.2.4). (c) The data used for the rational blow-up in the proof of Theorem 2.5.2.

(ii) For $t \in (0,1]$ such that $t\alpha \leq \varepsilon$, the rational blow-up is visible and the piece of the toric boundary connecting the end of the Wahl chain $C_{p,q}$ to the end of the compactifying divisor $\mathcal{D}_{p,q}$ is a symplectically embedded (-1)-sphere E_t. This sphere E_t intersects $C_{p,q}$ in exactly one point on C_m and similarly for $\mathcal{D}_{p,q}$. See Figure 2.17 (b).

As in Section 2.3, we write

$$E_t = \sum a_i D_i + \sum c_j C_j \in H_2\big(\widetilde{X}_{p,q}(t), \mathbb{Q}\big).$$

It is well known that (-1)-spheres are stable under deformations of the symplectic form as explained in [Wen18, Chapter 5].[38] This means that we have for all $t \in (0,1]$

$$0 < \int_{E_t} \widetilde{\omega}(t) = \sum a_i \int_{D_i} \widetilde{\omega}(t) + \sum c_j \int_{C_j} \widetilde{\omega}(t).$$

Since the compactifying divisor is disjoint from the embedding, the sum involving the terms connected to $\mathcal{D}_{p,q}$ will be constant. Considering the limit $t \to 0$ we see that the sum is, in fact, equal to $\lambda + \delta$. Since we assumed that the symplectic area of all but the last sphere in the Wahl chain vanishes, the inequality reads

$$0 < (\lambda + \delta) + c_m \int_{C_m} \widetilde{\omega}(t).$$

We can calculate the coefficient c_m via the equation $\boldsymbol{c} = M_C^{-1}\boldsymbol{e}_m$, where $\boldsymbol{e}_m$ is the intersection

[38] In the proof of [Ada+25, Theorem 1.5.2] this was justified by invoking spherical Gromov invariants.

profile of E_t with $C_{p,q}$ mentioned in (ii), i.e. the mth canonical basis vector. From (2.2.9) we deduce that

$$c_m = -\frac{1}{p^2}[pq-1]^{-1} = -\frac{1}{p^2}(p^2-(pq+1)),$$

which follows from Lemma 2.2.33 because $(M_C^{-1})_{mm} = -(e_m f_m)/p^2 = -e_m/p^2$. Using (i) the inequality therefore reads

$$t\alpha \le \lambda + \delta,$$

where $\delta > 0$ can be chosen arbitrarily small. This means $\alpha \le \lambda$, which is what we wanted to show. □

Remark 2.5.4. The curve that is used to obstruct the embedding in this proof is exactly the one that was used in the proof of the non-squeezing theorem in the $(p,q) = (n,1)$-case by the first and third authors in [AH25, Theorem D]. In the proof above the geometric situation becomes much more transparent, because the uniqueness theorem for Lagrangian pinwheels is available.

Theorem 2.5.5. *If there exists a symplectic embedding $B_{p,q}(\alpha) \overset{s}{\hookrightarrow} E_{p,q}(\infty,\lambda)$, then $\alpha \le \lambda$.*

Proof. The proof is analogous to the proof of Theorem 2.5.2 and is therefore omitted. □

Remark 2.5.6. Note that the proof of Theorem 2.5.2 also shows that for $p \ge 3$, $E_{p,q}(\alpha,1)$ embeds into $E_{p,q}(1,\infty)$ if and only if $\alpha \le 1$. Since $E_{p,q}(\alpha,1)$ embeds trivially into $E_{p,q}(\infty,1)$ for all $\alpha > 1$, it follows that the two pin-cylinders cannot be symplectomorphic.

Having established the non-squeezing theorem for pin-balls and recalling that in [Ada+25, Theorem 1.5.2] it was proven that the set

$$\mathcal{A}_{p,q} := \left\{(\alpha,\beta) \in \mathbb{R}^2_{>0} \;\middle|\; E_{p,q}(\alpha,\beta) \overset{s}{\hookrightarrow} \mathbb{C}P^2(1)\right\}$$

is either empty or carries a symplectic staircase structure, a natural question emerges.

Question. *Do the symplectic problems of embedding pin-ellipsoids into pin-balls carry the structure of symplectic staircases?*

In the forthcoming paper [AH] this question will be answered in the positive for $(p,q) = (2,1)$ and the corresponding ellipsoid domains in $B_{2;1,1} \cong T^*S^2$, by showing that the set

$$\mathcal{E}_{2,1} := \left\{(\alpha,\beta) \in \mathbb{R}^2_{>0} \;\middle|\; E_{2,1}(\alpha,\beta) \overset{s}{\hookrightarrow} B_{2,1}(1)\right\}$$

coincides with $\mathcal{A}_{2,1}$ and similarly for the $(2;1,1)$-case. These two sets are particularly interesting to compute since $E_{2,1}(\alpha,\beta)$ and $E_{2;1,1}(\alpha,\beta)$ are domains in $T^*\mathbb{R}P^2$ and T^*S^2, i.e. they carry inherent physical meaning and, moreover, they are easy to describe in terms of Finsler structures on the respective cotangent bundles.

2.5.2 Eliashberg–Polterovich's theorem

We illustrate the methods of this paper in a slightly different situation by reproving an old theorem of Eliashberg and Polterovich:

Theorem 2.5.7 ([EP96, Theorem 1.1.A])**.** *Suppose that $\Delta \subseteq (\mathbb{R}^4, \omega_{\mathrm{st}})$ is an embedded Lagrangian plane that coincides with a flat Lagrangian plane Δ_0 at infinity. Then Δ and Δ_0 are isotopic via a compactly supported Hamiltonian isotopy.*

Remark 2.5.8. Within the context of this paper, this theorem should be understood to be the $(p, q) = (1, 1)$ case of Theorem A. The strategy of proof is very similar to that of Theorem A. The main difference from the rest of this paper is that a "(1, 1)-pinwheel" is an embedded Lagrangian disc and therefore we have to carefully control the boundary behaviour of the embedded disc.

Remark 2.5.9. Theorem 2.5.7 is a foundational result in symplectic topology, as it shows that there are no "local" Lagrangian knots, i.e. there does not exist a local operation that knots a Lagrangian submanifold in dimension four. Eliashberg and Polterovich's proof uses the technique of filling by holomorphic discs. Since one cannot fill a Lagrangian plane by discs, they introduce an auxiliary object called a "plug", see also the survey [PS24]. In our proof the auxiliary object will be a compactifying divisor, which allows us to use rational closed J-curves, i.e. regulations, in order to unknot the Lagrangian plane.

The model of a flat Lagrangian plane that will make our proof most transparent is the Lagrangian plane that is given by the fixed locus of the anti-symplectic involution τ on $\mathbb{C}^2$ defined by $\tau(z_1, z_2) = (\bar{z}_2, \bar{z}_1)$. Denote this fixed point set by Δ_τ. Then

$$\Delta_\tau = \{(z, \bar{z}) \in \mathbb{C}^2 \mid z \in \mathbb{C}\}$$

and in polar coordinates $z_j = r_j e^{i\theta_j}$ on $\mathbb{C}^2$ we have

$$\Delta_\tau = \{(z_1, z_2) \mid r_1 = r_2 \text{ and } \theta_1 = -\theta_2\},$$

which means that Δ_τ is the visible Lagrangian that projects to the diagonal ray under the standard toric moment map on $\mathbb{C}^2$.

Proof of Theorem 2.5.7. Without loss of generality we assume that Δ coincides with Δ_τ outside $B^4(1 - \varepsilon)$. The proof of Theorem 2.5.7 proceeds in four steps.

Compactifying Lagrangian planes

The first step is to transfer the problem to a compact situation. Performing a symplectic cut along the standard $S^3 \subseteq \mathbb{C}^2$, with respect to the standard diagonal Hamiltonian action, we obtain $(\mathbb{C}P^2, \omega_{\mathrm{FS}})$. Moreover, the Lagrangian planes Δ and Δ_τ correspond to two embedded Lagrangian $\mathbb{R}P^2$s, which we denote by L and L_τ in the following, that coincide in a neighbourhood of the line at infinity $\mathbb{C}P^1_\infty \subseteq \mathbb{C}P^2$.[39]

Constructing a symplectomorphism of the complements of the Lagrangian $\mathbb{R}P^2$s in $\mathbb{C}P^2$

Recall that the rational blow-up of a Lagrangian $\mathbb{R}P^2$ in a symplectic manifold (X, ω) can be viewed as an operation induced by a symplectic cut. Pick a Weinstein neighbourhood of the Lagrangian $\mathbb{R}P^2$, which, in particular, defines a symplectic embedding of the unit cotangent disc bundle $D^*_\varepsilon\mathbb{R}P^2$ for some small $\varepsilon > 0$, where $D^*_\varepsilon\mathbb{R}P^2$ is defined with respect to the round metric on $\mathbb{R}P^2$. Then a symplectic cut can be performed along the periodic geodesic flow on $\partial D^*_\varepsilon\mathbb{R}P^2 := S^*_\varepsilon\mathbb{R}P^2$, since it is induced by a Hamiltonian function. See [BLW14, Section 2.2.1] for details on this construction and [Aud07; Ler95] for the original references.

In our situation we can investigate the effect of rationally blowing up L and L_τ on $\mathbb{C}P^2$ explicitly. Note that since L and L_τ coincide in a neighbourhood of $\mathbb{C}P^1_\infty$ the effect of the rational

[39] Since the involution that defines Δ_τ descends to $\mathbb{C}P^2$, it follows that the Lagrangians L and L_τ are obtained by gluing Lagrangian discs to a Lagrangian Möbius strip along their boundaries. This is also clear from the toric moment map, as illustrated in Figure 2.18.

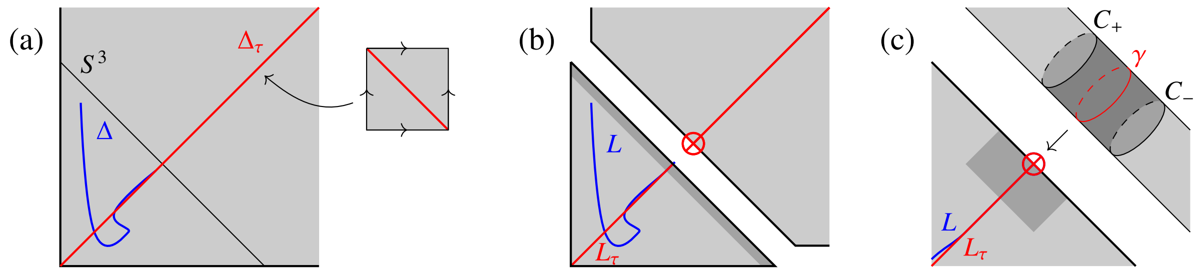

Figure 2.18: (a) Schematic illustration of the initial situation of Theorem 2.5.7. The plane Δ coincides with Δ_τ outside $B^4(1-\varepsilon)$. (b) The situation after the symplectic cut as described in Section 2.5.2. The Lagrangian $\mathbb{R}P^2$s L and L_τ coincide in a small neighbourhood of the line at infinity $\mathbb{C}P^1_\infty \subseteq \mathbb{C}P^2$, which is shaded darker in the figure. On the complementary piece obtained by the symplectic cut we obtain two Lagrangian Möbius strips that coincide. (c) The explicit Weinstein neighbourhood of the embedded, shared, Möbius strip shaded in darker gray and the configuration in $\mathbb{C}P^1_\infty$, which projects to the toric boundary.

blow-up on $\mathbb{C}P^1_\infty$ will be the same if the Weinstein neighbourhoods of L and L_τ coincide near $\mathbb{C}P^1_\infty$. Hence, we will consider L for the moment.

Lemma 2.5.10. *L can be rationally blown up in such a way that the rational blow-up (X, ω) contains a chain of embedded symplectic spheres $\mathcal{T} = (F^+, \Sigma, F_-)$ of self-intersection $[F^\pm]^2 = 0$ and $[\Sigma]^2 = -4$, where $F^+ \cup F^-$ is the "proper transform" of the line $\mathbb{C}P^1_\infty$.*

Proof. Consider the sphere $S^2 \subseteq \mathbb{R}^3$, equipped with the round metric g, and parametrize the equator by $\gamma(\theta) = (\cos(\theta), \sin(\theta), 0)$. Let s be the signed geodesic distance from the equator. Then

$$\gamma_s(\theta) = (\cos s \cos\theta, \cos s \sin\theta, \sin s)$$

is the displacement of $\gamma(\theta)$ by distance s along the normal geodesic to the equator. The round metric in these coordinates is given by $g = ds^2 + \cos^2(s)d\theta^2$. Taking the quotient $S^2/\!\sim\, = \mathbb{R}P^2$ by the antipodal map we find a tubular neighbourhood of the equator:

$$M_\rho := (S^1 \times (-\rho, \rho))/\!\sim \quad \text{where } (\theta, s) \sim (\theta + \pi, -s) \text{ and } \rho < \frac{\pi}{2},$$

which is a Möbius strip. The metric g descends to the quotient because it is invariant under the equivalence relation. Consider the cotangent bundle of the Möbius strip T^*M_ρ with coordinates $(s, \theta, p_s, p_\theta)$. Then taking the dual metric on T^*M_ρ, the codisc bundle of radius $\varepsilon > 0$ and its boundary are given by

$$D^*_\varepsilon M_\rho = \left\{ p_s^2 + \frac{p_\theta^2}{\cos^2(s)} \le \varepsilon^2 \right\} \quad \text{and} \quad S^*_\varepsilon M_\rho = \left\{ p_s^2 + \frac{p_\theta^2}{\cos^2(s)} = \varepsilon^2 \right\}.$$

Near the equator γ of $\mathbb{C}P^1_\infty$, which is the core circle of L, we have coordinates of the form $T^*S^1_{(\tau, p_\tau)} \times \mathbb{C}_{z=u+iv}$, in which the symplectic form ω_{FS} reads $dp_\tau \wedge d\tau \oplus du \wedge dv$, such that $\gamma = \{p_\tau = z = 0\}$ and the part of $\mathbb{C}P^1_\infty$ contained in this chart is given by $\{z = 0\}$. These coordinates come directly from the toric geometry as shown in Figure 2.18 (c). In these coordinates define the embedding $\iota : D^*_\varepsilon M_\rho \hookrightarrow \mathbb{C}P^2$ via

$$\iota\left([\theta, s, p_\theta, p_s]\right) := \left(2\theta, \frac{1}{4}(s^2+p_s^2)+\frac{1}{2}p_\theta, e^{-i\theta}(s-ip_s)\right), \tag{2.5.28}$$

which is easily seen to be well-defined and symplectic. Under the usual moment map on $T^*S^1 \times \mathbb{C}$, given by $\mu(\tau, p_\tau, z) := (p_\tau, |z|^2/2)$, the codisc bundle $D^*_\varepsilon M_\rho$ projects to a thin strip around the ray $\{(\lambda, 2\lambda) \mid \lambda \in \mathbb{R}_{\geq 0}\} \subseteq \mathbb{R}\times\mathbb{R}_{\geq 0}$.[40] The intersection of $\mathbb{C}P^1_\infty$ with $\iota(S^*_\varepsilon M)$ in this chart is given by

$$\iota(S^*_\varepsilon M) \cap \mathbb{C}P^1_\infty = \{z=0, \quad p_\tau = \pm\varepsilon/2\} \tag{2.5.29}$$

meaning that the intersection $\iota(S^*_\varepsilon M_\rho)\cap\mathbb{C}P^1_\infty$ consists of the images of the two circles

$$C_\pm := \{s=0, p_s=0, p_\theta=\pm\varepsilon\} \subseteq S^*_\varepsilon M_\rho,$$

see Figure 2.18 (c). The geodesic flow on $S^*_\varepsilon M_\rho$ is induced by the Hamiltonian flow of the standard quadratic Hamiltonian $H = \frac{1}{2}\|\cdot\|^2$ and the geodesic equations are easily computed from the Hamiltonian equations:

$$\dot\theta = \frac{p_\theta}{\cos^2(s)}, \qquad \dot s = p_s, \qquad \dot p_\theta = 0, \qquad \dot p_s = -p_\theta^2\frac{\sin(s)}{\cos^3(s)}. \tag{2.5.30}$$

This implies that $C_\pm$ are genuine geodesics, because the geodesic equations (2.5.30) on $C_\pm$ are $\dot\theta = \pm\varepsilon$, $\dot s = 0$, $\dot p_\theta = 0$, $\dot p_s = 0$, i.e. the geodesic flow just rotates $C_\pm$.

Now, consider a symplectic embedding $\psi : D^*_\varepsilon\mathbb{R}P^2 \hookrightarrow \mathbb{C}P^2$ obtained by restricting a Weinstein neighbourhood of $L \subseteq \mathbb{C}P^2$. Choosing $\varepsilon > 0$ and $\rho > 0$ small enough we can assume that $\psi|_{D^*_\varepsilon M_\rho}$ coincides with ι, defined in (2.5.28), and that $\psi|_{D^*_\varepsilon\mathbb{R}P^2\setminus D^*_\varepsilon M_\rho}$ is disjoint from $\mathbb{C}P^1_\infty$. Denote the two symplectic discs in $\mathbb{C}P^1_\infty$ that form the complement of the annulus $\mathbb{C}P^1_\infty \cap \psi(D^*_\varepsilon M_\rho)$ by $D^\pm$, as shown in Figure 2.18. Perform the rational blow-up of L by symplectically cutting along the embedded copy of $S^*_\varepsilon\mathbb{R}P^2$ and denote the resulting symplectic manifold by (X,ω). Since the boundaries of $D^\pm$ are geodesics, these two discs are two embedded symplectic spheres in (X,ω), which we denote by $F^\pm$.

The goal is now to compute the self-intersection number of these spheres. The computation is completely symmetric for the two spheres, and hence we carry it out for F^+. In order to do so we compute the Euler number of its normal bundle. The normal bundle $N_{\mathbb{C}P^1_\infty/\mathbb{C}P^2}$ is trivialized over a neighbourhood of C^+ by ι, and we can define the local section $\partial_u = \partial_{\Re(z)}$ in this trivialization. Moreover, because the normal bundle $N_{\mathbb{C}P^1_\infty/\mathbb{C}P^2}$ is trivial over the disc D_+ we can extend this section to a section ν_+ over D^+. Since the geodesic flow leaves the normal coordinate invariant, the section ν_+ descends to give a nowhere vanishing section of $N_{F^+/X}$, which implies:

$$F^+\cdot F^+ = e\left(N_{F^+/X}\right) = 0. \qquad \square$$

Remark 2.5.11. Lemma 2.5.10 can also be proven in another way, by working in complex coordinates. Oakley and Usher [OU16, Section 3] construct an explicit Weinstein neighbourhood of L_τ and a Hamiltonian torus action on $\mathbb{C}P^2\backslash L_\tau$.[41] This construction then allows one to perform the rational blow-up along L_τ in a visible manner. Since L and L_τ coincide near the line at

40 This discussion is completely analogous to the immersion of Lagrangian cylinders in [Eva23, Section 5.3].
41 This construction was also described by Wu [Wu15, Section 3].

infinity, this will yield an alternative proof of Lemma 2.5.10.

We view $T^*\mathbb{R}P^2$ as the quotient of $\{(p,q) \in \mathbb{R}^3 \oplus \mathbb{R}^3 \mid |q| = 1, \langle p, q\rangle = 0\}$ by the usual $\mathbb{Z}_2$-action and denote by $D_1^*\mathbb{R}P^2$ the unit codisc bundle. We view $\mathbb{C}P^2$ as the symplectic reduction of the standard sphere with radius $\sqrt{2}$, i.e. $S^5(\sqrt{2}) \subset \mathbb{C}^3$, and thus we understand elements of $\mathbb{C}P^2$ as tuples $[z_1 : z_2 : z_3]$ satisfying $|z_1|^2 + |z_2|^2 + |z_3|^2 = 2$, up to multiplication by $e^{i\theta}$.

Oakley and Usher construct an explicit Weinstein neighbourhood of L_τ via a map $\varphi : D_1^*\mathbb{R}P^2 \to \mathbb{C}P^2$, which is a symplectomorphism onto its image $\mathbb{C}P^2 \backslash Q$, where Q is the Fermat quadric $\{z_1^2 + z_2^2 + z_3^2 = 0\} \subset \mathbb{C}P^2$. Moreover, there is an explicit Hamiltonian torus action on $\mathbb{C}P^2 \backslash L_\tau$ given by

$$G([z_1, z_2, z_3]) := \mathrm{Im}(\bar{z}_1 z_2) \quad \text{and} \quad H([z_1, z_2, z_3]) := \frac{1}{4}\sqrt{4 - \left|\sum z_i^2\right|^2}. \tag{2.5.31}$$

The moment image is shown in Figure 2.19. The line at infinity $\mathbb{C}P^1_\infty$ is given by the equation

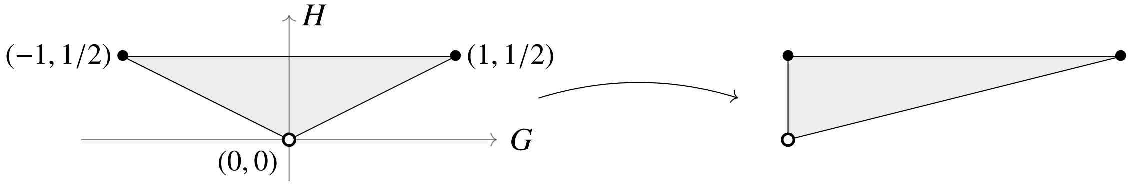

Figure 2.19: On the left, the moment image of (G, H), with the origin missing. Note that it is clear from (2.5.31) that L_τ sits over the origin. On the right, the moment image after an affine $SL_2(\mathbb{Z})$ transformation.

$\{z_3 = 0\}$ and, since it has real coefficients, intersects L_τ in a circle (an algebraic line of $\mathbb{R}P^2$). Thus $\mathbb{C}P^1_\infty$ splits into two discs in $\mathbb{C}P^2 \backslash L_\tau$, related by the involution τ. It is straightforward to check that each of these discs projects to one of the slanted boundary segments of the moment image.

The rational blow-up is then given by performing a symplectic cut along the visible hypersurface $H^{-1}(\alpha)$, for some $0 < \alpha < 1/2$, and the resulting moment image is the usual moment image of the Hirzebruch surface $\mathbb{F}_4$.

To recapitulate: performing the rational blow-up along the embedding ψ, introduced in the proof of Lemma 2.5.10, transforms $\mathbb{C}P^1_\infty$ into two embedded symplectic (0)-spheres $F^\pm$, which intersect the (-4)-sphere Σ, introduced by the rational blow-up, each in a single point, i.e. X contains the linear chain of symplectic spheres $\mathcal{T} := (F^+, \Sigma, F^-)$. Since the rational blow-up preserves rationality, see for example [BLW14, Lemma 2.3] or [AH25, Corollary 3.10], we know that (X, ω) is diffeomorphic to $S^2 \times S^2$ and that $[F^\pm] \in H_2(X; \mathbb{Z})$ is a fibre.

Since the argument that (X, ω) is rational and ruled can be carried out quite easily, we give it for the sake of completeness. Define $F := [F^\pm] \in H_2(X; \mathbb{Z})$ and for $J \in \mathcal{J}(X, \omega)$ consider the moduli space $\mathcal{M}_{0,1}(F, J)$ of rational J-holomorphic curves representing F with one marked point.[42] Moreover, define the following subspace of the space of ω-tame almost complex struc-

[42]Note that a priori we only know that $[\Sigma]$ and $[F]$ form an integral basis of the free part of $H_2(X; \mathbb{Z})$ because of an argument involving the intersection matrix of $[\Sigma]$ and $[F]$ and using the fact that $b_2(X; \mathbb{Q}) = 2$. However, since $F^\pm$ are symplectic spheres, they cannot represent a torsion element in $H_2(X; \mathbb{Z})$, so the integer homology classes they represent coincide.

tures
$$\mathcal{J}(\mathcal{T}) := \{J \in \mathcal{J}(X,\omega) \mid \text{each component of } \mathcal{T} \text{ is } J\text{-holomorphic}\}.$$
The space $\mathcal{J}(\mathcal{T})$ is non-empty since $\mathcal{T}$ is a symplectic normal crossing divisor.

Lemma 2.5.12. *For every $J \in \mathcal{J}(\mathcal{T})$, the moduli space $\mathcal{M}_{0,1}(F,J)$ is a compact manifold of dimension* $\dim(\mathcal{M}_{0,1}(F,J)) = 4$*. Moreover, (X,ω) is regulated by the rational curves in $\mathcal{M}_{0,0}(F,J)$. In particular, the regulation has no broken fibre and (X,ω) is rational and ruled.*

Proof. Note that F is a primitive homology class, because $\Sigma \cdot F = 1$, which in particular implies that all curves in $\mathcal{M}_{0,1}(F,J)$ are simple. Therefore, it follows by automatic transversality, see for example [Wen18, Chapter 2] or [HLS97], that $\mathcal{M}_{0,1}(F,J)$ is a smooth manifold of dimension $\dim(\mathcal{M}_{0,1}(F,J)) = 4$, because $c_1(F) = F \cdot F + 2 = 2$ and hence the index of a curve $u \in \mathcal{M}_{0,1}(F,J)$ is
$$\operatorname{ind}(u) = -2 + 2c_1([u]) = 2 > -2.$$
This implies that the evaluation map ev has degree one. The statement about the compactness of $\mathcal{M}_{0,1}(F,J)$ follows directly from the fact that F is primitive and from positivity of intersection. Moreover, Σ defines a section of the regulation and so (X,ω) is a rational and ruled surface, diffeomorphic to $S^2 \times S^2$. □

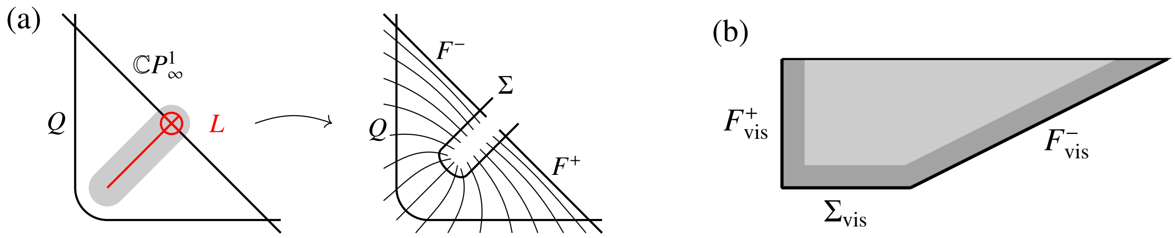

Figure 2.20: (a) A schematic illustration of the effect of the rational blow-up of L. Shaded in gray is the Weinstein neighbourhood discussed in Section 2.5.2. The "proper transform" of the line $\mathbb{C}P^1_\infty$ is given by the two fibres F^+ and F^-. (b) The reference model, namely the standard toric fibration on $\mathbb{F}_4$, with the configuration $\mathcal{T}_{\text{vis}} = (F^+_{\text{vis}}, \Sigma_{\text{vis}}, F^-_{\text{vis}})$ and the neighbourhood ν_{vis} shaded in darker gray.

Remark 2.5.13. In Figure 2.20 we illustrate how the regulation of (X,ω) is related to the initial situation in $\mathbb{C}P^2$. While it is general knowledge that in the complement of an embedded Lagrangian $\mathbb{R}P^2$ in $\mathbb{C}P^2$ there is always a smooth quadric, as can be shown, for example, by neck-stretching, in our case this follows a posteriori: in a normal neighbourhood of the configuration $\mathcal{T}$ we find a positive section, of self-intersection $+4$, of the regulation in the class F, which therefore also exists in $\mathbb{C}P^2 \setminus L$.

Constructing a compactly supported symplectomorphism of $\mathbb{C}^2$ that identifies the Lagrangian planes.

By Section 2.5.2, we find symplectic embeddings $\psi, \psi_\tau : D^*_\varepsilon\mathbb{R}P^2 \hookrightarrow \mathbb{C}P^2$, obtained by restricting Weinstein neighbourhoods of $L, L_\tau \subseteq \mathbb{C}P^2$, that coincide on $D^*_\varepsilon M_\rho$. Performing the rational blow-up along the embedded copies of $S^*_\varepsilon\mathbb{R}P^2$ yields two symplectic manifolds (X,ω) and (X_τ,ω_τ).

Moreover, these rational blow-ups contain linear chains of symplectic spheres $\mathcal{T} := (F^+, \Sigma, F^-)$ and $\mathcal{T}_\tau := (F_\tau^+, \Sigma_\tau, F_\tau^-)$. By the symplectic neighbourhood theorem there are symplectic embeddings $\phi : \nu \hookrightarrow (\mathbb{F}_4, \omega_4)$ and $\phi_\tau : \nu_\tau \hookrightarrow (\mathbb{F}_4, \omega_4)$ of normal neighbourhoods of these chains, such that $\phi_\tau(\nu_\tau) = \nu_{\text{vis}}$ and $\phi_\tau(\mathcal{T}_\tau) = \mathcal{T}_{\text{vis}}$ and similarly for ϕ, as shown in Figure 2.20. Here $(\mathbb{F}_4, \omega_4)$ denotes the 4th Hirzebruch surface equipped with a symplectic form ω_4 that realises the symplectic data of $\mathcal{T}$ and $\mathcal{T}_\tau$. Since X is rational and ruled, the complement $X \setminus \mathcal{T}$ is a disc bundle over an annulus and hence minimal. Therefore, Theorem 2.3.27 yields a symplectomorphism

$$\Phi : ((X, \omega), \mathcal{T}) \to ((\mathbb{F}_4, \omega_4), \mathcal{T}_{\text{vis}}),$$

as in the proof of Theorem 2.3.24, and similarly Φ_τ. Using these symplectomorphisms we obtain a symplectomorphism $\Psi : ((\mathbb{C}P^2, \omega_{\text{FS}}), L) \to ((\mathbb{C}P^2, \omega_{\text{FS}}), L_\tau)$ that fixes a neighbourhood of the line at infinity $\mathbb{C}P^1_\infty$ and hence lifts to a compactly supported symplectomorphism of $(\mathbb{C}^2, \omega_{\text{st}})$ that takes Δ to Δ_τ.

Upgrading the symplectomorphism to a compactly supported Hamiltonian diffeomorphism

To conclude the argument, it is now enough to observe that Gromov proved that $\text{Symp}_c(B^4, \omega_{\text{st}})$ is contractible. In particular, this implies that the symplectomorphism $\Psi \in \text{Symp}_c(\mathbb{C}^2, \omega_{\text{st}})$ that takes Δ to Δ_τ and is compactly supported in B^4, is actually a Hamiltonian diffeomorphism $\Psi^H \in \text{Ham}_c(\mathbb{C}^2, \omega_{\text{st}})$, with compact support in B^4. It follows that Δ and Δ_τ are isotopic via an ambient compactly supported Hamiltonian isotopy. □

Remark 2.5.14. This compactification is not the only choice that makes this argument work. One could equally well have chosen a monotone $S^2 \times S^2$ as compactification. The Lagrangian planes then transform to embedded Lagrangian spheres in $S^2 \times S^2$ and the theorem follows along similar lines.

2.5.3 Lagrangian pinwheels in $B_{p,q}$

The goal of this subsection is to show that apart from (p, q)- and $(p, p - q)$-pinwheels, no other Lagrangian pinwheel embeds into $B_{p,q}$, i.e. we prove the following theorem.

Theorem 2.5.15. *If $L_{m,n} \subseteq B_{p,q}$ with $0 < n < m$ coprime is a Lagrangian pinwheel, then*

$$(m, n) \in \{(p, q), (p, p - q)\}.$$

Proof. By the neighbourhood theorem for Lagrangian pinwheels, proven in [Kho13b, Section 3] or Theorem 2.6.6, we find a symplectic embedding $\iota : B_{m,n}(\varepsilon) \hookrightarrow B_{p,q}$. In particular, we have that $\iota^* c_1(B_{p,q}) = c_1(B_{m,n})$. Moreover, the Chern classes are primitive by [ES18, Lemma 2.13] and therefore we have $m \mid p$.[43] We will now run the argument carried out in Section 2.3.2. For the sake of brevity we will skip some steps in order to arrive at the important part of the argument quickly. The arguments are completely analogous to those in Section 2.3.2 and are therefore mainly a matter of bookkeeping.

A few points are crucial to observe at this point: the compactification $X_{p,q}$ of $B_{p,q}$ is as before, i.e. the intersection matrix of the compactifying divisor $\mathcal{D}_{p,q}$, denoted by $M_{\mathcal{D}}$, has the same numerology as before, whereas the numerology of the intersection matrix M_C, associated with

[43] Recall that $H^2(B_{p,q}, \mathbb{Z}) \cong \mathbb{Z}_p$.

the Wahl chain $C_{m,n}$, is governed by the numerology derived from (m,n). The fibre class F also has non-trivial intersection number with $L_{m,n}$.[44]

As before, we stretch the neck along the contact boundary of $\iota(B_{m,n}(\varepsilon))$ and consider the limit of curves in the class F. We again consider the component C of the holomorphic limit building that intersects the divisor D_0 and rationally blow up $X_{p,q}$ along ι to obtain a symplectic manifold $\widetilde{X}$. We denote the curve that corresponds to the component of the limit building that intersects the compactifying divisor in D_0 under this procedure by $\widetilde{C}$. Expressing the canonical class K of $\widetilde{X}$ and $\widetilde{C}$ in the basis defined by $\mathcal{D}_{p,q}$ and $C_{m,n}$ we obtain, in analogy to (2.3.18):

$$K = -F - \sum D_i + \sum k_j C_j, \qquad \widetilde{C} = \sum a_i D_i + \sum c_j C_j.$$

Repeating verbatim the proof of Lemma 2.3.17 we see that $\widetilde{C}$ is embedded and a sphere of non-positive self-intersection. We can assume that $\widetilde{C}$ is a (0)-sphere. We want to investigate the adjunction formula for this curve in order to determine (m,n) in terms of (p,q). The adjunction formula for $\widetilde{C}$ reads, in analogy to (2.3.21):

$$1 = \frac{p^2+1}{2p^2} - \frac{1}{2}\left(\boldsymbol{k}^T\chi + \chi^T M_C^{-1}\chi\right).$$

As in the proof of Lemma 2.3.18 we find that $\widetilde{C}$ intersects the Wahl chain $C_{m,n}$ precisely once, which means $\chi_j = 1$ for exactly one j in the notation from before. Then the adjunction formula takes the form

$$1 = \frac{p^2+1}{2p^2} - \frac{1}{2}\left(-1 + \frac{e_j + f_j}{m^2} - \frac{e_j f_j}{m^2}\right)$$

which is, by writing $p = lm$ for a positive integer l, equivalent to

$$0 = \frac{p^2+1}{2p^2} + \frac{1}{2m^2}(m^2 + e_j f_j - (e_j + f_j)) - 1 = \frac{1}{2p^2}(1 + l^2 e_j f_j - l^2(e_j + f_j))$$

meaning

$$0 = \left(\frac{1}{l^2} - 1 + (e_j - 1)(f_j - 1)\right).$$

This implies that $l^2 = 1$, i.e. $p = m$ because the term $(e_j - 1)(f_j - 1)$ is an integer, and therefore that $e_j = 1$ or $f_j = 1$, which means that $\widetilde{C}$ intersects $C_{m,n}$ at one of its ends. Because $\widetilde{C}$ is a (0)-curve we have

$$0 = a_0 + \chi^T M_C^{-1}\chi = \frac{pq-1}{p^2} - \frac{1}{p^2}e_j f_j,$$

which implies $pq - 1 = e_j f_j$. Since $\widetilde{C}$ intersects the Wahl chain at one of its ends we have $e_j f_j \in \{mn-1, [mn-1]^{-1}\} = \{pn-1, [pn-1]^{-1}\}$ and because $[pn-1]^{-1} = p^2-(pn+1) = p(p-n)-1$ this means $n = q$ or $n = p - q$. □

Corollary 2.5.16. *If there is a symplectic embedding of a star-shaped subset of $B_{m,n}$ into $B_{p,q}$, then $B_{m,n}$ is symplectomorphic to $B_{p,q}$.*

Proof. The star-shaped subset contains an (m,n)-pinwheel; therefore, by Theorem 2.5.15, $(m,n) \in$

[44] Assume that $[L_{m,n}] \cdot F = 0$. Then, as explained in [AH25, Appendix C], the class F can be represented by an embedded submanifold disjoint from $L_{m,n}$. However, this implies that $\iota^* c_1(B_{p,q}) = 0$, which is a contradiction.

