Algebraic & Geometric Topology 15 (2015) 3323–3353

Equivalence classes of augmentations and Morse complex sequences of Legendrian knots

MICHAEL B HENRY
DAN RUTHERFORD

Let L be a Legendrian knot in $\mathbb{R}^3$ with the standard contact structure. In earlier work of Henry, a map was constructed from equivalence classes of Morse complex sequences for L, which are combinatorial objects motivated by generating families, to homotopy classes of augmentations of the Legendrian contact homology algebra of L. Moreover, this map was shown to be a surjection. We show that this correspondence is, in fact, a bijection. As a corollary, homotopic augmentations determine the same graded normal ruling of L and have isomorphic linearized contact homology groups. A second corollary states that the count of equivalence classes of Morse complex sequences of a Legendrian knot is a Legendrian isotopy invariant.

57R17; 57M25, 53D40

1 Introduction

The symplectic techniques of holomorphic curves and generating families provide two effective classes of invariants of Legendrian knots in standard contact $\mathbb{R}^3$. The holomorphic curve approach, which in this low-dimensional setting takes on a combinatorial flavor, can be used to define a differential graded algebra (DGA). The DGA is known alternatively as the Legendrian contact homology DGA or the Chekanov–Eliashberg DGA and was originally defined by Chekanov [1] and Eliashberg, Givental and Hofer [5]. Generating families of Legendrian submanifolds in 1–jet spaces, including $\mathbb{R}^3$, have also been used to produce homological Legendrian invariants; see, for instance, Jordan and Traynor [13], Sabloff and Traynor [18] and Traynor [19; 20]. In addition to distinguishing Legendrian isotopy classes of knots, both the holomorphic and generating family invariants carry useful information about Lagrangian cobordisms, see Ekholm, Honda and Kálmán [4] and Sabloff and Traynor [18].

For Legendrian knots in $\mathbb{R}^3$, several close connections have been discovered between holomorphic curve and generating family invariants, although many questions remain. For example, the existence of a linear at infinity generating family for a Legendrian knot is known to be equivalent to the existence of a certain DGA morphism, called an

Published: 12 January 2016 DOI: 10.2140/agt.2015.15.3323

augmentation, from the Chekanov–Eliashberg DGA to its ground ring; see Pushkar' and Chekanov [3], Fuchs [7], Fuchs and Ishkhanov [8], Fuchs and Rutherford [9] and Sabloff [17]. However, it is unknown if this statement can be strengthened to a bijective correspondence between appropriate equivalence classes of generating families and augmentations. In this article, we approach this question using a discrete analog of a generating family called a Morse complex sequence, abbreviated MCS. MCSs have proven to be more tractable for explicit construction and computation; see, for example, Henry [10] and Henry and Rutherford [11; 12]. Section 2.2 sketches the connection between generating families and Morse complex sequences; a more complete description can be found in [11].

The concept of a Morse complex sequence originally appeared in unpublished work of Petya Pushkar, and first appears in print in the work of the first author [10] where MCSs are studied in connection with augmentations. In [10], a surjective map is defined from MCSs of L to augmentations of the Chekanov–Eliashberg DGA of L. Moreover, equivalent MCSs are mapped to homotopic augmentations. In the present article, we complement the results of [10] by showing in Lemma 3.1 that two MCSs mapped to homotopic augmentations must, in fact, be equivalent as MCSs. Combined with [10] this gives the following.

Theorem 1.1 *For any Legendrian knot $L \subset \mathbb{R}^3$ with generic front diagram, there is a bijection between equivalence classes of Morse complex sequences for L and homotopy classes of augmentations of the Chekanov–Eliashberg DGA of L.*

As a consequence, the number of MCS equivalence classes is a Legendrian isotopy invariant; see Corollary 4.1. The less immediate Corollary 4.2 combines Theorem 1.1 with previous work of the authors from [11] to deduce that homotopic augmentations must have isomorphic linearized homology groups. The set of linearized homology groups is a Legendrian isotopy invariant. Corollary 4.2 allows for a refinement of this invariant by considering multiplicities.

The remainder of the article is organized as follows. Section 2 recalls background concerning augmentations and Morse complex sequences. Section 3 contains the proof of Theorem 1.1 and Section 4 includes three corollaries.

Acknowledgments We gratefully acknowledge The Royal Academy of Belgium where conversations between the authors on this project began in August 2013 at a workshop on Legendrian submanifolds, holomorphic curves, and generating families. We also thank the workshop organizer Frédéric Bourgeois. In addition, we thank the American Institute of Mathematics for supporting a SQuaRE research group on augmentations and related topics, and we thank our fellow SQuaRE participants Dmitry Fuchs, Paul Melvin, Josh Sabloff, and Lisa Traynor.

2 Background

A *Legendrian knot* in the standard contact structure on $\mathbb{R}^3$ is a smooth knot $L\colon S^1 \to \mathbb{R}^3$ satisfying $L'(t) \in \ker(dz - y\,dx)$ for all $t \in S^1$. A smooth one-parameter family L_t, $0 \leq t \leq 1$, of Legendrian knots is a *Legendrian isotopy* between L_0 and L_1. The *front diagram* of L is the projection of L to the xz–plane. Every Legendrian knot is Legendrian isotopic, by an arbitrarily small Legendrian isotopy, to a Legendrian knot whose front diagram is embedded except at transverse self-intersections, called *crossings*, and semi-cubical cusps such that, in addition, all of these exceptional points have distinct x–coordinates. A Legendrian knot with such a front diagram is said to have a *σ–generic* front diagram; see, for example, the front diagram in Figure 1. In a neighborhood of an x value that is not the x–coordinate of a crossing or cusp, the front diagram looks like a collection of non-intersecting line segments commonly called the *strands* of D at x. Orient L. The *rotation number* $r(L)$ is $(d-u)/2$, where d (resp. u) is the number of cusps at which the orientation travels downward (resp. upward) with respect to the z–axis.

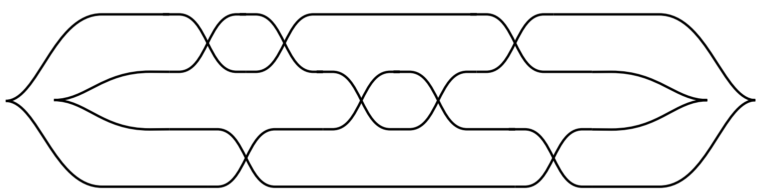

Figure 1: A σ–generic front diagram of a Legendrian knot with rotation number 0

2.1 Chekanov–Eliashberg algebra

Fix a Legendrian knot L with σ–generic front diagram D and rotation number 0. A *Maslov potential* is a map $\mu\colon L \to \mathbb{Z}$ that is constant except at cusp points of L where the Maslov potential of the lower strand of the cusp is one less than the upper strand. Let $A(D)$ be the $\mathbb{Z}/2\mathbb{Z}$ vector space generated by the labels $Q = \{q_1, \dots, q_n\}$ assigned to the crossings and right cusps of D. A generator $q \in Q$ is assigned a *grading* $|q|$, also called a *degree*, so that $|q|$ is 1 if q is a right cusp and, otherwise, $|q|$ is $\mu(T) - \mu(B)$ where T and B are the strands of D crossing at q and T has smaller slope. The graded algebra $\mathcal{A}(D)$ is the unital tensor algebra $TA(D)$. The *Chekanov–Eliashberg algebra*, written $(\mathcal{A}(D), \partial)$, is the algebra $\mathcal{A}(D)$ along with a degree -1 differential $\partial\colon \mathcal{A}(D) \to \mathcal{A}(D)$ that, in the case of the front diagram description from [15], is defined by counting certain admissible maps of the two-disk D^2 into the xz–plane.

Definition 2.1 below defines only those admissible maps needed in this article; we refer the reader to [15] for a complete definition of ∂.

An *augmentation* is an algebra homomorphism $\epsilon\colon \mathcal{A}(D) \to \mathbb{Z}/2\mathbb{Z}$ satisfying $\epsilon \circ \partial = 0$, $\epsilon(1) = 1$, and $\epsilon(q) = 1$ only if $|q| = 0$. The set $\mathrm{Aug}(D)$ is the set of all augmentations of $(\mathcal{A}(D), \partial)$. We say a crossing q is *augmented* by ϵ if $\epsilon(q) = 1$. An augmentation can be thought of as a morphism between the differential graded algebra $(\mathcal{A}(D), \partial)$ and the differential graded algebra $(\mathbb{Z}/2\mathbb{Z}, \partial')$ whose only non-zero element is in degree 0 and where $\partial' = 0$. From this perspective, there is a natural algebraic equivalence relation on $\mathrm{Aug}(D)$. Given ϵ and ϵ' in $\mathrm{Aug}(D)$, a *chain homotopy* from ϵ to ϵ' is a degree 1 linear map $H\colon (\mathcal{A}(D), \partial) \to (\mathbb{Z}/2\mathbb{Z}, \partial')$ satisfying $\epsilon - \epsilon' = \partial' \circ H + H \circ \partial$ and $H(ab) = H(a)\epsilon'(b) + (-1)^{|a|}\epsilon(a)H(b)$ for all $a, b \in \mathcal{A}(D)$. Since we are working over $\mathbb{Z}/2\mathbb{Z}$ and $\partial' = 0$, these conditions simplify to

$$\epsilon - \epsilon' = H \circ \partial \quad \text{and} \quad H(ab) = H(a)\epsilon'(b) + \epsilon(a)H(b). \tag{1}$$

By [14, Lemma 2.18], a chain homotopy H is determined by the values it takes on the degree -1 crossings of D.

We say augmentations ϵ and ϵ' are *homotopic* and write $\epsilon \simeq \epsilon'$ if there exists a chain homotopy from ϵ to ϵ'. As the notation implies and as is proven in [6], chain homotopy provides an equivalence relation on the set $\mathrm{Aug}(D)$. We let $\mathrm{Aug}^{\mathrm{ch}}(D)$ be $\mathrm{Aug}(D)/\simeq$. By [10, Proposition 4.5], the count of homotopy classes of augmentations is a Legendrian isotopy invariant.

Suppose ϵ and ϵ' are augmentations in $\mathrm{Aug}(D)$ and there exists a chain homotopy H from ϵ to ϵ'. Suppose q is a degree 0 crossing and $\langle \partial q, \prod_{i=1}^m q_{k_i} \rangle$ is 1, where $\langle \partial q, \prod_{i=1}^m q_{k_i} \rangle$ is the coefficient of $\prod_{i=1}^m q_{k_i}$ in ∂q. Then, by Equation (1),

$$\begin{aligned}
(\epsilon - \epsilon')(q) = H \circ \partial(q) &= H\Big(\prod_{i=1}^m q_{k_i} + \cdots\Big) \\
&= H\Big(\prod_{i=1}^m q_{k_i}\Big) + H(\cdots) \\
&= \sum_{j=1}^m \Big[\Big(\prod_{i=1}^{j-1} \epsilon(q_{k_i})\Big) H(q_{k_j}) \Big(\prod_{i=j+1}^m \epsilon'(q_{k_i})\Big)\Big] + H(\cdots).
\end{aligned}$$

At most one term in the sum

$$\sum_{j=1}^m \Big[\Big(\prod_{i=1}^{j-1} \epsilon(q_{k_i})\Big) H(q_{k_j}) \Big(\prod_{i=j+1}^m \epsilon'(q_{k_i})\Big)\Big]$$

may be non-zero, since ϵ and ϵ' are non-zero only on generators of degree 0 and H is non-zero only on generators of degree -1. Note that, for a fixed $j \in \{1, \dots, m\}$, the term

$$\Big(\prod_{i=1}^{j-1} \epsilon(q_{k_i})\Big) H(q_{k_j}) \Big(\prod_{i=j+1}^{m} \epsilon'(q_{k_i})\Big)$$

is non-zero if and only if $H(q_{k_j}) = 1$ holds and for $1 \leq i \leq j-1$ (resp. $j+1 \leq i \leq m$), the crossing q_{k_i} is augmented by ϵ (resp. ϵ').

The monomials $\prod_{i=1}^{m} q_{k_i}$ appearing in $\partial(q)$ correspond to certain mappings of the two-disk D^2 into the xz–plane that are immersions except for allowable exceptions along ∂D^2. Only monomials containing generators of degree 0 or -1 are relevant for our purposes. Therefore, we present only the description of such disks in the following definitions. Note that this restriction allows us to rule out some additional behaviors of ∂D^2 near right cusps that appear in [15] and lead to monomials that contain generators of degree 1.

Let D^2 be the disk of radius 1 centered at the origin in $\mathbb{R}^2$. Choose m points from $\partial D^2 \setminus \{(1,0)\}$. Label the chosen points $\{b_1, \dots, b_m\}$ counter-clockwise with b_1 the first point counter-clockwise from $(1,0)$.

Definition 2.1 In terms of the notation above, a $(0,-1)$*–admissible disk* is a continuous map from D^2 into the xz–plane that maps ∂D^2 to the front diagram D and is a smooth orientation preserving immersion when restricted to the interior of D^2 satisfying the following conditions:

1. The mapping takes $(1,0)$ to a degree 0 crossing q and the image of f in a neighborhood of $(1,0)$ looks as in Figure 2(a). We say the $(0,-1)$–admissible disk *originates at* q.
2. For exactly one $1 \leq j \leq m$, $f(b_j)$ is a degree -1 crossing q_{k_j} and the image of f in a neighborhood of b_j looks as in Figure 2(d) or (e).
3. For all $i \neq j$, $f(b_i)$ is a degree 0 crossing q_{k_i} and the image of f in a neighborhood of b_i looks as in Figure 2(d) or (e).
4. Along ∂D^2 the mapping is smooth except at $\{b_1, \dots, b_m\} \cup \{(1,0)\}$ as described in (1)–(3) and at points in $\partial D^2 \setminus (\{b_1, \dots, b_m\} \cup \{(1,0)\})$ where the image of f looks like either Figure 2(b) or (c).

We say the $(0,-1)$–admissible disk has *convex corners* at $q_{k_1}, \dots, q_{k_m}$. The $(0,-1)$–admissible disk is assigned the monomial $\prod_{i=1}^{m} q_{k_i}$. We say a $(0,-1)$–admissible disk is an (ϵ, ϵ', H)*–admissible disk* if, for some $1 \leq j \leq m$, $H(q_{k_j}) = 1$ holds and

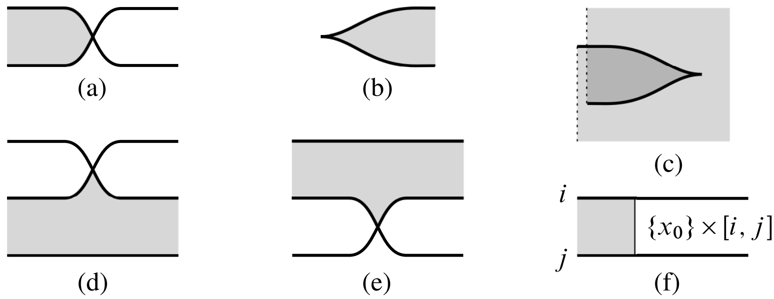

Figure 2: The possible singularities of the disk in Definition 2.1 and the half-disks in Definitions 3.3 and 3.2. The crossings in (d) and (e) are called convex corners. Near a boundary point that maps to a right cusp the image of a disk overlaps itself as indicated in (c) by the darkly shaded region.

for $1 \leq i \leq j-1$ (resp. $j+1 \leq i \leq m$), the crossing q_{k_i} is augmented by ϵ (resp. ϵ'); see Figure 3.

Henceforth, we consider admissible disks up to orientation preserving reparametrization of the domain (fixing $\{b_1, \ldots, b_m\} \cup \{(1,0)\}$), and all counts of disks are up to this equivalence relation.

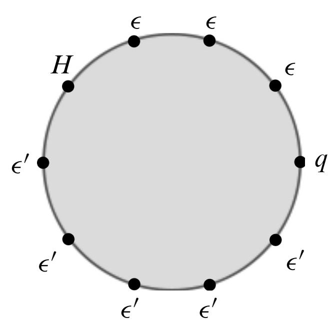

Figure 3: The domain of an (ϵ, ϵ', H)–admissible disk with labels indicating marked points mapped to crossings augmented by ϵ and ϵ' and the marked point mapped to the crossing satisfying $H(q_{k_j}) = 1$

The restrictions on the types of non-smooth points of an $(0,-1)$–admissible disk imply that q is the right-most point of the disk. From [15, Section 2], when a single q_{k_j} has degree -1 while q and all of the remaining q_{k_i} have degree 0, $\langle \partial q, \prod_{i=1}^m q_{k_i} \rangle = 1$ holds if and only if there are an odd number of $(0,-1)$–admissible disks originating at q and with monomial $\prod_{i=1}^m q_{k_i}$. Proposition 2.2 follows directly from the discussion above.

Proposition 2.2 *Suppose D is a σ–generic front diagram of a Legendrian knot and ϵ and ϵ' are augmentations in* $\mathrm{Aug}(D)$*. If q is a degree 0 crossing and H is a chain homotopy from ϵ to ϵ', then ϵ and ϵ' differ at q if and only if there are an odd number of (ϵ, ϵ', H)–admissible disks originating at q.*

2.2 Morse complex sequences

We briefly sketch the connection between generating families and Morse complex sequences and refer the reader to [11] for more details. A one-parameter family of smooth functions $f_x\colon \mathbb{R}^N \to \mathbb{R}$, parametrized by $x \in \mathbb{R}$, is a generating family for a Legendrian knot L with front diagram D if

$$D = \left\{(x, z) : z = f_x(e) \text{ for some } e \in \mathbb{R}^N \text{ satisfying } \frac{\partial f_x}{\partial e}(e) = 0\right\}.$$

With an appropriately chosen metric, a generic $x \in \mathbb{R}$ determines a Morse chain complex (C_x, d_x) on $\mathbb{R}^N$ and, as x varies, the evolution of the Morse complexes of f_x are well-understood; a cusp of D corresponds to the creation or elimination of a canceling pair of critical points and a crossing corresponds to two critical points exchanging critical values. As x varies, it is also possible for a fiberwise gradient flowline to momentarily flow between two critical points of the same index. Such an occurrence is called a handleslide and it determines an explicit chain isomorphism between successive Morse complexes. In summary, a generating family and choice of metric determine a one-parameter family of Morse chain complexes and the relationship between successive chain complexes is determined by the crossings and cusps of D and the handleslides. A Morse complex sequence on D is a finite sequence of chain complexes (C_m, d_m) and vertical marks on D that are meant to correspond to the Morse chain complexes and handleslides of a generating family and choice of metric. In addition, varying the choice of metric motivates an equivalence relation on MCSs.

Fix a Legendrian knot L with σ–generic front diagram D, rotation number 0, and Maslov potential μ. Theorem 1.1 proves that a certain surjective map in [10] from equivalence classes of Morse complex sequences to $\mathrm{Aug}^{\mathrm{ch}}(D)$ is, in fact, a bijection. We will use the definition of a Morse complex sequence given in [11]. This definition differs slightly from the definition in [10], but both definitions determine the same set of objects on L.

A *handleslide* on D is a vertical line segment disjoint from all crossings and cusps and with endpoints on strands of D that have the same Maslov potential.

Definition 2.3 A *Morse complex sequence* on a σ–generic front diagram D is the triple $\mathcal{C} = (\{(C_m, d_m)\}, \{x_m\}, H)$ satisfying:

(1) H is a set of handleslides on D.

(2) The real values $x_1 < x_2 < \cdots < x_M$ are x–coordinates distinct from the x–coordinates of crossings and cusps of D and handleslides of H. For each $1 \le m < M$, the set $\{(x,z) : x_m \le x \le x_{m+1}\}$ contains a single crossing, cusp or handleslide. The set $\{(x,z) : -\infty < x \le x_1\}$ contains the left-most left cusp and the set $\{(x,z) : x_M \le x < \infty\}$ contains the right-most right cusp.

(3) For each $1 \le m \le M$, the points of intersection of the vertical line $\{x_m\} \times \mathbb{R}$ and D are labeled $e_1, e_2, \dots, e_{s_m}$ from top to bottom. The vector space C_m is the $\mathbb{Z}$–graded $\mathbb{Z}/2\mathbb{Z}$ vector space generated by $e_1, e_2, \dots, e_{s_m}$, where the degree of each generator is the value of the Maslov potential on the corresponding strand of D, $|e_i| = \mu(e_i)$. The map $d_m \colon C_m \to C_m$ is a degree -1 differential that is triangular in the sense that

$$d_m e_i = \sum_{i<j} c_{ij} e_j, \quad c_{ij} \in \mathbb{Z}/2\mathbb{Z}.$$

(4) The coefficients $\langle d_1 e_1, e_2 \rangle$ and $\langle d_M e_1, e_2 \rangle$ are both 1. Suppose $1 \le m < M$ and let T be the tangle $D \cap \{(x,z) : x_m \le x \le x_{m+1}\}$. If T contains a left (resp. right) cusp between strands k and $k+1$, then $\langle d_{m+1} e_k, e_{k+1} \rangle$ is 1 (resp. $\langle d_m e_k, e_{k+1} \rangle$ is 1). If T contains a crossing between strands k and $k+1$, then $\langle d_m e_k, e_{k+1} \rangle$ is 0.

(5) For $1 \le m < M$, the crossing, cusp, or handleslide mark in the tangle $T = D \cap \{(x,z) : x_m \le x \le x_{m+1}\}$ determines an algebraic relationship between the chain complexes (C_m, d_m) and (C_{m+1}, d_{m+1}) as follows:

   (a) **Crossing** If the crossing is between strands k and $k+1$, then the map $\phi\colon (C_m, d_m) \to (C_{m+1}, d_{m+1})$ defined by

$$\phi(e_i) = \begin{cases} e_i & \text{if } i \notin \{k, k+1\}, \\ e_{k+1} & \text{if } i = k, \\ e_k & \text{if } i = k+1 \end{cases}$$

   is an isomorphism of chain complexes.

   (b) **Right cusp** If the right cusp is between strands k and $k+1$, then the linear map

$$\phi(e_i) = \begin{cases} [e_i] & \text{if } i < k, \\ [e_{i+2}] & \text{if } i \ge k \end{cases}$$

   is an isomorphism of chain complexes from (C_{m+1}, d_{m+1}) to the quotient of (C_m, d_m) by the acyclic subcomplex generated by $\{e_k, d_m e_k\}$.

   (c) **Left cusp** The case of a left cusp is the same as the case of a right cusp, though the roles of (C_m, d_m) and (C_{m+1}, d_{m+1}) are reversed.

(d) **Handleslide** If the handleslide mark has endpoints on strands k and l with $k < l$, then the map $h_{k,l}\colon (C_m, d_m) \to (C_{m+1}, d_{m+1})$ defined by

$$h_{k,l}(e_i) = \begin{cases} e_i & \text{if } i \neq k, \\ e_k + e_l & \text{if } i = k \end{cases}$$

is an isomorphism of chain complexes.

The set $\mathrm{MCS}(D)$ is the set of all Morse complex sequences on D.

Remark 2.4 Morse complex sequences may be defined over more general coefficient rings than $\mathbb{Z}/2\mathbb{Z}$; see [12]. We restrict attention to $\mathbb{Z}/2\mathbb{Z}$ coefficients as this is also done in [10].

Definition 2.5 An MCS $\mathcal{C} = (\{(C_m, d_m)\}, \{x_m\}, H)$ in $\mathrm{MCS}(D)$ has *simple left cusps* if, for each tangle $T = \{(x,z) : x_m \leq x \leq x_{m+1}\}$ containing a left cusp between strands k and $k+1$, the chain complex (C_{m+1}, d_{m+1}) satisfies $\langle d_{m+1} e_k, e_i \rangle = \langle d_{m+1} e_{k+1}, e_i \rangle = 0$ for all $k+1 < i$ and $\langle d_{m+1} e_j, e_{k+1} \rangle = \langle d_{m+1} e_j, e_k \rangle = 0$ for all $j < k$.