$\{(p,q),(p,p-q)\}$, and [Ada+25, Remark 2.8] shows that $B_{p,q}$ and $B_{p,p-q}$ are symplectomorphic since their base diagrams are related by an integral affine transformation. □

Remark 2.5.17. Note that a weaker version of this result could be obtained without evoking pseudoholomorphic curve techniques, just by considering the soft obstructions of Lagrangian pinwheels, namely that $c_1(L_{m,n}) = n \mod m$ and $L^2_{m,n} = -1 \mod m$ (see [AH25, Remark 2.8]). However, these considerations do not exclude the case of a symplectic embedding $B_{4,3} \hookrightarrow B_{8,1}$ for example.

2.6 Appendix

Pinwheels and their neighbourhoods

In this appendix, we elaborate on some aspects regarding pinwheels and their neighbourhoods. We introduce a weaker (perhaps more natural) definition of a Lagrangian pinwheel, show in Lemma 2.6.5 that it is equivalent to the one given in the main text, and demonstrate how these admit a “Weinstein” neighbourhood symplectomorphic to $B_{p,q}$ after completion. The neighbourhood theorem is originally due to Khodorovskiy [Kho13b, Section 3], and we improve upon it by showing that the symplectic embedding into $B_{p,q}$ takes the pinwheel to the visible one with respect to the standard ATF on $B_{p,q}$.

Definition 2.6.1. Define a *topological p-pinwheel* P_p to be the topological space $D^2/\sim$, where the equivalence relation is given by $z \sim z'$ if and only if $z, z' \in \partial D^2$ and $z = e^{2\pi i k/p} z'$ for some $k \in \mathbb{N}$.

Recall that we are always assuming, as in the rest of this paper, that $0 < q < p$ are two coprime integers. We now define a Lagrangian (p,q)-pinwheel differently from Definition 2.2.8.

Definition 2.6.2. A *Lagrangian* (p,q)*-pinwheel* in a symplectic 4-manifold (X,ω) is an immersion $f : D^2 \subseteq \mathbb{C} \looparrowright X$ such that

i) The restriction of f onto the interior of the disc $f|_{D^2\setminus\partial D^2}$ is a Lagrangian embedding and there exists an embedding $\gamma : \partial D^2 \hookrightarrow X$ such that $C := f(\partial D^2) = \gamma(\partial D^2)$, which we will call the *core circle*;

ii) The map $f_\partial := f|_{\partial D^2}$ factors through a continuous embedding of P_p;

iii) Furthermore, let $\Lambda \to C$ be the S^1-bundle whose fibre over $x \in C$ consists of Lagrangian 2-planes in T_xX that contain T_xC. Choose a trivialization $\Phi : \Lambda \to C \times S^1$. Pulling this trivialization back by f_∂ gives a trivialization of the pullback bundle $(f|_\partial)^*\Lambda \to \partial D^2$. The immersion defines a section $s := f_*(TD^2)$ of $(f|_\partial)^*\Lambda$, which in the trivialization defines a map $\vartheta : S^1 \to S^1$. We then require that this section satisfies

$$\vartheta(t+\delta) - \vartheta(t) = q\delta, \tag{2.6.32}$$

where $\delta := \frac{2\pi}{p}$.

Remark 2.6.3. Before we move on to investigate this definition in more detail, let us elaborate on the definition itself and its connections to earlier definitions of Lagrangian pinwheels. Both Khodorovskiy [Kho13b, Definition 3.1] and Evans–Smith [ES18, Definition 2.3] give a definition of a Lagrangian (p, q)-pinwheel that exchanges the last part of our definition with the weaker condition that the relative winding number, which is only defined modulo p, of the flanges around the core is equal to q.

i) For the purposes of Khodorovskiy and Evans–Smith the weaker definition suffices because in an arbitrarily small C^0-neighbourhood of a pinwheel satisfying this weaker definition there is a (p, q)-pinwheel that admits a standard neighbourhood. See [Kho13b, Definition 3.3] and [ES18, Definition 2.10].

ii) For us, however, the weaker definition is not sufficient: to show Hamiltonian uniqueness there is an obstruction coming from the configuration space of p-lines in the plane. Indeed, observe that, by definition, at a point on the core circle the immersion f defines p lines in the symplectic normal bundle. These lines cannot, in general, be taken to another configuration of p lines in the symplectic normal bundle by a linear symplectic (or even just smooth) map. Hence, it is absolutely crucial to fix the geometric tangent data along the core circle, which is exactly the last part of our Definition 2.6.2 (see also the proof of Lemma 2.6.5).

Theorem A shows uniqueness up to Lagrangian isotopy in $B_{p,q}$ of these weakly defined pinwheels. This follows from Theorem A together with the fact that weakly defined Lagrangian (p, q)-pinwheels are isotopic through weak Lagrangian (p, q)-pinwheels to a Lagrangian (p, q)-pinwheel, as just defined, and this isotopy can be localised around the core circle; cf. [Kho13b, Lemma 3.3].

Let $f : D^2 \looparrowright X$ define a Lagrangian (p, q)-pinwheel in (X, ω) and $\gamma : \partial D^2 \hookrightarrow X$ the parametrization of the core circle given in Definition 2.6.2. Since γ is isotropic we find coordinates around this core circle of the form $T^*S^1 \times \mathbb{C}$, with coordinates (τ, x) on T^*S^1, where τ is the base coordinate, and $z = u + iv$ on $\mathbb{C}$, such that the core circle is given by $S^1 \times \{0\}$. Moreover, the symplectic form is identified with $d\tau \wedge dx + du \wedge dv$. This means that, if we write $\Sigma : S^1 \times [0, \varepsilon) \to T^*S^1 \times \mathbb{C}$ for f in these coordinates, we can assume that

$$\Sigma(t, s) = (pt, x(t, s), z(t, s)). \tag{2.6.33}$$

Note that there is a natural choice of Lagrangian immersion of this collar that satisfies the conditions on the collar part in Definition 2.6.2 iii):

$$\Sigma^g(t, s) = \left(pt, \frac{q}{2p}s^2, se^{iqt}\right) \tag{2.6.34}$$

This immersion is exactly the immersion that naturally appears in toric geometry. See for example [Eva23, Section 5.3]. This parametrization warrants a definition.

Definition 2.6.4 ([Kho13b, Lemma 3.3])**.** A *good Lagrangian (p, q)-pinwheel* is a Lagrangian immersion $f : D^2 \looparrowright (X, \omega)$ such that we find coordinates around the core circle in which f is expressed as Σ^g, defined in (2.6.34).

In the main text, we have defined Lagrangian (p, q)-pinwheels as good Lagrangian (p, q)-pinwheels; this is allowed by the following lemma.

Lemma 2.6.5. *Definition 2.6.2 and the good Definition 2.6.4 of Lagrangian (p, q)-pinwheels agree.*

Proof. Our first claim is that coordinates around the core of any Lagrangian (p, q)-pinwheel can be chosen so that it is a good Lagrangian (p, q)-pinwheel to first order.

Consider the local expression $\Sigma(t, s) = (pt, x(t, s), z(t, s))$ obtained in (2.6.33). Along the core circle we have $\Sigma(t, 0) = (pt, 0, 0)$ and hence $\partial_t\Sigma(t, 0) = (p, 0, 0)$. Moreover,

$$\partial_s\Sigma(t, 0) = (0, \partial_s x(t, 0), V(t)), \quad \text{where } V(t) := \partial_s z(t, 0) \in \mathbb{C}.$$

We first show that $\partial_s x(t, 0) = 0$. Indeed, since Σ is Lagrangian,

$$0 = \omega(\partial_t\Sigma, \partial_s\Sigma) = p\partial_s x + du \wedge dv(\partial_t z, \partial_s z).$$

Using

$$du \wedge dv(X, Y) = \mathrm{Im}(\overline{dz(X)}dz(Y)),$$

this gives

$$p\partial_s x = -\mathrm{Im}(\overline{\partial_t z}\,\partial_s z).$$

Restricting to $s = 0$, the right hand side vanishes because $z(t, 0) = 0$, and hence $\partial_t z(t, 0) = 0$. Therefore,

$$\partial_s x(t, 0) = 0.$$

Thus $\partial_s\Sigma(t, 0) = (0, 0, V(t))$, where $V(t) \neq 0$. Using the polar coordinates we write $V(t) = \lambda(t)e^{i\psi(t)}$ for $\lambda(t) = |V(t)| > 0$. If we apply the diffeomorphism of the domain that is given by scaling the radial directions by $1/\lambda(t)$ and work in this new coordinate system we obtain $\partial_s\Sigma(t, 0) = (0, 0, e^{i\psi(t)})$.

By the definition of a Lagrangian (p, q)-pinwheel, in particular (2.6.32), we must have $\psi(t + \delta) = \psi(t) + q\delta$, where $\delta := 2\pi/p$. Define $\eta(t) := \psi(t) - qt$ and observe:

$$\eta(t + \delta) = \psi(t + \delta) - q(t + \delta) = \psi(t) + q\delta - q(t + \delta) = \eta(t).$$

This precisely means that η descends to a smooth function under the $p : 1$-projection, i.e. we can define $\eta_p(pt) := \eta(t)$. Now consider the symplectomorphism defined by

$$\Phi_{\eta_p}(\tau, x, z) = \left(\tau, x - \frac{1}{2}\eta_p'(\tau)|z|^2, e^{-i\eta_p(\tau)}z\right). \tag{2.6.35}$$

That Φ_{η_p} indeed defines a symplectic map follows from

$$d\tau \wedge d\left(x - \frac{1}{2}\eta_p'(\tau)|z|^2\right) = d\tau \wedge dx - \frac{1}{2}\eta_p'(\tau)d\tau \wedge d|z|^2$$

and, defining $w := e^{-i\eta_p(\tau)}z$,

$$\frac{i}{2}dw \wedge d\overline{w} = \frac{i}{2}dz \wedge d\bar{z} + \frac{1}{2}\eta_p'(\tau)d\tau \wedge d|z|^2.$$

The symplectomorphism Φ_{η_p} fixes the core pointwise and its differential rotates the complex normal direction by $e^{-i\eta_p(\tau)}$. In particular, this means that if we apply the coordinate change Φ_{η_p}

on the ambient coordinates and express Σ in these new coordinates we obtain

$$\partial_t\Sigma(t,0) = (p,0,0) \quad \text{and} \quad \partial_s\Sigma(t,0) = (0,0,e^{iqt}).$$

So the tangent planes of Σ^g and Σ along the core circle coincide. To sum up: we are now working in a coordinate system in which Σ agrees to first order with the model Σ^g along the core circle.

Our second claim is that now Σ *can be Hamiltonian isotoped to* Σ^g. Consider the p-fold cover of S^1 given by $\pi : S^1 \to S^1$, where $\pi(t) = pt$ and the p-fold covering this induces on $T^*S^1 \times \mathbb{C}$. Lifting the Lagrangian immersions Σ and Σ^g gives Lagrangian embeddings

$$\widetilde{\Sigma}(t,s) = (t, x(t,s), z(t,s)) \quad \text{and} \quad \widetilde{\Sigma}^g(t,s) = \left(t, \frac{q}{2p}s^2, se^{iqt}\right).$$

Since $\widetilde{\Sigma}$ and $\widetilde{\Sigma}^g$ agree along the core to first order, the Weinstein neighbourhood theorem implies that, after possibly shrinking the collar, $\widetilde{\Sigma}$ can be written as the graph of a closed one-form $\alpha \in \Omega^1(\widetilde{\Sigma}^g)$ on $\widetilde{\Sigma}^g$ that vanishes on the boundary $S^1 \times \{0\}$. Because $H^1(S^1 \times [0,\varepsilon), S^1 \times \{0\}) = 0$ it follows that α is exact, meaning that there exists a smooth function H such that $\alpha = dH$. Since $\widetilde{\Sigma}$ and $\widetilde{\Sigma}^g$ agree along the core circle to first order, H vanishes along the core circle to first order. Hence, the Hamiltonian vector field of H vanishes along the core circle.

Viewing H as a function on the Weinstein neighbourhood of $\widetilde{\Sigma}^g$, the Hamiltonian diffeomorphism generated by H maps $\widetilde{\Sigma}^g$ to $\widetilde{\Sigma}$. The crucial observation is now that the lifted core circle and the first order data along the core circle are invariant under the deck transformation. Consequently, the one-form α and its primitive H can be chosen to be deck-invariant, which ensures that the Hamiltonian flow generated by H is equivariant with respect to the deck transformation. Therefore, the Hamiltonian flow descends to the quotient, showing that in a neighbourhood of the core circle the immersion Σ is Hamiltonian isotopic to Σ^g. □

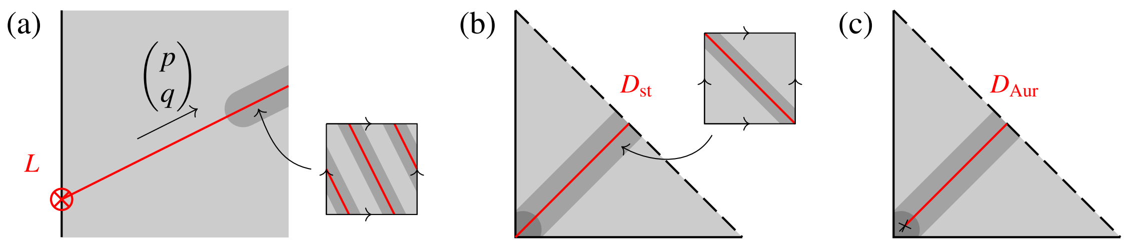

Figure 2.21: (a) A neighbourhood U_Σ of the core of a Lagrangian (p,q)-pinwheel L admits a Lagrangian torus fibration as shown. Shaded in darker gray is a Weinstein neighbourhood of the annulus Γ. (b) The standard toric fibration of the ball and the visible Lagrangian disc $D_{\rm st}$. A Weinstein neighbourhood of the annulus is shown in darker gray. Note that a neighbourhood of the centre of the disc contains, shown in even darker gray, a saturated neighbourhood of the origin, i.e. the fibre tori close to the origin are completely contained in the neighbourhood, while away from the origin the Weinstein neighbourhood is fibred by Lagrangian annuli. (c) The neighbourhood after a nodal trade, which yields a Weinstein neighbourhood of the disc $D_{\rm Aur}$.

Theorem 2.6.6. *Suppose that $L_{p,q} \subseteq (X, \omega)$ is a Lagrangian pinwheel in a symplectic manifold. Then for some $\varepsilon > 0$ there exists a symplectic embedding $\psi : B_{p,q}(\varepsilon) \hookrightarrow (X, \omega)$ such that $\psi(L^{vis}_{p,q}) = L_{p,q}$, i.e. there exists a neighbourhood U_L of $L_{p,q}$ that admits an ATF such that $L_{p,q}$ is a visible pinwheel.*

Proof. The idea of the proof is to produce ATFs in the neighbourhoods of the core circle and the capping disc and glue them appropriately. By Lemma 2.6.5 there is a neighbourhood U_Σ of the core circle of $L_{p,q}$ that is symplectomorphic to $S^1 \times D^3$ such that $L_{p,q}$ is visible in this neighbourhood as shown in Figure 2.21 (a). Consider the annulus Γ defined in the figure (to be the gluing region); the ATF defines an embedding of a small cotangent disc bundle of an annulus such that the zero section is identified with Γ. This embedding extends to an embed-

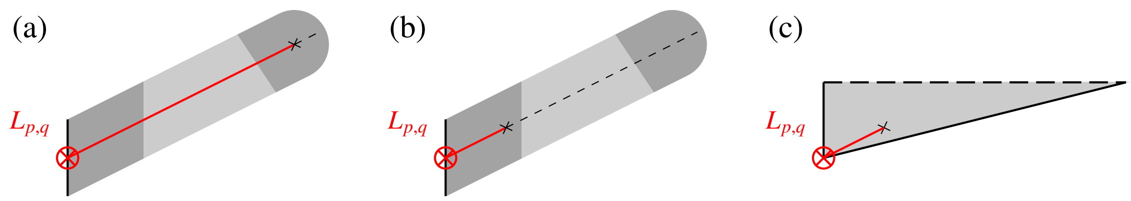

Figure 2.22: (a) The base diagram of the Lagrangian fibration constructed on a neighbourhood U of a Lagrangian pinwheel $L_{p,q} \subseteq (X, \omega)$. (b) The Lagrangian fibration on the same subset U after a nodal slide and (c) a restriction of the neighbourhood U given by taking the preimage of $\mathfrak{A}_{p,q}(\varepsilon)$, which is contained in (b). To see that (b) contains a base of the form $\mathfrak{A}_{p,q}(\varepsilon)$ change the branch cut in (b).

ding of a small cotangent disc bundle of the disc such that the zero section is mapped to the embedded capping disc Δ. We denote the image of this embedding by U_Δ. The union of the neighbourhoods $U := U_\Sigma \cup U_\Delta$ defines a neighbourhood of $L_{p,q}$. Moreover, this construction carries a Lagrangian fibration, which we now explain. On U_Δ we have a fibration as shown in Figure 2.21 (b) and after a nodal trade we see that U_Δ admits a fibration as shown in Figure 2.21 (c).[45] On U_Σ we have a Lagrangian fibration too, as explained in the beginning of the proof. Now, after shrinking U, we obtain a fibration as shown in Figure 2.22 (a) by gluing the two fibrations on the overlapping part.

We are not done yet: the fibres over the lightest-gray shaded regions in Figure 2.22 (a) and (b) are annuli and not tori. The goal is now to perform a nodal slide in order to pull the node over the region over which the Lagrangian fibres are only annuli and not full tori: this corresponds to the change of base diagram from Figure 2.22 (a) to (b). As this is a non-standard situation we have to use the interpretation of a nodal slide that Groman and Varolgunes give in [GV23, Sections 7.3 and 7.4]. The idea is to consider the preimage of the eigenline in Figure 2.22. This preimage carries a Hamiltonian S^1-action with stabilizer given by the focus-focus singularity. If we now quotient the preimage by this Hamiltonian S^1-action, we obtain a situation as shown in Figure 2.23 (a). Then changing the foliation as illustrated in the change from Figure 2.23 (a) to (b) changes the fibration on U from Figure 2.22 (a) to (b). Then after shrinking U again we obtain a neighbourhood U_L of $L_{p,q}$ as shown in Figure 2.22 (c). □

[45] Note that this is a non-trivial operation. A priori it is not clear that D_{Aur} can be assumed to coincide with D_{st}. However, this is guaranteed to be the case after a Hamiltonian isotopy by Theorem 2.5.7.

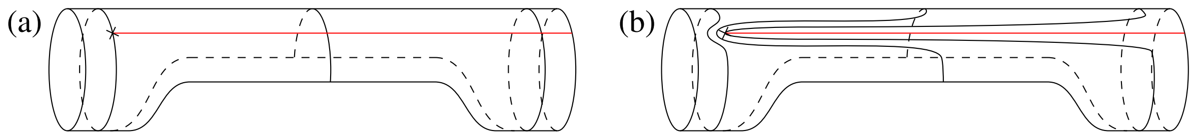

Figure 2.23: (a) The reduced space. (b) The reduced space with the changed foliation that induces the nodal slide.

Remark 2.6.7. Note that this is a refinement of the Weinstein-type theorem that Khodorovskiy proves in [Kho13b, Section 3]. We emphasize that Theorem 2.6.6 allows us to pass freely from embeddings of Lagrangian pinwheels to embeddings of small rational homology balls.

3 Markov Staircases

3.1 Introduction

3.1.1 Old questions for new domains

Let (X,ω) be a symplectic 4-manifold. By Darboux's theorem, small enough patches of X are symplectomorphic to subsets of the standard symplectic vector space $(\mathbb{C}^2,\omega_0)$. In particular one can always symplectically embed small enough symplectic ellipsoids

$$E(\alpha,\beta) = \left\{ (z_1,z_2) \in \mathbb{C}^2 \; : \; \frac{\pi|z_1|^2}{\alpha} + \frac{\pi|z_2|^2}{\beta} \leq 1 \right\} \tag{3.1.1}$$

into (X,ω). This raises several natural questions:

(Q1) **Isotopy problem.** Given two such embeddings of the same size, are they ambiently isotopic?

(Q2) **Embedding problem.** For which values of the parameters $\alpha,\beta > 0$ do symplectic embeddings of $E(\alpha,\beta)$ exist?

(Q3) **Ellipsoid packing problem.** Given several embeddings of differently-sized ellipsoids, when can they be made disjoint?

In the case of symplectic balls $B(\alpha) = E(\alpha,\alpha)$, these questions were raised in Gromov's foundational work [Gro85], and, ever since, the idea of using such embeddings to measure the symplectic sizes and quantitative properties of (X,ω) has been a central theme in symplectic topology. In the work of McDuff and Schlenk [MS12], it was shown that the set of pairs (α,β) for which there is a symplectic embedding of $E(\alpha,\beta)$ into the ball takes the shape of an infinite staircase, whose step-sizes are related to the sequence of odd-indexed Fibonacci numbers. Since then, the ellipsoid embedding question into various target spaces X has attracted a lot of attention, see for instance the survey [Sch18]. In this paper, instead of looking at new target spaces, we consider new domains.

Our goal is to study symplectic embedding questions for a family of domains we call *pin-ellipsoids* $E_{p,q}(\alpha,\beta)$, where p and q are positive coprime integers with $q \leq p$. See Definition 3.1.1 for their definition. These domains arise naturally (1) in Milnor fibres of smoothings of certain surface singularities called *cyclic quotient T-singularities*, and (2) as building blocks of *almost toric fibrations*. The name pin-ellipsoid is derived from the fact that every pin-ellipsoid contains a so-called *Lagrangian pin-wheel*. The case $p = q = 1$ corresponds to the case of the *standard ellipsoid* in Equation (3.1.1), whereas $E_{2,1}(\alpha,\beta)$ can be identified with a neighbourhood of the zero-section of $T^*\mathbb{RP}^2$. We address the above questions (Q1)–(Q3) for all pin-ellipsoid embeddings into $X = \mathbb{CP}^2$. Theorem 3.1.8 affirmatively answers the isotopy question (Q1), Theorem 3.1.13 solves the embedding problem (Q2) in the "visible range" (see Remark 3.1.2 (d)), and Theorem 3.1.18 solves the multiple pin-ellipsoid problem (Q3) for the special case of *pin-balls* where $\alpha = \beta$, analogous to Gromov's Two Ball Theorem [Gro85,

§ 0.3.B]. We begin by defining pin-ellipsoids and also quickly recall the Markov numbers, which will play a key role in the statements of all of our theorems.

3.1.2 Pin-ellipsoids, pin-balls and pin-wheels

Our definition of the pin-ellipsoid $E_{p,q}(\alpha,\beta)$ uses the theory of *almost toric fibrations* (ATFs), where the symplectic geometry of a 4-manifold is encoded in a 2-dimensional *almost toric base diagram*. These diagrams were introduced by Symington [Sym98]; see the book [Eva23] for more details about how to decode these diagrams.

Definition 3.1.1. Given positive integers $1 \le q \le p$ with $\gcd(p,q) = 1$ and positive real numbers α,β, let the *pin-ellipsoid* $E_{p,q}(\alpha,\beta)$ be the compact symplectic manifold corresponding to the almost toric base diagram $\mathfrak{A}_{p,q}(\alpha,\beta)$ shown in Figure 3.1. This base diagram has a slanted edge pointing in the $(p^2, pq-1)$-direction, having integral affine length α, and a vertical edge having integral affine length β; it has a single focus-focus fibre joined to the origin by a branch cut parallel to the (p,q)-direction. The top side, drawn — — is included in the diagram but it is not part of the toric boundary; the preimage of this side under the almost toric fibration is the boundary of $E_{p,q}(\alpha,\beta)$, which is a 3-manifold diffeomorphic to the lens space $L(p^2, pq-1)$.

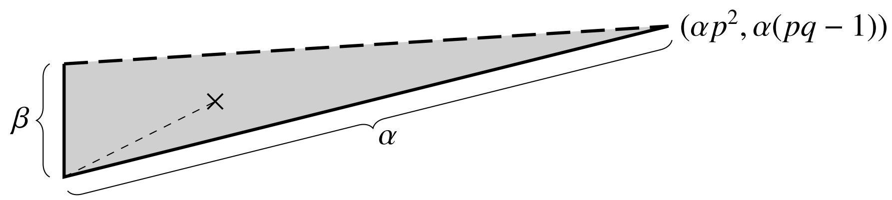

Figure 3.1: The almost toric base diagram $\mathfrak{A}_{p,q}(\alpha,\beta)$ for $E_{p,q}(\alpha,\beta)$. The branch cut is parallel to the (p,q)-direction. The toric boundary comprises the two solid edges (not the top side), whose affine lengths are α and β; these are two segments of an unbroken straight line with respect to the integral affine structure on the base diagram.

Remark 3.1.2. (a) We distinguish between *edges* of an almost toric base diagram (which we draw as thick, unbroken lines ——) which are part of the toric boundary (that is, the fibre over a point of an edge is a circle or a point), and *sides* (drawn as thick, long-dashed lines — — -) which are not part of the toric boundary (that is, the fibre over a point of a side is a 2-torus). Thin, short-dashed lines - - - - - on the interior of the polygon will represent branch cuts.

(b) Note that $E_{p,q}(\alpha,\beta)$ is symplectomorphic to $E_{p,p-q}(\beta,\alpha)$: the almost toric base diagrams are related by an integral affine transformation with determinant -1, which lifts to a symplectomorphism of the total spaces. It follows that pin-ellipsoid embedding problems in the case $(p,q) = (2,1)$ will be symmetric with respect to swapping α and β, just like for standard ellipsoids. For other (p,q) this will no longer hold true, as will be evident from Theorem 3.1.13. Compare also Remark 3.1.14.

(c) We call the pin-ellipsoid $E_{p,q}(\alpha,\alpha)$ the *pin-ball* $B_{p,q}(\alpha)$.

(d) Given an almost toric base diagram which contains a subset integral-affine equivalent to $\mathfrak{A}_{p,q}(\alpha,\beta)$, its preimage under the almost toric fibration is an embedded copy of $E_{p,q}(\alpha,\beta)$ in the associated almost toric manifold. We refer to these as *visible pin-ellipsoids*.

Definition 3.1.3. Inside $E_{p,q}(\alpha,\beta)$ there is an immersed Lagrangian disc $L_{p,q}$ which lives over the branch cut. This immersion fails to be an embedding along its boundary circle which meets the toric boundary; here it is p-to-1 and modelled on a standard model immersion, described in [Kho13b, Equations (3.4) and (3.5)] or [Eva23, Example 5.14]. This immersed disc is called a *Lagrangian* (p,q)*-pin-wheel*.

Remark 3.1.4. More generally, we call any immersed Lagrangian disc which is embedded away from its boundary and modelled on $L_{p,q}$ along its boundary a Lagrangian pin-wheel. Khodorovskiy [Kho13b, Lemma 3.4] shows that any Lagrangian pin-wheel admits a neighbourhood symplectomorphic to $E_{p,q}(\alpha,\beta)$ for any sufficiently small $\alpha,\beta > 0$. We will call a Lagrangian pin-wheel *visible* if it projects to an arc in the base of some almost toric fibration.

3.1.3 Markov numbers and their companions

Definition 3.1.5. Integer solutions to the Markov equation

$$p_1^2 + p_2^2 + p_3^2 = 3p_1p_2p_3$$

are called *Markov triples*. One can produce new Markov triples from a given one by three *mutations*, each of which substitutes a member of the triple by another Markov number. For example the mutation $(p_1, p_2, p_3) \to (p_1, p_2, 3p_1p_2 - p_3)$ replaces p_3 by $p_3' = 3p_1p_2 - p_3$. All Markov triples are obtained in this way from iterated mutations of $(1, 1, 1)$. Thus they are naturally arranged according to the graph depicted in Figure 3.2.

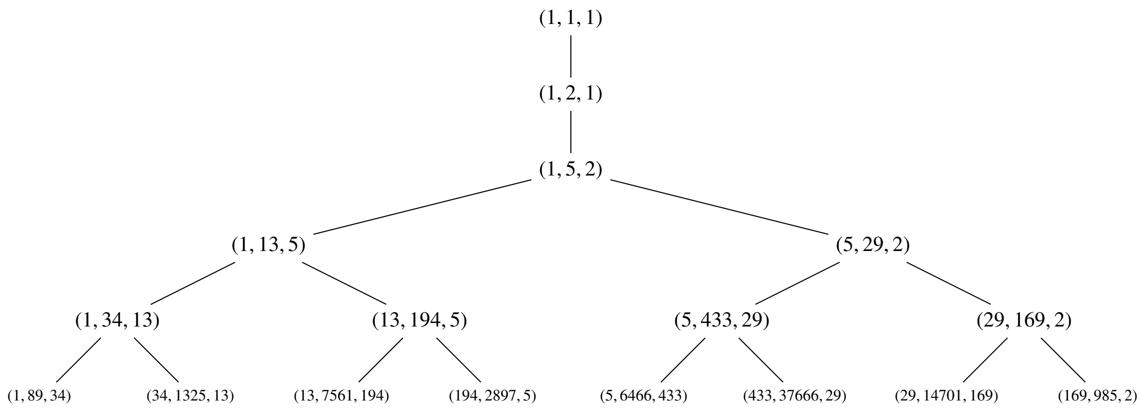

Figure 3.2: The beginning of the Markov graph, where we have omitted the repetitions at the first two Markov triples.

Definition 3.1.6 (Companion numbers)**.** Let p be a Markov number. A number q is called *companion number of* p if there exists a Markov triple (p, p_2, p_3) such that

$$q \equiv \pm 3p_2p_3^{-1} \mod p. \tag{3.1.2}$$

Remarks 3.1.7. *(a) For a fixed triple, Equation* (3.1.2) *defines two companion numbers with* $1 \leq q \leq p$ *which are related to one another by* $q \to p - q$*. If* $p \in \{1, 2\}$*, then the two companion numbers are both equal to* 1*, and for all other Markov numbers they differ and are strictly less than p. Mutations of Markov triples containing p preserve its companion numbers.*

(b) *Conversely, there is a unique mutation branch in the Markov tree in which p appears with companion number q, see for example [Aig13, Proposition 3.15]. It is conjectured that each Markov number has only one pair of companion numbers (see [Aig13, §2.3]), but since this conjecture is famously open, we need to specify both p and q in everything that follows.*

As was noticed by Vianna [Via14], the almost toric base diagrams of $\mathbb{CP}^2$ (up to integral affine equivalence and nodal slides) are in bijection with unordered Markov triples. This is closely related to the fact – due to Hacking–Prokhorov [HP10] – that the singular fibre of any $\mathbb{Q}$-Gorenstein degeneration of $\mathbb{CP}^2$ is (a partial smoothing of) a weighted projective plane $\mathbb{P}(p_1^2, p_2^2, p_3^2)$ where (p_1, p_2, p_3) is a Markov triple. Given a Markov triple, denote by $\mathfrak{D}(p_1, p_2, p_3)$ the corresponding almost toric base diagram. We call these almost toric base diagrams *Vianna triangles*: see Section 3.2.3 for a full description. For now, the important feature is that Vianna triangles allow us to see *visible* pin-ellipsoids

$$E_{p_i,q_i}(\alpha_i,\beta_i) \subset \mathbb{CP}^2 \quad \text{for any} \quad \alpha_i < \frac{p_{i+2}}{p_i p_{i+1}}, \quad \beta_i < \frac{p_{i+1}}{p_i p_{i+2}}, \tag{3.1.3}$$

where q_i is the companion number $3p_{i+1}p_{i+2}^{-1} \mod p_i$. Note that reordering the Markov triple by swapping p_{i+1} and p_{i+2} will swap the order of the bounds on α_i and β_i and also swap the companion numbers; this is consistent with Remark 3.1.2(b).

3.1.4 The isotopy problem

Evans and Smith [ES18] showed that $E_{p,q}(\alpha,\beta)$ admits a symplectic embedding into $\mathbb{CP}^2$ for some $\alpha,\beta > 0$ if and only if p is a Markov number and q is one of its companion numbers. Provided that $\alpha < \frac{p_3}{pp_2}$ and $\beta < \frac{p_2}{pp_3}$ for some Markov triple (p, p_2, p_3) with $q = 3p_2p_3^{-1} \mod p$, there is a visible embedding $E_{p,q}(\alpha,\beta) \hookrightarrow \mathbb{CP}^2$ coming from the Vianna triangle $\mathfrak{D}(p, p_2, p_3)$. But there exist symplectic embeddings also outside the visible range; for instance, symplectic folding can be used to show that for any $\alpha > 0$ there exists $\beta > 0$ such that $E_{p,q}(\alpha,\beta)$ symplectically embeds into $\mathbb{CP}^2$.

Theorem 3.1.8 (Isotopy Theorem)**.** *Any two symplectic embeddings of a pin-ellipsoid $E_{p,q}(\alpha,\beta)$ into $\mathbb{CP}^2$ are isotopic through symplectic embeddings.*

In other words, the space of symplectic embeddings of a pin-ellipsoid into $\mathbb{CP}^2$, if non-empty, is path-connected.

Remarks 3.1.9. (a) *The problem whether the space of symplectic embeddings of $E_{p,q}(\alpha,\beta)$ into $\mathbb{CP}^2$ is non-empty is addressed in § 3.1.5.*

(b) *Since $H^1(E_{p,q}(\alpha,\beta); \mathbb{R})$ vanishes, it follows that any two embeddings which are symplectically isotopic are related by a Hamiltonian diffeomorphism of $\mathbb{CP}^2$.*

(c) *In fact, we will reprove the results of Evans and Smith [ES18] along the way, avoiding the use of orbifold holomorphic curves. See Section 3.5.4.*

Corollary 3.1.10. *Lagrangian (p, q)-pin-wheels in $\mathbb{CP}^2$ are unique up to Hamiltonian isotopy.*

Proof. Indeed, Khodorovskiy [Kho13b, Lemma 3.4] shows that a Lagrangian pin-wheel has a standard neighbourhood symplectomorphic to an embedded pin-ellipsoid. Therefore, two pin-wheels yield two embedded pin-ellipsoids. If we choose them sufficiently small and of equal sizes α, β, Theorem 3.1.8 shows that these pin-ellipsoids are Hamiltonian isotopic, and the restriction of this isotopy to the pin-wheel gives the desired isotopy of pin-wheels. □

Since the $(2,1)$-pin-wheel is $\mathbb{RP}^2$, this generalises the known result about the uniqueness of Lagrangian $\mathbb{RP}^2$s in $\mathbb{CP}^2$ due to Hind [Hin12] and Li–Wu [LW12, Section 6.4.1], see also Borman–Li–Wu [BLW14, Theorem 1.3] and Adaloglou [Ada25] for an approach closer to ours.