The subset $\mathrm{MCS}_b(D) \subset \mathrm{MCS}(D)$ denotes the set of MCSs with simple left cusps. We use the letter b to be consistent with the notation of [10], where a left cusp is also called a "birth". This language is meant to draw a connection to the creation of a canceling pair of critical points, often called a birth, in a one-parameter family of Morse functions on a manifold.

Given an MCS $\mathcal{C} = (\{(C_m, d_m)\}, \{x_m\}, H)$ with simple left cusps, the chain complexes $\{(C_m, d_m)\}$ are uniquely determined by the crossings and cusps of D, the handleslides H, and requirements (5) (a)–(d) of Definition 2.3. Consequently, $\mathcal{C}$ may be represented visually by placing the handleslide marks H on the front diagram D; see Figure 4.

In [10] an equivalence relation on the set $\mathrm{MCS}(D)$ is defined that is motivated by a corresponding equivalence for generating families; see also [11]. Here we denote the set of equivalence classes of this relation by $\widehat{\mathrm{MCS}}(D) = \mathrm{MCS}(D)/\simeq$. We recall a version of this equivalence relation that applies to the more restricted set of MCSs with simple left cusps, $\mathrm{MCS}_b(D)$. We denote equivalence classes with respect to this relation by $\widehat{\mathrm{MCS}}_b(D)$. By [10, Proposition 3.17], the map from $\widehat{\mathrm{MCS}}_b(D)$ to $\widehat{\mathrm{MCS}}(D)$ induced by the inclusion $\mathrm{MCS}_b(D) \subset \mathrm{MCS}(D)$ is a bijection. Therefore, to prove Theorem 1.1, we need only consider MCSs in $\mathrm{MCS}_b(D)$ and MCS classes in $\widehat{\mathrm{MCS}}_b(D)$.

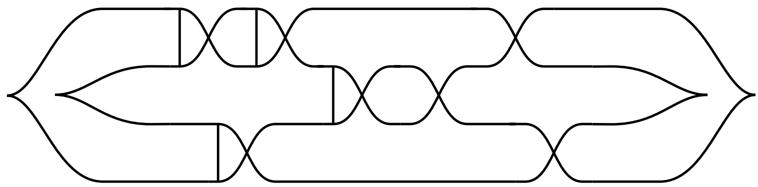

Figure 4: An MCS with simple left cusps. This MCS is also in A–form.

The equivalence relation on $\mathrm{MCS}_b(D)$ is generated by the *MCS moves* pictured in Figures 5 and 6. The numbering indicated will be used throughout this article. Additional moves result from reflecting each of the two figures in (3), (7), (9), (10) and (12) of Figure 5 about a horizontal axis and reflecting each of the two figures in (4), (9), (11) and (12) of Figure 5 about a vertical axis. The handleslide modification that results from reflecting Figure 5 (10) about a vertical axis is *not* an MCS move for MCSs with simple left cusps. (The absence of this reflected move is the only difference between the definitions of the equivalence relations on $\mathrm{MCS}_b(D)$ and $\mathrm{MCS}(D)$ discussed in the previous paragraph.) MCS move (13) requires explanation. Suppose $\mathcal{C} = (\{(C_m, d_m)\}, \{x_m\}, H)$ is an MCS on D and suppose there exist x_m and $1 \le k < l \le s_m$ such that $\mu(e_k) = \mu(e_l) - 1$. Then MCS move (13) introduces the collection of handleslides K defined as follows. The handleslides in K are of two types. First, if $i < k$ and $\langle d_m e_i, e_k \rangle = 1$ holds, then K contains a handleslide with endpoints on i and l. Second, if $l < j$ and $\langle d_m e_l, e_j \rangle = 1$ holds, then K contains a handleslide with endpoints on k and j.

By [10, Proposition 3.8], modifying the handleslide set of an MCS in $\mathrm{MCS}_b(D)$ as in one of the cases in Figures 5 and 6 results in another MCS in $\mathrm{MCS}_b(D)$. Therefore, the notion of equivalence in the following definition is well-defined. In addition, if an MCS move is applied to an MCS, then only those chain complexes near the location of the MCS move are affected. In other words, the MCS moves are local in the sense that they change both the handleslides and chain complexes of an MCS only in a local neighborhood.

Definition 2.6 Two MCSs $\mathcal{C}$ and $\mathcal{C}'$ in $\mathrm{MCS}_b(D)$ are *equivalent*, written $\mathcal{C} \simeq \mathcal{C}'$, if there exists a sequence $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_s$ in $\mathrm{MCS}_b(D)$ so that $\mathcal{C} = \mathcal{C}_1$, $\mathcal{C}' = \mathcal{C}_s$, and, for all $1 \le i < s$, the set of handleslide marks of $\mathcal{C}_i$ and $\mathcal{C}_{i+1}$ differ by exactly one MCS move. The set $\widehat{\mathrm{MCS}}_b(D)$ is the set $\mathrm{MCS}_b(D)/\simeq$.

MCSs of the following type have a standard form that makes their relationship with augmentations particularly simple to describe.

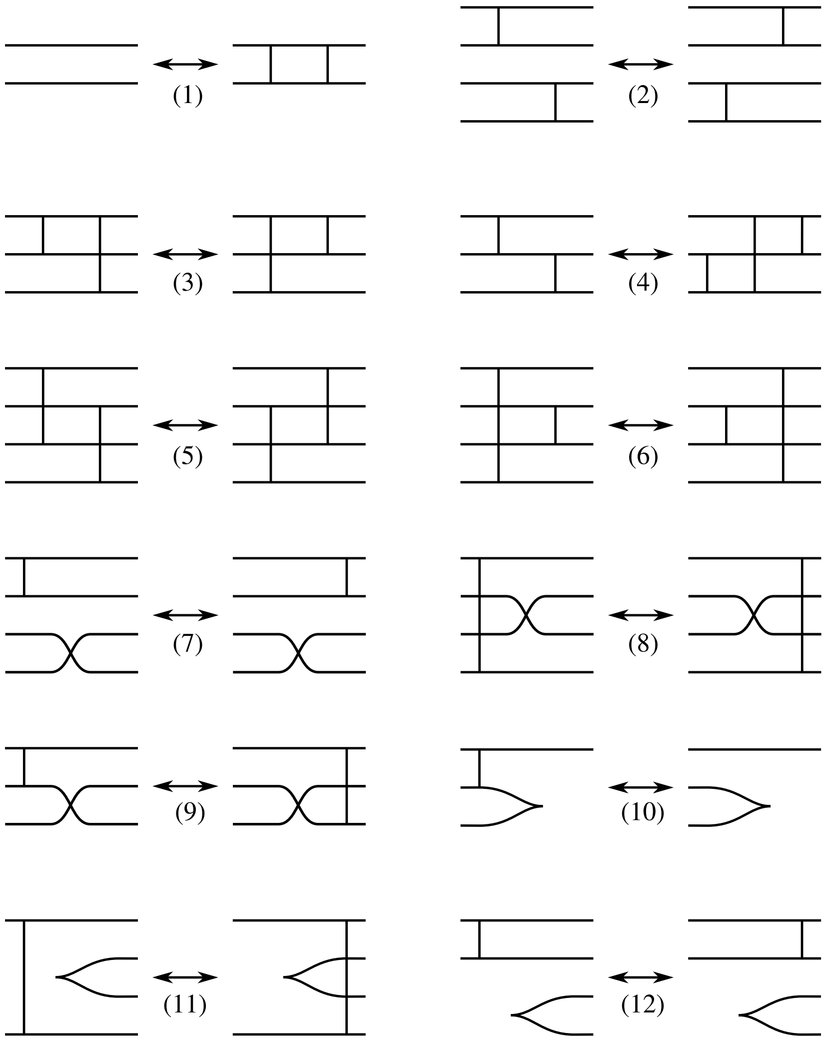

Figure 5: Handleslide modifications, called MCS moves, that result in an equivalent MCS

Definition 2.7 An MCS $\mathcal{C}$ in $\mathrm{MCS}_b(D)$ is in *A–form* if there exists a set R of degree 0 crossings so that just to the left of each q in R there is a handleslide with endpoints on the strands crossing at q and $\mathcal{C}$ has no other handleslides. A crossing q in R is said to be *marked*.

Figure 4 shows an MCS in A–form where R is the four left-most crossings. The subset $\mathrm{MCS}_A(D) \subset \mathrm{MCS}_b(D)$ consists of all A–form MCSs on D.

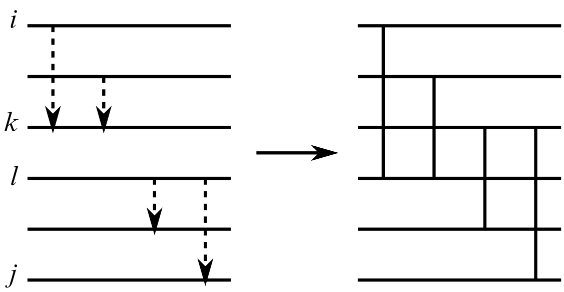

Figure 6: MCS move (13). On the left, a dotted arrow from strand α to strand β indicates that $\langle d_m e_\alpha, e_\beta \rangle$ is 1.

3 The main result

Suppose L is a Legendrian knot with σ–generic front diagram D, rotation number 0, and Maslov potential μ. The proof of Theorem 1.1 depends upon the following technical lemma, whose proof comprises most of this section.

Lemma 3.1 *Suppose D is a σ–generic front diagram and $\mathcal{C}$ and $\mathcal{C}'$ are in $\mathrm{MCS}_A(D)$ with corresponding augmentations $\epsilon_{\mathcal{C}}$ and $\epsilon_{\mathcal{C}'}$, respectively. If $\epsilon_{\mathcal{C}}$ and $\epsilon_{\mathcal{C}'}$ are homotopic, then $\mathcal{C}$ and $\mathcal{C}'$ are equivalent as MCSs.*

We now prove Theorem 1.1, assuming Lemma 3.1. Section 4 includes three corollaries of Theorem 1.1.

Proof of Theorem 1.1 By [10, Proposition 3.17], the natural inclusion of $\mathrm{MCS}_b(D)$ into $\mathrm{MCS}(D)$ induces a bijection from $\widehat{\mathrm{MCS}}_b(D)$ to $\widehat{\mathrm{MCS}}(D)$. Therefore, it suffices to construct a bijection from $\widehat{\mathrm{MCS}}_b(D)$ to $\mathrm{Aug}^{\mathrm{ch}}(D)$. In [10, Section 6], a surjective map $\widehat{\Psi}$ is constructed from $\widehat{\mathrm{MCS}}_b(D)$ to $\mathrm{Aug}^{\mathrm{ch}}(D)$. We will prove this map is injective. By [10, Theorem 1.6], every MCS is equivalent to an A–form MCS. Therefore, every MCS equivalence class contains an A–form representative. We give the definition of $\widehat{\Psi}$ in terms of A–form representatives and, in so doing, avoid most of the technical details of [10]. By [10, Corollary 6.21], given an MCS class $[\mathcal{C}]$ with A–form representative $\mathcal{C}$, $\widehat{\Psi}([\mathcal{C}])$ is the augmentation homotopy class $[\epsilon_{\mathcal{C}}]$, where a degree 0 crossing q is augmented by $\epsilon_{\mathcal{C}}$ if and only if q is marked by $\mathcal{C}$. Lemma 3.1 shows that if $\epsilon_{\mathcal{C}_1}$ is homotopic to $\epsilon_{\mathcal{C}_2}$, then $\mathcal{C}_1$ is equivalent to $\mathcal{C}_2$. It follows that $\widehat{\Psi}$ is injective. □

Before proving Lemma 3.1, we require two definitions and a lemma. Let D^2 be the disk of radius 1 centered at the origin in $\mathbb{R}^2$. Choose $m+2$ points on ∂D^2. Label the chosen points $\{b_0, \dots, b_{m+1}\}$ counter-clockwise. Let γ be the arc of ∂D^2 with endpoints b_{m+1} and b_0 and so that $b_1, \dots, b_m$ are not in γ. Given $x_0 \in \mathbb{R}$ that is not the x–coordinate of any crossing or cusp of D, we let $\{x_0\} \times [i, j]$ denote the vertical line segment with x–coordinate x_0 and endpoints on strands i and j of D, where the strands of D above $x = x_0$ are numbered $1, 2, \dots$ from top to bottom and $i < j$.

Definition 3.2 Let ϵ and ϵ' be homotopic augmentations in $\mathrm{Aug}(D)$ and let H be a chain homotopy from ϵ to ϵ'. An (ϵ, ϵ', H)*–half-disk* is a mapping of the two-disk D^2 into the xz–plane as in Definition 2.1, except for the following variations along the boundary:

(1) The arc γ maps to a vertical line $\{x_0\} \times [i, j]$; see Figure 2(f). We say the (ϵ, ϵ', H)–half-disk *originates at* $\{x_0\} \times [i, j]$.

(2) For exactly one $1 \leq j \leq m$, $f(b_j)$ is a degree -1 crossing q_{k_j}, $H(q_{k_j}) = 1$ holds, and f has a convex corner at $f(b_j)$; see Figure 2(d) or (e).

(3) If $1 \leq i < j$ (resp. $j < i \leq m$), $f(b_i)$ is a degree 0 crossing augmented by ϵ (resp. ϵ') and f has a convex corner at $f(b_i)$.

(4) The restriction of f to ∂D^2 is smooth except at $\{b_0, \dots, b_{m+1}\}$ as described in (1) and (2) and at points in $\partial D^2 \setminus (\{b_0, \dots, b_{m+1}\})$ where the image of f looks like Figure 2(b) or (c).

The set $\mathcal{H}(x_0, [i, j])$ consists of all (ϵ, ϵ', H)–half-disks originating at $\{x_0\} \times [i, j]$ up to reparametrization, and $\#\mathcal{H}(x_0, [i, j])$ is the mod 2 count of elements in $\mathcal{H}(x_0, [i, j])$.

Definition 3.3 Let ϵ be an augmentation in $\mathrm{Aug}(D)$. An ϵ*–half-disk* is a mapping f of the two-disk D^2 into the xz–plane as in Definition 3.2 except that conditions (2) and (3) are replaced with the requirement that all convex corners are at crossings that are augmented by ϵ.

The set $\mathcal{G}^\epsilon(x_0, [i, j])$ consists of all ϵ–half-disks originating at $\{x_0\} \times [i, j]$ up to reparametrization, and $\#\mathcal{G}^\epsilon(x_0, [i, j])$ is the mod 2 count of elements in $\mathcal{G}^\epsilon(x_0, [i, j])$.

As in Definition 2.1, the points in the vertical line $\{x_0\} \times [i, j]$ are the right-most points of either an (ϵ, ϵ', H)–half-disk or an ϵ–half-disk. It follows from the definitions that $\mu(i) = \mu(j)$ in the case of an (ϵ, ϵ', H)–half-disk and $\mu(i) = \mu(j) + 1$ in the case of an ϵ–half-disk

By [10, Corollary 6.21], the map Φ: $\mathrm{MCS}_A(D) \to \mathrm{Aug}(D)$ defined as follows is a bijection. Given $\mathcal{C} \in \mathrm{MCS}_A(D)$ and a generator q of $\mathcal{A}(D)$, $\Phi(\mathcal{C})(q) = 1$ holds if and only if q is a marked crossing of $\mathcal{C}$. We let $\epsilon_{\mathcal{C}}$ be the augmentation $\Phi(\mathcal{C})$.

Lemma 3.4 below generalizes [11, Lemma 7.10] and [12, Lemma 5.4] by removing the assumption that the front diagram D is nearly plat. Note that "gradient paths" from [11, Lemma 7.10] correspond to $\epsilon_{\mathcal{C}}$–half-disks in our terminology, and that [12, Lemma 5.4] allows more general coefficients.

Lemma 3.4 *Suppose D is a σ–generic front diagram and $\mathcal{C} = (\{C_m, d_m\}, \{x_m\}, H)$ is in $\mathrm{MCS}_A(D)$. Suppose $\mathcal{C}$ has $M \in \mathbb{N}$ chain complexes, $p \in \{1, \dots, M\}$, and x_p is to the immediate right of a crossing or cusp. Then, for all $i < j$,*

$$\langle d_p e_i, e_j \rangle = \#\mathcal{G}^{\epsilon_{\mathcal{C}}}(x_p, [i, j]). \tag{2}$$

Proof We induct on p. The base case, $p = 1$, follows since there is a unique disk in $\mathcal{G}^{\epsilon_{\mathcal{C}}}(x_1, [1, 2])$, as in Figure 2(b), while $\langle d_1 e_1, e_2 \rangle = 1$ holds according to Definition 2.3(4) .

Assume now that x_p sits to the immediate right of a crossing or cusp and that the result is known for smaller values of p. We complete the inductive step by considering cases.

Left cusp Suppose x_p is to the right of a left cusp with the two strands that meet at the cusp labeled k and $k+1$ at x_p. Define $\tau\colon \{1, \dots, s_{p-1}\} \to \{1, \dots, s_p\}$ by

$$\tau(i) = \begin{cases} i & \text{if } i < k, \\ i+2 & \text{if } i \geq k. \end{cases}$$

(Note that $s_{p-1} = s_p - 2$.) For any $1 \leq i' < j' \leq s_{p-1}$ there is a bijection between $\mathcal{G}^{\epsilon_{\mathcal{C}}}(x_{p-1}, [i', j'])$ and $\mathcal{G}^{\epsilon_{\mathcal{C}}}(x_p, [\tau(i'), \tau(j')])$; see, for example, Figure 7. Moreover, Definition 2.3(5)(c) together with the requirement that $\mathcal{C}$ has simple left cusps give

$$\langle d_{p-1} e_{i'}, e_{j'} \rangle = \langle d_p e_{\tau(i')}, e_{\tau(j')} \rangle,$$

so (2) follows when $i = \tau(i')$ and $j = \tau(j')$.

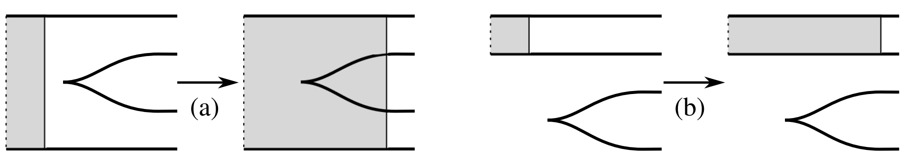

Figure 7: Possible extensions of an ϵ–half-disk or (ϵ, ϵ', H)–half-disk past a left cusp

It remains to consider those cases where $\{i, j\} \cap \{k, k+1\} \neq \varnothing$. Suppose that precisely one of i or j belongs to $\{k, k+1\}$. As $\mathcal{C}$ has simple left cusps, we have $\langle d_p e_i, e_j \rangle = 0$. In addition, the restriction on the behavior of an ϵ–half disk near a left cusp from Figure 2(b) gives that $\mathcal{G}^{\epsilon_{\mathcal{C}}}(x_p, [i, j]) = \varnothing$, so (2) holds. Finally, when $i = k$ and

$j = k+1$, there is a unique ϵ–half disk in $\mathcal{G}^{\epsilon c}(x_p, [k, k+1])$. (This disk has no convex corners, so Definition 3.3(2) is vacuously satisfied.) Therefore, (2) follows in view of Definition 2.3(4).

Crossing When x_p sits immediately to the right of a crossing, the inductive step is achieved precisely as in [11, Lemma 7.10] or [12, Lemma 5.4] (the signs in the latter reference may be ignored). The arguments in these references apply regardless of whether or not the crossing is marked.

Right cusp Suppose a right cusp sits between x_p and x_{p-1} with the strands that meet at the cusp labeled k and $k+1$ at x_{p-1}. Let $a_{i,j}$ be $\langle d_{p-1}e_i, e_j\rangle$. In the quotient of (C_{p-1}, d_{p-1}) by the subcomplex spanned by e_k and $d_{p-1}e_k$, we have

$$0 = [d_{p-1}e_k] = [e_{k+1}] + \sum_{k+1<j} a_{k,j}[e_j],$$

so

$$d_{p-1}[e_i] = \sum_{i<j} a_{i,j}[e_j] = \sum_{i<j<k} a_{i,j}[e_j] + \sum_{k+1<j} (a_{i,j} + a_{i,k+1}\cdot a_{k,j})[e_j].$$

By Definition 2.3 (5) (b), this gives the computation of the differential in (C_p, d_p) as

$$\langle d_p e_i, e_j\rangle = \langle d_{p-1}e_{\pi(i)}, e_{\pi(j)}\rangle + \langle d_{p-1}e_{\pi(i)}, e_{k+1}\rangle \cdot \langle d_{p-1}e_k, e_{\pi(j)}\rangle, \tag{3}$$

where $\pi\colon \{1, \dots, s_p\} \to \{1, \dots, s_{p-1}\}$ is defined by

$$\pi(i) = \begin{cases} i & \text{if } i < k, \\ i+2 & \text{if } i \geq k. \end{cases}$$

We note that the second term on the right can be non-zero only if $i < k \leq j$; see Figure 8.

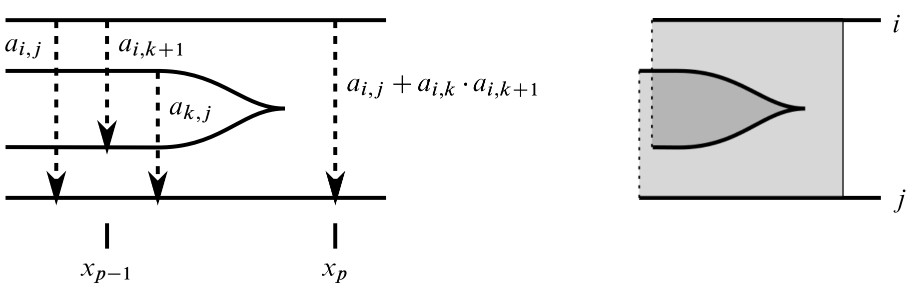

Figure 8: (Left) The relation between differentials at a right cusp. A dotted arrow at x_l pointing from strand i to strand j indicates the matrix coefficient $\langle d_l e_i, e_j\rangle$. (Right) The appearance of disks in $\mathcal{G}^{\epsilon c}(x_p, [i, j])$ with a boundary point at the right cusp between x_{p-1} and x_p.