Remark 3.1.11. Every Lagrangian pin-wheel $L \subset \mathbb{CP}^2$ is a non-trivial Lagrangian barrier in the sense of Biran [Bir01]: the Gromov width of the complement $\mathbb{CP}^2 \setminus L$ is strictly less than the Gromov with of $\mathbb{CP}^2$. Indeed, by Corollary 3.1.10 any pin-wheel in $\mathbb{CP}^2$ is Hamiltonian isotopic to one of the *visible* Lagrangian pin-wheels studied by Brendel and Schlenk [BS24], who proved that these are barriers and moreover computed the Gromov width of the complement.

3.1.5 The embedding problem

Notation 3.1.12. *Answering the embedding problem (Q2) for pin-ellipsoid embeddings into $\mathbb{CP}^2$ means determining the set*

$$\mathcal{A}_{p,q} := \left\{(\alpha,\beta) \in \mathbb{R}^2_{>0} \;:\; E_{p,q}(\alpha,\beta) \overset{s}{\hookrightarrow} \mathbb{CP}^2\right\}, \tag{3.1.4}$$

where p is a Markov number and q one of its companion numbers. By the definition of symplectic embedding, the set $\mathcal{A}_{p,q}$ is open in $\mathbb{R}^2_{>0}$. Let (p, m_0, m_1) be a Markov triple with $q \equiv 3m_0m_1^{-1} \mod p$ and consider the recursion

$$m_{i+2} = 3pm_{i+1} - m_i, \tag{3.1.5}$$

defining a sequence $\{m_i\}_{i\in\mathbb{Z}}$. The set $\{(p, m_i, m_{i+1}) \mid i \in \mathbb{Z}\}$ consists of all Markov triples which can be obtained from (p, m_0, m_1) by mutations preserving p: in the Markov tree in Figure 3.2 this set forms a $\bigwedge$-shaped tree. Let

$$\square_i(p,q) := \left(0, \frac{m_{i+1}}{pm_i}\right) \times \left(0, \frac{m_i}{pm_{i+1}}\right), \quad \mathrm{Stair}(p,q) := \bigcup_{i\in\mathbb{Z}} \square_i(p,q), \tag{3.1.6}$$

and

$$\sigma_p := \frac{1}{2}\left(3 + \sqrt{9 - \frac{4}{p^2}}\right).$$

Theorem 3.1.13 (Staircase Theorem)**.** *Let p be a Markov number and q one of its companion numbers. Then*

$$\mathcal{A}_{p,q} \cap (0,\sigma_p)^2 = \mathrm{Stair}(p,q). \tag{3.1.7}$$

Remarks 3.1.14. *(a) Since $\mathrm{vol}(E_{p,q}(\alpha,\beta)) = \frac{1}{2}p^2\alpha\beta$, the volume constraint for the symplectic embedding problem yields that $\mathcal{A}_{p,q}$ lies below the curve $\{p^2\alpha\beta = 1\}$.*

(b) All the pin-ellipsoids in this staircase can be realised as visible pin-ellipsoids coming from Vianna triangles (compare with Equation (3.1.3)) and by Theorem 3.1.8 are isotopic to a visible pin-ellipsoid. Conversely, all visible embeddings of $E_{p,q}(\alpha,\beta)$ satisfy $\max\{\alpha,\beta\} < \sigma_p$. It is striking that visible constructions are often sharp for such quantitative problems.

(c) The sets in Equation (3.1.7) are infinite staircases whose outer corners lie on the volume constraint and accumulate at the boundary of the square $(0, \sigma_p)^2 \subset \mathbb{R}^2$, see Figures 3.4 and 3.5. The first such staircase, which we recover as the case $(p, q) = (1, 1)$, was found by McDuff and Schlenk [MS12]. Since then, similar staircases have been found for many different target spaces – we refer to [Sch18] and the references therein for more on this and related problems. The usual convention is to use the normalisation freedom to reduce to the case $E(1, \beta) \hookrightarrow \mathbb{CP}^2(A)$, where $A > 0$ is the symplectic area of the line. In our case, we prefer the normalisation convention $E(\alpha, \beta) \hookrightarrow \mathbb{CP}^2(1)$. This is in part justified by the fact that the staircases we find in Theorem 3.1.13 are not symmetric with respect to swapping α and β, making it more natural to keep α and β on equal footing. In fact, the set obtained by swapping coordinates in Stair(p, q) *describes $\mathcal{A}_{p,p-q}$, the staircase associated to the second companion number $p - q$, by Remark 3.1.2(b).*

(d) We do not know the set $\mathcal{A}_{p,q}$ for $\alpha \geq \sigma_p$ or $\beta \geq \sigma_p$, except for $p = q = 1$ where it can be derived from [MS12, Sections 4 and 5]. For $p \geq 2$ we even do not know whether the intersection of the closure $\overline{\mathcal{A}_{p,q}}$ with the vertical segment at $\alpha = \sigma_p$ is a point or an interval of positive length. For $p = 1$ this intersection is a point, since the upper boundary of $\mathcal{A}_{1,1}$ on $\{\tau^2 \leq \alpha \leq \frac{21}{8}\}$ is given by $\beta = 3 - \alpha$.

For many (though not all, [DR23, Remark 5.1]) symplectic 4-manifolds (X, ω), there are full symplectic fillings $E(\alpha, \beta) \hookrightarrow (X, \omega)$ by standard ellipsoids whenever $\frac{\alpha}{\beta}$ is sufficiently large or sufficiently small. For instance, the boundary of $\mathcal{A}_{1,1}$ over $\{\frac{17}{6} \leq \alpha < \infty\}$ is given by $\beta = \frac{1}{\alpha}$. Is this so for all Markov numbers p? More precisely:

Question. *Let (p, q) be as above. Are there full symplectic fillings $E_{p,q}(\alpha, \beta) \hookrightarrow \mathbb{CP}^2$ for all sufficiently small and all sufficiently large $\frac{\alpha}{\beta}$?*

Example 3.1.15. Let us illustrate Theorem 3.1.13 by the example $p = 2, q = 1$, see Figures 3.3 and 3.4. We start with the Vianna triangle $\mathfrak{D}(1, 1, 1)$ appearing as the image of $\mathbb{CP}^2$ under the standard toric moment map. After one mutation, we obtain the triangle $\mathfrak{D}(2, 1, 1)$, containing a vertex which is integral affine equivalent to the cone spanned by the vectors $(0, -1)$ and $(4, -1)$ adjacent to edges which both have integral affine length $\frac{1}{2}$. Thus we deduce the existence of a volume filling embedding $E_{2,1}(\frac{1}{2}, \frac{1}{2}) \hookrightarrow \mathbb{CP}^2$. This domain is symplectomorphic to the codisc bundle $D^*_{1/2\pi}\mathbb{RP}^2$ of radius $\frac{1}{2\pi}$ in the round metric of curvature 1 on $\mathbb{RP}^2$. One more mutation at another vertex yields the triangles $\mathfrak{D}(2, 1, 5)$ and $\mathfrak{D}(2, 5, 1)$ and hence volume filling embeddings of $E_{2,1}(\frac{5}{2}, \frac{1}{10})$ and $E_{2,1}(\frac{1}{10}, \frac{5}{2})$ into $\mathbb{CP}^2$, respectively. This generates three steps of the $(2, 1)$-staircase, as illustrated in Figure 3.4.

The rest of the $p = 2$, $q = 1$ staircase comes from the sequence:

$$\{m_i\}_{i \in \mathbb{Z}} = \{\ldots, 29, 5, 1, 1, 5, 29, 169, \ldots\}$$

of odd-index Pell numbers, which yields the staircase depicted in Figure 3.4. This staircase is symmetric with respect to swapping the coordinate axes.

Example 3.1.16. For $(p, q) = (5, 1)$, we obtain

$$\{m_i\}_{i \in \mathbb{Z}} = \{\ldots, 433, 29, 2, 1, 13, 194, 2897, \ldots\},$$

yielding the staircase depicted in Figure 3.5. This staircase is not symmetric with respect to swapping the coordinate axes.

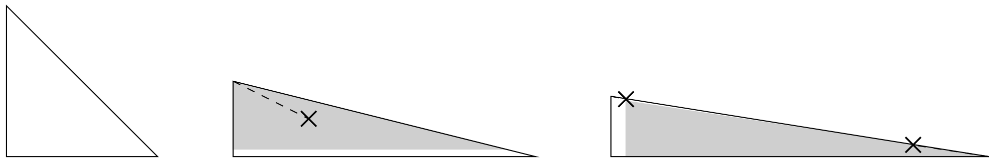

(a) $\mathfrak{D}(1,1,1)$ (b) $\mathfrak{D}(2,1,1)$ (c) $\mathfrak{D}(2,5,1)$

Figure 3.3: The first three Vianna triangles (a) $\mathfrak{D}(1,1,1)$, (b) $\mathfrak{D}(2,1,1)$ and (c) $\mathfrak{D}(2,5,1)$ together with the first two visible pin-ellipsoids (shaded in (b) and (c)) for the $p = 2, q = 1$ staircase. These visible embeddings can be made to fill an arbitrarily large portion of the triangle; in case (c) you need to use a nodal slide to shrink the other branch cut.

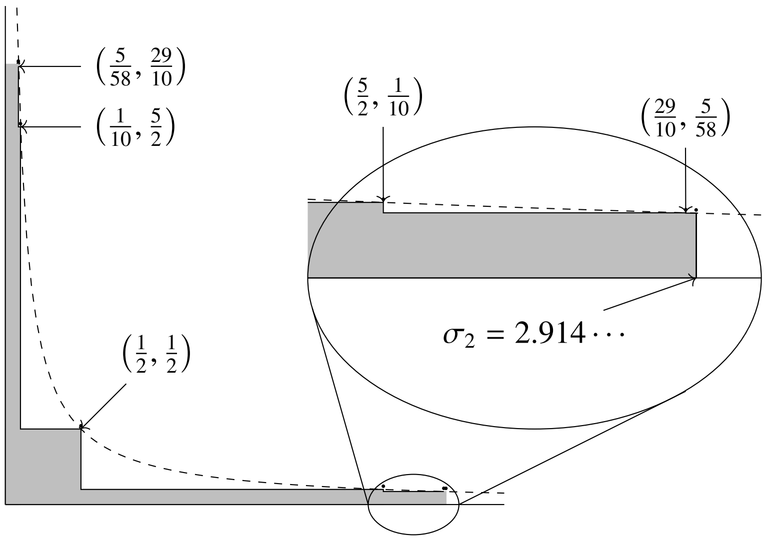

Figure 3.4: The staircase for $p = 2$, $q = 1$.

Remark 3.1.17. In particular, for a Markov number p with companion number q, Theorem 3.1.13 computes the largest (p,q)-pin-ball that symplectically embeds into $\mathbb{CP}^2$:

$$\sup\{\alpha > 0 \,:\, B_{p,q}(\alpha) \overset{s}{\hookrightarrow} X\} = \min\left\{\frac{p_2}{pp_3}, \frac{p_3}{pp_2}\right\}$$

where (p, p_2, p_3) denotes the Markov triple in which p appears as largest entry with companion number q. The only pin-balls admitting volume-filling embeddings into $\mathbb{CP}^2$ are therefore those with $(p,q) = (1,1)$ and $(p,q) = (2,1)$.

3.1.6 Pin-ball packing problem

The following result generalizes Gromov's Two Ball Theorem in dimension four.

Theorem 3.1.18 (Two Pin-Ball Theorem)**.** *Let p_1, p_2 be Markov numbers appearing in the same Markov triple*[1]*, and let p_3 be the smaller of the two Markov numbers completing the triple. Let*

[1]Whenever such a triple exists, there are precisely two such triples, and they are related by a single mutation.

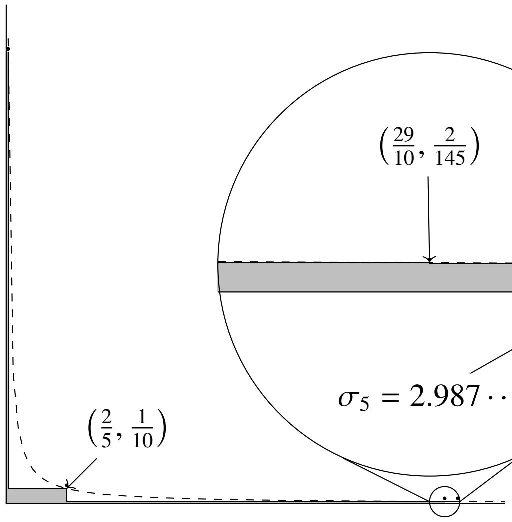

Figure 3.5: The staircase for $p = 5$, $q = 1$.

$q_i = \pm 3p_{i+1}p_{i+2}^{-1} \mod p_i$ *for* $i = 1, 2$ *be companion numbers. Then there exists a symplectic embedding*

$$B_{p_1,q_1}(\alpha_1) \sqcup B_{p_2,q_2}(\alpha_2) \hookrightarrow \mathbb{CP}^2, \tag{3.1.8}$$

if and only if $\alpha_1 < \frac{p_2}{p_1 p_3}$ *and* $\alpha_2 < \frac{p_1}{p_2 p_3}$ *and* $\alpha_1 + \alpha_2 < \frac{p_3}{p_1 p_2}$.

Remark 3.1.19. Note that one of the inequalities on the α_js is implied by the bound on $\alpha_1 + \alpha_2$. To see this, assume that $p_2 = \max\{p_1, p_2, p_3\}$; the case $p_1 = \max\{p_1, p_2, p_3\}$ follows by exactly the same argument. Then, assuming that $\alpha_1 + \alpha_2 < p_3/(p_1 p_2)$, we get

$$\alpha_1 < \alpha_1 + \alpha_2 < \frac{p_3}{p_1 p_2} \leq \frac{p_2}{p_1 p_3},$$

which proves the inequality for α_1. Moreover, Theorem 3.1.13 proves that $\alpha_2 < \frac{p_1}{p_2 p_3}$. This means that to establish the necessity of the inequalities in Theorem 3.1.18, we only need to prove the bound on $\alpha_1 + \alpha_2$.

Remark 3.1.20. The embeddings in Theorem 3.1.18 again appear as visible embeddings in Vianna triangles, see Figure 3.6. In general, if an embedding of the form (3.1.8) exists and $p_1, p_2 \geq 2$, then p_1 and p_2 are necessarily Markov numbers appearing in the same Markov triple and q_1 and q_2 are companion numbers of p_1 and p_2, respectively, see [ES18, Theorem 1.2]. However if, say, $p_1 = 1$, then $B_{p_1,q_1}(\alpha_1)$ is a standard ball and embeddings of the type (3.1.8) exist even if the Markov numbers $1, p_2$ do not appear in the same triple. In that case, constraints on α_1 and α_2 can be derived by the same methods as in [BS24].

Corollary 3.1.21 (Three Pin-Ball Theorem)**.** *Let* (p_1, p_2, p_3) *be a Markov triple with companion numbers* q_1, q_2, q_3. *Then there exists a symplectic embedding*

$$B_{p_1,q_1}(\alpha_1) \sqcup B_{p_2,q_2}(\alpha_2) \sqcup B_{p_3,q_3}(\alpha_3) \hookrightarrow \mathbb{CP}^2, \tag{3.1.9}$$

if and only if $\alpha_1 + \alpha_2 < \frac{p_3}{p_1 p_2}$, $\alpha_1 + \alpha_3 < \frac{p_2}{p_1 p_3}$ *and* $\alpha_2 + \alpha_3 < \frac{p_1}{p_2 p_3}$.

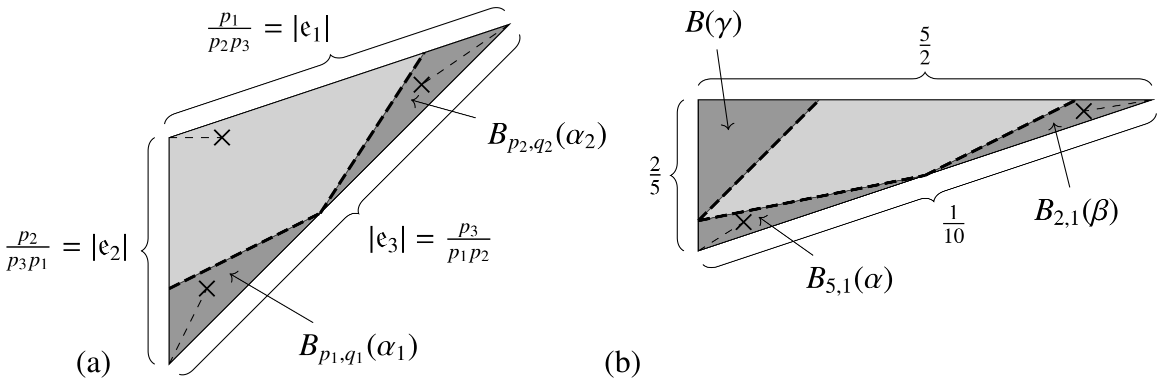

Figure 3.6: (a) The Vianna triangle $\mathfrak{D}(p_1, p_2, p_3)$ and an illustration of the constructive side of the Two Pin-Ball Theorem 3.1.18. (b) The Vianna triangle $\mathfrak{D}(5, 2, 1)$ and the pin-balls for the constructive side of the Two Pin-Ball Theorem for $\{p_1, p_2\} = \{1, 5\}$ and for $\{p_1, p_2\} = \{2, 5\}$ and of the Three Pin-Ball Theorem 3.1.21.

Proof. If $p_1 = p_2 = p_3 = 1$, the claim follows directly from Theorem 3.1.18 (or also from Gromov's 2-ball theorem). We can therefore assume that $p_3 \leq p_2 < p_1$. By Theorem 3.1.18,

$$\alpha_1 + \alpha_2 < \frac{p_3}{p_1 p_2}.$$

If we mutate (p_1, p_2, p_3) at p_2, then $p_2' = 3p_1 p_3 - p_2 > p_2$. Hence Theorem 3.1.18 applied to p_1, p_3 yields

$$\alpha_1 + \alpha_3 < \frac{p_2}{p_1 p_3}.$$

Finally, if we mutate (p_1, p_2, p_3) at p_1, then $p_1' = 3p_2 p_3 - p_1 < p_1$. Hence Theorem 3.1.18 applied to p_2, p_3 yields

$$\alpha_2 + \alpha_3 < \frac{p_1'}{p_2 p_3} < \frac{p_1}{p_2 p_3}. \qquad \square$$

Remark 3.1.22. Note that, by [ES18, Theorem 1.2], it is impossible to pack more than three pin-balls with $p_i \geq 2$.

3.1.7 Idea of proof for Theorem 3.1.8

Our proof builds on the idea of Hind [Hin12] and Li–Wu [LW12, § 6.1.1] who showed the Lagrangian uniqueness of $\mathbb{RP}^2$ in $\mathbb{CP}^2$ by turning this problem into the uniqueness problem of symplectic -4-spheres in $S^2 \times S^2$ obtained by a symplectic cut on the boundary of small codisc-bundle of an embedded $\mathbb{RP}^2$ in $\mathbb{CP}^2$. Suppose we have two symplectic embeddings $\iota_1, \iota_2 \colon E_{p,q}(\alpha, \beta) \hookrightarrow \mathbb{CP}^2$. We shall perform a rational blow-up along each of these pin-ellipsoids; this yields two 4-manifolds $\widetilde{X}_1$ and $\widetilde{X}_2$ containing configurations C_1 and C_2 of embedded symplectic spheres. We show that, for $i = 1, 2$, $\widetilde{X}_i$ is ruled, and that:

- in each case, one of the curves from C_i is a section of this ruled surface,
- the other curves of C_i (together with some additional curves) form one or two broken rulings.

- the configuration $\mathcal{T}_i$ consisting of C_i and these additional curves (and possibly one or two additional smooth rulings) is completely determined by p, q and the choice of data of the rational blow-up. In other words $\mathcal{T}_1$ and $\mathcal{T}_2$ have symplectomorphic neighbourhoods $\nu_1 \cong \nu_2$.
- the common boundary $\partial\nu_i$ is contactomorphic to $S^1 \times S^2$ and the complement of ν_i in $\widetilde{X}_i$ is minimal. In particular, uniqueness of symplectic fillings of these manifolds tells us that the symplectic completion of $\widetilde{X}_i \setminus \nu_i$ is symplectomorphic to either $\mathbb{R}^4$ or $T^*S^1 \times \mathbb{C}$.

Since $\widetilde{X}_1$ and $\widetilde{X}_2$ are essentially obtained by gluing together symplectomorphic pieces, and there is only one way to glue these together, we find that $\widetilde{X}_1$ and $\widetilde{X}_2$ are symplectomorphic via a symplectomorphism taking C_1 to C_2. This descends to give a symplectomorphism of the rational blow-down $\mathbb{CP}^2 \to \mathbb{CP}^2$ taking one pin-ellipsoid to the other. Finally, this symplectomorphism is Hamiltonian due to Gromov's classical result stating that the symplectomorphism group of $\mathbb{CP}^2$ is connected.

3.1.8 Idea of proof for Theorems 3.1.13 and 3.1.18

To explain the proof of the Staircase Theorem (Theorem 3.1.13), let p be a Markov number and q a companion number. As mentioned in Remark 3.1.14 (b), the points in $\mathrm{Stair}(p,q)$ are realised by visible embeddings coming from Vianna triangles. It is therefore enough to show that for all $i \in \mathbb{Z}$, $E_{p,q}(\alpha,\beta)$ does not symplectically embed into $\mathbb{CP}^2$ whenever $\alpha > \frac{m_i}{pm_{i-1}}$ and $\beta > \frac{m_i}{pm_{i+1}}$. Here m_i denote the members of the sequence defined in (3.1.5). Note that these points correspond to the inner corners of the staircase $\mathrm{Stair}(p,q)$.

Assuming that such an embedding $\iota : E_{p,q}(\alpha,\beta) \hookrightarrow \mathbb{CP}^2$ exists, we construct a one-parameter family of symplectic manifolds $(\widetilde{X}, \widetilde{\omega}_t)$ by rational symplectic blow-up[2] of a small neighbourhood of $\iota\left(tE_{p,q}(\alpha,\beta)\right)$, where the parameter $t \in (0,1]$ measures the size of the blow-up. For small enough t, we can assume, by Theorem 3.1.8, that the embedding coincides with the visible embedding coming from the Vianna triangle $\mathfrak{D}(p, m_{i-1}, m_i)$. In that case, we construct a *visible* symplectic -1-sphere in the base diagram of $(\widetilde{X}, \widetilde{\omega}_t)$, see Figure 3.22.[3] Hence we find a class $S \in H_2(\widetilde{X};\mathbb{Z})$ whose Gromov invariant is one. Furthermore, homological computations show that the symplectic area of S as measured by $\widetilde{\omega}_t$ becomes negative for t close enough to 1. However, the Gromov invariant is a deformation invariant, meaning that S should be represented by a holomorphic curve for all $\widetilde{\omega}_t$. If $\alpha > \frac{m_i}{pm_{i-1}}$ and $\beta > \frac{m_i}{pm_{i+1}}$ for some i then we show that t can be chosen large enough that the symplectic area of S is negative. Since holomorphic curves have positive area, we obtain a contradiction.

This idea of using visible curves to obstruct embeddings of standard ellipsoids can be found in McDuff–Siegel [MS25b, Figure 4.2.1 and Corollary 6.1.4], so the main new input here is the Isotopy Theorem. Note that in the proof we only use a weak version of the Isotopy Theorem 3.1.8, namely that two pin-ellipsoid embeddings can be made to agree on an arbitrarily small pin-subellipsoid. For symplectic embeddings of balls ($p = 1$, where the pin-wheel becomes a point) this is an elementary fact; as a consequence, the proof of the Fibonacci staircase (in the visible range, below the aspect ratio ϕ^4) is immediate from almost toric geometry: the optimal embeddings are visible and the obstructions come from visible curves. This is implicit in [MS25b].

[2] Technically, we need to take a blow-up of the rational blow-up which we refer to as a *pavilion blow-up*.

[3] In the case $p = 1$ this curve corresponds to the unicuspidal curve in $\mathbb{CP}^2$ used in [MS25b, Figure 4.2.1], but here it is found directly in $\widetilde{X}$ and hence it is a smoothly embedded -1-curve.

The proof of the Two Pin-Ball Theorem (Theorem 3.1.18) also comes from positivity of area for a pseudoholomorphic curve, and can be seen as an elaborate version of Gromov's proof of his Two Ball Theorem. Suppose we have disjoint symplectically embedded pin-balls $B_{p_1,q_1}(\alpha_1)$ and $B_{p_2,q_2}(\alpha_2)$ in $\mathbb{CP}^2$. In this case, we rationally blow-up both pin-balls; this produces a manifold $\widetilde{X}$ containing two chains of symplectic spheres. We find a symplectic -1-sphere $\widetilde{C}$ in $\widetilde{X}$ which connects these chains and whose symplectic area is bounded by $\frac{p_3}{p_1 p_2}$ where p_3 is the smaller Markov number making (p_1, p_2, p_3) a Markov triple. Contracting the chains yields an orbifold with two orbifold points, and the image of $\widetilde{C}$ passes through both of them. We then lift to the uniformising cover of each orbifold chart and apply the monotonicity formula for areas of minimal surfaces to get a lower bound of $\alpha_1 + \alpha_2$ on the symplectic area of $\widetilde{C}$. In the case when the pin-balls are simultaneously visible, the sphere $\widetilde{C}$ is also visible as part of the toric boundary; see Figure 3.7. One could prove the existence of $\widetilde{C}$ in general by proving an isotopy theorem for pairs of pin-balls, but to save space we instead appeal to the paper of Evans and Smith [ES18, Theorem 4.16].

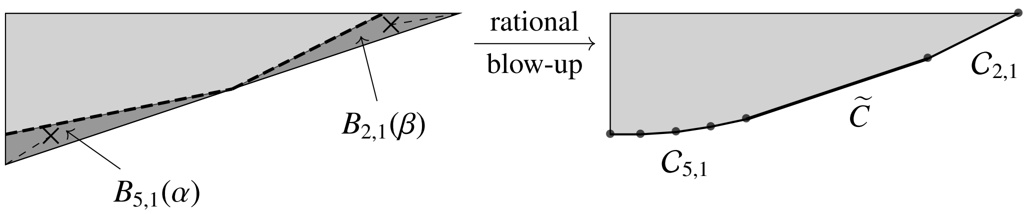

Figure 3.7: $\widetilde{C}$ is the exceptional curve used to prove the obstructive side of the Two Pin-Ball Theorem 3.1.18. The situation above is the visible situation for the $(5, 2)$-case. The Wahl chains $C_{5,1}$ and $C_{2,1}$ are connected by the exceptional curve $\widetilde{C}$. In general we will find the curve $\widetilde{C}$ by a neck stretching procedure.

3.1.9 Outline of the paper

In Section 3.2, we introduce the basic constructions from almost toric geometry which are needed for the rest of the paper, including the Vianna triangles (almost toric base diagrams for $\mathbb{CP}^2$) which allow us to see the visible pin-ellipsoids, the *pavilion blow-up* (a mild generalisation of rational blow-up) and the *culet curve* (a distinguished curve in the pavilion blow-up of a pin-ellipsoid in $\mathbb{CP}^2$). In Section 3.3, we explain in detail how to equip the pavilion blow-up with a symplectic structure and compatible almost complex structure and how to perform computations in the homology after blowing up. In Section 3.4, we collect some basic results about ruled symplectic 4-manifolds which will be used in the main argument. Section 3.5 contains the proofs of the main results:

- The main technical results are: the existence of a square zero holomorphic sphere in the rational blow-up (Theorem 3.5.1, proved in Section 3.5.4) and the description of how this sphere can degenerate to give broken rulings (Theorem 3.5.4, proved in Section 3.5.6).
- The proofs of the Isotopy Theorem 3.1.8 and the Staircase Theorem 3.1.13, assuming these technical results, are given in Section 3.5.2 and 3.5.3 respectively.
- The Two Pin-Ball Theorem (Theorem 3.1.18) is proved in Section 3.5.5.

Acknowledgements

The authors would like to thank Gerard Bargalló i Gómez, Marco Golla, Richard Hind, Paul Levy, Grisha Mikhalkin, Dusa McDuff, Federica Pasquotto, George Politopoulos, Joel Schmitz, Kyler Siegel, Laura Starkston, Giancarlo Urzúa, Renato Vianna and Weiyi Zhang for helpful conversations.

JB is supported by SNSF Ambizione Grant PZ00P2-223460.

JE is supported by EPSRC Standard Grant EP/W015749/1.

3.2 Notation and definitions

3.2.1 Wahl singularities and their resolutions

Definition 3.2.1 (Triangle $\Delta_{p,q}(\alpha,\beta)$)**.** If we ignore the node and branch cut for the almost toric base diagram $\mathfrak{A}_{p,q}(\alpha,\beta)$ from Definition 3.1.1 then we get a triangle (see Figure 3.8):

$$\Delta_{p,q}(\alpha,\beta) := \left\{ x \in \mathbb{R}^2 \mid \rho_0 \cdot x \geq 0, \quad \rho_{m+1} \cdot x \geq 0, \quad \gamma \cdot x \geq -\alpha\beta p^2 \right\} \tag{3.2.10}$$

where ρ_0, ρ_{m+1} and γ are the vectors

$$\rho_0 = (1,0), \quad \rho_{m+1} = (1-pq, p^2), \quad \text{and} \quad \gamma = (\alpha pq - \alpha - \beta, -\alpha p^2).$$

We use the index $m+1$ here because we will soon add further edges with inward normals $\rho_1, \ldots, \rho_m$; see Definition 3.2.7.

Definition 3.2.2 (Girdled orbifold $O_{p,q}(\alpha,\beta)$)**.** This triangle is the moment map image of a compact toric orbifold $O_{p,q}(\alpha,\beta)$ with boundary. As in Remark 3.1.2(a), the top side of $\Delta_{p,q}(\alpha,\beta)$ with inward normal γ is considered to be part of the boundary, not the toric boundary[4]. We refer to the top side as the *girdle*[5] and we call this bounded toric orbifold the *girdled orbifold*.

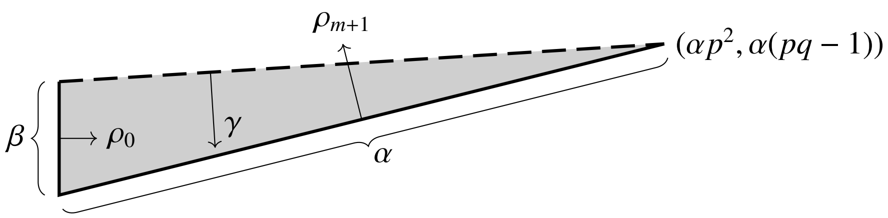

Figure 3.8: The moment image $\Delta_{p,q}(\alpha,\beta)$ of a toric orbifold. The toric boundary comprises the two solid edges, *not* the top side; the orbifold point is at the vertex of the polygon. We have also labelled the inward normals ρ_0 and ρ_{m+1} of the toric boundary and γ of the top side.

The girdled orbifold $O_{p,q}(\alpha,\beta)$ has a singularity x living over the vertex of $\Delta_{p,q}(\alpha,\beta)$, locally modelled on the cyclic quotient singularity $\frac{1}{p^2}(1, pq-1)$. This is the quotient of $\mathbb{C}^2$ by the action of $\boldsymbol{\mu}_{p^2} = \{\xi \in \mathbb{C}^\times \mid \xi^{p^2} = 1\}$ where $\xi \cdot (x,y) = (\xi x, \xi^{pq-1} y)$. Singularities of this form

[4]The boundary has real codimension 1; the toric boundary is a divisor, with complex codimension 1.

[5]The terminology we introduce in this paper (girdle, pavilion, culet) is standard for cut gems [Bri90].

are called *Wahl singularities*. These are precisely the cyclic quotient singularities which admit a $\mathbb{Q}$-Gorenstein smoothing of Milnor number zero. The Milnor fibre of this smoothing is a pin-ellipsoid, so whenever a smooth surface admits a $\mathbb{Q}$-Gorenstein degeneration with Wahl singularities, we can find symplectically embedded pin-ellipsoids in the smooth fibres.

Definition 3.2.3 (Resolutions)**.** A *resolution* of $\mathcal{O}_{p,q}(\alpha,\beta)$ is a smooth complex manifold $(\widetilde{U},\widetilde{J})$ together with a holomorphic map $\pi\colon \widetilde{U} \to \mathcal{O}_{p,q}(\alpha,\beta)$ satisfying the following property. There exists a configuration $\mathcal{C}$ of complex curves $C \subset \widetilde{U}$ such that $\pi(C) = x$ for all $C \in \mathcal{C}$ and such that π restricts to a biholomorphism $\widetilde{U} \setminus \bigcup_{C\in\mathcal{C}} C \to \mathcal{O}_{p,q}(\alpha,\beta) \setminus \{x\}$. We say that the resolution is *chain-shaped* if the curves $C \in \mathcal{C}$ are embedded spheres and form a chain, that is, $\mathcal{C} = \{C_1, C_2, \ldots, C_m\}$ with

$$C_i \cdot C_j = \begin{cases} 1 & \text{if } |i-j| = 1 \\ 0 & \text{otherwise.} \end{cases}$$

Definition 3.2.4 (Girdled resolution $\widetilde{U}_\rho$)**.** There is a standard family of chain-shaped resolutions of $\mathcal{O}_{p,q}(\alpha,\beta)$ which are themselves toric. They arise from subdivisions of the inward normal fan for the girdled orbifold. Recall that the inward normals to the two edges of $\Delta_{p,q}(\alpha,\beta)$ are ρ_0 and ρ_{m+1}. A subdivision $\boldsymbol{\rho}$ is a sequence of primitive integer vectors $\rho_1, \ldots, \rho_m$ in the convex sector bounded by ρ_0 and ρ_{m+1} with the property that $\rho_i \wedge \rho_{i+1} \geq 1$ for $i = 0, 1, \ldots, m$. Here, $v \wedge w$ denotes the determinant of the 2-by-2 matrix with columns v and w. If the subdivision $\boldsymbol{\rho}$ satisfies $\rho_i \wedge \rho_{i+1} = 1$ the subdivided fan $\boldsymbol{\rho}$ determines the resolution as a normal toric variety, see for example [Ful93, Section 2.6]. Since we are resolving a girdled orbifold, we will call this the *girdled resolution* $\widetilde{U}_\rho$ associated to $\boldsymbol{\rho}$.