To complete the proof, we combine Equation (3) with the observation that ϵ–half-disks satisfy a bijection

$$\mathcal{G}^{\epsilon c}(x_p,[i,j]) \cong \mathcal{G}^{\epsilon c}(x_{p-1},[\pi(i),\pi(j)]) \\ \cup\big(\mathcal{G}^{\epsilon c}(x_{p-1},[\pi(i),k+1]) \times \mathcal{G}^{\epsilon c}(x_{p-1},[k,\pi(j)])\big)$$

explained as follows. Those disks in $\mathcal{G}^{\epsilon c}(x_p,[i,j])$ whose boundaries do not intersect the cusp point are in bijection with $\mathcal{G}^{\epsilon c}(x_{p-1},[\pi(i),\pi(j)])$; see Figure 9. Disks in $\mathcal{G}^{\epsilon c}(x_p,[i,j])$ whose boundaries do intersect the cusp point appear between x_{p-1} and x_p as pictured in Figure 8. Removing the portion of the disk between x_{p-1} and x_p leaves a pair of initially overlapping disks from $\mathcal{G}^{\epsilon c}(x_{p-1},[\pi(i),k+1]) \times \mathcal{G}^{\epsilon c}(x_{p-1},[k,\pi(j)])$, and this correspondence is bijective. □

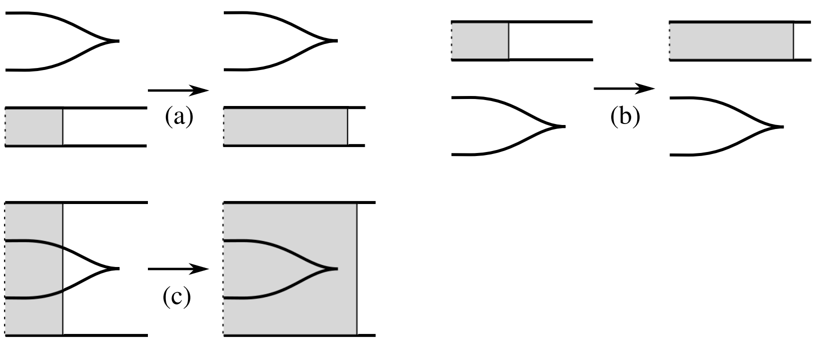

Figure 9: Possible extensions of an ϵ–half-disk or (ϵ,ϵ',H)–half-disk past a right cusp

We outline the central idea of Lemma 3.1 before proceeding to the proof. Recall that an augmentation ϵ has an associated A–form MCS $\mathcal{C}$ where a degree 0 crossing q is marked by $\mathcal{C}$ if and only if $\epsilon(q)$ is 1. The proof of Theorem 1.1 reduced to showing that if augmentations ϵ and ϵ' are homotopic, then their associated A–form MCSs $\mathcal{C}$ and $\mathcal{C}'$ are equivalent. This is accomplished in Lemma 3.1 where an algorithm is given to translate a chain homotopy H from ϵ to ϵ' into a sequence of MCS moves from $\mathcal{C}$ to $\mathcal{C}'$. In particular, for each degree -1 crossing p sent to 1 by H, we employ MCS move (13) just to the left of p to introduce new handleslides. We prove that these handleslides give the mod 2 count of certain (ϵ,ϵ',H)–half disks. Moving these handleslides to the right in the front diagram D, we find that a degree 0 crossing q is changed from marked to unmarked or from unmarked to marked if and only if there exists an odd number of (ϵ,ϵ',H)–half disks originating at q. Therefore, by Proposition 2.2 and the definition of $\mathcal{C}$ and $\mathcal{C}'$, q is changed from marked to unmarked or from unmarked to marked if and only if $\mathcal{C}$ and $\mathcal{C}'$ differ at q. We may therefore conclude that $\mathcal{C}$ and $\mathcal{C}'$ are equivalent.

Proof of Lemma 3.1 Suppose $\mathcal{C}$ and $\mathcal{C}'$ are A–form MCSs and $\epsilon_{\mathcal{C}}$ is homotopic to $\epsilon_{\mathcal{C}'}$. We simplify notation by letting ϵ be $\epsilon_{\mathcal{C}}$ and ϵ' be $\epsilon_{\mathcal{C}'}$. Since ϵ and ϵ' are homotopic, there exists a chain homotopy $H\colon \mathcal{A}(D) \to \mathbb{Z}/2\mathbb{Z}$. Label the degree -1 crossings sent to 1 by H, from left to right, $p_1, \dots, p_m$.

To prove the lemma, we will construct a sequence of MCSs $\mathcal{C}_0, \dots, \mathcal{C}_s$ so that $\mathcal{C}_0$ is $\mathcal{C}$ and $\mathcal{C}_s$ is $\mathcal{C}'$, and, for all $0 \leq r < s$, $\mathcal{C}_r \simeq \mathcal{C}_{r+1}$ holds. The construction of the $\mathcal{C}_r$ is inductive, and each of the $\mathcal{C}_r$ will contain a (possibly empty) collection of handleslides V_r that are grouped together immediately to the right of a particular crossing or cusp.

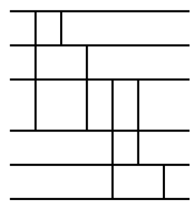

Figure 10: An ordered collection of handleslides

For our purposes it will be convenient to require that the handleslides in each of the V_r are ordered in the following sense. We say a collection of handleslides is *ordered* if, given two handleslides h and h' in the collection with endpoints on strands $i < j$ and $i' < j'$ respectively, h is right of h' if and only if $i > i'$ holds, or $i = i'$ and $j < j'$ hold; see Figure 10. We let $v_r^{i,j}$ be 1 if there exists a handleslide in V_r with endpoints on strands i and j, where $i < j$. Otherwise, $v_r^{i,j}$ is defined to be 0. In a slight abuse of notation, we also let $v_r^{i,j}$ refer to the handleslide in V_r with endpoints on i and j, if such a handleslide exists.

We will verify that Property 1 below holds for all $0 \leq r \leq s$ as we inductively construct MCSs $\mathcal{C}_r$ with ordered handleslide collections V_r.

Property 1 (a) The MCS $\mathcal{C}_r$ agrees with $\mathcal{C}'$ to the left of V_r and $\mathcal{C}$ to the right of V_r.

(b) For all $i < j$,

$$v_r^{i,j} = \#\mathcal{H}(x_r, [i, j]),$$

where x_r is the x–coordinate of the left-most handleslide in V_r.

Each time r increases, the collection of handleslides V_r is pushed to the right past one cusp or crossing. We continue this inductive process until we arrive at an MCS $\mathcal{C}_s$ with

V_s located just to the left of the right-most right cusp of D. Since the two strands of this cusp do not have the same Maslov potential, V_s must be empty. Then, Property 1 (a) shows that $\mathcal{C}_s$ is $\mathcal{C}'$. As $\mathcal{C} = \mathcal{C}_0$, and for $0 \leq i < s$, $\mathcal{C}_i \simeq \mathcal{C}_{i+1}$ it will then follow that $\mathcal{C} \simeq \mathcal{C}'$ holds, as desired.

In the remainder of the proof we construct the sequence of MCSs $\mathcal{C}_0, \dots, \mathcal{C}_s$. Since a crossing q is augmented by ϵ (resp. ϵ') if and only if q is marked by $\mathcal{C}$ (resp. $\mathcal{C}'$), Proposition 2.2 implies $\mathcal{C}$ and $\mathcal{C}'$ differ at q if and only if there exists an odd number of (ϵ, ϵ', H)–admissible disks originating at q. Since q is the right-most point of an (ϵ, ϵ', H)–admissible disk originating at q, it follows that there are no admissible disks originating to the left of p_1 (which is the first crossing sent to 1 by H). Therefore, $\mathcal{C}$ and $\mathcal{C}'$ are identical to the left of p_1. We can then set $\mathcal{C}_0 = \mathcal{C}$ and define V_0 to be empty, but located just to the left of p_1. It follows that Property 1 holds for $\mathcal{C}_0$.

Given $\mathcal{C}_r$ and V_r, we will construct $\mathcal{C}_{r+1}$ and V_{r+1} by applying MCS moves to $\mathcal{C}_r$. We will prove that if Property 1 holds for $\mathcal{C}_r$, then it holds for $\mathcal{C}_{r+1}$ as well. We consider five cases depending on the type of crossing or cusp just to the right of V_r. Let q be the crossing or cusp point to the immediate right of V_r. Let x_r (resp. x_{r+1}) be an x–coordinate to the immediate left (resp. right) of q. In each of the five cases, we first analyze the (ϵ, ϵ', H)–half-disks in $\mathcal{H}(x_{r+1}, [i, j])$ before describing the sequence of MCS moves used to construct $\mathcal{C}_{r+1}$ from $\mathcal{C}_r$ and proving Property 1 holds for $\mathcal{C}_{r+1}$ and V_{r+1}.

In the first three cases considered, q is a crossing between strands k and $k+1$ where the strands of D have been numbered $1, \dots, s_r$, from top to bottom, just to the left of q. Let ρ: $\{1, \dots, s_r\} \to \{1, \dots, s_r\}$ be the permutation that transposes k and $k+1$.

Crossing q such that $|q| \neq 0$ and $H(q) \neq 1$ Since $|q|$ is non-zero, $\mathcal{H}(x_r, [k, k+1])$ and $\mathcal{H}(x_{r+1}, [k, k+1])$ are both empty. Given $1 \leq i < j \leq s_r$ such that $(i, j) \neq (k, k+1)$, one has that (ϵ, ϵ', H)–half-disks in $\mathcal{H}(x_{r+1}, [i, j])$ cannot have a convex corner at q, since $|q| \neq 0$ and $H(q) \neq 1$ hold. In fact, $\#\mathcal{H}(x_{r+1}, [i, j]) = \#\mathcal{H}(x_r, [\rho(i), \rho(j)])$ holds, since there is a natural bijection between $\mathcal{H}(x_{r+1}, [i, j])$ and $\mathcal{H}(x_r, [\rho(i), \rho(j)])$; see, for example, Figure 11.

We now define the sequence of MCS moves that create $\mathcal{C}_{r+1}$ from $\mathcal{C}_r$ and prove Property 1 holds for $\mathcal{C}_{r+1}$. Move all handleslides of V_r to the right of q using MCS moves (7)–(9). Since $|q|$ is non-zero, $v_r^{k,k+1}$ is 0, and therefore all handleslides of V_r can be moved past q and no new handleslides are created by doing so. The resulting collection can be ordered, using MCS moves, without creating new handleslides. The reordering requires rearranging handleslides with one endpoint on either strand k or $k+1$. Since $|q|$ is non-zero, there is no handleslide between k and $k+1$, and therefore the rearrangement can be done without using MCS move (4). The resulting

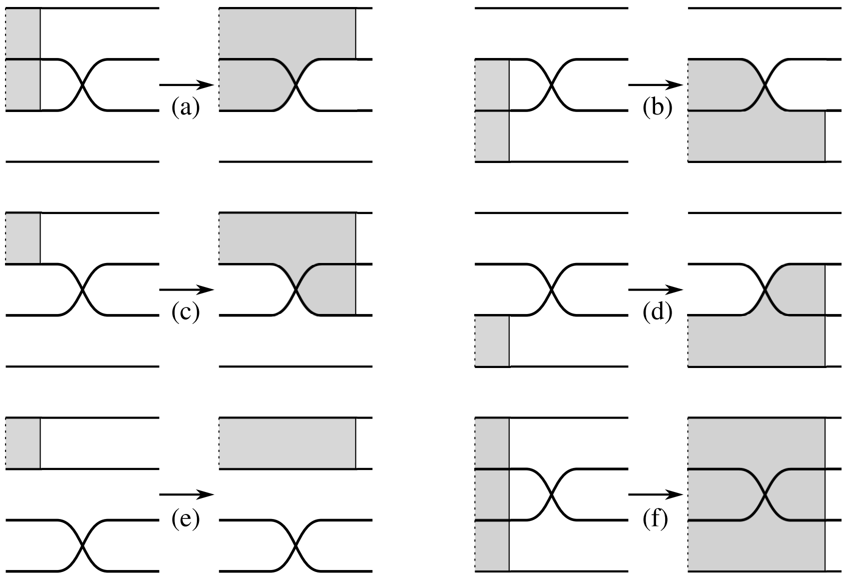

Figure 11: Possible extensions of an (ϵ, ϵ', H)–half-disk past a crossing

ordered collection is V_{r+1} and the MCS is $\mathcal{C}_{r+1}$, and $v_{r+1}^{i,j} = v_r^{\rho(i),\rho(j)}$ holds for all $1 \le i < j \le s_r$. By Property 1(b), $v_r^{\rho(i),\rho(j)} = \#\mathcal{H}(x_r, [\rho(i), \rho(j)])$ holds and as shown above, $\#\mathcal{H}(x_{r+1}, [i, j]) = \#\mathcal{H}(x_r, [\rho(i), \rho(j)])$ holds. Therefore, Property 1(b) holds for $\mathcal{C}_{r+1}$. Property 1(a) holds for $\mathcal{C}_r$ and since $|q|$ is non-zero, q is not marked by either $\mathcal{C}_{r+1}$ or $\mathcal{C}'$. Therefore, Property 1(a) holds for $\mathcal{C}_{r+1}$.