A general prescription for finding resolutions of a cyclic quotient singularity $\frac{1}{n}(a,1)$ works as follows. The exceptional set of the minimal resolution is a chain of holomorphic spheres $C_1, \ldots, C_m$ with $C_i^2 = -b_i$ where $[b_1, \ldots, b_m]$ is the Hirzebruch–Jung continued fraction expansion of n/a, that is:

$$[b_1, \ldots, b_m] := b_1 - \cfrac{1}{b_2 - \cfrac{1}{\ddots - \frac{1}{b_m}}}.$$

Remark 3.2.5. Girdled resolutions all *dominate* a unique *minimal resolution*, meaning that they are obtained from it by a sequence of complex blow-ups. For the minimal resolution, the curves C_i are embedded spheres with self-intersection $C_i^2 = -b_i$ where $[b_1, \ldots, b_m]$ is the Hirzebruch–Jung continued fraction expansion of $p^2/(pq-1)$. Non-minimal resolutions are obtained from the minimal one by blowing up at intersection points between the exceptional spheres C_i.

Definition 3.2.6 (Wahl chains and duals)**.** In the special case of Wahl singularities, we call $[b_1, \ldots, b_m]$ a *Wahl chain*. Recall that the *dual* of the singularity $\frac{1}{n}(1,a)$ is the singularity $\frac{1}{n}(1,\overline{a})$ where $a\overline{a} = 1 \mod n$; the continued fraction of $n/\overline{a}$ is $[b_m, \ldots, b_1]$, the *dual* of the chain for n/a.

Definition 3.2.7 (Pavilion)**.** Let $\boldsymbol{\rho}$ be a subdivision and, for each $\rho_i \in \boldsymbol{\rho}$, choose a positive real number $\lambda_i > 0$. We then consider the polygon

$$\mathrm{Pav}_{\rho,\lambda}\left(\Delta_{p,q}(\alpha,\beta)\right) := \left\{x \in \Delta_{p,q}(\alpha,\beta) \mid \rho_i \cdot x \geq \lambda_i \text{ for } i = 1, \ldots, m\right\}, \tag{3.2.11}$$

see Figure 3.9. We require that $\rho_i \cdot x > \lambda_i$ for all points x on the girdle (recall that this is the top side of the moment triangle). A choice of $\boldsymbol{\rho}$ and $\boldsymbol{\lambda}$ satisfying this requirement is called a *pavilion*. If, moreover, for every $i = 1, 2, \dots, m$ we have that (a) $\rho_i \wedge \rho_{i+1} = 1$ and (b) there is an edge $\mathfrak{f}_i$ of the polygon with inward normal ρ_i and positive length, then we call the pavilion *Delzant*. We call a pavilion *minimal* if the corresponding toric girdled resolution $\widetilde{U}_\rho$ is the minimal resolution.

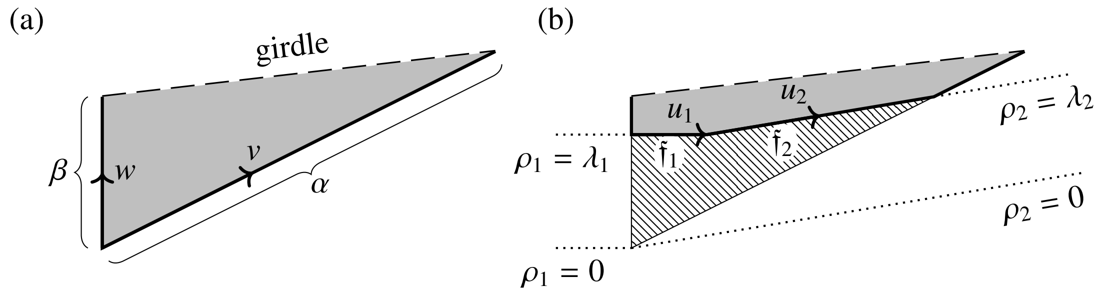

Figure 3.9: (a): $\Delta_{p,q}(\alpha,\beta)$ with its girdle. (b): A pavilion $\mathrm{Pav}_{\rho,\lambda}\left(\Delta_{p,q}(\alpha,\beta)\right)$ for $\Delta_{p,q}(\alpha,\beta)$.

If we choose a Delzant pavilion and a strictly convex function ψ on it (with suitable behaviour along the edges of the polygon) then we get a toric Kähler metric on the resolution constructed using the Hessian of ψ (see [Abr03, Eq. (2.1)]; the boundary behaviour along edges is obtained from [Abr03, Eq. (2.8)] by taking Legendre transform). Therefore, $\mathrm{Pav}_{\rho,\lambda}\left(\Delta_{p,q}(\alpha,\beta)\right)$ is the moment image for a toric Kähler structure on the girdled resolution $\widetilde{U}_\rho$. We write $\widetilde{\omega}_{\rho,\lambda,\psi}$ for the Kähler form.

Definition 3.2.8 (Pavilion blow-up)**.** We call the symplectic manifold $\left(\widetilde{U}_\rho, \widetilde{\omega}_{\rho,\lambda,\psi}\right)$ the *pavilion blow-up of* $E_{p,q}(\alpha,\beta)$.

Remark 3.2.9. The specific choices of α, β and $\boldsymbol{\lambda}$ will be important for some of the quantitative questions we consider. The freedom to choose non-minimal resolutions will help with finding holomorphic curves which give obstructions to pin-ellipsoid embeddings.

Remark 3.2.10. Although we construct the symplectic form in this way, one could also think of the pavilion blow-up as coming from a sequence of symplectic cuts. This would have the disadvantage that different choices of $\boldsymbol{\lambda}$ would yield symplectic forms *on different manifolds* (all diffeomorphic, but still different) making it harder to compare almost complex structures. Nonetheless, we will sometimes abuse language and talk about "making a pavilion cut"; we just mean picking a particular $\boldsymbol{\rho}$ and $\boldsymbol{\lambda}$.

3.2.2 Offcuts

Definition 3.2.11. Given a pavilion $\boldsymbol{\rho}, \boldsymbol{\lambda}$ for $\Delta_{p,q}(\alpha,\beta)$, we can think of the pavilion as a subset of the almost toric base diagram $\mathfrak{A}_{p,q}(\alpha,\beta)$, and nodally slide so that the focus-focus singularity lies below $\mathrm{Pav}_{\rho,\lambda}\left(\Delta_{p,q}(\alpha,\beta)\right)$. Define the *offcut*

$$\mathrm{Off}_{\rho,\lambda}\left(\mathfrak{A}_{p,q}(\alpha,\beta)\right) := \mathfrak{A}_{p,q}(\alpha,\beta) \setminus \mathrm{Pav}_{\rho,\lambda}\left(\Delta_{p,q}(\alpha,\beta)\right).$$

We call the associated almost toric domain the *almost toric offcut* $\mathrm{Off}_{\rho,\lambda}\left(E_{p,q}(\alpha,\beta)\right)$; see Figure 3.10.

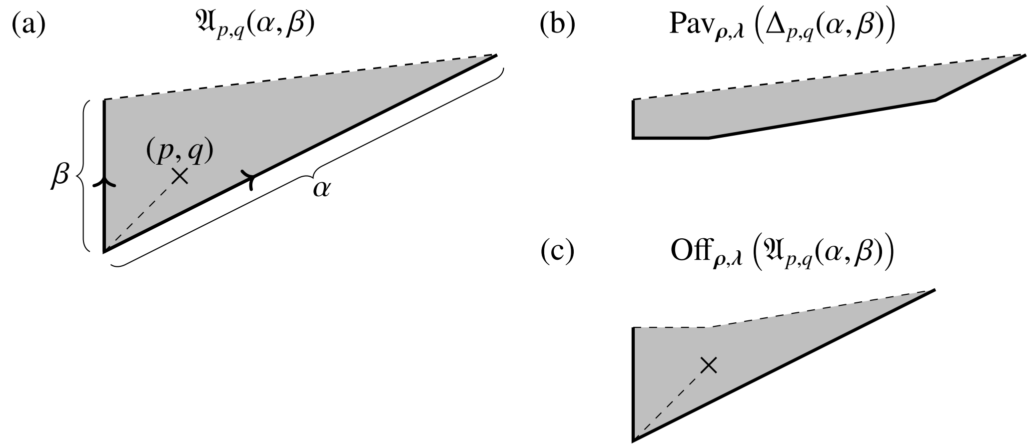

Figure 3.10: (a): The almost toric base diagram $\mathfrak{A}_{p,q}(\alpha,\beta)$. (b): The moment image of the pavilion blow-up $\widetilde{U}_\rho$. (c): The almost toric offcut $\mathrm{Off}_{\rho,\lambda}\left(E_{p,q}(\alpha,\beta)\right)$.

3.2.3 Vianna triangles

Let $\mathfrak{D}(1,1,1)$ be the triangle with vertices $\mathfrak{v}_1 = (0,0)$, $\mathfrak{v}_2 = (1,0)$ and $\mathfrak{v}_3 = (0,1)$. This is the moment image of a Hamiltonian torus action on $\mathbb{CP}^2$. We can make nodal trades at the three corners to get an almost toric base diagram with three nodes $\mathfrak{n}_1, \mathfrak{n}_2, \mathfrak{n}_3$ where $\mathfrak{n}_i$ is connected to the vertex $\mathfrak{v}_i$ by a branch cut $\mathfrak{b}_i$. The result is the base of an almost toric fibration on $\mathbb{CP}^2$. We can perform mutations of this almost toric base diagram and we obtain a sequence of almost toric base diagrams called *Vianna triangles*. For a general introduction to Vianna triangles and their elementary properties see [Eva23, Appendix I]. We continue to write $\mathfrak{v}_1, \mathfrak{v}_2, \mathfrak{v}_3$ for the vertices (ordered anticlockwise) and $\mathfrak{b}_i$ for the branch cut connecting v_i to the node $\mathfrak{n}_i$. Recall that the determinant[6] at the vertex $\mathfrak{v}_i$ is equal to p_i^2 for some positive integer p_i, and together, (p_1, p_2, p_3) form a *Markov triple*. Recall that (p_1, p_2, p_3) is called a Markov triple if

$$p_1^2 + p_2^2 + p_3^2 = 3p_1p_2p_3. \tag{3.2.12}$$

We will write $\mathfrak{D}(p_1, p_2, p_3)$ for the Vianna triangle corresponding to (p_1, p_2, p_3). If we mutate $\mathfrak{D}(p_1, p_2, p_3)$ at the vertex p_2, say, then the result is $\mathfrak{D}(p_1', p_2', p_3')$ where

$$p_1' = p_1, \quad p_2' = p_3, \quad p_3' = 3p_1p_3 - p_2. \tag{3.2.13}$$

Label the edges of $\mathfrak{D}(p_1, p_2, p_3)$ as $\mathfrak{e}_1, \mathfrak{e}_2, \mathfrak{e}_3$, where $\mathfrak{e}_i$ is opposite vertex $\mathfrak{v}_i$. The affine length of $\mathfrak{e}_i$ is $\frac{p_i}{p_{i+1}p_{i+2}}$, see [Eva23, Corollary I.13].

If we ignore the nodes and branch cuts, then $\mathfrak{D}(p_1, p_2, p_3)$ is the moment image for the toric orbifold surface $\mathbb{P}(p_1^2, p_2^2, p_3^2)$ which has a $\frac{1}{p_i^2}(p_{i+2}^2, p_{i+1}^2)$ cyclic quotient singularity[7] living over

[6] If the primitive integer vectors along the edges emanating from $\mathfrak{v}_i$ are u and v, then the determinant of $\mathfrak{v}_i$ is defined to be $|u \wedge v|$, that is the absolute value of the determinant of the matrix whose columns are u and v.

[7] Note that there is a confusing ordering issue here. The conventions we are using for the moment polygon ensure that the Hirzebruch–Jung chain appears in order if you read the self-intersections off anticlockwise. However, with this convention, the moment-preimage of the vertical edge in $\Delta_{p,q}(\alpha,\beta)$ is contained in the x-axis and the moment-preimage of the slanted edge is contained in the y-axis. Indeed, the affine change of coordinates in [Eva23, Example 3.21] which finds the moment polygon for an orbifold singularity has negative determinant.

the vertex $\mathfrak{v}_i$. Note that (taking indices modulo 3) $\frac{1}{p_i^2}(p_{i+2}^2, p_{i+1}^2) \cong \frac{1}{p_i^2}(p_{i+1}^2 p_{i+2}^{-2}, 1)$, where the inverse is computed modulo p_i^2. Equation (3.2.12) tells us that $p_{i+1}^2 p_{i+2}^{-2} = 3p_i p_{i+2}^{-1} p_{i+1} - 1 \mod p_i^2$. Furthermore, $p_{i+1}^2 p_{i+2}^{-2} = p_i q_i - 1 \mod p_i^2$ since the singularity is Wahl. Therefore,

$$q_i = +3p_{i+1}p_{i+2}^{-1} \mod p_i. \tag{3.2.14}$$

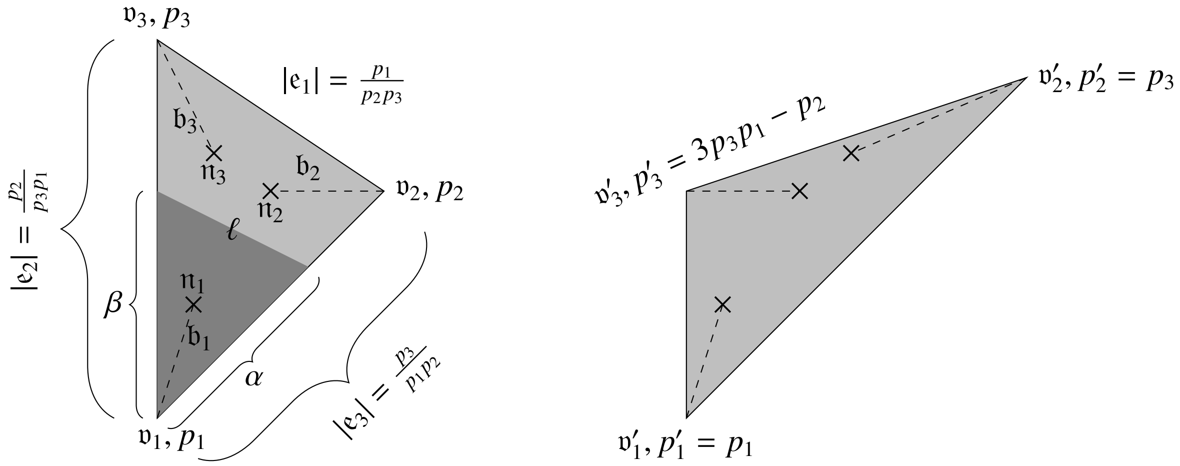

Figure 3.11: A Vianna triangle and its mutation at $\mathfrak{v}_2$; quantities for the mutated triangle are written with primes. The darker shaded region bounded by the line ℓ is the visible $\mathfrak{A}_{p_1,q_1}(\alpha,\beta)$. Next to each vertex $\mathfrak{v}_i$ we also give the Markov number p_i with $\det(\mathfrak{v}_i) = p_i^2$.

Lemma 3.2.12. *Given a Markov triple (p_1, p_2, p_3), take the companion number $q_i = 3p_{i+1}p_{i+2}^{-1} \mod p_i$ for $i = 1, 2, 3$ (indices taken modulo 3). Then $\mathbb{CP}^2$ contains a visible symplectic $E_{p_i,q_i}(\alpha,\beta)$ for any $\alpha < \frac{p_{i+2}}{p_i p_{i+1}}$ and $\beta < \frac{p_{i+1}}{p_i p_{i+2}}$.*

Proof. The constraints on α and β mean that we can pick the points $\mathfrak{p}_{i+2}$ on $\mathfrak{e}_{i+2}$ an affine length α from v_i and $\mathfrak{p}_{i+1}$ on $\mathfrak{e}_{i+1}$ an integral affine length β from v_i. Connect them with a straight line ℓ; nodally slide $\mathfrak{n}_i$ so that it lives on one side of ℓ and $\mathfrak{n}_{i+1}$ and $\mathfrak{n}_{i+2}$ so that they live on the far side of ℓ. The line ℓ divides $\mathfrak{D}(p_1, p_2, p_3)$ into two pieces; the piece containing $\mathfrak{n}_i$ is integral-affine equivalent to $\mathfrak{A}_{p_i,q_i}(\alpha,\beta)$, so we obtain a visible embedding of $E_{p_i,q_i}(\alpha,\beta)$ as required; see Figure 3.11. □

3.2.4 The culet

Once again, ignore the almost toric data and consider $\mathfrak{D}(p_1, p_2, p_3)$ as the moment image of the toric orbifold $\mathbb{P}(p_1^2, p_2^2, p_3^2)$. Suppose that p_1 is the biggest number in the Markov triple (p_1, p_2, p_3) and use an integral affine transformation to make the edge $\mathfrak{e}_1$ of $\mathfrak{D}(p_1, p_2, p_3)$ horizontal, with the rest of the triangle below.

Definition 3.2.13. Let us subdivide the fan of inward normals for the Vianna triangle by adding a ray pointing vertically upwards. This gives partial resolution of the $\frac{1}{p_1^2}(p_2^2, p_3^2)$ singularity

This is why we write $\frac{1}{p_i^2}(p_{i+2}^2, p_{i+1}^2)$ rather than $\frac{1}{p_i^2}(p_{i+1}^2, p_{i+2}^2)$.

introducing a single exceptional curve which we call the *culet curve*.[8] If we fix a sufficiently small λ then we get a (non-Delzant) pavilion with a new horizontal edge at the bottom, which we call the *culet*; see Figure 3.12.

There are still (up to) two orbifold singularities on this edge at the new vertices $\mathfrak{v}_1'$ and $\mathfrak{v}_1''$; these are toric, with their moment images modelled on the *alternate angles* for the corners $\mathfrak{v}_2$ and $\mathfrak{v}_3$ and so are *dual* to these Wahl singularities [MS23, Proposition 6]. We can further (minimally) resolve these singularities by completing our subdivision to a Delzant pavilion $\boldsymbol{\rho}$. Corresponding to these new rays in the fan, we have a chain of exceptional curves of the form $[W_0^\vee, w, W_1^\vee]$ where $W_i^\vee$ are dual Wahl chains and $-w$ is the self-intersection of the culet curve. By a theorem of Urzúa and Zúñiga [UZ25, Proposition 4.1], this is only possible if:

- both $W_i^\vee$ are nonempty chains and $w = 10$,
- precisely one of the $W_i^\vee$ is nonempty and $w = 7$, or
- both of the $W_i^\vee$ are empty and $w = 4$.

Hence the chain $[W_0^\vee, w, W_1^\vee]$ does not contain any (-1)-spheres and is therefore minimal.

Remark 3.2.14. If we pick a non-minimal resolution, it dominates the minimal resolution. We will still refer to the proper transform of the culet curve as the culet curve, so the notion makes sense for arbitrary resolutions.

Definition 3.2.15. Given a chain-shaped resolution of the $\frac{1}{p_1^2}(p_2^2, p_3^2)$ cyclic quotient singularity, we write $C_1, \dots, C_m$ for the exceptional curves of this resolution and write i_{culet} for the index of the culet curve.

Definition 3.2.16. Let w be $-C_{i_{\text{culet}}}^2$ for the minimal resolution. We call this the *Manetti weight* of the pin-ellipsoid (note that it only makes sense when p is a Markov number and q a companion of p).

Remark 3.2.17. Recall that $w \in \{4, 7, 10\}$ by the aforementioned result of Urzúa and Zúñiga [UZ25, Proposition 4.1]. The Manetti weights depend on the pairs (p, q) as follows. The pair $(p, q) = (2, 1)$ is the only one which yields $w = 4$. The Manetti weight is $w = 7$ if and only if (p, q) comes from a Markov triple containing a 1. In every other case it is $w = 10$. Note that the square of the culet curve in a non-minimal resolution will be bigger than the weight if our non-minimal resolution is obtained from the minimal one by blow up of points on the culet curve.

Since the edge $\mathfrak{e}_1$ is parallel to the culet, there is a visible symplectic sphere in the sense of Symington [Sym03, Definition 7.2] having square zero and living over a vertical arc connecting $\mathfrak{e}_1$ and the culet. One of the main ideas in what follows will be to find the analogue of this square zero sphere for an arbitrary symplectic embedding $E_{p,q}(\alpha, \beta) \hookrightarrow \mathbb{CP}^2$.

3.3 Geometry and topology of pavilion blow-ups

In Section 3.3.1, we explain in detail the almost complex and symplectic structures with which we will equip pavilion blow-ups. In Section 3.3.2, we explain how to compute in the cohomology ring of the pavilion blow-up.

[8] We will continue to refer to the proper transform of this curve as the culet curve in any resolution dominating this one.

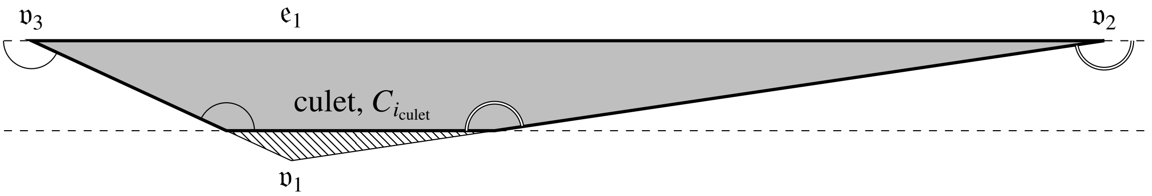

Figure 3.12: The culet cut is horizontal, parallel to the edge $\mathfrak{e}_1$ of affine length $p_1/(p_2p_3)$. The new corners are integral-affine dual to the corners at $\mathfrak{v}_2$ and $\mathfrak{v}_3$.

3.3.1 Pavilion blow up and its geometric structures

Notation 3.3.1. *Given a symplectic manifold* (W, ω) *with a convex (respectively concave) boundary, let* $\overline{W}$ *denote its symplectic completion, obtained by gluing on copies of the positive (respectively negative) half of the symplectisation of the boundary. If* J *is an* ω*-compatible almost complex structure on* W *which is cylindrical near the ends then let* $\overline{J}$ *denote the unique almost complex structure on* $\overline{W}$ *extending* J *cylindrically to the ends. If* W *has both convex and concave ends then define* $\overline{W}_+$ *(respectively* $\overline{W}_-$*) to be the result of completing only the convex (respectively concave) ends.*

Definition 3.3.2. A neighbourhood N of the boundary of $E_{p,q}(\alpha,\beta)$ is symplectomorphic to a neighbourhood of the boundary of the girdled orbifold $O_{p,q}(\alpha,\beta)$, namely the moment-preimage of the shaded region in $\Delta_{p,q}(\alpha,\beta)$ shown in Figure 3.13. Since this toric orbifold is a finite quotient of a compact set in $\mathbb{C}^2$, it has a complex structure. Let J_N be the pullback of this complex structure to N.

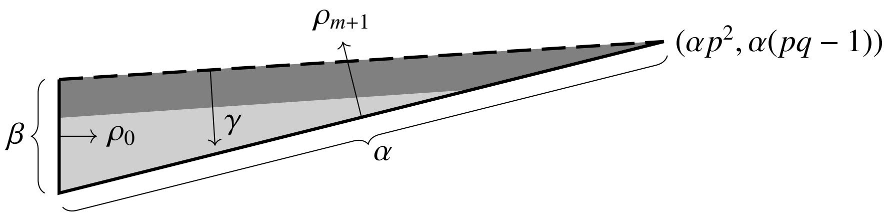

Figure 3.13: The moment image $\Delta_{p,q}(\alpha,\beta)$ of a toric orbifold. The preimage under the moment map of the shaded region is symplectomorphic to a neighbourhood of $\partial E_{p,q}(\alpha,\beta)$ in $E_{p,q}(\alpha,\beta)$. The toric boundary comprises the two solid edges, *not* the top side; the orbifold point is at the vertex of the polygon. We have also labelled the inward normals ρ_0 and ρ_{m+1} of the toric boundary and γ of the top side (girdle).

Definition 3.3.3. Let (X, ω) be a symplectic manifold and suppose that we have a symplectic embedding $\iota\colon E_{p,q}(\alpha,\beta) \hookrightarrow X$. Let $\mathcal{J}_N$ be the space of compatible almost complex structures on X which coincide with $\iota_* J_N$ on $\iota(N)$.

These almost complex structures are *adjusted* to the neck region N in the sense of [Bou+03, Section 2.2] and so amenable to *neck-stretching*. Neck-stretching (or splitting) is a specific one-parameter deformation of the almost complex structure introduced in [EGH00, Section 1.3]; neck-stretching for pin-ellipsoids is discussed in detail in [Fer16, Sections 3.2–3.3] and [ES18, Section 3.1].

Definition 3.3.4. Given a $J \in \mathcal{J}_N$, we get a neck-stretching sequence of compatible almost complex structures J_t on X with $J_0 = J$; let us write $U = \iota(E_{p,q}(\alpha,\beta))$ and $V = X \setminus U$ and $\overline{U}$, $\overline{V}$ for the symplectic completions. Since $(\overline{N}_-, \overline{J_N})$ is isomorphic to a punctured neighbourhood of the origin in $\mathbb{C}^2/G$, we can compactify $(\overline{V}, \overline{J})$ by adding in a point x to obtain an almost complex orbifold $(\widehat{X}, \widehat{J})$. Let $\widehat{N} := \overline{N}_- \cup \{x\}$; this is an embedded copy of the girdled orbifold $O_{p,q}(\alpha,\beta)$.

We will now resolve $\widehat{X}$ using a pavilion as in Definition 3.2.8. We will need to choose a Delzant pavilion such that the image of the neck N under the moment map is contained in $\mathrm{Pav}_{\rho,\lambda}\left(\Delta_{p,q}(\alpha,\beta)\right)$; in practice, we will fix the choice of pavilion first, and then choose N sufficiently close to the girdle that it satisfies this condition.

Definition 3.3.5 (Pavilion blow-up)**.** Given a Delzant pavilion $\boldsymbol{\rho}, \boldsymbol{\lambda}$ for $E_{p,q}(\alpha,\beta)$ and a symplectic embedding $\iota\colon E_{p,q}(\alpha,\beta) \hookrightarrow X$, we obtain a resolution $\widetilde{X}$ of $\widehat{X}$ by replacing $\widehat{N} \subset \widehat{X}$ with the girdled resolution $\widetilde{U}_\rho$ of $O_{p,q}(\alpha,\beta)$.

(i) Let $\widetilde{J}$ be the almost complex structure on $\widetilde{X}$ which coincides with J on V and agrees with the complex structure on $\widetilde{U}_\rho$.

(ii) Note that if $J \in \mathcal{J}_N$, the associated almost Kähler metric $\omega(-, J-)$ is flat on N. Let ψ_0 be the convex function on the moment-image of N which encodes the Kähler potential of $\omega|_N$. Given a convex function ψ on $\mathrm{Pav}_{\rho,\lambda}\left(\Delta_{p,q}(\alpha,\beta)\right)$ which coincides with ψ_0 on the shaded region (moment-image of N) and has the correct behaviour along the edges, we obtain a symplectic form $\widetilde{\omega}$ on $\widetilde{X}$ which coincides with ω on V and $\widetilde{\omega}_{\rho,\lambda,\psi}$ on $\widetilde{U}_\rho$; this is well-defined because ψ is chosen to coincide with ψ_0 on the moment-image of N.

We call the manifold $\widetilde{X} := V \cup \widetilde{U}_\rho$, equipped with the almost complex structure $\widetilde{J}$ and the symplectic form $\widetilde{\omega}$ the *pavilion blow-up* of X along the embedded pin-ellipsoid U. If we want to keep the dependence of $\widetilde{X}$ on the various choices explicit, we will write $\mathrm{Pav}^\iota_{\rho,\lambda}(X)$ for $\widetilde{X}$. We will sometimes abusively drop the subscript on $\widetilde{U}$.

Remark 3.3.6. A pavilion blow-up with minimal Delzant pavilion is just the usual *symplectic rational blow-up* established by Symington [Sym01; Sym98], which was elaborated upon by Khodorovskiy [Kho13b]. It is the inverse operation of the *rational blow-down*, introduced in the smooth category by Fintushel and Stern [FS97]; see also [Eva23, Section 9.2] for an exposition in the symplectic category. In the last decade the symplectic rational blow-up was used by many authors to study Lagrangian pin-wheels and pin-ellipsoids. See for example the work of Borman–Li–Wu [BLW14], Shevchishin–Smirnov [SS20] and Adaloglou [Ada26; AH25] on Lagrangian embeddings of $\mathbb{RP}^2$; the work of Khodorovskiy [Kho13a] and Evans–Smith [ES20] on bounding Wahl singularities; the work of Brendel–Schlenk on Lagrangian barriers [BS24] generalising Biran's work [Bir01]; and the work of Adaloglou [AH25] and Adaloglou–Hauber [AH25] on symplectic embeddings of $B_{n,1}$.

Figure 2.6 illustrates the spaces $X, \overline{V}, \widehat{X}, \widetilde{X}$, the subspaces $U, V, \widetilde{U}, N, \overline{N}_-, \widehat{N}$, and the relations between them all.

Lemma 3.3.7. *There are bijections between the following collections of objects:*

- *irreducible finite energy punctured $\overline{J}$-holomorphic curves C in $\overline{V}$,*
- *irreducible orbifold $\widehat{J}$-holomorphic curves $\widehat{C}$ in $\widehat{X}$,*

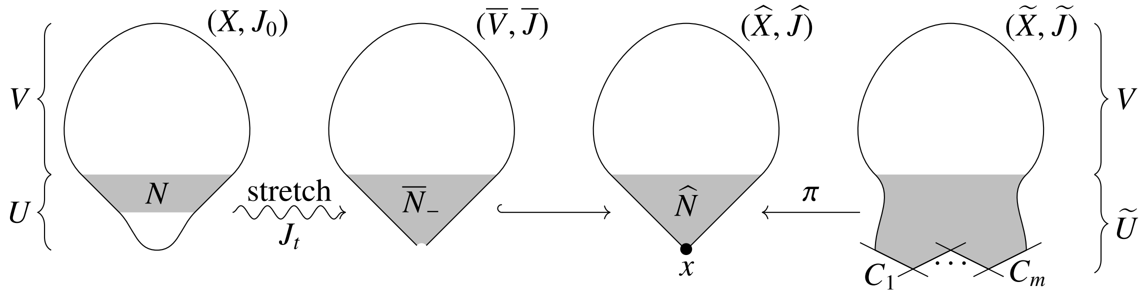

Figure 3.14: The spaces defined in Section 3.3.1

- *irreducible $\widetilde{J}$-holomorphic curves $\widetilde{C}$ in $\widetilde{X}$ other than $C_1, \ldots, C_m$.*

Proof. Given a finite energy punctured $\overline{J}$-holomorphic curve $C \subset \overline{V}$, we write $\widehat{C}$ for its closure in $\widehat{X}$. The inverse operation is puncturing (removing x). Given an orbifold curve $\widehat{C}$ in $\widehat{X}$, write $\widetilde{C}$ for the proper transform in $\widetilde{X}$. The inverse operation is projection along π, which sends any curve other than $C_1, \ldots, C_m$ to an orbifold curve. □

Lemma 3.3.8. *(a) Any irreducible $\widetilde{J}$-holomorphic curve in $\widetilde{X}$ which is not one of the C_i must enter the interior of V.*

(b) By choosing J generically on V, we can ensure that all curves which enter V are regular[9], so that the only irregular $\widetilde{J}$-holomorphic curves are amongst[10] the $C_1, \ldots, C_m$.

Proof. (a) If $\widetilde{C} \subset \widetilde{X}$ is a $\widetilde{J}$-curve other than $C_1, \ldots, C_m$ which does not enter the interior of V then the associated punctured curve C is contained in $\overline{N}_-$. Since the positive boundary of $\overline{N}_-$ is convex, this contradicts the maximum principle.

(b) This follows from [MS04, Proposition 3.2.1 and Remark 3.2.3]. □

Remark 3.3.9. Note that everything that was introduced and discussed within this section makes sense in the case where there are several disjoint pin-ellipsoids.

3.3.2 Topology of the pavilion blow-up

Given a symplectic embedding $\iota\colon E_{p,q}(\alpha,\beta) \hookrightarrow X$ and a choice of pavilion, we will need to understand how to compute intersection numbers of curves in the pavilion blow-up $\widetilde{X}$ by computing locally in $\widetilde{U}$ and $V = X \setminus \widetilde{U}$ separately, and adding. In what follows, if A is a class in integral homology, we will write $A_{\mathbb{Q}}$ for the corresponding class in rational homology. Denote the interface between $\widetilde{U}$ and V by Σ and recall that Σ is topologically a lens space $L(p^2, pq-1)$.

Lemma 3.3.10. *We have $H_1(U, \Sigma; \mathbb{Z}) = H_3(U, \Sigma; \mathbb{Z}) = 0$ and $H_2(U, \Sigma; \mathbb{Z}) \cong \mathbb{Z}_p$.*

Proof. This follows immediately from the long exact sequence of the pair (U, Σ):

[9] A curve is called *regular* if its linearised Cauchy–Riemann operator is surjective; see [MS04, Section 3.1]. For an embedded holomorphic sphere C in a 4-manifold, [MS04, Lemma 3.3.3] implies that regularity is equivalent to the condition $C^2 \geq -1$.

[10] If the resolution is minimal then *all* the C_i are irregular ($C_i^2 \leq -2$). Otherwise there are also some -1-spheres amongst them.

$$
\begin{tikzcd}
\cdots \arrow[r] & 0 = H_3(U;\mathbb{Z}) \arrow[r] & H_3(U,\Sigma;\mathbb{Z}) \arrow[dll] \\
0 = H_2(\Sigma;\mathbb{Z}) \arrow[r] & 0 = H_2(U;\mathbb{Z}) \arrow[r] & H_2(U,\Sigma;\mathbb{Z}) \arrow[dll] \\
\mathbb{Z}_{p^2} = H_1(\Sigma;\mathbb{Z}) \arrow[r, "\bmod p"'] & H_1(U;\mathbb{Z}) \cong \mathbb{Z}_p \arrow[r] & H_1(U,\Sigma;\mathbb{Z}) \arrow[r] & 0.
\end{tikzcd}
$$

□

Lemma 3.3.11. *Given a class $A \in H_2(\widetilde{X};\mathbb{Z})$, there exist rational numbers $a_1, \dots, a_m$ and a class $A_X \in H_2(X;\mathbb{Z})$ such that*

$$A_{\mathbb{Q}} = \frac{1}{p} A_X + \sum_{i=1}^{m} a_i [C_i]. \tag{3.3.15}$$

Here we are identifying $H_2(X;\mathbb{Q})$ as a subset of $H_2(\widetilde{X};\mathbb{Q})$ using the fact that $H_2(X;\mathbb{Q}) \cong H_2(V;\mathbb{Q})$. Moreover, the class pA_X is an integral class which can be represented by a cycle contained in V.