Crossing q such that $|q| = 0$ Let v_q be 1 if q is marked by $\mathcal{C}$ and 0 otherwise. Since Property 1(a) holds for $\mathcal{C}_r$, if q is marked by $\mathcal{C}$, then q is marked by $\mathcal{C}_r$ as well. We slightly abuse notation and also let v_q be the handleslide at q in $\mathcal{C}_r$ in the case such exists.

Suppose $i \neq k+1$ and $j \neq k$. Half-disks in $\mathcal{H}(x_{r+1}, [i, j])$ cannot have a convex corner at q, and therefore there is a bijection from $\mathcal{H}(x_r, [i, j])$ to $\mathcal{H}(x_{r+1}, [\rho(i), \rho(j)])$; see, for example, Figure 11(c)–(f). Since Property 1(b) holds for $\mathcal{C}_r$, $\#\mathcal{H}(x_{r+1}, [i, j]) = v_r^{\rho(i),\rho(j)}$ holds.

Note that $\mathcal{H}(x_{r+1}, [k, k+1])$ is empty. Suppose one of $i = k+1$ or $j = k$ holds. Half-disks in $\mathcal{H}(x_r, [\rho(i), \rho(j)])$ may be smoothly extended past q as in Figure 11(a) and (b). Therefore, there exists an injection from $\mathcal{H}(x_r, [\rho(i), \rho(j)])$ to $\mathcal{H}(x_{r+1}, [i, j])$. However, there may be half-disks in $\mathcal{H}(x_{r+1}, [i, j])$ that have a convex corner at q. If $j = k$ (resp. $i = k+1$) and q is marked by $\mathcal{C}'$ (resp. $\mathcal{C}$), then a half-disk in $\mathcal{H}(x_{r+1}, [i, j])$ can have a convex corner at q; see Figure 12(a) and (b) respectively.

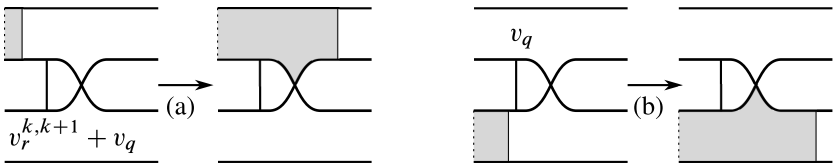

Figure 12: Extending an (ϵ, ϵ', H)–half-disk past a degree 0 crossing so as to have a convex corner at the crossing

Such half-disks are in bijection with half-disks in $\mathcal{H}(x_r, [i, j])$, as can be seen in Figure 12(a) and (b), and by Property 1(b), are counted by $v_r^{i,j}$. Since $v_r^{k,k+1}$ is 1 if and only if $\mathcal{C}'$ and $\mathcal{C}$ differ at q, q is marked by $\mathcal{C}'$ (resp. $\mathcal{C}$) if and only if $v_r^{k,k+1} + v_q$ is 1 (resp. v_q is 1). Therefore, if $j = k$ (resp. $i = k+1$), then the mod 2 count of half-disks in $\mathcal{H}(x_{r+1}, [i, j])$ with a convex corner at q is $(v_r^{k,k+1} + v_q) \cdot v_r^{i,k}$ (resp. $v_q \cdot v_r^{k+1,j}$). In summary,

$$\#\mathcal{H}(x_{r+1}, [i, j]) = \begin{cases} v_r^{i,k+1} + (v_r^{k,k+1} + v_q) \cdot v_r^{i,k} & \text{if } j = k, \\ v_r^{k,j} + v_q \cdot v_r^{k+1,j} & \text{if } i = k+1, \\ 0 & \text{if } i = k \text{ and } j = k+1, \\ v_r^{\rho(i),\rho(j)} & \text{otherwise.} \end{cases} \tag{4}$$

We now define the sequence of MCS moves that create $\mathcal{C}_{r+1}$ from $\mathcal{C}_r$ and prove Property 1 holds for $\mathcal{C}_{r+1}$. We move each handleslide $v_r^{i,j}$ of V_r past q beginning with the right-most handleslide in V_r. If $i \geq k$ and $j \neq k+1$ hold, use MCS moves (2)–(9) to move $v_r^{i,j}$ past v_q, if v_q is 1, and then past the crossing q. If v_q is 1 and $i = k+1$, a new handleslide with endpoints on strands k and j is created when MCS move (4) is used to move $v_r^{i,j}$ past v_q. Move this handleslide to the right of q as well. It is not possible to move $v_r^{k,k+1}$ past q and so, for now, we simply leave $v_r^{k,k+1}$ to the left of q. If $i < k$ holds, use MCS moves (2)–(9) to move $v_r^{i,j}$ past $v_r^{k,k+1}$, if $v_r^{k,k+1}$ is 1, then past v_q, if v_q is 1, and then past the crossing q. If v_q is 1 or $v_r^{k,k+1}$ is 1, and $j = k$, a new handleslide with endpoints on strands i and $k+1$ is created when MCS move (4) is used to move $v_r^{i,j}$ past v_q or $v_r^{k,k+1}$. Move this handleslide to the right of q as well. Once all $v_r^{i,j}$, except $v_r^{k,k+1}$, have been moved past q, use MCS moves (5) and (1) to order the collection of handleslides just to the right of q and remove pairs of handleslides that have the same endpoints. The resulting collection is V_{r+1}. From our work above, we have:

$$v_{r+1}^{i,j} = \begin{cases} v_r^{i,k+1} + (v_r^{k,k+1} + v_q) \cdot v_r^{i,k} & \text{if } j = k, \\ v_r^{k,j} + v_q \cdot v_r^{k+1,j} & \text{if } i = k+1, \\ 0 & \text{if } i = k \text{ and } j = k+1, \\ v_r^{\rho(i),\rho(j)} & \text{otherwise.} \end{cases} \tag{5}$$

Use MCS move (1) to remove both $v_r^{k,k+1}$ and v_q, in the case that they both exist. The resulting MCS is $\mathcal{C}_{r+1}$. By Property 1, $v_r^{k,k+1}$ is 1 if and only if there is an odd number of (ϵ, ϵ', H)–half-disks originating at $\{x_r\} \times [k, k+1]$. There is a bijection between such disks and the (ϵ, ϵ', H)–admissible disks originating at q; see Figure 13. Therefore, by Proposition 2.2, $v_r^{k,k+1}$ is 1 if and only if $\mathcal{C}$ and $\mathcal{C}'$ differ at q. Therefore, Property 1(a) holds for $\mathcal{C}_{r+1}$. Finally, Equations (5) and (4) imply Property 1(b) holds for $\mathcal{C}_{r+1}$.

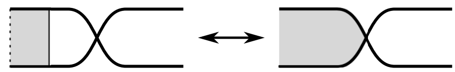

Figure 13: The bijection between (ϵ, ϵ', H)–half-disks originating at $\{x_r\} \times [k, k+1]$ and (ϵ, ϵ', H)–half-disks originating at a crossing q between strands k and $k+1$

Crossing p_i where $1 < i \leq m$ Suppose p_i is a degree -1 crossing between strands k and $k+1$ and $H(p_i) = 1$ holds. Suppose $i \neq k+1$ and $j \neq k$. Half-disks in $\mathcal{H}(x_{r+1}, [i, j])$ cannot have a convex corner at p_i and, therefore, there is a bijection from $\mathcal{H}(x_r, [i, j])$ to $\mathcal{H}(x_{r+1}, [\rho(i), \rho(j)])$; see, for example, Figure 11(c)–(f). Since Property 1(b) holds for $\mathcal{C}_r$, $\#\mathcal{H}(x_{r+1}, [i, j]) = v_r^{\rho(i), \rho(j)}$ holds.

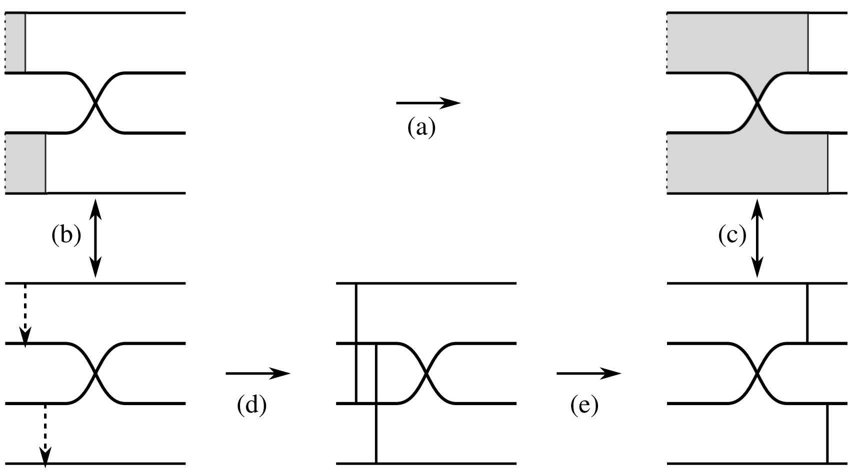

Figure 14: The correspondence between (ϵ, ϵ', H)–half-disks with a convex corner at the degree -1 crossing in the figure and handleslides introduced by MCS move (13) to the left of the crossing. In step (d), two handleslides are created by MCS move (13). These handleslides correspond to the two (ϵ, ϵ', H)–half-disks in the top right figure, each of which has a convex corner at the crossing.

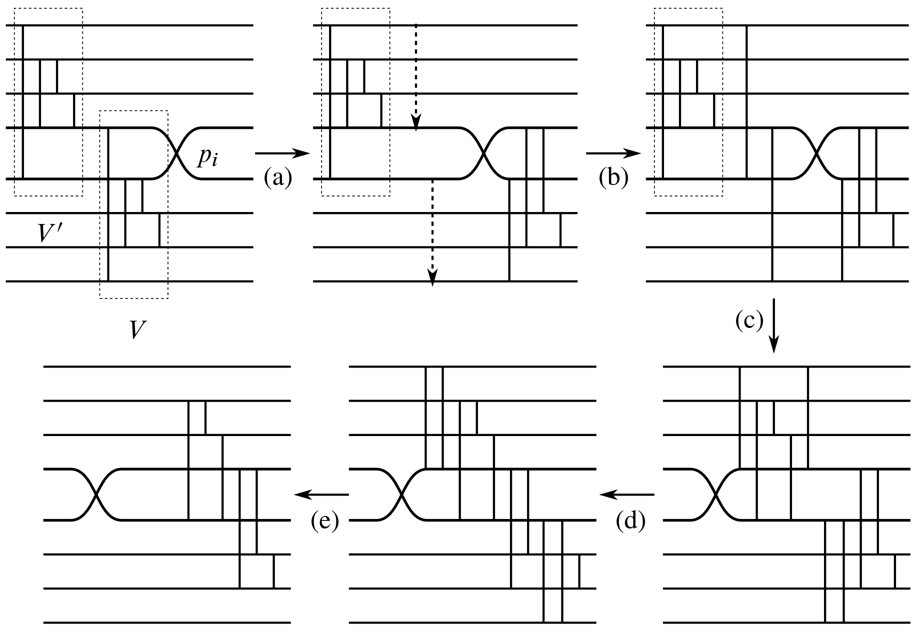

Figure 15: The sequence of MCS moves at a crossing p_i where $|p_i| = -1$ and $H(p_i) = 1$ both hold.

Note that $\mathcal{H}(x_{r+1}, [k, k+1])$ is empty. Suppose one of $i = k+1$ or $j = k$ holds. Half-disks in $\mathcal{H}(x_r, [\rho(i), \rho(j)])$ may be smoothly extended past p_i as in Figure 11(a) and (b). Therefore, there exists an injection from $\mathcal{H}(x_r, [\rho(i), \rho(j)])$ to $\mathcal{H}(x_{r+1}, [i, j])$. However, there may be half-disks in $\mathcal{H}(x_{r+1}, [i, j])$ that have a convex corner at q. In the case that $j = k$ (resp. $i = k+1$), such disks correspond to ϵ–half-disks (resp. ϵ'–half-disks) that have been extended past p_i so as to have a convex corner at p_i; see Figure 14(a)–(e). Let (C, d) (resp. (C', d')) be the chain complex of $\mathcal{C}$ (resp. $\mathcal{C}'$) just to the left of p_i. Property 1(a) implies that, in $\mathcal{C}_r$, (C, d) (resp. (C', d')) is the chain complex to the immediate right (resp. left) of V_r. By Lemma 3.4, the mod 2 count of such half-disks is $\langle de_i, e_j \rangle$ and $\langle d'e_i, e_j \rangle$ respectively. Since Property 1(b) holds for $\mathcal{C}_r$, we may summarize the work of the previous two paragraphs as follows:

$$\#\mathcal{H}(x_{r+1}, [i, j]) = \begin{cases} v_r^{\rho(i),\rho(j)} + \langle de_i, e_j \rangle & \text{if } j = k, \\ v_r^{\rho(i),\rho(j)} + \langle d'e_i, e_j \rangle & \text{if } i = k+1, \\ 0 & \text{if } i = k \text{ and } j = k+1, \\ v_r^{\rho(i),\rho(j)} & \text{otherwise.} \end{cases} \tag{6}$$

We now define the sequence of MCS moves that create $\mathcal{C}_{r+1}$ from $\mathcal{C}_r$ and prove Property 1 holds for $\mathcal{C}_{r+1}$. Let $V \subset V_r$ (resp. $V' \subset V_r$) be the handleslides $v_r^{i,j}$ in V_r

satisfying $i \geq k$ (resp. $i < k$). Since V_r is ordered, V is right of V'; see Figure 15. Let $(\overline{C}, \overline{d})$ be the chain complex of $\mathcal{C}_r$ between V' and V. Use MCS moves (7) and (9) to move the handleslides in V past p_i; see Figure 15(a). Since p_i has degree -1, $v_r^{k,k+1}$ is 0 and $\mu(k) = \mu(k+1) - 1$ holds. Therefore, strands k and $k+1$ satisfy the conditions of MCS move (13). Use MCS move (13) to introduce new handleslides between V' and p_i; see Figure 15(b). MCS move (13) introduces a handleslide with endpoints i and j if and only if either $j = k+1$ and $\langle \overline{d} e_i, e_k \rangle$ is 1, or $i = k$ and $\langle \overline{d} e_{k+1}, e_j \rangle$ is 1. Recall that (C, d) (resp. (C', d')) is the chain complex of $\mathcal{C}_r$ to the immediate right (resp. left) of V_r. Since V_r is ordered, the handleslides between (C, d) and $(\overline{C}, \overline{d})$ have upper endpoints on $k, \dots, s_r$ and the handleslides between (C', d') and $(\overline{C}, \overline{d})$ have upper endpoints on $1, \dots, k-1$. Because of the ordering of handleslides in V_r, the coefficient $\langle \overline{d} e_i, e_k \rangle$ (resp. $\langle \overline{d} e_{k+1}, e_j \rangle$) is unaffected by handleslides in V (resp. V'). As a consequence $\langle \overline{d} e_i, e_k \rangle = \langle d e_i, e_k \rangle$ holds for all $i < k$ and $\langle \overline{d} e_{k+1}, e_j \rangle = \langle d' e_{k+1}, e_j \rangle$ holds for all $k+1 < j$. Therefore, MCS move (13) introduces a handleslide with endpoints i and j if and only if either $j = k+1$ and $\langle d e_i, e_k \rangle$ is 1, or $i = k$ and $\langle d' e_{k+1}, e_j \rangle$ is 1. Move the handleslides created by MCS move (13) and the handleslides in V' past p_i using MCS moves (7)–(9); see Figure 15(c). Use MCS moves (1), (3), (5) and (6) to order the collection of handleslides to the right of p_i and remove pairs of handleslides with identical endpoints; see Figure 15 (d) and (e). In particular, this can be done without creating any new handleslides. The resulting ordered collection of handleslides is V_{r+1} and the MCS is $\mathcal{C}_{r+1}$. Since the only new handleslides created were those created by the single application of MCS move (13),

$$(7) \qquad v_{r+1}^{i,j} = \begin{cases} v_r^{\rho(i),\rho(j)} + \langle d e_i, e_j \rangle & \text{if } j = k, \\ v_r^{\rho(i),\rho(j)} + \langle d' e_i, e_j \rangle & \text{if } i = k+1, \\ 0 & \text{if } i = k \text{ and } j = k+1, \\ v_r^{\rho(i),\rho(j)} & \text{otherwise.} \end{cases}$$

Equations (6) and (7) imply Property 1(b) holds for $\mathcal{C}_{r+1}$. Finally, Property 1(a) holds for $\mathcal{C}_r$ and $|q| \neq 0$ implies q is not marked by either $\mathcal{C}_{r+1}$ or $\mathcal{C}'$. Therefore, Property 1(a) holds for $\mathcal{C}_{r+1}$.

Left cusp Suppose q is a left cusp. Number the strands of D, from top to bottom, $1, \dots, s_r$ (resp. $1, \dots, s_{r+1}$) just to the left (resp. right) of q. Define $\tau \colon \{1, \dots, s_r\} \to \{1, \dots, s_{r+1}\}$ by

$$\tau(i) = \begin{cases} i & \text{if } i < k, \\ i+2 & \text{if } i \geq k. \end{cases}$$

(Note that $s_r = s_{r+1} - 2$.) For any $1 \leq i' < j' \leq s_r$, there is a bijection between $\mathcal{H}(x_{r+1}, [\tau(i'), \tau(j')])$ and $\mathcal{H}(x_r, [i', j'])$; see, for example, Figure 7(a) and (b). If

$\{i,j\} \cap \{k,k+1\}$ is non-empty, then $\mathcal{H}(x_{r+1},[i,j])$ is empty. Therefore, since Property 1(b) holds for $\mathcal{C}_r$,

$$\#\mathcal{H}(x_{r+1},[i,j]) = \begin{cases} v_r^{\tau^{-1}(i),\tau^{-1}(j)} & \text{if } \{i,j\}\cap\{k,k+1\}=\varnothing, \\ 0 & \text{otherwise.} \end{cases} \tag{8}$$

Use MCS moves (11) and (12) to move each handleslide in V_r past q. The resulting collection V_{r+1} is ordered and the resulting MCS is $\mathcal{C}_{r+1}$. The endpoints of a handleslide remain on the same strands of D as it is moved past q. Therefore, we have

$$v_{r+1}^{i,j} = \begin{cases} v_r^{\tau^{-1}(i),\tau^{-1}(j)} & \text{if } \{i,j\}\cap\{k,k+1\}=\varnothing, \\ 0 & \text{otherwise.} \end{cases} \tag{9}$$

Equations (8) and (9) imply Property 1(b) holds for $\mathcal{C}_{r+1}$. Since q is not a crossing and Property 1(a) holds for $\mathcal{C}_r$, it must hold for $\mathcal{C}_{r+1}$ as well.

Right cusp Suppose q is a right cusp between strands k and $k+1$. Let (C,d) (resp. (C',d')) be the chain complex of $\mathcal{C}$ (resp. $\mathcal{C}'$) just to the left of q. Property 1(a) implies that, in $\mathcal{C}_r$, (C,d) (resp. (C',d')) is the chain complex to the immediate right (resp. left) of V_r. Number the strands of D, from top to bottom, by $1,\dots,s_{r+1}$ (resp. by $1,\dots,s_r$) just to the right (resp. left) of q. Define $\pi\colon \{1,\dots,s_{r+1}\} \to \{1,\dots,s_r\}$ by

$$\pi(i) = \begin{cases} i & \text{if } i<k, \\ i+2 & \text{if } i \ge k. \end{cases}$$

(Note that $s_r = s_{r+1}+2$.)

If $j<k$ or $i \ge k$, then

$$\#\mathcal{H}(x_{r+1},[i,j]) = v_r^{\pi(i),\pi(j)} \tag{10}$$

holds, since Property 1(b) holds for $\mathcal{C}_r$ and there is a bijection from $\mathcal{H}(x_r,[\pi(i),\pi(j)])$ to $\mathcal{H}(x_{r+1},[i,j])$; see Figure 9(a) and (b).

When $j \ge k$ and $i<k$, we claim that there is a bijection

$$\begin{aligned} \mathcal{H}(x_{r+1},[i,j]) \cong \mathcal{H}(x_r,[\pi(i),\pi(j)]) & \\ \cup (\mathcal{G}^{\epsilon}(x_r,[\pi(i),k+1]) \times \mathcal{H}(x_r,[k,\pi(j)])) & \\ \cup (\mathcal{H}(x_r,[\pi(i),k+1]) \times \mathcal{G}^{\epsilon'}(x_r,[k,\pi(j)])). & \end{aligned} \tag{11}$$

Suppose $j \ge k$ and $i<k$. Half-disks in $\mathcal{H}(x_r,[\pi(i),\pi(j)])$ may be smoothly extended past q as in Figure 9(c). Therefore, there exists an injection from $\mathcal{H}(x_r,[\pi(i),\pi(j)])$ to $\mathcal{H}(x_{r+1},[i,j])$. However, there may be half-disks in $\mathcal{H}(x_{r+1},[i,j])$ whose boundary intersects the cusp point; see Figure 16(a) and Figure 17(a).

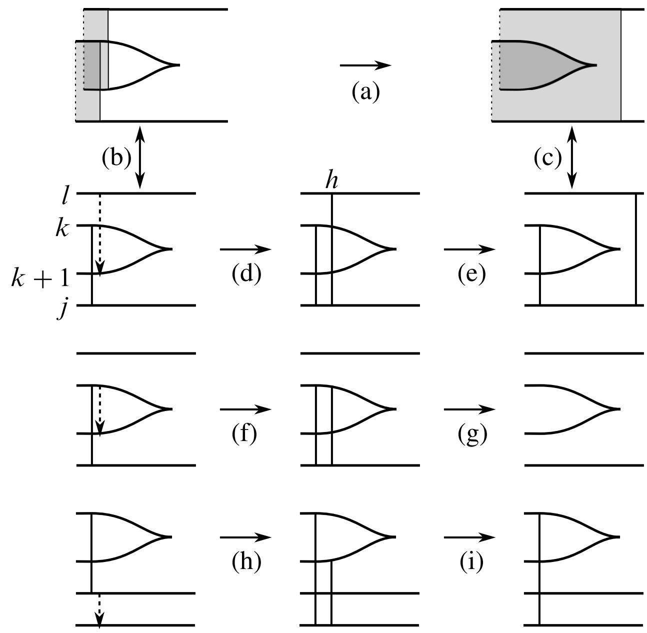

Figure 16: (a)–(e): The correspondence between Type 1 (ϵ, ϵ', H)–half-disks whose boundary intersects the right cusp in the figure and handleslides introduced by MCS move (13) to the left of the crossing. (f), (g): A handleslide introduced by MCS move (13) that is removed, along with $v_r^{k,j}$, by an MCS (1) move. (h), (i): A handleslide introduced by MCS move (13) that is removed by an MCS (10) move.