Proof. The Mayer–Vietoris sequence says:

$$0 = H_2(\Sigma;\mathbb{Z}) \to H_2(\widetilde{U};\mathbb{Z}) \oplus H_2(V;\mathbb{Z}) \to H_2(\widetilde{X};\mathbb{Z}) \to H_1(\Sigma;\mathbb{Z}) = \mathbb{Z}_{p^2} \to \cdots,$$

so if $A \in H_2(\widetilde{X};\mathbb{Z})$ then $p^2 A$ can be written as $A'_{\widetilde{U}} + A'_V$. Therefore $A_{\mathbb{Q}} = \frac{1}{p^2} A'_{\widetilde{U}} + \frac{1}{p^2} A'_V$. We will take $A_{\widetilde{U}} := \frac{1}{p^2} A'_{\widetilde{U}} = \sum_{i=1}^{m} a_i [C_i]$ for some rational numbers $a_i \in \frac{1}{p^2}\mathbb{Z}$. To understand A'_V, consider the exact sequence of the pair (X, V):

$$0 = H_3(U,\Sigma;\mathbb{Z}) \to H_2(V;\mathbb{Z}) \to H_2(X;\mathbb{Z}) \to H_2(X,V;\mathbb{Z}) = H_2(U,\Sigma;\mathbb{Z}) = \mathbb{Z}_p,$$

where we used excision to identify $H_2(X, V)$ with $H_2(U, \Sigma)$. This means that $A'_V = pA_X$ for some $A_X \in H_2(X;\mathbb{Z})$. □

Remark 3.3.12. We can calculate the numbers a_j as follows if we know how A intersects the curves C_j. Let $\alpha_j = A \cdot [C_j] \in \mathbb{Z}$ and let $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_m)^T$ and $\boldsymbol{a} = (a_1, \dots, a_m)$. Then

$$\alpha_j = \sum_{i=1}^{m} a_i C_i \cdot C_j, \text{ or } \boldsymbol{\alpha} = M\boldsymbol{a},$$

where M is the m-by-m matrix with entries $M_{ij} = C_i \cdot C_j$. Therefore $\boldsymbol{a} = M^{-1}\boldsymbol{\alpha}$. It will be useful to have formulas for the entries of M^{-1}.

Lemma 3.3.13. *Suppose that the Wahl chain for the minimal resolution of our singularity is*

$$[b_1, \dots, b_m], \qquad \text{with } b_i = -C_i^2,$$

and fix an index i. Consider the continued fractions $e_i/e_{i-1} = [b_{i-1}, \dots, b_1]$ and $f_i/f_{i+1} =$

$[b_{i+1}, \dots, b_m]$. *The entries of* M^{-1}_{ij} *are*

$$M^{-1}_{ij} = \begin{cases} -(e_i f_j)/p^2 \text{ if } i \leq j \\ -(e_j f_i)/p^2 \text{ if } i > j \end{cases} .$$

Proof. The continued fractions are computed as

$$\frac{e_{i+1}}{e_i} = b_i - \frac{1}{[b_{i-1}, \dots, b_1]} = \frac{b_i e_i - e_{i-1}}{e_i}, \qquad \frac{f_{i-1}}{f_i} = b_i - \frac{1}{[b_{i+1}, \dots, b_m]} = \frac{b_i f_i - f_{i+1}}{f_i},$$

so

$$e_{i+1} + e_{i-1} = b_i e_i, \qquad f_{i+1} + f_{i-1} = b_i f_i,$$

which are Fibonacci-like recursions with initial conditions

$$e_0 = 0, \quad e_1 = 1, \qquad f_0 = p^2, \quad f_1 = pq - 1.$$

For the matrix M^{-1} as postulated in the lemma, the off-diagonal entries $(M^{-1}M)_{ij}$ above the diagonal are[11]

$$-\frac{1}{p^2}\left(e_i f_{j-1} - e_i b_j f_j + e_i f_{j+1}\right) = 0.$$

Those below the diagonal are

$$-\frac{1}{p^2}\left(e_{j-1} f_i - e_j b_j f_i + e_{j+1} f_i\right) = 0.$$

Those on the diagonal are

$$-\frac{1}{p^2}\left(e_{i-1} f_i - b_i e_i f_i + e_i f_{i+1}\right) = \frac{1}{p^2}\left(e_i f_{i-1} - e_{i-1} f_i\right).$$

We can show that $e_i f_{i-1} - e_{i-1} f_i = p^2$ by induction: it holds when $i = 1$ and

$$e_{i+1} f_i - e_i f_{i+1} = \begin{vmatrix} e_{i+1} & f_{i+1} \\ e_i & f_i \end{vmatrix} = \begin{vmatrix} b_i e_i - e_{i-1} & b_i f_i - f_{i-1} \\ e_i & f_i \end{vmatrix} = \begin{vmatrix} e_i & f_i \\ e_{i-1} & f_{i-1} \end{vmatrix}$$

by subtracting b_i times row 2 from row 1 and swapping the rows.

This shows that $M^{-1}M$ is the identity matrix, as required. □

Remark 3.3.14. We will only be concerned with the case when $X = \mathbb{CP}^2$. In this case, we write $\mathcal{E}$ for the rational class $\frac{1}{p}H \in H_2(X; \mathbb{Q})$, where H is the class of the line, so $p^2\mathcal{E}$ can be

[11] The entries in the first and last column have only two terms, but the formula still holds true if we decide to set $e_0 = f_{m+1} = 0$.

represented by a cycle in V. We can write any class $A \in H_2(\widetilde{X}; \mathbb{Q})$ in the form

$$A = a_0 \mathcal{E} + \sum_{i=1}^{m} a_i [C_i]. \tag{3.3.16}$$

The square of A is then given by the formula

$$A^2 = \frac{a_0^2}{p^2} + \boldsymbol{a}^T M \boldsymbol{a}. \tag{3.3.17}$$

Remark 3.3.15. In the case of multiple embeddings $\iota\colon \bigsqcup_{j=1}^{\ell} E_{p_j,q_j}(\alpha_j, \beta_j) \hookrightarrow X$ we define $\Delta := \prod_{j=1}^{\ell} p_j$. Then, by performing a pavilion blow-up on all the ℓ embedded pin-ellipsoids, we obtain spaces $\widetilde{X}$, $\widetilde{V}$ and V, and Equation (3.3.15) becomes

$$A_{\mathbb{Q}} = \frac{1}{\Delta} A_X + \sum_{j=1}^{\ell} \sum_{i=1}^{m_j} a_{j,i} [C_{j,i}].$$

Similarly, in the case that $X = \mathbb{CP}^2$, Equations (3.3.16) and (3.3.17) become:

$$A = a_0 \mathcal{E} + \sum_{j=1}^{\ell} \sum_{i=1}^{m_j} a_{j,i} [C_{j,i}] \quad \text{and} \quad A^2 = \frac{a_0^2}{\Delta^2} + \sum_{j=1}^{\ell} \boldsymbol{a}_j^T M_j \boldsymbol{a}_j. \tag{3.3.18}$$

3.4 Ruled 4-manifolds

The results of this section are well-known to experts; we thank Weiyi Zhang for informing us that the methods of his paper [Zha21] for irrational ruled surfaces can be adapted to the rational case.

3.4.1 Regulations and rulings

Definition 3.4.1. Given a stable J-holomorphic curve S, write $S = \sum_{i \in I} m_i S_i$ where S_i are the non-constant irreducible components of S, indexed by a set I, and m_i their covering multiplicities. The *dual graph* of S is the labelled graph Γ_S:

- whose vertex set is I,
- whose vertex i is labelled by the self-intersection number S_i^2,
- and where there are $S_i \cdot S_j$ edges between the vertices $i \neq j$.

Definition 3.4.2. Let (Y, J) be a tamed almost complex 4-manifold. Given a homology class $A \in H_2(Y; \mathbb{Z})$, let $\overline{\mathcal{M}}_{0,n}(A, J)$ denote the moduli space of J-holomorphic stable maps in the class A having n marked points. We say that Y *admits a J-holomorphic regulation*[12] *in the class* A if all the following conditions hold:

(a) $A^2 = 0$ and $c_1(Y) \cdot A = 2$

[12] This would ordinarily be called a *ruling*, but ruling can also mean one of the curves in the regulation and this is how we will use it. Regulation is to *regula* (Latin for rule) as a triangulation is to *triangle*.

(b) the evaluation map $\mathrm{ev}\colon \overline{\mathcal{M}}_{0,1}(A, J) \to Y$ has degree 1.

The curves in the moduli space $\overline{\mathcal{M}}_{0,0}(A, J)$ are called *rulings*. Smooth curves corresponding to points of the (possibly empty) subspace $\mathcal{M}_{0,0}(A, J)$ we call *smooth rulings*; by the adjunction formula, these are embedded spheres. Curves corresponding to points of $\partial\overline{\mathcal{M}}_{0,0}(A, J) = \overline{\mathcal{M}}_{0,0}(A, J) \setminus \mathcal{M}_{0,0}(A, J)$ we call *broken rulings*. We call the regulation of Y *non-degenerate* if there is at least one smooth ruling.

Example 3.4.3. (1) The simplest example of a ruled surface is $\mathbb{CP}^1 \times \mathbb{CP}^1$: this admits two non-degenerate regulations, one in the class $[\mathbb{CP}^1 \times \{\mathfrak{p}\}]$ and one in the class $[\{\mathfrak{p}\} \times \mathbb{CP}^1]$. This ruled surface arises as a quadric surface in $\mathbb{CP}^3$ and the rulings in both regulations are lines.

(2) If we degenerate the quadric to a nodal quadric then there is still a regulation (but now only one) consisting of lines $\{F_t \mid t \in \mathbb{CP}^1\}$ passing through the node: the two regulations of the smooth surface degenerate onto this regulation. If we resolve the node (introducing a -2-curve σ_∞) then the result is the Hirzebruch surface $\mathbb{F}_2$. The proper transforms $\tilde{F}_t$ of the rulings of the nodal quadric yield the standard non-degenerate regulation of $\mathbb{F}_2$; there is another, degenerate, regulation whose rulings are $\tilde{F}_t + \sigma_\infty$.

Theorem 3.4.4. *If Y contains a smoothly embedded J-holomorphic curve C with $[C]^2 = 0$ then C is part of a non-degenerate J-holomorphic regulation in the class $[C]$.*

Proof. We get that $c_1(Y) \cdot [C] = 2$ by adjunction. Since $c_1(Y) \cdot [C] > 0$, automatic transversality guarantees that the moduli space $\mathcal{M}_{0,1}([C], J)$ is a smooth manifold of dimension 4 and the evaluation map $\mathrm{ev}\colon \mathcal{M}_{0,1}([C], J) \to Y$ is a pseudocycle. The degree of this pseudocycle must be 1 because there is at most one smooth J-holomorphic sphere in the class $[C]$ passing through any point (by positivity of intersections because $[C]^2 = 0$) and for all points in a neighbourhood of C there is precisely one (namely C or its small deformations, which exist by automatic transversality). In particular, $\mathrm{ev}\colon \overline{\mathcal{M}}_{0,1}([C], J) \to Y$ has degree 1. □

3.4.2 Broken rulings

Proposition 3.4.5. *Let (Y, J) be a tamed almost complex manifold admitting a non-degenerate regulation in the class $A \in H_2(Y; \mathbb{Z})$. Let C be a smooth ruling, and let $S = \sum m_i S_i$ be a broken ruling of the regulation.*

(a) The curves S_i are embedded rational curves disjoint from C.

(b) The dual graph Γ_S is a tree.

(c) The irreducible components S_i have $[S_i]^2 < 0$.

(d) At least one irreducible component is a sphere of square -1.

Proof. (a) and (b). The existence of a smooth ruling of the regulation implies that the class A is J-nef in the sense of Li and Zhang [LZ15, Lemma 2.8]. This implies that the dual graph is a tree and all the components are embedded spheres [LZ15, Theorem 1.5]. Note that the components S_i are disjoint from any smooth ruling C, since $0 = A^2 = C \cdot \sum m_i S_i = \sum m_i (C \cdot S_i)$ and $C \cdot S_i \geq 0$ for all i, so $C \cdot S_i = 0$ for all components S_i.

(c) To see why $S_i^2 < 0$ for all i, we adapt the proof of [MO15, Lemma 3.3.1]. Suppose that S_i is a component which appears with multiplicity m_i in S. Since the curves $\sum_{j\neq i} m_j S_j$ and $C + m_i S_i$ are geometrically distinct, they intersect non-negatively; in fact, since the dual graph is connected with at least two vertices, S_i must intersect at least one S_j with $j \neq i$, so they have positive intersection. Therefore

$$0 < \sum_{j\neq i} m_j S_j \cdot (A + m_i S_i) = (A - m_i S_i) \cdot (A + m_i S_i) = -m_i^2 S_i^2,$$

which implies that $S_i^2 < 0$.

(d) Since $c_1(A) = 2$, there must be at least one S_i with $c_1(S_i) > 0$; since $c_1(S_i) = S_i^2 + 2$ by adjunction and $S_i^2 < 0$, this must be a -1-sphere. □

Remark 3.4.6. We thank Weiyi Zhang for pointing out that this proposition can also be deduced from [LZ15, Lemma 4.10 and Corollary 4.11].

Lemma 3.4.7. *In a non-degenerate regulation, two broken rulings are either identical or disjoint.*

Proof. Suppose that S and S' are broken rulings of a non-degenerate regulation. Since $[S] = [S']$ and $[S] \cdot [S'] = [S]^2 = 0$, the only way for S and S' to intersect is if they share a common component. Suppose that they are not identical (but share a common component). Take a component T of S that is adjacent to a shared component, but is not contained in S'. Then $T \cdot S' > 0$. But non-degeneracy means that S' is homologous to a smooth ruling C, and $T \cdot C = 0$ by Proposition 3.4.5(a), so $T \cdot S' = 0$. This yields a contradiction. □

Remark 3.4.8. The existence of a smooth ruling is necessary (see Example 3.4.3(2)). We thank Weiyi Zhang for pointing out [LLW22, Theorem 1.7] which guarantees the existence of regulations with at least one smooth ruling on any rational symplectic 4-manifold other than $\mathbb{CP}^2$.

In the sequel, we look at spheres without marked point. Note that by Gromov compactness, there is a finite number of broken rulings of a non-degenerate regulation.

Definition 3.4.9. Given a broken ruling $S \in \overline{\mathcal{M}}_{0,0}(A, J)$, pick a contractible neighbourhood of S in $\overline{\mathcal{M}}_{0,0}(A, J)$ which contains no other broken rulings and let $N \subset Y$ be the union of all rulings from this neighbourhood in the moduli space. We call such an N a *tubular neighbourhood* of S.

Lemma 3.4.10. *Let S be a broken ruling of a non-degenerate regulation in a class A and let N be a tubular neighbourhood of S. Let J' be a new tame almost complex structure which coincides with J outside N and for which S is still J'-holomorphic. Then there is still a non-degenerate J'-holomorphic regulation in the class A and it has no new broken rulings.*

Proof. Since the smooth rulings outside N are unaffected by the modification of the almost complex structure, they are still part of a non-degenerate J'-holomorphic regulation and so is S (because it is still part of the moduli space $\overline{\mathcal{M}}_{0,0}(A, J')$). If there are any new broken rulings then they must be disjoint from both S and from the rulings outside N, and so they must be contained in $N \setminus S$. But $N \setminus S$ is a sphere bundle over a non-compact surface (the punctured neighbourhood of $S \in \overline{\mathcal{M}}_{0,0}(A, J')$) and hence $H_2(N \setminus S; \mathbb{Z}) = \mathbb{Z}$, generated by the class A. Therefore there are no homology classes into which the rulings can break (since A has minimal area amongst multiples of A). □

Proposition 3.4.11. *Let (Y, J) be an almost complex 4-manifold tamed by a symplectic form ω. Suppose Y admits a non-degenerate J-holomorphic regulation in the class A. Let $U \subset Y$ be a subset on which J is integrable. We can find another almost complex structure J' tamed by ω satisfying:*

(i) $J'|_U = J|_U$.

(ii) Y still admits a J'-holomorphic regulation, and has exactly the same broken rulings (identical as point sets) as the original regulation.

(iii) J' is integrable on a neighbourhood of all the broken rulings.

Proof. Sikorav [Sik97, Theorem 3] shows that one can make J' integrable on a neighbourhood of any given holomorphic curve; Chen and Zhang [CZ24, Appendix A] modified this argument to produce a tamed J'. Since we are modifying the almost complex structure only on a tubular neighbourhood of the broken rulings, Lemma 3.4.10 guarantees that we do not produce any new broken rulings. □

Corollary 3.4.12. *The dual graph of a broken ruling in a non-degenerate regulation is obtained from the graph* $\overset{0}{\bullet}$ *by a sequence of the following moves (combinatorial blow-ups):*

(1) Add a new vertex labelled -1 connected by a single edge to an existing vertex i, whose label drops by 1.

(2) Add a new vertex labelled -1 at the middle of an edge connecting two vertices i and j, whose labels both drop by 1.

Proof. We know from Proposition 3.4.5(d) that any broken ruling S contains an embedded -1-sphere E. Once we have made the almost complex structure integrable near S, we can blow down E holomorphically to obtain a new almost complex manifold with a non-degenerate regulation. Iterate this procedure until the ruling becomes smooth; then its dual graph is a single vertex labelled 0. Inverting this process, at each step we blow up a point which either lives at a smooth point of the ruling or a node of the ruling, which changes the dual graph by a combinatorial blow-up of type 1 or 2 respectively. □

The possible dual graphs with at most three vertices are shown in Figure 3.15.

Remark 3.4.13. Note that by [Lis08, Lemma 2.2] chain-shaped dual graphs obtained from iterated combinatorial blow-up are precisely the *zero continued fractions*; for example

$$2 - \frac{1}{1 - \frac{1}{2}} = 0.$$

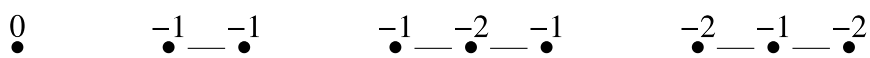

Figure 3.15: The dual graphs for broken rulings with at most three irreducible components.

Corollary 3.4.14. *If S is a broken ruling of a non-degenerate regulation and S has precisely one -1-sphere amongst its irreducible components, then the same is true of all its blow-downs until we reach* $\overset{-1}{\bullet}\!-\!\!-\overset{-1}{\bullet}$ *. Blowing back up, the next dual graph must be* $\overset{-2}{\bullet}\!-\!\!-\overset{-1}{\bullet}\!-\!\!-\overset{-2}{\bullet}$*, and all subsequent ones are obtained from this by combinatorially blowing-up a point on the -1-sphere; for example the next step might be* $\overset{-2}{\bullet}\!-\!\!-\overset{-2}{\bullet}\!-\!\!-\overset{-1}{\bullet}\!-\!\!-\overset{-3}{\bullet}$ *or* $\overset{-2}{\bullet}\!-\!\!-\underset{-2}{\overset{-1\,\bullet\atop |}{\bullet}}\!-\!\!-\overset{-2}{\bullet}$.

3.5 Pin-ellipsoids in the complex projective plane

3.5.1 Goal and strategy

In what follows, we specialise to the case where $X = \mathbb{CP}^2$ equipped with the Fubini–Study form giving a line area 1. In this preliminary section, we will outline the key technical results needed to prove Theorems 3.1.8, 3.1.13, 3.1.18 and Corollary 3.1.21.

Assume that we have a symplectic embedding $\iota\colon E_{p,q}(\alpha,\beta) \hookrightarrow X$ for some p, q, α, β, write U for the image of this embedding and $V = X \setminus U$. Choose a Delzant pavilion $\boldsymbol{\rho}, \boldsymbol{\lambda}$ and equip the pavilion blow-up $\widetilde{X} := \mathrm{Pav}^{\iota}_{\rho,\lambda}(X)$ with an almost complex structure $\widetilde{J}$ and symplectic form $\widetilde{\omega}$ as in Definition 3.3.5. Let $C_1, \ldots, C_m$ be the exceptional curves of the pavilion.

Theorem 3.5.1. *Assume that $\widetilde{J}$ is chosen generically on V and that the Delzant pavilion is minimal. Then there exists a smoothly embedded $\widetilde{J}$-holomorphic sphere $\widetilde{C} \subset \widetilde{X}$ with*

$$\widetilde{C}^2 = 0 \quad \text{and} \quad \widetilde{C} \cdot C_j = \begin{cases} 0 \text{ unless } j = i_{\mathrm{culet}}, \\ 1 \text{ if } j = i_{\mathrm{culet}}. \end{cases}$$

Remark 3.5.2. Note that Theorem 3.5.1 is obvious for a visible embedding ι^{vis}: one can then find $\widetilde{C}$ over a segment perpendicular to the line of $C_{i_{\mathrm{culet}}}$ and to the edge $\mathfrak{e}_1$, cf. Figure 3.12. Theorem 3.5.1 follows easily from the results of Evans and Smith [ES18, Theorem 4.15], as we will explain in Appendix 3.6, but we will give an alternative proof below which avoids working with orbifold holomorphic curves. As a consequence we will give a more direct, orbifold-free proof of the fact that p is a Markov number and q is one of its companion numbers (though, upon comparison with the orbifold proof from [ES18], it is clear that the two proofs are doing the same thing in different language).

Remark 3.5.3. In the case of two disjointly embedded pin-ellipsoids, we perform a pavilion blow-up as above at both of them to obtain a space $\widetilde{X}$. By similar arguments, we obtain the existence of a holomorphic sphere $\widetilde{C}$ satisfying:

- $\widetilde{C}^2 = -1$,
- $\widetilde{C} = c_0\mathcal{E} + \sum c_{1,i}C_{1,i} + \sum c_{2,i}C_{2,i}$ where c_0 is the smaller of the two solutions to
$$c_0^2 + p_1^2 + p_2^2 = 3c_0p_1p_2.$$
- $\widetilde{C}$ intersects each Wahl chain in the pavilion blow-up once transversely at a point belonging to one of the curves at the end of the chain.

This curve will be used in Section 3.5.5 to prove Theorem 3.1.18 and Corollary 3.1.21. Again, this curve is visible (living over an edge) if both pin-ellipsoids are visible.

Theorem 3.5.1 implies that $\widetilde{X}$ admits a regulation. The next result gives some properties of this regulation.

Theorem 3.5.4. *Assume that $\widetilde{J}$ is chosen generically on V and that the Delzant pavilion is minimal. Then the curve $\widetilde{C}$ from Theorem 3.5.1 is part of a $\widetilde{J}$-holomorphic regulation with the following properties:*

(a) The culet curve $C_{i_{\text{culet}}}$ is a section.

(b) (i) If $C^2_{i_{\text{culet}}} = -4$ then there are no broken rulings.

(ii) If $C^2_{i_{\text{culet}}} = -7$ then $i_{\text{culet}} = 1$ (or m) and there is one broken ruling consisting of $C_2 \cup \cdots \cup C_m \cup E$ (respectively $C_1 \cup \cdots \cup C_{m-1} \cup E$) where E is an embedded -1-sphere intersecting precisely one of the curves C_i in the broken ruling once transversely.

(iii) If $C^2_{i_{\text{culet}}} = -10$ then there are two broken rulings:

$$C_1 \cup \cdots \cup C_{i_{\text{culet}}-1} \cup E_1 \quad \text{and} \quad C_{i_{\text{culet}}+1} \cup \cdots \cup C_m \cup E_2$$

where E_1 and E_2 are embedded -1-spheres which intersect precisely one of the curves C_i in the broken ruling once transversely.

In all cases, the curves C_i which hit the -1-curves E, E_1 and E_2 are completely determined by the numbers p and q.

Corollary 3.5.5. *If the pavilion is not minimal then we will still find a regulation with at most two broken rulings. The possible structures of the broken rulings is described in the proof and illustrated in Figure 3.18.*

Proof. Recall from Remark 3.2.5 that a non-minimal pavilion is obtained from a minimal one by a sequence of toric blow-ups, corresponding to adding further rays to the fan $\boldsymbol{\rho}$. Since these blow-ups occur at intersection points between the curves $C_1, \ldots, C_m$, they are disjoint from a general ruling of the regulation, so the proper transform of a general ruling is still a holomorphic curve of square zero, so the result is still ruled and the culet curve[13] is still a section.

The broken rulings are given by taking the total transforms of either an existing broken ruling (we call this Type I) or of a smooth ruling (which we call Type II). Since the blow-ups are toric, the Type II broken rulings are chain-shaped and their dual graphs correspond to zero continued fractions as in Remark 3.4.13. The dual graphs for Type I broken rulings still have the form of a tree comprising a chain of spheres from the pavilion together with a -1-sphere attached somewhere along the chain, but now the resolution curves from the pavilion can be -1-curves. See Figure 3.18 for an example of a non-minimal pavilion blow-up with one Type I and one Type II broken ruling. □

Definition 3.5.6. We will define a tree $\mathcal{T}$ of symplectic curves in $\widetilde{X}$ as follows.

[13] Recall from Remark 3.2.14 that the culet curve for a non-minimal resolution is the proper transform of the culet curve from the minimal resolution.

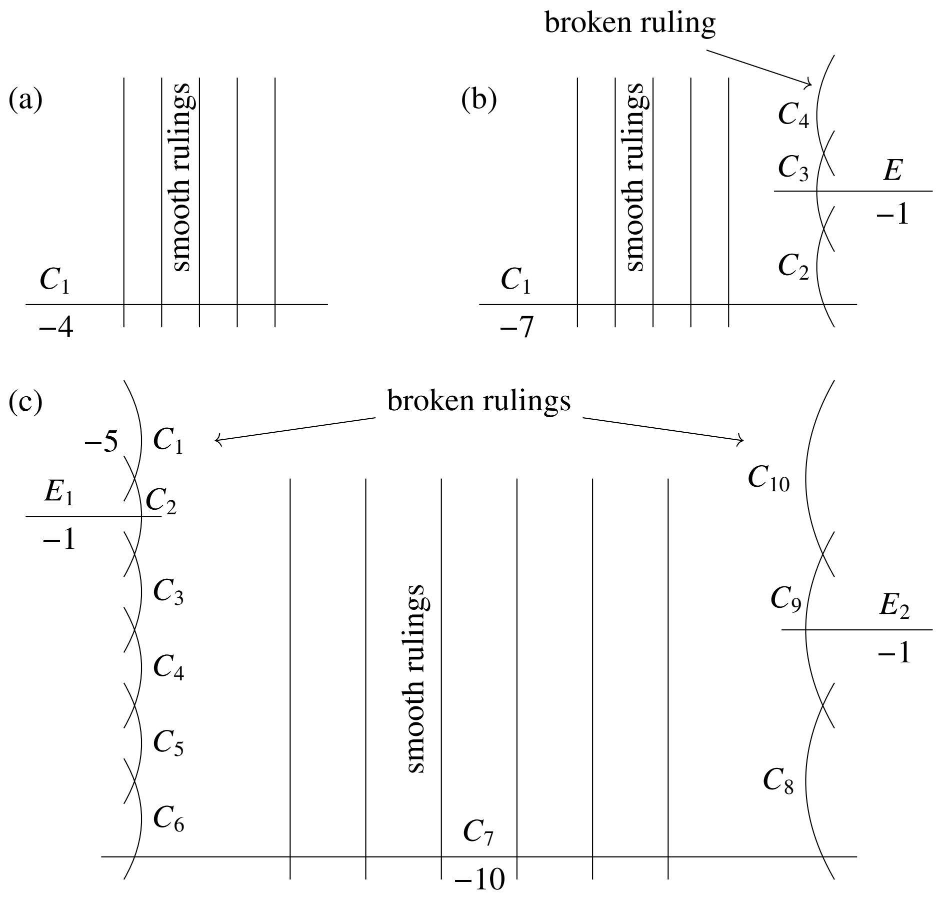

Figure 3.16: Examples illustrating the structure of the regulation described in Theorem 3.5.4 for the cases (a) weight 4 ($p = 2, q = 1$), (b) weight 7 ($p = 5, q = 1$) and (c) weight 10 ($p = 29, q = 7$). Curves are labelled with their name and self-intersection, except vertical straight lines (which represent smooth rulings with square 0) and -2-spheres whose self-intersection is omitted.

- If there are two broken rulings then $\mathcal{T}$ consists of $C_{i_{\text{culet}}}$ and the broken rulings.
- If there is precisely one broken ruling then we choose an unbroken ruling F and take $\mathcal{T}$ to be the union of F, $C_{i_{\text{culet}}}$ and the broken ruling.
- If there are no broken rulings (which can happen only if $w = 4$ and the pavilion is minimal) then we choose two unbroken rulings F_1 and F_2 and take $\mathcal{T}$ to be the union $F_1 \cup F_2 \cup C_1$.

Denote this splitting of $\mathcal{T}$ into three parts by $\mathcal{T} = \mathcal{F}_1 \cup C_{i_{\text{culet}}} \cup \mathcal{F}_2$.

Corollary 3.5.7. *Let $\mathcal{T}$ be the tree of curves from Definition 3.5.6. The complement $\widetilde{X} \setminus \mathcal{T}$ is a minimal symplectic manifold.*

Proof. The space $\widetilde{X}$ is obtained from a Hirzebruch surface by a sequence of blow-ups, since by Corollary 3.4.14 we can consecutively blow down the -1-curves contained in the broken rulings until we arrive at a situation in which there are no broken rulings. Then, $\widetilde{X} \setminus \mathcal{T}$ is identified with

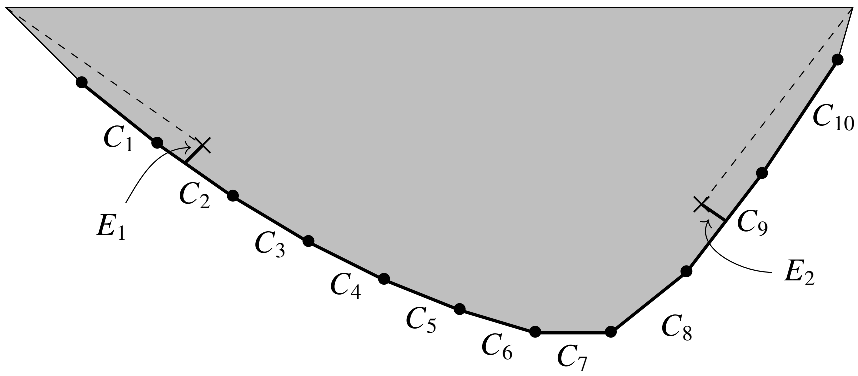

Figure 3.17: An almost toric base diagram for a minimal pavilion blow-up of a visible $E_{29,7}(\alpha,\beta) \subset \mathbb{CP}^2$. This is almost toric representation of Figure 3.16 (c).

a subset in the complement of a section in a Hirzebruch surface. Since we can always arrange the consecutive blow-downs to happen away from the culet curve, $\widetilde{X} \setminus \mathcal{T}$ lives in the normal neighbourhood of a symplectic sphere with positive self-intersection number. Therefore, $\widetilde{X} \setminus \mathcal{T}$ is minimal. □

Remark 3.5.8. It is possible to compute the minimal model that we mentioned in the proof of Corollary 3.5.7: it is diffeomorphic to the Hirzebruch surface $\mathbb{F}_w$ whose negative section has square $-w$. This is diffeomorphic to either $S^2 \times S^2$ (if w is even) or $\mathbb{CP}^2 \# \overline{\mathbb{CP}}^2$ (if w is odd).

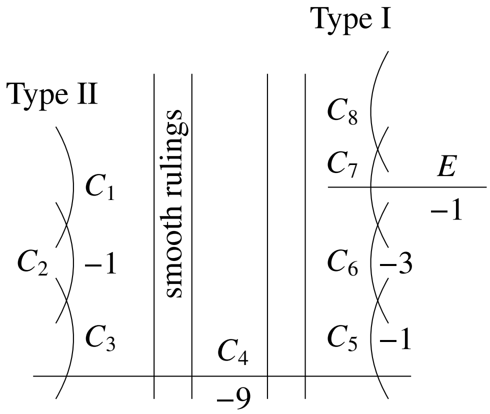

Figure 3.18: An example illustrating the structure of the regulation for a non-minimal pavilion blow-up, as described in Corollary 3.5.5. This example is obtained from Figure 3.16(b) by blowing up four times: once where the original broken ruling intersects the culet curve (yielding the Type I broken ruling) and three more times on a general smooth ruling (yielding the Type II broken ruling). Curves are labelled with their name and self-intersection, except vertical straight lines (which represent smooth rulings with square 0) and -2-spheres whose self-intersection is omitted.

In the next few sections, we will show how Corollary 3.5.5 implies the Isotopy and Staircase Theorems 3.1.8 and 3.1.13 from the introduction. In Section 3.5.4, we will prove Theorem 3.5.1.

In Section 3.5.5, we will modify the proof to handle multiple pin-balls and also give a proof of the Two Pin-Ball Theorem. Finally, in Section 3.5.6 we will prove Theorem 3.5.4.

3.5.2 Proof of the Isotopy Theorem (Theorem 3.1.8)

In this section, we explain how Corollary 3.5.5 solves the isotopy problem for pin-ellipsoids in $\mathbb{CP}^2$. We begin by stating a further corollary of Corollary 3.5.5.

Corollary 3.5.9. *Let $\iota\colon E_{p,q}(\alpha,\beta) \hookrightarrow \mathbb{CP}^2$ be a symplectic embedding, let $\boldsymbol{\rho}, \boldsymbol{\lambda}$ be a (possibly non-minimal) choice of pavilion and let $\widetilde{X}$ be the pavilion blow-up along ι. Let $\mathcal{T}$ be the tree of symplectic spheres from Definition 3.5.6. Then the complement $\widetilde{X} \setminus \mathcal{T}$ is diffeomorphic to a disc bundle over the annulus and the self-intersections and symplectic areas of the spheres in $\mathcal{T}$ depend only on p, q and the choice of pavilion.*

Proof. Since $\mathcal{T}$ consists of a section and two rulings of the regulation on $\widetilde{X}$, including all broken rulings, the restriction of the regulation to $\widetilde{X} \setminus \mathcal{T}$ exhibits this complement as a disc bundle over the annulus.

The tree comprises the linear chain $C_1 \cup \ldots \cup C_m$ together with some additional curves. These additional curves are of the following types:

- a proper transform of a smooth ruling in the minimal model, coming from a broken ruling of Type II;
- a -1-curve attached to one of the curves C_i in the chain, coming from a broken ruling of Type I;
- a smooth ruling, if there are fewer than two broken rulings.

The dual graph of the linear chain is determined by the choice of pavilion, as are the positions in the dual graph where the additional curves attach. The symplectic areas of the curves C_i are also determined by the pavilion. Writing the additional curves in the form given by Equation (3.3.16), we can find the coefficients by computing their intersection numbers with the curves C_i and with a smooth ruling; then their symplectic areas can be computed using the fact that $\int_{\mathcal{E}} \widetilde{\omega} = 1/p^2$ and the knowledge of the symplectic areas of the C_i coming from the pavilion. □

For the rest of this subsection we assume that $\iota, \iota'\colon E_{p,q}(\alpha,\beta) \hookrightarrow \mathbb{CP}^2$ are symplectic embeddings. Let $\boldsymbol{\rho}, \boldsymbol{\lambda}$ be a choice of pavilion for $E_{p,q}(\alpha,\beta)$. As before we denote the symplectic manifolds obtained by performing the pavilion blow-up along these embeddings by $\widetilde{X} = \mathrm{Pav}^{\iota}_{\boldsymbol{\rho},\boldsymbol{\lambda}}(\mathbb{CP}^2)$ and $\widetilde{X}' = \mathrm{Pav}^{\iota'}_{\boldsymbol{\rho},\boldsymbol{\lambda}}(\mathbb{CP}^2)$.

Corollary 3.5.10. *There exists a symplectomorphism*

$$\Psi\colon \big((\widetilde{X}, \omega), \mathcal{T}\big) \to \big((\widetilde{X}', \omega'), \mathcal{T}'\big)$$

which carries the tree $\mathcal{T}$ of symplectic spheres in $\widetilde{X}$ coming from Corollary 3.5.9 to the tree $\mathcal{T}'$ of symplectic spheres in $\widetilde{X}'$.

Proof. Since the symplectomorphism type of the trees $\mathcal{T}, \mathcal{T}'$ of symplectic spheres depends only on the pavilion, the symplectic neighbourhood theorem gives symplectic embeddings $\psi\colon \nu \hookrightarrow \widetilde{X}$ and $\psi'\colon \nu \hookrightarrow \widetilde{X}'$ of a plumbing ν of symplectic disk bundles over symplectic spheres into both pavilion blow-ups $\widetilde{X}, \widetilde{X}'$. Our goal is to extend the symplectomorphism

$\psi' \circ \psi^{-1}\colon \psi(\nu) \to \psi'(\nu)$ to a symplectomorphism $\Psi\colon (\widetilde{X}, \omega) \to (\widetilde{X}', \omega')$. Recall from Definition 3.5.6 that the trees split as a union $\mathcal{T} = \mathcal{F}_1 \cup C_{i_{\mathrm{culet}}} \cup \mathcal{F}_2$ and $\mathcal{T}' = \mathcal{F}_1' \cup C'_{i_{\mathrm{culet}}} \cup \mathcal{F}_2'$ of two rulings and the culet curve (a section). By the proof of Corollary 3.5.5, not only the two trees but also these two decompositions are homeomorphic. In particular, $\psi' \circ \psi^{-1}$ takes $C := C_{i_{\mathrm{culet}}}$ to $C' := C'_{i_{\mathrm{culet}}}$.

Step 1: Choice of regulations. Recall that the regulation of $\widetilde{X}$ was constructed with respect to an almost complex structure J. Denote the restriction of this almost complex structure to $\psi(\nu)$ by j. Extend $(\psi' \circ \psi^{-1})_* j$ to an almost complex structure J' on $\widetilde{X}'$ and consider the regulation of $\widetilde{X}'$ associated to J'. Note that $\psi' \circ \psi^{-1}$ takes rulings to rulings in a neighbourhood of $\mathcal{F}_1$ and $\mathcal{F}_2$.

Step 2: Splitting $\widetilde{X}$ by splitting C. Since C is a section, the regulation determines a continuous projection map $\pi\colon \widetilde{X} \to C$ that is smooth away from the rulings $\mathcal{F}_1$, $\mathcal{F}_2$. Let p_1, p_2 be the images under π of $\mathcal{F}_1$ and $\mathcal{F}_2$. Choose small disjoint embedded closed curves γ_1, γ_2 in C around p_1, p_2. They decompose C into a closed disc D_1 around p_1, an annulus A, and a closed disc D_2 around p_2. Since π is continuous, we can choose the γ_i so close to p_i that both $\pi^{-1}(D_i)$ are contained in $\psi(\nu)$. In this way we obtain a splitting

$$\widetilde{X} = \pi^{-1}(D_1) \cup \pi^{-1}(A) \cup \pi^{-1}(D_2).$$

Choosing γ_i' to be the curves $(\psi' \circ \psi^{-1})(\gamma_i)$ in the section C' of $\widetilde{X}'$, we obtain analogous splittings $C' = D_1' \cup A' \cup D_2'$ and

$$\widetilde{X}' = (\pi')^{-1}(D_1') \cup (\pi')^{-1}(A) \cup (\pi')^{-1}(D_2).$$

Note that $\psi' \circ \psi^{-1}$ restricts to a fibre-preserving map $\pi^{-1}(D_i) \to \pi^{-1}(D_i')$ and in particular to a fibre-preserving map $\pi^{-1}(\gamma_i) \to (\pi')^{-1}(\gamma_i')$.

Step 3: A good chart for ω near $\pi^{-1}(\gamma_i)$. Fix $i \in \{1, 2\}$. The S^2-bundle $\pi\colon \pi^{-1}(\gamma_i) \to \gamma_i$ is the trivial S^2-bundle over the circle γ_i, because the monodromy around γ_i is symplectic and hence preserves the orientation of the fibre. The same argument proves the first of the following two assertions.

(i) The monodromy of the bundle $\gamma_i \times S^2 \to \gamma_i$ in $\pi_0(\mathrm{Diff}(S^2))$ is trivial.

(ii) The symplectic normal bundle $\mathcal{N} \to \gamma_i \times S^2$ with fibre $\{\theta\} \times (T_q S^2)^{\perp\omega}$ over $(\theta, q) \in \gamma_i \times S^2$ is trivial.

Proof of (ii): If $\mathcal{F}_i$ is a broken ruling we successively blow down the exceptional divisors of the broken ruling $\pi^{-1}(p_i)$ as in § 3.4.2. The blow-downs are supported away from the wall $\pi^{-1}(\gamma_i)$. Furthermore, the broken ruling $\pi^{-1}(p_i)$ becomes one smooth fibre, and so the regulation over $\pi^{-1}(D_i)$ becomes a trivial symplectic sphere bundle over a disc. This implies assertion (ii).

The following lemma gives a one-sided normal form neighbourhood for the above loops of symplectic spheres.

Lemma 3.5.11. *Consider the closed disc D with symplectic form $\omega_D = dx \wedge dy$, the sphere S^2 with an area form ω_{S^2} whose total area agrees with the ω-area of a smooth fibre of π, and take the product symplectic form $\omega_D \oplus \omega_{S^2}$ on $D \times S^2$. Then there exists a neighbourhood U of*

$\partial D \times S^2$ in $D \times S^2$ and a diffeomorphism α from U to a neighbourhood of $\pi^{-1}(\gamma_i)$ in $\pi^{-1}(D_i)$ that pulls back ω to $\omega_D \oplus \omega_{S^2}$.

Proof. The claim would be well-known if we only had to deal with one sphere $\{\theta\} \times S^2$ instead of $\gamma_i \times S^2$. In our case, we use assertions (i) and (ii) above to construct a smooth embedding $\underline{\alpha}\colon \partial D \times S^2 \to X$ onto $\pi^{-1}(\gamma_i)$ such that $\underline{\alpha}^*\omega = (\omega_D \oplus \omega_{S^2})|_{T(\partial D \times S^2)}$. The neighbourhood theorem for hypersurfaces from [Got82] (or also Exercise 3.4.17 in [MS16]) now implies that $\underline{\alpha}$ extends to the claimed diffeomorphism α on a neighbourhood U. □

Step 4: Extending $\psi' \circ \psi^{-1}$. We first extend the symplectomorphism

$$\psi' \circ \psi^{-1}\colon \pi^{-1}(D_1) \cup \nu(C) \cup \pi^{-1}(D_2) \;\to\; \pi^{-1}(D_1') \cup \nu(C') \cup \pi^{-1}(D_2')$$

to a diffeomorphism $\Psi_0\colon \widetilde{X} \to \widetilde{X}'$. (This is possible since $\widetilde{X} \setminus \mathcal{T}$ and $\widetilde{X}' \setminus \mathcal{T}'$ are both disc bundles over an annulus by Corollary 3.5.9. We could assume that Ψ_0 preserves the fibres of π and π', but this is not needed in the sequel.) We then have the two symplectic forms ω and $\Psi_0^*\omega'$ on $\widetilde{X}$, that agree on $\pi^{-1}(D_1) \cup \nu(C) \cup \pi^{-1}(D_2)$. We are left with showing that there exists a diffeomorphism ρ of $\widetilde{X}$ that maps $\mathcal{T}$ to $\mathcal{T}'$ and satisfies $\rho^*\Psi_0^*\omega' = \omega$.

View $S^2 \times S^2$ as $C \times S^2 = (D_1 \cup A \cup D_2) \times S^2$ where C is equipped with the symplectic form pulled back from ω. By Lemma 3.5.11 we can construct two symplectic forms ω_1 and ω_2 on $S^2 \times S^2$ by taking on $A \times S^2$ the forms ω and $\Psi_0^*\omega'$, but taking on $(D_1 \cup D_2) \times S^2$ the split form $\omega_D \times \omega_{S^2}$. By a Moser argument, we can find smaller closed discs $D_i^< \subset D_i$ and a diffeomorphism μ of $S^2 \times S^2$ such that:

- μ is the identity on $(D_1^< \cup D_2^<) \times S^2$,
- μ preserves $C \subset S^2 \times S^2$, and
- $\mu^*\omega_1$ is a split form near C. Since ω_1 and ω_2 coincide near C, this is also true of $\mu^*\omega_2$.

Choose a closed neighbourhood K of C that is contained in this neighbourhood, and choose smaller closed discs $D_i^{\lll} \subset D_i^<$ around p_i. By Theorem 3.5.12 below, there exists a diffeomorphism Φ of $S^2 \times S^2$ that is the identity on $K \cup ((D_1^{\lll} \cup D_2^{\lll}) \times S^2)$ and pulls back $\mu^*\omega_2$ to $\mu^*\omega_1$. We now take the diffeomorphism ρ of $\widetilde{X}$ to be $\mu \circ \Phi \circ \mu^{-1}$ on $\widetilde{X} \setminus \pi^{-1}(p_1 \cup p_2) = (S^2 \times S^2) \setminus (\{p_1 \cup p_2\} \times S^2)$ and equal to the identity on the broken rulings $\pi^{-1}(p_1 \cup p_2)$. Then $\Psi_0 \circ \rho$ still takes $\mathcal{T}$ to $\mathcal{T}'$, and it pulls back ω' to ω.

We are left with proving the following improvement of the Gromov–McDuff Theorem [MS04, Theorem 9.4.7 (ii)].

Theorem 3.5.12. *Let $\omega_{a,b}$ be the usual split symplectic form on $S^2 \times S^2$ that gives the first factor area a and the second factor area b. Represent $(S^2 \times S^2, \omega_{a,b})$ by its moment image as in Figure 3.19. Let p_N and p_S the North and South poles of S^2, and consider the three spheres*

$$S_1 = S^2 \times \{p_S\}, \quad S_2^N = \{p_N\} \times S^2, \quad S_2^S = \{p_S\} \times S^2$$

in $S^2 \times S^2$. Let ω be a symplectic form on $S^2 \times S^2$ that agrees with $\omega_{a,b}$ on a neighbourhood of $S_1 \cup S_2^N \cup S_2^S$. Then there exists a symplectomorphism $\Phi\colon (S^2 \times S^2, \omega_{a,b}) \to (S^2 \times S^2, \omega)$ that is the identity on a neighbourhood of $S_1 \cup S_2^N \cup S_2^S$.

We note that this theorem (and in fact also a stronger version with four spheres taken out) has been proved independently in [CMM25, Lemma 3.1.3].

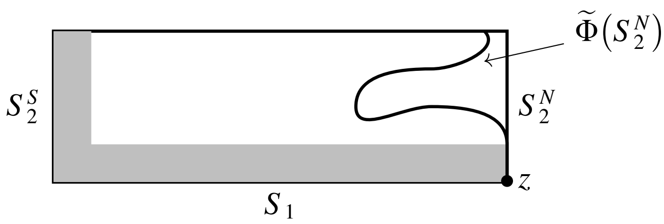

Figure 3.19: The moment rectangle for the manifold $S^2 \times S^2$, illustrating the proof of Theorem 3.5.12. The symplectomorphism $\widetilde{\Phi}$ coming from the Gromov–McDuff Theorem, fixes the shaded region pointwise. We must isotope it to something fixing a neighbourhood of S_2^N by first isotoping $\widetilde{\Phi}\left(S_2^N\right)$ back to S_2^N, then fixing the symplectic normal bundle of S_2^N, and finally straightening the symplectomorphism in a neighbourhood of S_2^N.

Proof of Theorem 3.5.12. According to [MS04, Theorem 9.4.7] there exists a symplectomorphism $\widetilde{\Phi}\colon (S^2 \times S^2, \omega_{a,b}) \to (S^2 \times S^2, \omega)$ that is the identity on a neighbourhood of $S_1 \cup S_2^S$. Set $z = (p_N, p_S) = S_1 \cap S_2^N$. For $t \in [0, 1]$ choose a smooth path J_t of ω-compatible almost complex structures such that

(i) S_1 and S_2^S are J_t-holomorphic for all t,

(ii) $\widetilde{\Phi}(S_2^N)$ is J_0-holomorphic,

(iii) S_2^N is J_1-holomorphic.

Such a path exists since $S_1, S_2^S, S_2^N, \widetilde{\Phi}(S_2^N)$ are embedded ω-symplectic spheres and since S_2^S is disjoint from $\widetilde{\Phi}(S_2^N)$ and S_1 intersects S_2^N and $\widetilde{\Phi}(S_2^N)$ transversely, positively, and only in z.

Let S_2^t be the unique J_t-holomorphic sphere in class $A := [S_2^N] \in H_2(S^2 \times S^2; \mathbb{Z})$ passing through z. This exists because $A^2 = 0$ so the Gromov–Taubes invariant counting holomorphic spheres in this class with a single point constraint is equal to 1 and curves in this class cannot bubble, S_2^S being J_t-holomorphic.[14]

The spheres S_2^t depend smoothly on t, because the path J_t is smooth in t and each S_2^t is automatically regular (being a square zero sphere). Now $\{S_2^t\}$ is a smooth path of embedded ω-symplectic spheres such that $S_2^0 = \widetilde{\Phi}(S_2^N)$ and $S_2^1 = S_2^N$. By (i), each sphere S_2^t intersects S_1 transversally at z and is disjoint from S_2^S. We can therefore modify the path $\{S_2^t\}$ near z such that all S_2^t coincide with S_2^N near z, and such that the S_2^t are still disjoint from S_2^S and intersect S_1 only in z.

By [ST05, Proposition 0.3] or [McL12, Lemma 5.16] there exists a Hamiltonian isotopy φ^t of $(S^2 \times S^2, \omega)$ whose support is contained in any given neighbourhood of $\bigcup_{t\in[0,1]} S_2^t$ such that $\varphi^t(S_2^0) = S_2^t$. In particular, we can assume that the support of φ^t is disjoint from S_2^S. Since the S_2^t already agree near z, we can also assume that the support is disjoint from S_1. Hence $\varphi^1 \circ \widetilde{\Phi}$ is still the identity near $S_1 \cup S_2^S$ and takes S_2^N to itself. Let σ be the restriction of $\varphi^1 \circ \widetilde{\Phi}$ to S_2^N. Since σ fixes a neighbourhood of z, it is the time-1 map of a Hamiltonian isotopy σ^t of S_2^N that fixes a neighbourhood of z. Extend σ^t to a Hamiltonian isotopy $\hat{\sigma}_t$ of $(S^2 \times S^2, \omega)$ that has

[14] This is because a J_t-holomorphic bubble curve will necessarily remain disjoint from S_2^S, because z is not contained in S_2^S and by positivity of intersection, and must therefore live in a homology class which is a positive multiple of A. But the area cannot be larger than that of S_2^N, so there can be only one irreducible component.

support near S_2^N and is the identity near z. Then $\widehat{\Phi} := \hat{\sigma}_1^{-1} \circ \varphi^1 \circ \widetilde{\Phi}$ is a symplectomorphism $(S^2 \times S^2, \omega_{a,b}) \to (S^2 \times S^2, \omega)$ that is the identity near $S_2^S \cup S_1$ and is the identity along S_2^N.

The differential of $\widehat{\Phi}$ at a point $p \in S_2^N$ with respect to the symplectic splitting of $T_p(S^2 \times S^2)$ is of the form

$$d_p\widehat{\Phi} = \begin{pmatrix} \mathrm{id}_2 & B \\ 0_2 & D \end{pmatrix} : T_pS_2^N \oplus (T_pS_2^N)^\omega \to T_pS_2^N \oplus (T_pS_2^N)^\omega.$$

Since $d_p\widehat{\Phi}$ is symplectic, it leaves $(T_pS_2^N)^\omega$ invariant. Hence $B = 0_2$ and D is symplectic. The space $\mathrm{Symp}(2;\mathbb{R})$ deformation retracts onto $U(1) = S^1$. Since the space of pointed maps from the 2-sphere to the circle is connected, we can isotopy $d\widehat{\Phi}|_{S_2^N}$ at the bundle level to the identity along S_2^N. Using Moser's method as in the proof of the symplectic neighbourhood theorem, we can then find a symplectic isotopy from $\widehat{\Phi}$ to the identity *near* S_2^N. Since the tubular neighbourhood of S_2^N has trivial first homology, we can make this isotopy Hamiltonian. We finally multiply the Hamiltonian function H_t, that we may choose to vanish near z, with a function χ that is 1 on a small tubular neighbourhood of S_2^N and vanishes outside a slightly larger tubular neighbourhood. The composition $\Phi := \phi_{\chi H} \circ \widehat{\Phi}$ is then a symplectomorphism $(S^2 \times S^2, \omega_{a,b}) \to (S^2 \times S^2, \omega)$ that is the identity *near* S_2^N. The proof of Theorem 3.5.12 is complete. □

We are now ready to prove Theorem 3.1.8, that we restate as:

Corollary 3.5.13 (Hamiltonian uniqueness of pin-ellipsoids)**.** *Given two symplectic embeddings $\iota, \iota' \colon E_{p,q}(\alpha,\beta) \hookrightarrow \mathbb{CP}^2$ there exists a Hamiltonian diffeomorphism Φ of $\mathbb{CP}^2$ such that $\iota = \Phi \circ \iota'$.*

Proof. We first give the idea of the proof. Two embeddings ι, ι' as in the corollary yield splittings of $\mathbb{CP}^2$, namely

$$\mathbb{CP}^2 = \left(\mathbb{CP}^2 \setminus U\right) \cup U$$

for $U = \iota(E_{p,q}(\alpha,\beta))$, and analogously for ι'. By Corollary 3.5.10 the two complements can be symplectically identified, and the embedded ellipsoids are trivially symplectically identified through the embeddings. We will adjust the first identification near the boundaries so that it agrees with the second one, and hence obtain a symplectomorphism Φ of $(\mathbb{CP}^2, \omega_{\mathrm{FS}})$ that intertwines the ellipsoid embeddings. This symplectomorphism is Hamiltonian thanks to Gromov's result from [Gro85] that the symplectomorphim group of $(\mathbb{CP}^2, \omega_{\mathrm{FS}})$ is connected.

We now come to the actual proof. By definition of "symplectic embedding", there exist $\hat{\alpha} > \alpha, \hat{\beta} > \beta$ such that the maps ι, ι' are restrictions of symplectic embeddings

$$\hat{\iota}, \hat{\iota}' \colon E_{p,q}(\hat{\alpha},\hat{\beta}) \hookrightarrow \mathbb{CP}^2.$$

Choose the almost toric offcut $\mathrm{Off}_{\rho,\lambda}(E_{p,q}(\hat{\alpha},\hat{\beta}))$ such that

$$E_{p,q}(\alpha,\beta) \subset \mathrm{Off}_{\rho,\lambda}(E_{p,q}(\hat{\alpha},\hat{\beta})) \subset E_{p,q}(\hat{\alpha},\hat{\beta})$$

and such that $\overline{\mathrm{Off}_{\rho,\lambda}}(E_{p,q}(\hat{\alpha},\hat{\beta})) \subset \mathrm{Int}\left(E_{p,q}(\hat{\alpha},\hat{\beta})\right)$.[15] Doing the pavilion blow-up, as introduced in Definition 3.3.5, we obtain two closed symplectic manifolds, which we denote by $(\widetilde{X}, \widetilde{\omega})$

[15] Even though the edges of the offcut have rational slopes, we can clearly choose ρ, λ such that the offcut has these properties.

and $(\widetilde{X}', \widetilde{\omega}')$. Denote by C_{vis} the linear chain of symplectic spheres in the pavilion blow-up $\mathrm{Pav}_{\rho,\lambda}\big(E_{p,q}(\hat{\alpha},\hat{\beta})\big)$, and let $\mathcal{N}(C_{\mathrm{vis}})$ be a neighbourhood of C_{vis} in $\mathrm{Pav}_{\rho,\lambda}\big(E_{p,q}(\hat{\alpha},\hat{\beta})\big)$.[16] In the following we denote the linear subchains of $\mathcal{T}$ and $\mathcal{T}'$ corresponding to C_{vis} by C and C'. Trivially, the complements of the trees $\mathcal{T}$ and $\mathcal{T}'$ of symplectic spheres in these pavilion blow-ups are identified with the complement of the offcuts "downstairs" in $\mathbb{CP}^2$. We write

$$\jmath\colon \mathbb{CP}^2 \setminus \hat{\imath}\big(\overline{\mathrm{Off}}_{\rho,\lambda}(E_{p,q}(\hat{\alpha},\hat{\beta}))\big) \to (\widetilde{X} \setminus C, \widetilde{\omega})$$

and analogously $\jmath'$ for these identifications. By Corollary 3.5.10, there exists a symplectomorphism $\psi\colon (\widetilde{X}', \widetilde{\omega}') \to (\widetilde{X}, \widetilde{\omega})$ that takes $\mathcal{T}'$ to $\mathcal{T}$. In particular, ψ takes C to C'. Now consider the diagram of symplectic maps shown in Figure 3.20. Since the embedding $\hat{\imath}$ is defined on all of $E_{p,q}(\hat{\alpha},\hat{\beta})$, the embedding

$$\jmath \circ \hat{\imath}\colon \mathcal{N}(C_{\mathrm{vis}}) \setminus C_{\mathrm{vis}} \hookrightarrow \widetilde{X}$$

uniquely extends to a smooth embedding $\sigma\colon \mathcal{N}(C_{\mathrm{vis}}) \hookrightarrow \widetilde{X}$. Since $\jmath \circ \hat{\imath}$ is symplectic, σ is also symplectic, and $\sigma(C_{\mathrm{vis}}) = C$. Using in addition that ψ is a symplectomorphism C' to C, we see in the same way that

$$\psi \circ \jmath' \circ \hat{\imath}'\colon \mathcal{N}(C_{\mathrm{vis}}) \setminus C_{\mathrm{vis}} \hookrightarrow \widetilde{X} \tag{3.5.19}$$

extends to a symplectic embedding $\sigma'\colon \mathcal{N}(C_{\mathrm{vis}}) \hookrightarrow \widetilde{X}$ that also takes C_{vis} to C.

$$\begin{tikzcd}
\big((\widetilde{X}', \widetilde{\omega}'), C'\big) \arrow[rr, "\psi"] & & \big((\widetilde{X}, \widetilde{\omega}), C\big) \\
\mathbb{CP}^2 \setminus \hat{\imath}'\big(\overline{\mathrm{Off}}_{\rho,\lambda}(E_{p,q}(\hat{\alpha},\hat{\beta}))\big) \arrow[u, "\jmath'"] & & \mathbb{CP}^2 \setminus \hat{\imath}\big(\overline{\mathrm{Off}}_{\rho,\lambda}(E_{p,q}(\hat{\alpha},\hat{\beta}))\big) \arrow[u, "\jmath"'] \\
\cap & & \cap \\
\mathbb{CP}^2 & & \mathbb{CP}^2 \\
\cup & & \cup \\
\hat{\imath}'(E_{p,q}(\hat{\alpha},\hat{\beta})) & & \hat{\imath}(E_{p,q}(\hat{\alpha},\hat{\beta})) \\
& E_{p,q}(\hat{\alpha},\hat{\beta}) \arrow[ul, "\hat{\imath}'"] \arrow[ur, "\hat{\imath}"'] &
\end{tikzcd}$$

Figure 3.20: The diagram of symplectic maps used in the proof of Corollary 3.5.13.

Lemma 3.5.14. *There exists a compactly supported Hamiltonian diffeomorphism h of a neighbourhood $\mathcal{N}(C)$ such that near C,*

$$h \circ \sigma' = \sigma. \tag{3.5.20}$$

[16] The chain of spheres C_{vis} is visible in the almost toric diagram as in Figure 3.10 (b). The combinatorial data of C_{vis} is exactly that of $\mathcal{T}$ and $\mathcal{T}'$, except that $\mathcal{T}$ and $\mathcal{T}'$ may in addition contain one or two (-1)-spheres, depending on the Manetti weight. Recall that these (-1)-spheres stem from the "ambient" topology of $\mathbb{CP}^2$ and are not contained in the linear chain corresponding to the choice of pavilion. Therefore, the "interface" between the embeddings of the pin-ellipsoids and the symplectomorphism of their complements will exactly be C_{vis}.

Proof. Set $\tau = \sigma \circ (\sigma')^{-1}|_C$. Then $\tau(C_j) = C_j$ for all spheres C_j in C, and τ fixes the transverse intersection points $C_j \cap C_{j+1}$. The group of orientation preserving diffeomorphisms of a sphere that fix one or two points is connected. We therefore find a smooth path τ_t from id_C to τ. By Moser's deformation argument, we can assume that the τ_t are symplectic. Using Weinstein's symplectic neighbourhood theorem, we can extend τ_t to a smooth path of symplectic embeddings $\hat{\tau}_t \colon \mathcal{N}_t \to \mathcal{N}(C)$ of tubular neighbourhoods $\mathcal{N}_t$ of C that starts at the inclusion. Since $H_1(\mathcal{N}_t) = 0$, this path is generated by a Hamiltonian vector field. Multiplying the Hamiltonian function by a cut-off function that is 1 near C and 0 outside a small neighbourhood of C, we obtain a Hamiltonian diffeomorphism h supported near $\mathcal{T}$ that agrees with $\sigma \circ (\sigma')^{-1}$ near C. ◻

In view of Equations we can define (3.5.19) and (3.5.20) define $\Phi \in \mathrm{Symp}(\mathbb{CP}^2, \omega_{\mathrm{FS}})$ by

$$\mathbb{CP}^2 \setminus \hat{\iota}'\left(\overline{\mathrm{Off}}_{\rho,\lambda}(E_{p,q}(\hat{\alpha},\hat{\beta}))\right) \xrightarrow{j^{-1}\circ h\circ\psi\circ j'} \mathbb{CP}^2 \setminus \hat{\iota}\left(\overline{\mathrm{Off}}_{\rho,\lambda}(E_{p,q}(\hat{\alpha},\hat{\beta}))\right)$$

and

$$\hat{\iota}'(\mathcal{N}) \xrightarrow{\hat{\iota}\circ(\hat{\iota}')^{-1}} \hat{\iota}(\mathcal{N}),$$

where $\mathcal{N}$ is a sufficiently small neighbourhood of $\mathrm{Off}_{\rho,\lambda}(E_{p,q}(\hat{\alpha},\hat{\beta}))$ in $E_{p,q}(\hat{\alpha},\hat{\beta})$. Because of Lemma 3.5.14 this symplectomorphism is well-defined. Due to Gromov's theorem [Gro85] that $\mathrm{Symp}(\mathbb{CP}^2, \omega_{\mathrm{FS}})$ is connected, Φ is a Hamiltonian diffeomorphism, and hence the map we were looking for. ◻

3.5.3 Proof of the Staircase Theorem (Theorem 3.1.13)

In this section, we will use Corollary 3.5.13 to prove the Staircase Theorem, which we restate in the following form:

Theorem 3.5.15 (Markov staircases)**.** *Let (p_1, p_2, p_3) be a Markov triple with $q_1 := 3p_2p_3^{-1} \bmod p_1$ and let $(p_1', p_2', p_3') = (p_1, p_3, 3p_3p_1 - p_2)$ be its mutation at p_2. There is no symplectic embedding $E_{p_1,q_1}(\alpha,\beta) \hookrightarrow \mathbb{CP}^2$ if both $\alpha > \frac{p_3}{p_1p_2}$ and $\beta > \frac{p_2'}{p_3'p_1'}$.*

Fix the triple (p_1, p_2, p_3), companion q_1 and mutant triple p_1', p_2', p_3' as in the statement of Theorem 3.5.15. Consider the Vianna triangles $\mathfrak{D}(p_1, p_2, p_3)$ and $\mathfrak{D}(p_1', p_2', p_3')$ whose edges $\mathfrak{e}_i$ and $\mathfrak{e}_i'$ have affine integral lengths

$$|\mathfrak{e}_1| = \frac{p_1}{p_2p_3} \qquad |\mathfrak{e}_2| = \frac{p_2}{p_3p_1} \qquad |\mathfrak{e}_3| = \frac{p_3}{p_1p_2}$$
$$|\mathfrak{e}_1'| = \frac{p_1'}{p_2'p_3'} \qquad |\mathfrak{e}_2'| = \frac{p_2'}{p_3'p_1'} \qquad |\mathfrak{e}_3'| = \frac{p_3'}{p_1'p_2'}.$$

It suffices to show that if $E_{p_1,q_1}(\alpha,\beta)$ admits a symplectic embedding and $\alpha > |\mathfrak{e}_3|$ then $\beta \le |\mathfrak{e}_2'|$. If we draw the two Vianna triangles $\mathfrak{D}(p_1, p_2, p_3)$ and $\mathfrak{D}(p_1', p_2', p_3')$ superimposed then we can see they both contain $\mathfrak{A}_{p_1,q_1}(|\mathfrak{e}_3|, |\mathfrak{e}_2'|)$ (Figure 3.21).

Lemma 3.5.16. *Consider the triangle $\mathfrak{A}_{p_1,q_1}(|\mathfrak{e}_3|, |\mathfrak{e}_2'|)$ and recall that its girdle is the top side.*

(a) The primitive integer vector pointing along the girdle is $\begin{pmatrix} p_3' \\ (p_3'q_1 - 3p_2')/p_1 \end{pmatrix}$.

(b) The affine length of the girdle is equal to $\frac{p_1p_3}{p_2p_3'}$.

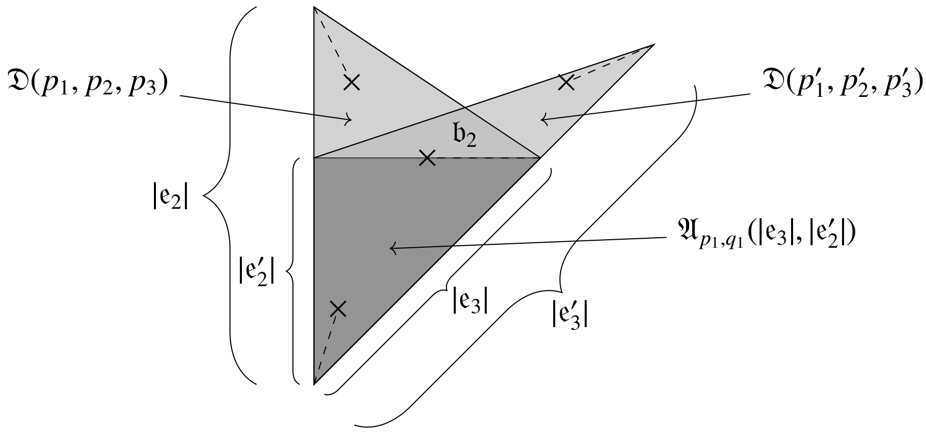

Figure 3.21: The Vianna triangles $\mathfrak{D}(p_1, p_2, p_3)$ and $\mathfrak{D}(p'_1, p'_2, p'_3)$ superimposed. The dark shaded region is a copy of $\mathfrak{A}_{p_1,q_1}(|\mathfrak{e}_3|, |\mathfrak{e}'_2|)$ in their intersection.

(c) The affine displacement between the girdle and the vertex $\mathfrak{v}_1$ is equal to p_3/p_1.

Proof. The vertices of $\mathfrak{A}_{p,q}(|\mathfrak{e}_3|, |\mathfrak{e}'_2|)$ are at

$$\begin{pmatrix}0\\0\end{pmatrix}, \quad |\mathfrak{e}_3|\begin{pmatrix}p_1^2\\p_1q_1-1\end{pmatrix}, \quad \begin{pmatrix}0\\|\mathfrak{e}'_2|\end{pmatrix},$$

so the vector pointing along the girdle is

$$\begin{pmatrix}|\mathfrak{e}_3|p_1^2\\|\mathfrak{e}_3|p_1q_1-|\mathfrak{e}_3|-|\mathfrak{e}'_2|\end{pmatrix} = \frac{p_3}{p_1p_2p'_3}\begin{pmatrix}p'_3p_1^2\\p'_3p_1q_1-p'_3-p_2\end{pmatrix} = \frac{p_3}{p_2p'_3}\begin{pmatrix}p'_3p_1\\p'_3q_1-3p_3\end{pmatrix},$$

where we used the facts that $p'_1 = p_1$, $p'_2 = p_3$ and $p'_3 + p_2 = 3p_1p_3$.

Since no Markov number is divisible by 3, and since the numbers in a Markov triple are always pairwise coprime, $\gcd(p'_3, p'_3q_1 - 3p_3) = \gcd(3p_1p_3 - p_2, 3p_3) = \gcd(p_2, p_3) = 1$.

We also have $\gcd(p_1, p'_3q_1 - 3p_3) = \gcd(p_1, (3p_1p_3 - p_2)q_1 - 3p_3) = \gcd(p_1, p_2q_1 + 3p_3) = p_1$ since[17] $q_1 = 3p_2p_3^{-1} = -3p_3p_2^{-1} \mod p_1$. Therefore the vector along the girdle is

$$\frac{p_1p_3}{p_2p'_3}\begin{pmatrix}p'_3\\(p'_3q_1-3p_3)/p_1\end{pmatrix}$$

where $\frac{p_1p_3}{p_2p'_3}$ is the integral affine length and $u := \begin{pmatrix}p'_3\\(p'_3q_1-3p_3)/p_1\end{pmatrix}$ is a primitive integer vector. The affine displacement between the girdle and the vertex at the origin is therefore

$$u^\perp \cdot \begin{pmatrix}0\\|\mathfrak{e}'_2|\end{pmatrix} = \begin{pmatrix}-(p'_3q_1-3p_3)/p_1\\p'_3\end{pmatrix} \cdot \begin{pmatrix}0\\p_3/(p_1p'_3)\end{pmatrix} = p_3/p_1. \qquad \square$$

Lemma 3.5.17. *If $\alpha > |\mathfrak{e}_3|$ and $\beta > |\mathfrak{e}'_2|$ then there exists a Delzant pavilion for $E_{p_1,q_1}(\alpha,\beta)$ (which introduces new edges $\mathfrak{f}_1, \dots, \mathfrak{f}_m$ with inward normals $\boldsymbol{\rho} = (\rho_1, \dots, \rho_m)$ corresponding to curves $C_1, \dots, C_m$) and an index i such that:*

[17]Note that $p_2^2 + p_3^2 = 3p_1p_2p_3 - p_1^2 = 0 \mod p_1$ so $p_2p_3^{-1} = -p_3p_2^{-1} \mod p_1$.

(a) the edge $\mathfrak{f}_i$ *is parallel to the branch cut* $\mathfrak{b}_2$ *in the Vianna triangle* $\mathfrak{D}(p_1, p_2, p_3)$*;*

(b) $\int_{C_i} \widetilde{\omega} > \frac{p_1 p_3}{p_2 p_3'}$ *;*

(c) if we define

$$b_i = -C_i^2, \quad e_i/e_{i-1} = [b_{i-1}, \ldots, b_1] \text{ and } f_i/f_{i+1} = [b_{i+1}, \ldots, b_m]$$

then $e_i = p_3'$ *and* $f_i = p_2$*.*

Proof. The fact that $\alpha > |\mathfrak{e}_3|$ and $\beta > |\mathfrak{e}_2'|$ means that the almost toric base diagram $\mathfrak{A}_{p_1,q_1}(|\mathfrak{e}_3|, |\mathfrak{e}_2'|)$ is strictly contained in $\mathfrak{A}_{p_1,q_1}(\alpha,\beta)$. The top side of this subdiagram is precisely where the branch cut $\mathfrak{b}_2$ would run in the Vianna triangle $\mathfrak{D}(p_1, p_2, p_3)$ (see Figure 3.21). Since the containment of the subdiagram is strict, this means that we can make a pavilion cut parallel to $\mathfrak{b}_2$ just above this top side (with further very small cuts near the ends of this side to ensure the pavilion is Delzant). The symplectic area of the corresponding curve C_i is given by the affine length of the side. Since our cut is slightly above the top side of $\mathfrak{A}_{p_1,q_1}(|\mathfrak{e}_3|, |\mathfrak{e}_2'|)$ (whose affine length is $\frac{p_1 p_3}{p_2 p_3'}$ by Lemma 3.5.16) and since we can take our remaining cuts to be arbitrarily small, we can ensure that $\int_{C_i} \widetilde{\omega} > \frac{p_1 p_3}{p_2 p_3'}$.

Finally, since the integer vector along the girdle is $\begin{pmatrix} p_3' \\ (p_3' q_1 - 3p_3)/p_1 \end{pmatrix}$, we can compute e_i and f_i as

$$\begin{aligned} e_i &= \begin{pmatrix} p_3' \\ (p_3' q_1 - 3p_3)/p_1 \end{pmatrix} \wedge \begin{pmatrix} 0 \\ 1 \end{pmatrix} = p_3', \\ f_i &= \begin{pmatrix} p_3' \\ (p_3' q_1 - 3p_3)/p_1 \end{pmatrix} \wedge \begin{pmatrix} p_1^2 \\ p_1 q_1 - 1 \end{pmatrix} \\ &= p_3'(p_1 q_1 - 1) - p_1(p_3' q_1 - 3p_3) = 3p_3 p_1 - p_3' = p_2. \end{aligned}$$ □

We now complete the proof. Let $\boldsymbol{\rho}$ be the normal vectors and $\boldsymbol{\lambda}$ the choice of constants giving the pavilion from Lemma 3.5.17. Let $\boldsymbol{\lambda}_t = (t\lambda_1, \ldots, t\lambda_m)$ for $t \in [0, 1]$ and consider the family of pavilion cuts $\left(\mathrm{Pav}^{\iota}_{\boldsymbol{\rho},\boldsymbol{\lambda}_t}(\mathbb{CP}^2), \widetilde{\omega}_t\right)$. When t is sufficiently small, there is a visible symplectic embedding $\iota_{\mathrm{vis}} \colon \mathrm{Off}_{\boldsymbol{\rho},\boldsymbol{\lambda}_t}(E_{p,q}(\alpha,\beta)) \hookrightarrow \mathbb{CP}^2$ given by the neighbourhood of a vertex of a Vianna triangle. By Corollary 3.5.13, the restriction

$$\iota|_{\mathrm{Off}_{\boldsymbol{\rho},\boldsymbol{\lambda}_t}(E_{p,q}(\alpha,\beta))} \colon \mathrm{Off}_{\boldsymbol{\rho},\boldsymbol{\lambda}_t}(E_{p,q}(\alpha,\beta)) \hookrightarrow \mathbb{CP}^2$$

is Hamiltonian isotopic to ι_{vis}, so $\mathrm{Pav}^{\iota}_{\boldsymbol{\rho},\boldsymbol{\lambda}_t}(\mathbb{CP}^2)$ is symplectomorphic to $\mathrm{Pav}^{\iota_{\mathrm{vis}}}_{\boldsymbol{\rho},\boldsymbol{\lambda}_t}(\mathbb{CP}^2)$. In the pavilion blow-up $\mathrm{Pav}^{\iota_{\mathrm{vis}}}_{\boldsymbol{\rho},\boldsymbol{\lambda}_t}(\mathbb{CP}^2)$ along the visible offcut, there is a visible symplectic -1-sphere S_t, living over an arc in the almost toric base diagram which connects the node coming out of v_2 to the edge $\mathfrak{f}_i$ corresponding to C_i. This is because $\mathfrak{f}_i$ is parallel to $\mathfrak{b}_2$; see [Eva23, Example 9.1 and Figure 9.1] and Figure 3.22.

This symplectic sphere S_t intersects C_i and none of the other curves C_j. As in Equation (3.3.16), we can write

$$[S_t] = s_0\,\mathcal{E} + \sum_{j=1}^{m} s_j [C_j] \in H_2(\mathrm{Pav}^{\iota}_{\boldsymbol{\rho},\boldsymbol{\lambda}_t}(\mathbb{CP}^2); \mathbb{Q})$$

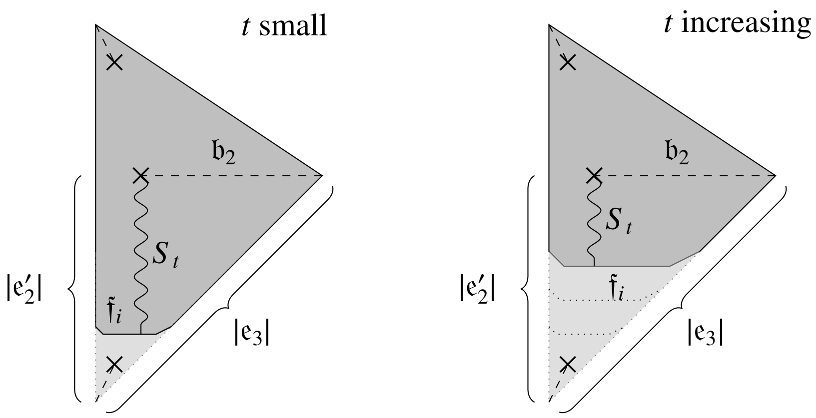

Figure 3.22: The pavilion blow-up $\mathrm{Pav}^{\iota_{\mathrm{vis}}}_{\rho,\lambda_t}(\mathbb{CP}^2)$ along a visible $E_{p,q}(\alpha,\beta)$ in $\mathbb{CP}^2$ for two different values of t. The visible sphere S_t lives over an arc connecting a node to the edge $\mathfrak{f}_i$. As t increases, the pavilion cuts move upwards and the symplectic area of S gets smaller. If $\alpha > |\mathfrak{e}_3|$ and $\beta > |\mathfrak{e}_2'|$, we could push $\mathfrak{f}_i$ up above the branch cut $\mathfrak{b}_2$ and S_t would end up with negative area.

and $\sigma_j := [S_t]\cdot[C_j]$, so $\boldsymbol{\sigma} = (\sigma_1,\ldots,\sigma_m) = (0,\ldots,0,1,0,\ldots,0)$ with the 1 in the ith place. Since S_t is a symplectic -1-sphere, the spherical Gromov–Taubes invariant $\mathrm{Gr}([S_t])$ counting holomorphic spheres in $\mathrm{Pav}^{\iota}_{\rho,\lambda_t}(\mathbb{CP}^2)$ in the class $[S_t]$ is equal to 1, see for instance [McD97, Proposition 4.1]. Since Gromov–Taubes invariants are deformation invariant, there is still a holomorphic sphere in this homology class for $\mathrm{Pav}^{\iota}_{\rho,\lambda_1}(\mathbb{CP}^2)$. This means that

$$0 < \int_{[S_1]} \widetilde{\omega}_1 = s_0 \int_{\mathcal{E}} \widetilde{\omega}_1 + \sum_{j=1}^{m} s_j \int_{C_j} \widetilde{\omega}_1.$$

Since $\boldsymbol{s} = M^{-1}\boldsymbol{\sigma}$, and since by Lemma 3.3.13 the entries of M^{-1} are all negative, we get

$$0 < s_0 \int_{\mathcal{E}} \widetilde{\omega}_1 + s_i \int_{C_i} \widetilde{\omega}_1.$$

Recall from Remark 3.3.14 that $p_1^2\mathcal{E}$ can be represented by a cycle in V (the complement of our embedding). Hence $s_0 \int_{\mathcal{E}} \widetilde{\omega}_t$ is constant in t; in particular, it is equal to the limit of the areas of the visible $\widetilde{\omega}_t$-symplectic spheres S_t as $t \to 0$; by [Sym03, Proposition 7.8], this equals the affine displacement $\frac{p_3}{p_1}$ between the girdle and the bottom vertex computed in Lemma 3.5.16(c). Therefore,

$$0 < \frac{p_3}{p_1} + s_i \int_{C_i} \widetilde{\omega}_1.$$

By Lemma 3.5.17(b), this tells us that

$$-s_i \frac{p_1 p_3}{p_2 p_3'} < \frac{p_3}{p_1}.$$

But $s_i = M_{ii}^{-1} = -\frac{e_i f_i}{p_1^2} = -\frac{p_2 p_3'}{p_1^2}$ by Lemma 3.3.13 and Lemma 3.5.17(c), so we get

$$\frac{p_3}{p_1} < \frac{p_3}{p_1},$$

which is a contradiction. This shows that we cannot have both $\alpha > \frac{p_3}{p_1 p_2}$ and $\beta > \frac{p_2'}{p_3' p_1'}$, as required.

3.5.4 Proof of Theorem 3.5.1 (Evans–Smith *sans* orbifolds)

Recall the notation from Definitions 3.3.4–3.3.5: given a symplectic embedding

$$\iota\colon E_{p,q}(\alpha,\beta) \hookrightarrow X := \mathbb{CP}^2,$$

we let $U = \iota(E_{p,q}(\alpha,\beta))$, let $V = X \setminus U$, and let $\overline{U}$ and $\overline{V}$ be the symplectic completions of U and V; we also write $\Sigma = \partial U$. We pick a neck-stretching sequence of almost complex structures J_t as in Definition 3.3.4.

Lemma 3.5.18. *For each t, pick a J_t-holomorphic line L_t in the class $H \in H_2(\mathbb{CP}^2;\mathbb{Z})$. As $t \to \infty$, there is a subsequence $t_i \to \infty$ such that the sequence L_{t_i} converges (in the sense of Gromov–Hofer convergence [Bou+03, Section 9.1]) to a holomorphic limit building with a nonempty component in $\overline{U}$ and a nonempty component in $\overline{V}$ (and potentially components in intermediate symplectisation levels).*

Proof. Convergence to a limit building is the content of the SFT compactness theorem [Bou+03; CM05]. There must be a component in $\overline{V}$ because $\overline{U}$ is exact and therefore unable to contain a closed holomorphic curve. There must be components in $\overline{U}$ because the anticanonical bundle of X restricted to the complement of a line can be trivialised, whilst the first Chern class of $\overline{U}$ is nontrivial (though it is torsion; see [ES18, Lemma 2.13]). □

Let C be one of the components of the SFT limit curve living in $\overline{V}$. Fix a minimal Delzant pavilion $\boldsymbol{\rho}, \boldsymbol{\lambda}$, let $\widetilde{X}$ be the pavilion blow-up and let $\widetilde{U} \subset \widetilde{X}$ be the girdled resolution with complement V. Let $\widetilde{C}$ be the corresponding closed curve in $\widetilde{X}$ given by Lemma 3.3.7. Our strategy will be to carefully analyse the adjunction formula for $\widetilde{C}$, using the discrepancies of the singularity to compute $K \cdot \widetilde{C}$ where $K = -\mathrm{PD}(c_1(\widetilde{X}))$ is the canonical class.

For that purpose, we write $K_{\mathbb{Q}}$ and $\widetilde{C}_{\mathbb{Q}}$, as in Equation (3.3.16), in the form

$$K_{\mathbb{Q}} = -3p\mathcal{E} + \sum k_i C_i, \qquad \widetilde{C}_{\mathbb{Q}} = c_0\mathcal{E} + \sum c_i C_i.$$

The coefficient of $\mathcal{E}$ in $K_{\mathbb{Q}}$ comes from the facts that $K_{\mathbb{CP}^2} = -3H$ and $H = p\mathcal{E}$. The numbers k_i are called the discrepancies of the singularity and can be computed to be $k_i = -1 + \frac{e_i + f_i}{p^2}$, see [HTU17, Section 2.1]. Here, e_i and f_i are defined as in Lemma 3.3.13. Note that these all lie in the interval $(-1, 0)$ since we have taken the minimal resolution; if we had taken a non-minimal resolution then some of the discrepancies would have been positive.

Lemma 3.5.19. *We have $0 < c_0 \le p$.*

Proof. The symplectic form on $\overline{V}$ extends to a closed 2-form ϖ on $\widetilde{X}$ which vanishes on $\bigcup_{i=1}^m C_i$ (this is essentially the form on the pavilion blow-up with $\boldsymbol{\lambda} = \mathbf{0}$). The integral $\int_{\widetilde{C}} \varpi$ agrees with the symplectic area of $C \subset \overline{V}$, which is bounded above by the symplectic area of a line in $\mathbb{CP}^2$

(which we are normalising to be 1). We have $\int_{\widetilde{C}} \varpi = c_0/p$ since the symplectic area of the line (which is homologous to $p\mathcal{E}$) is 1. Note that ϖ is non-negative on $\widetilde{J}$-complex lines in tangent spaces and vanishes only on points of $\bigcup_{i=1}^m C_i$, so $c_0/p = \int_{\widetilde{C}} \varpi > 0$. □

Proposition 3.5.20. *The curve $\widetilde{C}$ is an embedded sphere of self-intersection $\widetilde{C}^2 \le 0$.*

Proof. The curve $\widetilde{C}$ has arithmetic genus zero (since we obtained it by stretching a holomorphic sphere). The adjunction formula for $\widetilde{C}$ is

$$\sum_{x \in \mathrm{Sing}(\widetilde{C})} \delta_x = \frac{\widetilde{C}^2 + K \cdot \widetilde{C}}{2} + 1$$

where δ_x is the multiplicity of the singular point $x \in \widetilde{C}$. We have:

$$\widetilde{C}^2 = c_0^2/p^2 + \left(\sum_{i=1}^m c_i C_i\right)^2$$
$$\widetilde{C} \cdot K = \frac{-3c_0}{p} + \sum_{i=1}^m \chi_i k_i,$$

where $\chi_i = \widetilde{C} \cdot C_i$. Since the intersection matrix $M_{ij} = C_i \cdot C_j$ is negative definite, $\left(\sum_{i=1}^m c_i C_i\right)^2 \le 0$. By positivity of intersections, $\chi_i = \widetilde{C} \cdot C_i \ge 0$ for all i, so negativity of the discrepancies implies $\sum_{i=1}^m \chi_i k_i \le 0$. Therefore the adjunction formula becomes

$$\sum_{x \in \mathrm{Sing}(\widetilde{C})} \delta_x \le \frac{c_0^2 - 3pc_0}{2p^2} + 1.$$

Since c_0 satisfies $0 < c_0 \le p$ by Lemma 3.5.19, we find that $c_0^2 - 3pc_0$ is strictly negative. Since the numbers δ_x are positive integers, we deduce that $\mathrm{Sing}(\widetilde{C}) = \emptyset$. This shows that $\widetilde{C}$ is an embedded sphere.

The self-intersection of $\widetilde{C}$ is $c_0^2/p^2 + \left(\sum_{i=1}^m c_i C_i\right)^2$. Since M_{ij} is negative definite, and since $c_0 \le p$, this is ≤ 1. Therefore $\widetilde{C}^2 \le 1$ and $\widetilde{C}^2 \le 0$ unless $c_0 = p$ and $\boldsymbol{c} = 0$. In the latter case, $\widetilde{C}$ would be homologous to H. The following shows that this is impossible. Recall that $U = \iota(E_{p,q}(\alpha, \beta))$ deformation-retracts onto a CW complex called a pin-wheel, which has a fundamental class in homology with coefficients in $\mathbb{Z}_\ell$ for any prime ℓ dividing p (note that, despite the name, p is not assumed to be prime). The pin-wheel has nonzero self-intersection, so the fundamental class of this pin-wheel has nonzero $\mathbb{Z}_\ell$ intersection number with a line (which generates $\mathbb{Z}_\ell$-homology). This shows that $\widetilde{C}$ cannot be both homologous to H and disjoint from $\bigcup_{i=1}^m C_i$, which concludes the proof that $\widetilde{C}^2 \le 0$. □

Remark 3.5.21. Choosing the almost complex structure on V generically, we can ensure that the curve $\widetilde{C}$ has square 0 or -1, because the expected dimension of the moduli space is only positive in these two cases.

Proposition 3.5.22. *At least one of the components C of the SFT limit yields a curve $\widetilde{C} \subset \widetilde{X}$ which has square zero.*

Proof. We can impose an arbitrary point constraint on the line before we stretch, so that we find SFT-limit curves through every point of $\overline{V}$. If all of these were to lift to curves of negative square then, since there is only a countable number of such curves (at most one in each homology class), we could not find one through every point. This is a contradiction. □

To summarise, we have found an embedded $\widetilde{C} \subset \widetilde{X}$ with $\widetilde{C}^2 = 0$. Recall that we write $\widetilde{C} = c_0\mathcal{E} + \sum c_i C_i$ with $0 < c_0 \leq p$, and we define $\chi_i = \widetilde{C} \cdot C_i$, $\boldsymbol{c}^T = (c_1, \ldots, c_m)$, $\boldsymbol{\chi}^T = (\chi_1, \ldots, \chi_m)$ and $M_{ij} = C_i \cdot C_j$, so that $\boldsymbol{\chi} = M\boldsymbol{c}$. A formula for M^{-1}_{ij} was given in Lemma 3.3.13 in terms of the numbers e_i, f_i defined there. The fact that $\widetilde{C}^2 = 0$ implies

$$0 = \widetilde{C}^2 = \frac{c_0^2}{p^2} + \boldsymbol{\chi}^T M^{-1} \boldsymbol{\chi}.$$

Finally, the adjunction formula tells us that

$$0 = \frac{\widetilde{C}^2 + K \cdot \widetilde{C}}{2} + 1,$$

that is

$$1 = \frac{3pc_0 - c_0^2}{2p^2} - \frac{1}{2}\boldsymbol{\chi}^T M^{-1} \boldsymbol{\chi} - \frac{1}{2}\boldsymbol{k}^T \boldsymbol{\chi} \tag{3.5.21}$$

where $\boldsymbol{k}^T = (k_1, \ldots, k_m)$ is the vector of (negative) discrepancies defined by $K = -3p\mathcal{E} + \sum k_i C_i$.

Lemma 3.5.23. *We have (a) $\chi_i \leq 1$ for all i and (b) $\chi_i = 1$ for precisely one i.*

Proof. All the entries of $-M^{-1}$ are positive and also $\chi_i \geq 0$, so the second and third terms on the right-hand side of Equation (3.5.21) contribute at least

$$\sum_{i=1}^{m} \left(\frac{1}{2p^2}\chi_i^2 e_i f_i + \frac{1}{2}\chi_i \left(1 - \frac{e_i + f_i}{p^2}\right)\right) \tag{3.5.22}$$

to a sum whose total is 1.

If $\chi_i \geq 2$ then

$$\frac{1}{2p^2}\chi_i^2 e_i f_i + \frac{1}{2}\chi_i \left(1 - \frac{e_i + f_i}{p^2}\right) \geq 1 + \frac{1}{p^2}(2e_i f_i - e_i - f_i) \geq 1,$$

since $e_i, f_i \geq 1$. This is not possible, since the first term in Equation (3.5.21) is positive. Therefore $\chi_i \leq 1$ for all i. This proves (a).

If $\chi_i = 1$ then

$$\frac{1}{2p^2}\chi_i^2 e_i f_i + \frac{1}{2}\chi_i \left(1 - \frac{e_i + f_i}{p^2}\right) = \frac{1}{2}\left(1 + \frac{1}{p^2}(e_i f_i - e_i - f_i)\right) \geq \frac{1}{2}(1 - 1/p^2).$$

Let r be the number of indices i with $\chi_i \neq 0$. Then Equation (3.5.21) implies

$$1 \geq \frac{1}{2p^2}\left(3pc_0 - c_0^2\right) + \frac{r}{2}\left(1 - \frac{1}{p^2}\right).$$

The function $3pc_0 - c_0^2$ is an increasing function of c_0 on the interval $1 \leq c_0 \leq p$ and with minimum $3p-1$, so

$$1 \geq \frac{r}{2}\left(1 - \frac{1}{p^2}\right) + \frac{3p-1}{2p^2}.$$

If $r = 2$, the right hand side is > 1. If $r \geq 3$, the right hand side is also > 1, since then (using $p \geq 2$) we have $\frac{r}{2} - \frac{r}{2p^2} \geq \frac{3}{2} - \frac{3}{8} = \frac{9}{8} > 1$. The only possibility is that $r = 1$, which proves (b). □

Lemma 3.5.24. *There exist integers p_2, p_3 such that (p, p_2, p_3) is a Markov triple.*

Proof. We now know that the curve $\widetilde{C}$ intersects precisely one of the curves C_i (and does so once transversely), that is $\chi_i = 1$ and $\chi_j = 0$ for $j \neq i$. Therefore,

$$0 = \widetilde{C}^2 = \frac{c_0^2}{p^2} + \chi^T M^{-1} \chi = \frac{1}{p^2}\left(c_0^2 - e_i f_i\right),$$

so

$$c_0^2 = e_i f_i.$$

Equation (3.5.21) becomes

$$1 = \frac{3pc_0 - c_0^2}{2p^2} + \frac{e_i f_i}{2p^2} + \frac{1}{2} - \frac{e_i + f_i}{2p^2},$$

that is,

$$p^2 + e_i + f_i = 3p\sqrt{e_i f_i}.$$

Set $w_0 = p^2$, $w_1 = e_i$ and $w_2 = f_i$; we find

$$\frac{(w_0 + w_1 + w_2)^2}{w_0 w_1 w_2} = 9.$$

Now [HKW24, Lemma 2.8] implies that w_0, w_1, w_2 is a squared Markov triple, that is $e_i = p_2^2$ and $f_i = p_3^2$ for some Markov triple (p, p_2, p_3). □

Corollary 3.5.25. *The singularity $\frac{1}{p^2}(1, pq-1)$ has the form $\frac{1}{p^2}(p_3^2, p_2^2)$, so $q = \pm 3p_2p_3^{-1} \mod p$, and the curve C_i which $\widetilde{C}$ hits is the culet curve.*

Proof. Consider the toric singularity $\frac{1}{p^2}(1, pq-1)$, whose cone is bounded by the rays $\mathbb{R}_{\geq 0}(1, 0)$ and $\mathbb{R}_{\geq 0}(1 - pq, p^2)$. The curve C_i in the minimal resolution corresponds to a ray in this cone, say $\mathbb{R}_{\geq 0}(-t, s)$. Add the *opposite* ray $\mathbb{R}_{\geq 0}(t, -s)$ into the fan, and instead of getting a resolution, we get a compactification of the toric singularity (the fan is now complete). In terms of the moment polygon, this is equivalent to making a symplectic cut parallel to the (s, t)-direction whose inward normal points *downwards*. See Figure 3.23.

Since the partial resolution corresponding to the ray $\mathbb{R}_{\geq 0}(-t, s)$ would introduce the cyclic singularities $\frac{1}{e_i}(1, e_{i-1}^{-1} \mod e_i)$ and $\frac{1}{f_i}(1, f_{i+1} \mod f_i)$ (by definition of e_i and f_i from Lemma 3.3.13), the compactification will have dual singularities

$$\frac{1}{e_i}(1, e_i - e_{i-1}^{-1}), \qquad \frac{1}{f_i}(1, f_i - f_{i+1}).$$

Since $\gcd(p^2, e_i, f_i) = 1$, the result is a weighted projective space $\mathbb{P}(p^2, e_i, f_i)$. Since p^2, e_i, f_i is a Markov triple (p_1^2, p_2^2, p_3^2), the original singularity is $\frac{1}{p_1^2}(p_3^2, p_2^2)$ and the two new singularities are dual to the Wahl singularities $\frac{1}{p_2^2}(p_1^2, p_3^2)$ and $\frac{1}{p_3^2}(p_2^2, p_1^2)$, so by the discussion in Section 3.2.4, the curve C_i is the culet curve. The fact that $q = \pm 3p_2p_3^{-1} \mod p$ follows as in Equation (3.2.14). □

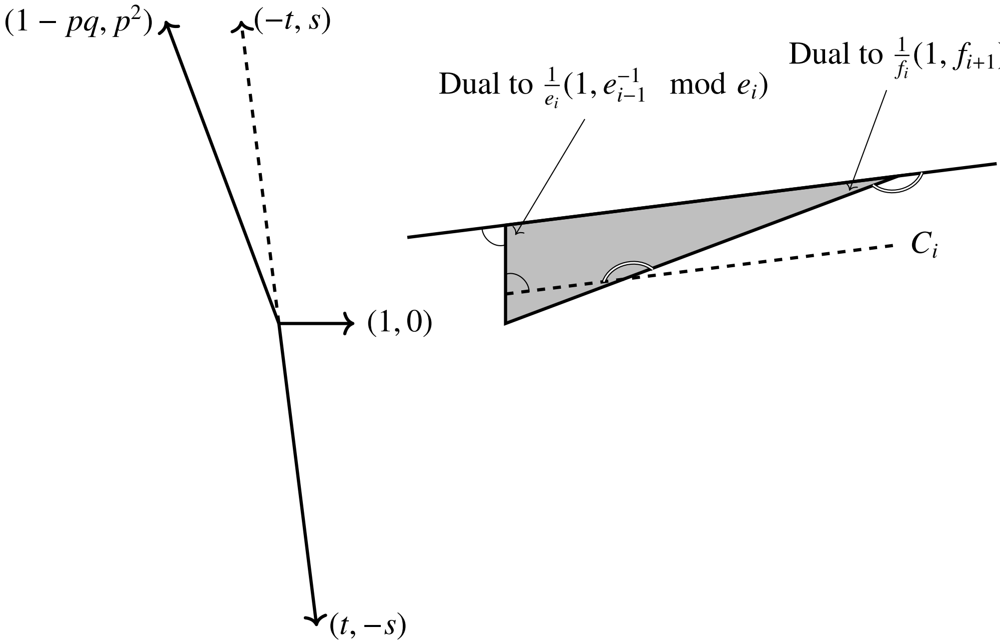

Figure 3.23: Left: The fan for the compactification of the toric $\frac{1}{p^2}(1, pq - 1)$ singularity. Right: The moment polygon for the compactification. The two new singularities (top left and top right corners of the triangle) are dual to the singularities $\frac{1}{e_i}(1, e_{i-1}^{-1} \mod e_i)$ and $\frac{1}{f_i}(1, f_{i+1})$ that would be introduced by the partial resolution with exceptional locus C_i.

This completes the proof of Theorem 3.5.1 and also reproves the theorem of Evans and Smith that $E_{p,q}(\alpha, \beta)$ can only embed in $\mathbb{CP}^2$ if p is a Markov number and q a companion.

3.5.5 Proof of the Two Pin-Ball Theorem (Theorem 3.1.18)

For this subsection we assume that we have a symplectic embedding of two pin-balls

$$\iota \colon B_{p_1,q_1}(\alpha_1) \sqcup B_{p_2,q_2}(\alpha_2) \hookrightarrow \mathbb{CP}^2,$$

where $1 < p_1 < p_2$, and we denote $\Delta = p_1p_2$. We take the same setup as in the beginning of Section 3.5.4, that is, we take C to be a component of the SFT limit curve in $\overline{V}$ obtained by the neck-stretching procedure and denote the corresponding closed curve in the minimal resolution $\widetilde{X}$ by $\widetilde{C}$. As explained in Remark 3.3.15, we write

$$K_{\mathbb{Q}} = -3\Delta\mathcal{E} + \sum_{i=1}^{m_1} k_{1,i}C_{1,i} + \sum_{i=1}^{m_2} k_{2,i}C_{2,i} \quad \text{and} \quad \widetilde{C}_{\mathbb{Q}} = c_0\mathcal{E} + \sum_{i=1}^{m_1} c_{1,i}C_{1,i} + \sum_{i=1}^{m_2} c_{2,i}C_{2,i}. \quad (3.5.23)$$

Repeating the proofs of Lemma 3.5.19 and Proposition 3.5.20 verbatim, we obtain $0 < c_0 \le \Delta$ and that $\widetilde{C}$ is an embedded sphere with $\widetilde{C}^2 \le 0$. Writing $\chi_j = M_j\left(c_{j,1}, \dots, c_{j,m_j}\right)$ and $\boldsymbol{k}_j = \left(k_{j,1}, \dots, k_{j,m_j}\right)$, the adjunction formula gives us an analogue of Equation (3.5.21):

$$1 = \frac{3\Delta c_0 - c_0^2}{2\Delta^2} - \frac{1}{2}\left(\chi_1^T M_1^{-1}\chi_1 + \boldsymbol{k}_1^T\chi_1\right) - \frac{1}{2}\left(\chi_2^T M_2^{-1}\chi_2 + \boldsymbol{k}_2^T\chi_2\right). \tag{3.5.24}$$

In their proof of [ES18, Theorem 4.16] Evans and Smith show the existence of an orbifold curve whose proper transform is the curve $\widetilde{C}$ we want; it hits each Wahl chain precisely once transversely: the first chain at a point of either $C_{1,1}$ or C_{1,m_1}, the second at a point of either $C_{2,1}$ or C_{2,m_2} (this translates into the property that $K_z = K_{z'} = 0$ in the notation of that paper). As in the previous section, one could repeat all their arguments in the minimal resolution, without referring to orbifolds, but we omit this derivation. Using this we find that Equation (3.5.24) for $\widetilde{C}$ is equivalent to

$$3c_0p_1p_2 = c_0^2 + p_1^2 + p_2^2, \tag{3.5.25}$$

which is just the Markov equation. This means that there are exactly two solutions for c_0, and because $c_0 < \frac{\Delta}{3p_1} = \frac{p_2}{3}$ by the proof of [ES18, Lemma 4.11.B], we know that c_0 is the smaller of the two solutions to (3.5.25).

Now the SFT-limit curve C has area $\int_{c_0\mathcal{E}}\widetilde{\omega} = c_0/(p_1p_2)$. For $i = 1, 2$, let $\overline{N}_i$ be the negative end of $\overline{V}$ coming from the neck around $B_{p_i,q_i}(\alpha_i)$ and let D_i be the unique non-compact connected component of $C \cap \overline{N}_i$. Each D_i is properly embedded.

Lemma 3.5.26. *We have $\alpha_i \le \mathrm{Area}(D_i)$ for $i = 1, 2$.*

Proof. Add in the orbifold points x_1 and x_2 to obtain girdled orbifolds $\widehat{N}_i$, which are locally modelled on $\mathbb{C}^2/\mathbb{Z}_{p_i^2}$. For each $i = 1, 2$, take the uniformising cover $\mathbb{C}^2 \to \mathbb{C}^2/\mathbb{Z}_{p_i^2}$ and take the preimage of D_i; this is a union of punctured curves in a Euclidean ball $B \subset \mathbb{C}^2$. In fact, the preimage consists of a single punctured curve D_i': since $\widetilde{C}$ intersects one of the end-curves of the Wahl chain, D_i is asymptotic to the quotient of either the punctured z_1 or z_2 axis in $\mathbb{C}^2$, so that $\pi_1(D_i)$ maps surjectively to $\mathbb{Z}_{p_i^2}$.

Since D_i' is a finite-energy curve, removal of singularities tells us that it compactifies to give a properly embedded holomorphic curve $D_i'' \subset B \subset \mathbb{C}^2$ passing through the origin. The monotonicity formula for minimal surfaces now tells us that the area of this holomorphic curve is at least the area of a complex line intersected with B. Therefore the area of D_i is bounded from below by the area of the projection of a complex line to the girdled orbifold $\widehat{N}_i$, which is α_i. □

As a corollary,

$$\alpha_1 + \alpha_2 < \mathrm{Area}(C) = \frac{c_0}{p_1p_2},$$

because the curve C passes through the region between the two necks. As explained in Remark 3.1.19, this proves the Two Pin-Ball Theorem (Theorem 3.1.18).

3.5.6 Proof of Theorem 3.5.4

The strategy of the proof is to mimic, as closely as possible, an argument of Manetti [Man91] but staying in the world of symplectic geometry. Given a symplectic embedding $\iota : E_{p,q}(\alpha,\beta) \hookrightarrow X$ for some p, q, α, β, choose a minimal pavilion in the sense of Definition 3.2.7.

By Theorem 3.5.1, there exists a square zero $\widetilde{J}$-holomorphic sphere $\widetilde{C} \subset \widetilde{X}$ which hits $C_{i_{\text{culet}}}$ once transversely and is disjoint from the other curves C_j, $j \neq i_{\text{culet}}$. By Theorem 3.4.4, $\widetilde{C}$ is part of a $\widetilde{J}$-holomorphic regulation on $\widetilde{X}$. By Lemma 3.3.8, we can choose J generically on V so as to ensure that the only irregular genus zero J-curves are $C_1, \ldots, C_m$: all the other embedded J-holomorphic spheres are either -1-spheres or have non-negative square. By Proposition 3.4.5(c), the broken rulings can only involve spheres of negative square, so the broken rulings are built out of the curves $C_1, \ldots, C_m$ and some -1-spheres.

Lemma 3.5.27. *Every curve C_j except the culet curve $C_{i_{\text{culet}}}$ appears as an irreducible component in one of the broken rulings.*

Proof. The curves C_j other than $C_{i_{\text{culet}}}$ are disjoint from $\widetilde{C}$. Since $\widetilde{C}^2 = 0$, we can impose an arbitrary point constraint $x \in \widetilde{X}$ and find a ruling passing through x. If $x \in C_j$ then we find a ruling S intersecting C_j. Since $[S] = [\widetilde{C}]$, we find $S \cdot C_j = 0$, and thus the ruling S must contain C_j as an irreducible component. □

We now use Proposition 3.4.11 to ensure that $\widetilde{J}$ is integrable in a neighbourhood of the -1-curves in the broken rulings, keeping it integrable along $C_1, \ldots, C_m$ (where it is already integrable). Recall from Proposition 3.4.5(d) that every broken ruling contains at least one -1-sphere. Now we can successively holomorphically contract -1-spheres in broken rulings until we reach an almost complex 4-manifold Y with a non-degenerate holomorphic regulation with no broken rulings. This must have second Betti number 2 (in fact, it must be diffeomorphic to either $S^2 \times S^2$ or $\mathbb{CP}^2 \# \overline{\mathbb{CP}}^2$, see Remark 3.5.8).

Lemma 3.5.28. *The contraction from our resolved surface $\widetilde{X}$ to the minimal model Y contracts $m - 1$ curves.*

Proof. Let k be the number of curves contracted in going from $\widetilde{X}$ to Y. The second Betti number of $\widetilde{X}$ is $k + 2$. But the second Betti number of $\widetilde{X}$ is equal to $1 + m$ since it is obtained from $\mathbb{CP}^2$ by pavilion blow-up, introducing the chain of m spheres $C_1, \ldots, C_m$. Therefore $k + 2 = m + 1$ and $k = m - 1$. □

Lemma 3.5.29. *There are at most two broken rulings on $\widetilde{X}$. More precisely,*

- *If $i_{\text{culet}} = 1 = m$ then there are no broken rulings.*
- *If $i_{\text{culet}} = 1 < m$ (respectively $1 < i_{\text{culet}} = m$) then there is precisely one broken ruling comprising $C_2, \ldots, C_m$ (respectively $C_1, \ldots, C_{m-1}$) and one additional $\widetilde{J}$-holomorphic -1-sphere E.*
- *If $1 < i_{\text{culet}} < m$ then there are precisely two broken rulings: one comprising $C_1, \ldots, C_{i_{\text{culet}}-1}$ and an additional -1-sphere E_1, the other comprising $C_{i_{\text{culet}}+1}, \ldots, C_m$ and an additional -1-sphere E_2.*

Proof. By Proposition 3.4.5, each broken ruling is a tree of embedded $\widetilde{J}$-holomorphic spheres with negative self-intersection, which are therefore either amongst the curves $C_1, \ldots, C_m$ or else -1-spheres. All of the curves C_i other than the culet curve $C_{i_{\text{culet}}}$ appear amongst the irreducible components of some broken rulings by Lemma 3.5.27.

If $m = 1$ then $i_{\text{culet}} = 1$ and by Lemma 3.5.28 we have $\widetilde{X} = Y$, so there are no broken rulings. This situation arises only for pin-ellipsoids with $p = 2$.

If $i_{\text{culet}} = 1 < m$ (respectively $1 < i_{\text{culet}} = m$) then by Lemma 3.4.7 the curves $C_2, \ldots, C_m$ (respectively $C_1, \ldots, C_{m-1}$) will belong to a single broken ruling. Since these curves all have $C_i^2 \leq -2$, Proposition 3.4.5(d) shows that this broken ruling must also contain a -1-sphere E. We need to contract all but one of the curves in this ruling to get to the minimal model Y. This is $m-1$ curves, so there can be no other broken rulings by Lemma 3.5.28. Moreover, the broken ruling is precisely $E \cup \bigcup_{i=2}^{m} C_i$. See Figure 3.16(b) for an example; this situation arises precisely for pin-ellipsoids with $p \geq 5$ along the Fibonacci branch of the Markov tree.

If $i_{\text{culet}} \neq 1, m$ then there will be two broken rulings containing the subchains $C_1, \ldots, C_{i_{\text{culet}}-1}$ and $C_{i_{\text{culet}}+1}, \ldots, C_m$. Each of these broken rulings will also contain a -1-sphere, say E_1 and E_2. In each broken ruling we need to contract all but one of the curves, which means we contract a total of $i_{\text{culet}} - 1$ curves in the first and $m - i_{\text{culet}}$ curves in the second ruling, making a total of $m - 1$ curves. Again, Lemma 3.5.28 tells us there can be no further rulings and the two rulings we have found contain no further curves. See Figure 3.16(c) for an example. □

Remark 3.5.30. Moreover, we can identify which of the curves $C_1, \ldots, C_m$ intersect the -1-spheres E_1 and E_2 in the broken rulings. Since the dual graph of the broken ruling is a tree, E_1 and E_2 can each intersect at most one of the curves in the chain, and blowing down E_1 (respectively E_2) will turn the subchain $C_1, \ldots, C_{i_{\text{culet}}-1}$ (respectively $C_{i_{\text{culet}}+1}, \ldots, C_m$) into a zero continued fraction by Remark 3.4.13. Recall that $C_1, \ldots, C_{i_{\text{culet}}-1}$ and $C_{i_{\text{culet}}+1}, \ldots, C_m$ form dual Wahl chains. It is well-known (see for example [UZ25, Proposition-Definition 2.3 and Proposition 2.4]) that a dual Wahl chain has a unique position where it can be combinatorially blown down to obtain a zero-continued fraction, so this gives the positions where E_1 and E_2 intersect the subchains.

This completes the proof of Theorem 3.5.4.

3.6 Appendix

3.6.1 Alternative proof of Theorem 3.5.1

In this appendix, we explain how Theorem 3.5.1 follows from [ES18, Theorem 4.15]. Given a symplectic embedding $\iota : E_{p,q}(\alpha,\beta) \hookrightarrow \mathbb{CP}^2$, they found a Markov triple (p, b, c) and constructed a $\widehat{J}$-holomorphic orbifold curve $\widehat{C}$ in the orbifold $\widehat{X}$ (see Figure 3.14 for the notation) satisfying the following properties:

(ES1) Near the orbifold point $p \in \widehat{X}$, $\widehat{C}$ has precisely one branch (in the terminology of [ES18], $|Z| = 1$).

(ES2) If we look at the lift of $\widehat{C}$ to the local uniformising cover $\mathbb{C}^2 \to \mathbb{C}^2/G$, there exist local coordinates (z_1, z_2) transforming under the action of $\zeta \in G$ as $(\zeta^{m_1} z_1, \zeta^{m_2} z_2)$ with respect to which the lift is given by $z \mapsto \left(z^{b^2}, a_1 z^{c^2}\right)$. In particular, the Puiseux series has length one and the link of the singularity of the curve in the cover is a (b^2, c^2)-torus knot.

(ES3) The real dimension of the moduli space containing $\widehat{C}$ is equal to its virtual dimension which equals 2.

Remark 3.6.1. We take this opportunity to correct [ES18, Remark 3.15]. Namely, as described in [ES18, Section 3.3], the germ of the orbifold curve $f : \widehat{C} \to \widehat{X}$ near an orbifold point $z \in \widehat{C}$ involves a choice of homomorphism $\rho_z : H_z \to G_z$, where H_z is the isotropy group of $\widehat{C}$ at z and

G_z is the isotropy group of $\widehat{X}$ near $f(z)$; a lift $\tilde{f}_\alpha$ of f to the local uniformising cover near $f(z)$ satisfies the equivariance condition [ES18, Eq. (2)]

$$\tilde{f}_\alpha(\zeta x) = \rho_z(\zeta)\tilde{f}_\alpha(x).$$

In [ES18, Remark 3.15], this was applied to $\tilde{f}_\alpha(x) = \left(x^Q, \sum a_i x^{R_i}\right)$ and the cyclic quotient singularity $\frac{1}{p^2}(1, pq-1)$ to say that

$$\left(\zeta^Q x^Q, \sum a_i \zeta^{R_i} x^{R_i}\right) = \left(\zeta x^Q, \zeta^{pq-1} \sum a_i x^{R_i}\right)$$

and deduce that $Q \equiv 1 \mod p^2$, $R_i \equiv pq-1 \mod p^2$ (or vice versa). However, this assumes that ρ_z is the identity. More generally, ρ_z could have the form $\rho_z(\zeta) = \zeta^n$ for some n. If $\gcd(n, p^2) = p^2/d$ then at best we can deduce that

$$R_i/Q = (pq-1)^{\pm 1} \mod d.$$

In fact, this is all that was used in what followed. It should have been clear that the stronger conclusion is false from the fact that eventually $Q = b^2$, $R_1 = c^2$, for example if $(p, b, c) = (2, 5, 29)$ then neither 2^2 nor 5^2 is congruent to 1 modulo 29^2.

The uniformising cover. Apply an integral affine transformation to $\mathfrak{D}(p_1, p_2, p_3)$ as in 3.2.4 to make the edge $\mathfrak{e}_1$ horizontal with the rest of the triangle below. Let $\boldsymbol{w}_2$ and $\boldsymbol{w}_3$ be the primitive integer vectors pointing out of $\mathfrak{v}_1$ along $\mathfrak{e}_2$ and $\mathfrak{e}_3$ respectively; note that $\boldsymbol{w}_2 = (-u_2, p_3^2)$ and $\boldsymbol{w}_3 = (u_3, p_2^2)$ for some u_2, u_3, since $(p_2\boldsymbol{w}_2)/(p_3p_1)$ and $(p_3\boldsymbol{w}_3)/(p_1p_2)$ must reach the same height p_2p_3/p_1 (compare with [BS24, Paragraph 2 of the proof of Proposition 2.5] for the computation of this height). Consider the uniformising cover $\mathbb{C}^2 \to \mathbb{C}^2/G$ for the orbifold singularity $\frac{1}{p_1^2}(p_2^2, p_3^2)$. The cover $\mathbb{C}^2$ is also toric, its moment polygon is the positive quadrant, and the uniformising cover is induced by the integral affine map given by the matrix whose columns are $\boldsymbol{w}_2$ and $\boldsymbol{w}_3$. Under this map, a symplectic cut along the $(p_2^2, -p_3^2)$ direction in the moment image of $\mathbb{C}^2$ maps to a horizontal cut in the moment image of $\frac{1}{p_1^2}(p_2^2, p_3^2)$, that is the culet cut.

Proposition 3.6.2. *If $\widehat{C} \subset \widehat{X}$ is a curve with Properties (ES1–3) then the proper transform $\widetilde{C} \subset \widetilde{X}$ satisfies the following properties:*

(a) $\widetilde{C}$ is a smooth sphere which intersects the culet curve $C_{i_{\text{culet}}}$ once transversely and is disjoint from the other curves C_j, $j \neq i_{\text{culet}}$.

(b) $\widetilde{C}$ lives in a moduli space of real dimension equal to its virtual dimension, which equals 2. In particular, $\widetilde{C}$ is a square zero $\widetilde{J}$-holomorphic sphere, and therefore appears as a smooth ruling in a $\widetilde{J}$-holomorphic regulation.

Proof. (a) By Property (ES2), the germ of the lifted curve in the uniformising cover $\mathbb{C}^2 \to \mathbb{C}^2/G$ is parametrised by $z \mapsto \left(z^{b^2}, a_1 z^{c^2}\right)$. We can resolve the singularity of the germ by performing a weighted blow-up. Consider the surface

$$\left\{(z_1, z_2, [w_1 : w_2]) \,:\, z_1^{c^2} w_2^{b^2} = w_1^{c^2} z_2^{b^2}\right\} \subset \mathbb{C}^2 \times \mathbb{P}^1(b^2, c^2).$$

The total transform of our germ under the projection of this surface to $\mathbb{C}^2$ is given by the equations $z_1^{c^2} w_2^{b^2} = w_1^{c^2} z_2^{b^2}$ and $z_1^{c^2} = z_2^{b^2}$. In the chart $w_2 = 1$, we have $z_1^{c^2} = w_1^{c^2} z_2^{b^2} = z_2^{b^2}$, so the proper transform becomes $w_1^{c^2} = 1$; this looks like a curve with c^2 irreducible components $w_1 = \mu$ as μ runs over the set $\boldsymbol{\mu}_{c^2}$ of c^2th roots of unity, but there is a residual action of $\boldsymbol{\mu}_{c^2}$ (the rescalings preserving the condition $w_2 = 1$) which relates all of these, so the proper transform of the germ has a single, smooth irreducible component in this chart which intersects the exceptional curve $z_2 = 0$ once transversely; we find the same component in the other chart $w_1 = 1$. This weighted blow-up corresponds to making a symplectic cut of $\mathbb{C}^2$ along the hypersurface living over the line parallel to the vector $(b^2, -c^2)$, which projects under the partial resolution of the uniformising cover to the culet curve, as explained in 3.6.1.

(b) The same analysis applies to the proper transforms of the other curves $\widehat{C}_t$ in the moduli space of $\widehat{C}$, so that the proper transforms $\widetilde{C}_t$ live in the same moduli space. Moreover, if $\widetilde{C}_s$ is in the same moduli space of $\widetilde{C}$ then its projection to $\widehat{X}$ lives in the same moduli space as $\widehat{C}$, so the moduli spaces are diffeomorphic and the dimension computation follows from Property (ES3). The fact that $\widetilde{C}$ is a sphere of square zero follows from the fact that a sphere of square n lives in a moduli space of virtual real dimension $2n + 2$ (which equals the actual dimension as long as $n \geq -1$ by automatic transversality). □

This completes the alternative proof of Theorem 3.5.1.

Bibliography

[Abo12] M. Abouzaid. "Nearby Lagrangians with vanishing Maslov class are homotopy equivalent". In: *Invent. Math.* 189.2 (2012), pp. 251–313. DOI: 10.1007/s00222-011-0365-0.

[Abo+25a] M. Abouzaid, D. Álvarez-Gavela, S. Courte, and T. Kragh. *Normal invariant of nearby Lagrangians via twisted derivative*. 2025. arXiv: 2505.12515 [math.SG].

[Abo+25b] M. Abouzaid, S. Courte, S. Guillermou, and T. Kragh. "Twisted generating functions and the nearby Lagrangian conjecture". In: *Duke Math. J.* 174.5 (2025), pp. 949–1011. DOI: 10.1215/00127094-2024-0052.

[AK18] M. Abouzaid and T. Kragh. "Simple homotopy equivalence of nearby Lagrangians". In: *Acta Math.* 220.2 (2018), pp. 207–237. DOI: 10.4310/ACTA.2018.v220.n2.a1.

[Abr98] M. Abreu. "Topology of symplectomorphism groups of $S^2 \times S^2$". In: *Invent. Math.* 131.1 (1998), pp. 1–23. DOI: 10.1007/s002220050196.

[Abr03] M. Abreu. "Kähler geometry of toric manifolds in symplectic coordinates". In: *Symplectic and contact topology: interactions and perspectives*. Vol. 35. Fields Inst. Commun. Amer. Math. Soc., Providence, RI, 2003, pp. 1–24. DOI: 10.1090/fic/035/01.

[AM00] M. Abreu and D. McDuff. "Topology of symplectomorphism groups of rational ruled surfaces". In: *J. Am. Math. Soc.* 13.4 (2000), pp. 971–1009. DOI: 10.1090/S0894-0347-00-00344-1.

[Ada25] N. Adaloglou. "Uniqueness of Lagrangians in $T^*\mathbb{R}P^2$". In: *Ann. Math. Qué.* 49.1 (2025), pp. 215–222. DOI: 10.1007/s40316-024-00238-3.

[Ada26] N. Adaloglou. "Embeddings and disjunction of Lagrangian pinwheels via rational blow-ups". In: *J. Symplectic Geom.* 24.1 (2026), pp. 105–129. DOI: 10.4310/JSG.260423234950.

[ABH26] N. Adaloglou, G. Bargalló i Gómez, and J. Hauber. *The nearby Lagrangian conjecture for pinwheels*. Preprint, arXiv:2605.22473 [math.SG] (2026). 2026.

[Ada+25] N. Adaloglou, J. Brendel, J. Evans, J. Hauber, and F. Schlenk. *Markov staircases*. 2025. arXiv: 2509.03224 [math.SG].

[AH25] N. Adaloglou and J. Hauber. *Pinwheels in symplectic rational and ruled surfaces and non-squeezing of rational homology balls*. 2025. arXiv: 2503.16250 [math.SG].

[AH] N. Adaloglou and J. Hauber. "Symplectic staircases for domains in cotangent bundles". Forthcoming.

[Aig13] M. Aigner. *Markov's theorem and 100 years of the uniqueness conjecture*. A mathematical journey from irrational numbers to perfect matchings. Springer, Cham, 2013, pp. x+257. DOI: 978-3-319-00888-2.

[Arn86] V. I. Arnol'd. "First steps in symplectic topology". In: *Russ. Math. Surv.* 41.6 (1986), pp. 1–21. DOI: 10.1070/RM1986v041n06ABEH004221.

[AMW25] M. Atallah, C. Y. Mak, and W. Wu. *C^0-rigidity of the Hamiltonian diffeomorphism group of symplectic rational surfaces*. Preprint, arXiv:2508.20285 [math.SG]. 2025.

[Aud07] M. Audin. "Lagrangian skeletons, periodic geodesic flows and symplectic cuttings". In: *Manuscr. Math.* 124.4 (2007), pp. 533–550. DOI: 10.1007/s00229-007-0134-y.

[Aur97] D. Auroux. "Asymptotically holomorphic families of symplectic submanifolds". In: *Geom. Funct. Anal.* 7.6 (1997), pp. 971–995. DOI: 10.1007/s000390050033.

[Ban78] A. Banyaga. "Sur la structure du groupe des difféomorphismes qui preservent une forme symplectique". In: *Comment. Math. Helv.* 53 (1978), pp. 174–227. DOI: 10.1007/BF02566074.

[BO12] M. Bhupal and K. Ono. "Symplectic fillings of links of quotient surface singularities". In: *Nagoya Math. J.* 207 (2012), pp. 1–45.

[Bir96] P. Biran. "Connectedness of spaces of symplectic embeddings". In: *Int. Math. Res. Not.* 1996.10 (1996), pp. 487–491. DOI: 10.1155/S1073792896000323.

[Bir97] P. Biran. "Symplectic packing in dimension 4". In: *Geom. Funct. Anal.* 7.3 (1997), pp. 420–437. DOI: 10.1007/s000390050014.

[Bir01] P. Biran. "Lagrangian barriers and symplectic embeddings". In: *Geom. Funct. Anal.* 11.3 (2001), pp. 407–464. DOI: 10.1007/PL00001678.

[BLW14] M. S. Borman, T.-J. Li, and W. Wu. "Spherical Lagrangians via ball packings and symplectic cutting". In: *Sel. Math., New Ser.* 20.1 (2014), pp. 261–283. DOI: 10.1007/s00029-013-0120-z.

[Bou+03] F. Bourgeois, Y. Eliashberg, H. Hofer, K. Wysocki, and E. Zehnder. "Compactness results in symplectic field theory". In: *Geom. Topol.* 7 (2003), pp. 799–888. DOI: 10.2140/gt.2003.7.799.

[BHS24] J. Brendel, J. Hauber, and J. Schmitz. *Semi-Local Exotic Lagrangian Tori in Dimension Four*. Preprint, arXiv:2403.00408 [math.SG] (2024). 2024.

[BS24] J. Brendel and F. Schlenk. "Pinwheels as Lagrangian barriers". In: *Commun. Contemp. Math.* 26.5 (2024). Id/No 2350020. DOI: 10.1142/S0219199723500207.

[Bri90] Encyclopaedia Britannica. ""facet"". In: *Micropedia*. 15th. Vol. 4. Encyclopaedia Britannica, Inc., 1990.

[Buc25] M. R. Buck. *Lagrangian Spheres and Cyclic Quotient T-singularities*. 2025. arXiv: 2509.18976 [math.SG].

[CZ24] H. Chen and W. Zhang. "Kodaira dimensions of almost complex manifolds II". In: *Comm. Anal. Geom.* 32.3 (2024), pp. 751–790. DOI: 10.4310/cag.241015030334.

[CM05] K. Cieliebak and K. Mohnke. “Compactness for punctured holomorphic curves”. In: *J. Symplectic Geom.* 3.4 (2005), pp. 589–654. DOI: 10.4310/JSG.2005.v3.n4.a5.

[CMM25] D. Cristofaro-Gardiner, N. Magill, and D. McDuff. *Curvy points, the perimeter, and the complexity of convex toric domains*. Preprint, arXiv:2506.23498 [math.SG]. 2025.

[DR23] Cristofaro-Gardiner D. and Hind R. “Boundaries of open symplectic manifolds and the failure of packing stability”. In: *arXiv:2307.01140* (2023).

[Dim19] G. Dimitroglou Rizell. “The classification of Lagrangians nearby the Whitney immersion”. In: *Geometry & Topology* 23 (2019), pp. 3367–3458. DOI: 10.2140/gt.2019.23.3367.

[DGI16] G. Dimitroglou Rizell, E. Goodman, and A. Ivrii. “Lagrangian isotopy of tori in $S^2 \times S^2$ and $\mathbb{C}P^2$”. In: *Geom. Funct. Anal.* 26.5 (2016), pp. 1297–1358. DOI: 10.1007/s00039-016-0388-1.

[EGH00] Y. Eliashberg, A. Givental, and H. Hofer. “Introduction to Symplectic Field Theory”. In: *GAFA 2000. Visions in mathematics—Towards 2000. Proceedings of a meeting, Tel Aviv, Israel, August 25–September 3, 1999. Part II.* Basel: Birkhäuser, 2000, pp. 560–673.

[EP96] Y. Eliashberg and L. Polterovich. “Local Lagrangian 2-knots are trivial”. In: *Ann. Math. (2)* 144.1 (1996), pp. 61–76. DOI: 10.2307/2118583.

[Etn+23] J. B. Etnyre, H. Min, L. Piccirillo, and A. Roy. *Small symplectic caps and embeddings of homology balls in the complex projective plane*. 2023. arXiv: 2305.16207 [math.GT].

[Eva11] J. Evans. “Symplectic mapping class groups of some Stein and rational surfaces”. In: *J. Symplectic Geom.* 9.1 (2011), pp. 45–82. DOI: 10.4310/JSG.2011.v9.n1.a4.

[Eva23] J. Evans. *Lectures on Lagrangian Torus Fibrations*. London Mathematical Society Student Texts. Cambridge University Press, 2023. DOI: 10.1017/9781009372671.

[Eva24] J. Evans. *KIAS Lectures on Symplectic Aspects of Degenerations*. arXiv:2403.03519 [math.SG]. 2024.

[ES18] J. Evans and I. Smith. “Markov numbers and Lagrangian cell complexes in the complex projective plane”. In: *Geom. Topol.* 22.2 (2018), pp. 1143–1180. DOI: 10.2140/gt.2018.22.1143.

[ES20] J. Evans and I. Smith. “Bounds on Wahl singularities from symplectic topology”. In: *Algebr. Geom.* 7.1 (2020), pp. 59–85. DOI: 10.14231/AG-2020-003.

[Fer16] R. Ferreira de Velloso Vianna. “Infinitely many exotic monotone Lagrangian tori in $\mathbb{CP}^2$”. In: *J. Topol.* 9.2 (2016), pp. 535–551. DOI: 10.1112/jtopol/jtw002.

[FS97] R. Fintushel and R. J. Stern. “Rational blowdowns of smooth 4-manifolds”. In: *J. Differential Geom.* 46.2 (1997), pp. 181–235.

[FSS08] K. Fukaya, P. Seidel, and I. Smith. “Exact Lagrangian submanifolds in simply-connected cotangent bundles”. In: *Invent. Math.* 172.1 (2008), pp. 1–27. DOI: 10.1007/s00222-007-0092-8.

[Ful93] W. Fulton. *Introduction to toric varieties*. Vol. 131. Annals of Mathematics Studies. The William H. Roever Lectures in Geometry. Princeton University Press, Princeton, NJ, 1993, pp. xii+157. DOI: 10.1515/9781400882526.

[GO25] M. Golla and B. Owens. *The Farey tree and embeddings of lens spaces and rational balls in* $\mathbb{CP}^2$. Preprint, arXiv:2512.09183 [math.GT]. 2025.

[Got82] Mark J. Gotay. “On coisotropic imbeddings of presymplectic manifolds”. In: *Proc. Amer. Math. Soc.* 84.1 (1982), pp. 111–114. DOI: 10.2307/2043821.

[GV23] Y. Groman and U. Varolgunes. “Locality of relative symplectic cohomology for complete embeddings”. In: *Compos. Math.* 159.12 (2023), pp. 2551–2637. DOI: 10.1112/S0010437X23007492.

[Gro85] M. Gromov. “Pseudo holomorphic curves in symplectic manifolds”. In: *Inventiones mathematicae* 82 (1985), pp. 307–347.

[HP10] P. Hacking and Y. Prokhorov. “Smoothable del Pezzo surfaces with quotient singularities”. In: *Compos. Math.* 146.1 (2010). DOI: 10.1112/S0010437X09004370.

[HTU17] P. Hacking, J. Tevelev, and G. Urzúa. “Flipping surfaces”. In: *J. Algebr. Geom.* 26.2 (2017), pp. 279–345. DOI: 10.1090/jag/682.

[HKW24] J. Hausen, K. Király, and M. Wrobel. “On degenerations of the projective plane”. In: *arXiv:2405.04862* (2024).

[Hin04] R. Hind. “Lagrangian spheres in $S^2 \times S^2$”. In: *Geom. Funct. Anal.* 14.2 (2004), pp. 303–318. DOI: 10.1007/s00039-004-0459-6.

[Hin12] R. Hind. “Lagrangian unknottedness in Stein surfaces”. In: *Asian J. Math.* 16.1 (2012), pp. 1–36. DOI: 10.4310/AJM.2012.v16.n1.a1.

[HPW16] R. Hind, M. Pinsonnault, and W. Wu. “Symplectomophism groups of non-compact manifolds, orbifold balls, and a space of Lagrangians”. In: *J. Symplectic Geom.* 14(1) (2016), pp. 203–226.

[HLS97] H. Hofer, V. Lizan, and J.-C. Sikorav. “On genericity for holomorphic curves in four-dimensional almost-complex manifolds”. In: *J. Geom. Anal.* 7.1 (1997), pp. 149–159. DOI: 10.1007/BF02921708.

[Kar18] D. Karabas. *Microlocal Sheaves on Pinwheels*. 2018. arXiv: 1810.09021 [math.SG].

[KSW26] A. Keating, I. Smith, and M. Wemyss. *Splitting symplectic monodromy*. 2026. arXiv: 2601.20438 [math.SG].

[Kho13a] T. Khodorovskiy. *Bounds on Embeddings of Rational Homology Balls in Symplectic 4-manifolds*. 2013. arXiv: 1307.4321 [math.SG].

[Kho13b] T. Khodorovskiy. *Symplectic Rational Blow-up*. 2013. arXiv: 1303.2581 [math.SG].

[KS88] J. Kollár and N. I. Shepherd-Barron. "Threefolds and deformations of surface singularities". In: *Invent. Math.* 91.2 (1988), pp. 299–338. DOI: 10.1007/BF01389370.

[Kra13] T. Kragh. "Parametrized ring-spectra and the nearby Lagrangian conjecture". In: *Geom. Topol.* 17.2 (2013), pp. 639–731. DOI: 10.2140/gt.2013.17.639.

[LM14] Y. Lekili and M. Maydanskiy. "The symplectic topology of some rational homology balls". In: *Comment. Math. Helv.* 89.3 (2014), pp. 571–596. DOI: 10.4171/CMH/327.

[Ler95] E. Lerman. "Symplectic cuts". In: *Math. Res. Lett.* 2.3 (1995), pp. 247–258. DOI: 10.4310/MRL.1995.v2.n3.a2.

[LLW22] J. Li, T.-J. Li, and W. Wu. "The space of tamed almost complex structures on symplectic 4-manifolds via symplectic spheres". In: *Riv. Math. Univ. Parma (N.S.)* 13.2 (2022), pp. 651–670.

[LW12] T.-J. Li and W. Wu. "Lagrangian spheres, symplectic surfaces and the symplectic mapping class group". In: *Geom. Topol.* 16.2 (2012), pp. 1121–1169. DOI: 10.2140/gt.2012.16.1121.

[LZ15] T.-J. Li and W. Zhang. "J-holomorphic curves in a nef class". In: *Int. Math. Res. Not. IMRN* 22 (2015), pp. 12070–12104. DOI: 10.1093/imrn/rnv056.

[Lis08] P. Lisca. "On symplectic fillings of lens spaces". In: *Trans. Amer. Math. Soc.* 360.2 (2008), pp. 765–799. DOI: 10.1090/S0002-9947-07-04228-6.

[LP22] P. Lisca and A. Parma. "On Stein rational balls smoothly but not symplectically embedded in $\mathbb{CP}^2$". In: *Bull. Lond. Math. Soc.* 54.3 (2022), pp. 949–960. DOI: 10.1112/blms.12607.

[LP24] P. Lisca and A. Parma. "On almost complex embeddings of rational homology balls". In: *Frontiers in geometry and topology. Summer school and research conference, The Abdus Salam International Centre for Theoretical Physics, Trieste, Italy, August 1–12, 2022*. Providence, RI: American Mathematical Society (AMS), 2024, pp. 183–193. DOI: 10.1090/pspum/109/01995.

[Man91] M. Manetti. "Normal degenerations of the complex projective plane". In: *J. Reine Angew. Math.* 419 (1991), pp. 89–118. DOI: 10.1515/crll.1991.419.89.

[McD91] D. McDuff. "Blow ups and symplectic embeddings in dimension 4". In: *Topology* 30.3 (1991), pp. 409–421. DOI: 10.1016/0040-9383(91)90021-U.

[McD93] D. McDuff. "Remarks on the uniqueness of symplectic blowing up". In: *Symplectic geometry. Based on lectures given at a workshop and conference held at the University of Warwick, Coventry, UK in August 1990*. Cambridge: Cambridge University Press, 1993, pp. 157–167.

[McD97] D. McDuff. "Lectures on Gromov invariants for symplectic 4-manifolds". In: *Gauge theory and symplectic geometry. Proceedings of the NATO Advanced Study Institute and séminaire de mathématiques supérieures, Montréal, Canada, July 3–14, 1995*. Dordrecht: Kluwer Academic Publishers, 1997, pp. 175–210.

[MO15] D. McDuff and E. Opshtein. “Nongeneric J-holomorphic curves and singular inflation”. In: *Algebr. Geom. Topol.* 15.1 (2015), pp. 231–286. DOI: 10.2140/agt.2015.15.231.

[MP94] D. McDuff and L. Polterovich. “Symplectic packings and algebraic geometry”. In: *Invent. Math.* 115.3 (1994), pp. 405–429. DOI: 10.1007/BF01231766.

[MS04] D. McDuff and D. Salamon. *J-holomorphic curves and symplectic topology*. Vol. 52. American Mathematical Society Colloquium Publications. American Mathematical Society, Providence, RI, 2004, pp. xii+669. DOI: 10.1090/coll/052.

[MS16] D. McDuff and D. Salamon. *Introduction to symplectic topology*. 3rd edition. Vol. 27. Oxf. Grad. Texts Math. Oxford: Oxford University Press, 2016.

[MS12] D. McDuff and F. Schlenk. “The embedding capacity of 4-dimensional symplectic ellipsoids”. In: *Ann. of Math. (2)* 175.3 (2012), pp. 1191–1282. DOI: 10.4007/annals.2012.175.3.5.

[MS25a] D. McDuff and K. Siegel. “Ellipsoidal superpotentials and singular curve counts”. In: *Int. Math. Res. Not.* 2025.21 (2025). Id/No rnaf285, p. 37. DOI: 10.1093/imrn/rnaf285.

[MS25b] D. McDuff and K. Siegel. “Singular algebraic curves and infinite symplectic staircases”. In: *Invent. Math.* 242.2 (2025), pp. 387–459. DOI: 10.1007/s00222-025-01359-4.

[McL12] M. McLean. “The growth rate of symplectic homology and affine varieties”. In: *Geom. Funct. Anal.* 22.2 (2012), pp. 369–442. DOI: 10.1007/s00039-012-0158-7.

[MS23] G. Mikhalkin and M. Shkolnikov. “Wave fronts and caustics in the tropical plane”. In: *arXiv:2310.17269* (2023).

[OU16] J. Oakley and M. Usher. “On certain Lagrangian submanifolds of $S^2 \times S^2$ and $\mathbb{C}P^n$”. In: *Algebr. Geom. Topol.* 16.1 (2016), pp. 149–209. DOI: 10.2140/agt.2016.16.149.

[Owe18] B. Owens. “Equivariant embeddings of rational homology balls”. In: *Q. J. Math.* 69.3 (2018), pp. 1101–1121. DOI: 10.1093/qmath/hay016.

[Owe20] B. Owens. “Smooth, nonsymplectic embeddings of rational balls in the complex projective plane”. In: *Q. J. Math.* 71.3 (2020), pp. 997–1007. DOI: 10.1093/qmathj/haaa013.

[Par+18] H. Park, J. Park, D. Shin, and G. Urzúa. “Milnor fibers and symplectic fillings of quotient surface singularities”. In: *Adv. Math.* 329 (2018), pp. 1156–1230. DOI: 10.1016/j.aim.2018.03.002.

[PS24] L. Polterovich and F. Schlenk. *Lagrangian knots and unknots – an essay.* With an appendix by G. Dimitroglou Rizell. 2024. arXiv: 2406.15967 [math.SG].

[Sch18] F. Schlenk. “Symplectic embedding problems, old and new”. In: *Bull. Amer. Math. Soc. (N.S.)* 55.2 (2018), pp. 139–182. DOI: 10.1090/bull/1587.

[Sei99] P. Seidel. “Lagrangian two-spheres can be symplectically knotted”. In: *J. Differ. Geom.* 52.1 (1999), pp. 145–171. DOI: 10.4310/jdg/1214425219.

[Sei00] P. Seidel. "Graded Lagrangian submanifolds". In: *Bull. Soc. Math. Fr.* 128.1 (2000), pp. 103–149. DOI: 10.24033/bsmf.2365.

[Sei] P. Seidel. "Floer homology and the symplectic isotopy problem". Ph.D. thesis, Oxford University, 1997.

[SS20] V. Shevchishin and G. Smirnov. "Symplectic triangle inequality". In: *Proc. Am. Math. Soc.* 148.4 (2020), pp. 1389–1397. DOI: 10.1090/proc/14842.

[ST05] B. Siebert and G. Tian. "On the holomorphicity of genus two Lefschetz fibrations". In: *Ann. of Math. (2)* 161.2 (2005), pp. 959–1020. DOI: 10.4007/annals.2005.161.959.

[Sik97] J.-C. Sikorav. "Singularities of J-holomorphic curves". In: *Math. Z.* 226.3 (1997), pp. 359–373. DOI: 10.1007/PL00004346.

[Sym98] M. Symington. "Symplectic rational blowdowns". In: *J. Differ. Geom.* 50.3 (1998), pp. 505–518. DOI: 10.4310/jdg/1214424968.

[Sym01] M. Symington. "Generalized symplectic rational blowdowns". In: *Algebr. Geom. Topol.* 1 (2001), pp. 503–518.

[Sym03] M. Symington. "Four dimensions from two in symplectic topology". In: *Topology and geometry of manifolds (Athens, GA, 2001)*. Vol. 71. Proc. Sympos. Pure Math. Amer. Math. Soc., Providence, RI, 2003, pp. 153–208. DOI: 10.1090/pspum/071/2024634.

[UZ25] G. Urzúa and J. P. Zúñiga. *The birational geometry of Markov numbers*. 2025. arXiv: 2310.17957 [math.AG].

[Via14] R. Vianna. "On exotic Lagrangian tori in $\mathbb{CP}^2$". In: *Geom. Topol.* 18.4 (2014), pp. 2419–2476. DOI: 10.2140/gt.2014.18.2419.

[Wen18] C. Wendl. *Holomorphic curves in low dimensions. From symplectic ruled surfaces to planar contact manifolds*. Vol. 2216. Lect. Notes Math. Berlin: Springer, 2018. DOI: 10.1007/978-3-319-91371-1.

[Wu14] W. Wu. "Exact Lagrangians in A_n-surface singularities". In: *Math. Ann.* 359.1-2 (2014), pp. 153–168. DOI: 10.1007/s00208-013-0993-3.

[Wu15] W. Wu. "On an exotic Lagrangian torus in $\mathbb{C}P^2$". In: *Compos. Math.* 151.7 (2015), pp. 1372–1394. DOI: 10.1112/S0010437X14007945.

[Zha21] W. Zhang. "Moduli space of J-holomorphic subvarieties". In: *Selecta Math. (N.S.)* 27.2 (2021), Paper No. 29, 44. DOI: 10.1007/s00029-021-00648-z.