We divide half-disks whose boundary intersects q into two types as follows. Any such half-disk has one convex corner at a degree -1 crossing, which we denote p. Trace the boundary of such a half-disk counter-clockwise beginning at the vertical line $\{x_{r+1}\} \times [i, j]$. In a Type 1 (resp. Type 2) half-disk, p appears after (resp. before) q. A Type 1 (resp. Type 2) half-disk can be uniquely decomposed into an (ϵ, ϵ', H)–half-disk and an ϵ–half-disk (resp. ϵ'–half-disk) as in Figure 16(a) and (b) (resp. Figure 17(a) and (b)). Therefore, the set in the second (resp. third) line of Equation (11) is in bijection with Type 1 (resp. Type 2) half-disks. Since Property 1(b) holds for $\mathcal{C}_r$ and Lemma 3.4 holds for both (C, d) and (C', d'), the mod 2 count of Type 1 half-disks is $\langle de_{\pi(i)}, e_{k+1}\rangle \cdot v_r^{k,\pi(j)}$ and the mod 2 count of Type 2 half-disks is $v_r^{\pi(i),k+1} \cdot \langle d'e_k, e_{\pi(j)}\rangle$. Therefore, for $j \geq k$ and $i < k$, we have the formula

$$(12) \quad \#\mathcal{H}(x_{r+1}, [i, j]) = v_r^{\pi(i),\pi(j)} + v_r^{\pi(i),k+1} \cdot \langle d'e_k, e_{\pi(j)}\rangle + \langle de_{\pi(i)}, e_{k+1}\rangle \cdot v_r^{k,\pi(j)}.$$

We now define the sequence of MCS moves that create $\mathcal{C}_{r+1}$ from $\mathcal{C}_r$ and prove Property 1 holds for $\mathcal{C}_{r+1}$. We move the handleslides of V_r past q iteratively beginning with the right-most handleslide. Suppose $v_r^{i,j}$ is the right-most handleslide of V_r that has yet to be moved past q. If $i > k+1$ or $j < k$, use MCS move (12) to move $v_r^{i,j}$ past q. If $i < k$ and $j > k+1$, use MCS move (11) to move $v_r^{i,j}$ past q. If $i = k+1$ or $j = k$, use MCS move (10) to remove $v_r^{i,j}$. Since $\mu(k) = \mu(k+1)+1$ and a handleslide has endpoints on strands with the same Maslov potential, $v_r^{k,k+1}$ must be 0. It remains to consider the two cases $i = k, j > k+1$ and $i < k, j = k+1$.

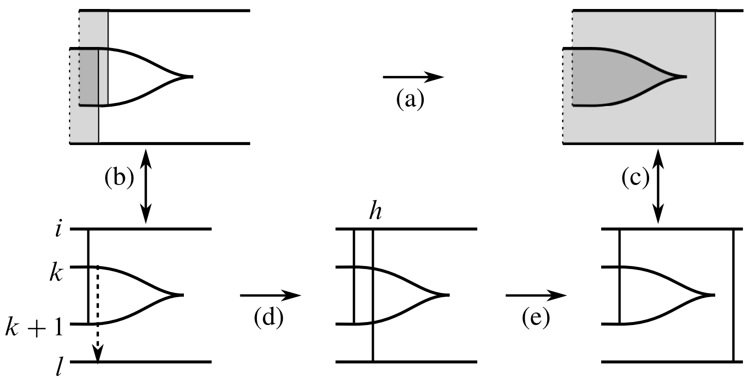

Figure 17: (a)–(e): The correspondence between Type 2 (ϵ, ϵ', H)–half-disks whose boundary intersects the right cusp in the figure and handleslides introduced by MCS move (13) to the left of the crossing.

Suppose $v_r^{i,j}$ is $v_r^{k,j}$ where $j > k+1$. Since $\mu(k) = \mu(j)$ and $\mu(k) = \mu(k+1)+1$ both hold, $\mu(k+1) = \mu(j) - 1$ holds and, thus, strands $k+1$ and j satisfy the conditions of MCS move (13). Use MCS move (13) to create new handleslide marks; see the arrow directed to the right in Figure 6. Let $(\overline{C}, \overline{d})$ be the chain complex of $\mathcal{C}_r$ just to the right of $v_r^{k,j}$. The handleslides created are of three types. By Definition 2.3(4), $\langle \overline{d} e_k, e_{k+1} \rangle$ is 1. Therefore, MCS move (13) introduces a handleslide with endpoints k and j; see Figure 16(f). Use MCS move (1) to remove this handleslide and $v_r^{k,j}$; see Figure 16(g). For each l such that $\langle \overline{d} e_j, e_l \rangle$ is 1, MCS move (13) introduces a handleslide with endpoints $k+1$ and l; see Figure 16(h). Use MCS move (10) to remove this handleslide; see Figure 16(i). Suppose $l < k$ and $\langle \overline{d} e_l, e_{k+1} \rangle$ is 1. The third type of handleslide introduced by MCS move (13) has endpoints l and j; see Figure 16(d). Let h be this handleslide. Use MCS move (11) to move h past q; see Figure 16(e). Recall that (C, d) is the chain complex of $\mathcal{C}_r$ to the immediate right of V_r. Since V_r is ordered, the handleslides between $(\overline{C}, \overline{d})$ and (C, d) have endpoints on strands $k+1, \dots, s_p$. The coefficient $\langle \overline{d} e_l, e_{k+1} \rangle$ is unaffected by such handleslides and, thus, $\langle \overline{d} e_l, e_{k+1} \rangle = \langle d e_l, e_{k+1} \rangle$ holds. Therefore, h exists if and only if $\langle de_l, e_{k+1}\rangle \cdot v_r^{k,j}$ is 1. As we noted earlier, $\langle de_l, e_{k+1}\rangle \cdot v_r^{k,j}$ is 1 if and only if the mod 2 count of Type 1 half-disks in $\mathcal{H}(x_{r+1}, [l, j+2])$ is 1. Therefore, h exists if and only if the mod 2 count of Type 1 half-disks in $\mathcal{H}(x_{r+1}, [l, j+2])$ is 1.

Suppose $v_r^{i,j}$ is $v_r^{i,k+1}$ where $i < k$. Note that $\mu(i) = \mu(k) - 1$ holds and, thus, strands i and k satisfy the conditions of MCS move (13). Use MCS move (13) to create new handleslides to the immediate right of $v_r^{i,k+1}$. Suppose $l > k+1$ and $\langle \overline{d} e_k, e_l\rangle$ is 1. MCS move (13) introduces a handleslide with endpoints i and l; see Figure 17(d). Let h be this handleslide. Use MCS move (11) to move h past q; see Figure 17(e). Following an analogous argument as was used in the case of a Type 1 half-disk, h exists if and only if the mod 2 count of Type 2 half-disks in $\mathcal{H}(x_{r+1}, [i, l+2])$ is 1. MCS move (13) also introduces handleslides analogous to those in Figure 16(f) and (h), which are removed in same manner as was done in Figure 16(g) and (i).

Once we have applied the above algorithm to each handleslide in V_r, we are left with a collection of handleslides V to the right of q. The ordering of V_r ensures the only new handleslides were those introduced by applications of MCS move (13). Therefore, given $1 \leq i < j \leq s_{r+1}$, there may be up to 3 handleslides in V with endpoints on i and j; one counts (ϵ, ϵ', H)–half-disks extended past q as in Figure 9, one counts Type 1 half-disks as in Figure 16(a)–(e), and the third counts Type 2 half-disks as in Figure 17(a)–(e). Use MCS moves (1), (3), (5), and (6) to remove pairs of handleslides with identical endpoints and order V. In particular, V can be ordered without creating new handleslides. The resulting ordered collection of handleslides is V_r and the MCS is $\mathcal{C}_{r+1}$. If $j < k$ or $i \geq k$, then

$$v_{r+1}^{i,j} = v_r^{\pi(i),\pi(j)}$$

holds and, if $j \geq k$ and $i < k$, then

$$v_{r+1}^{i,j} = v_r^{\pi(i),\pi(j)} + v_r^{\pi(i),k+1} \cdot \langle d' e_k, e_{\pi(j)}\rangle + \langle de_{\pi(i)}, e_{k+1}\rangle \cdot v_r^{k,\pi(j)}$$

holds. These equations, along with Equations (10) and (12), imply Property 1(b) holds for $\mathcal{C}_{r+1}$. Finally, since q is not a crossing and Property 1(a) holds for $\mathcal{C}_r$, it must hold for $\mathcal{C}_{r+1}$ as well.

This completes the construction of the MCSs $\mathcal{C}_0, \dots, \mathcal{C}_s$. □

4 Corollaries to Theorem 1.1

In the following corollaries to Theorem 1.1, D is the σ–generic front diagram of a Legendrian knot with rotation number 0. Recall that an augmentation ϵ in $\mathrm{Aug}(D)$ has a corresponding A–form MCS $\mathcal{C}$ where, for a degree 0 crossing q, $\epsilon(q) = 1$ holds if and only if q is marked by $\mathcal{C}$.

Corollary 4.1 *The count of MCS classes of a Legendrian knot is a Legendrian isotopy invariant.*

Corollary 4.1 follows from the fact that the count of homotopy classes of augmentations is a Legendrian isotopy invariant and every Legendrian knot is Legendrian isotopic to a Legendrian knot with σ–generic front diagram by an arbitrarily small Legendrian isotopy. Corollary 4.1 is stated and a proof is briefly sketched by Petya Pushkar in a letter to Dmitry Fuchs from 2000. The proposed proof investigates the effect of Legendrian Reidemeister moves on the number of MCS classes and is different from the approach in this article.

Given the Chekanov–Eliashberg algebra $(\mathcal{A}(D), \partial)$, the differential $\partial^\epsilon\colon \mathcal{A}(D) \to \mathcal{A}(D)$ is $\phi^\epsilon \circ \partial \circ (\phi^\epsilon)^{-1}$, where $\phi^\epsilon\colon \mathcal{A}(D) \to \mathcal{A}(D)$ is the algebra map defined on generators by $\phi^\epsilon(q) = q + \epsilon(q)$. The group LCH$(\epsilon)$, called the *linearized contact homology* of ϵ, is the homology of the chain complex $(A(D), \partial_1^\epsilon)$, where $\partial_1^\epsilon(q)$ is the length 1 monomials of $\partial^\epsilon(q)$. By [2], the set $\{\mathrm{LCH}(\epsilon)\}_{\epsilon \in \mathrm{Aug}(D)}$ is a Legendrian isotopy invariant, which we will call the *LCH invariant*.

Corollary 4.2 *If ϵ and ϵ' are homotopic as augmentations, then LCH(ϵ) and LCH(ϵ') are isomorphic as homology groups. Therefore, augmentation homotopy classes have well-defined linearized contact homology groups.*

Proof We will apply two theorems from [10]. In order to do so, the front diagram must be "nearly plat". A front diagram is *plat* if all left cusps have the same x–coordinate, all right cusps have the same x–coordinate, and no two crossings have the same x–coordinate. A front diagram is *nearly plat* if it is the result of perturbing a plat front diagram slightly so that no two cusps have the same x–coordinate.

We now deduce the corollary in the case that D is nearly plat. Suppose ϵ and ϵ' are homotopic. By Lemma 3.1, the A–form MCSs $\mathcal{C}$ and $\mathcal{C}'$ corresponding to ϵ and ϵ' are equivalent as MCSs. In [11], differential graded algebras $(\mathcal{A}_\mathcal{C}, d)$ and $(\mathcal{A}_{\mathcal{C}'}, d')$ are assigned to $\mathcal{C}$ and $\mathcal{C}'$, respectively. The linear level of each algebra is a chain complex $(A_\mathcal{C}, d_1)$ and $(A_{\mathcal{C}'}, d_1')$, respectively. By [11, Theorem 5.5], $(A_\mathcal{C}, d_1)$ and $(A_{\mathcal{C}'}, d_1')$ are isomorphic. By [11, Theorem 7.3], $(A(D), \partial_1^\epsilon)$ is isomorphic to $(A_\mathcal{C}, d_1)$ and $(A(D), \partial_1^{\epsilon'})$ is isomorphic to $(A_{\mathcal{C}'}, d_1')$. Therefore, LCH(ϵ) and LCH(ϵ') are isomorphic as homology groups.

For the general case of a Chekanov–Eliashberg algebra $(\mathcal{A}, \partial)$ assigned to a front (or Lagrangian) diagram that is not nearly plat, we argue as follows. By [1], the Chekanov–Eliashberg algebras assigned to Legendrian isotopic Legendrian knots are stable tame isomorphic. Any Legendrian knot is Legendrian isotopic to a knot with nearly plat front diagram, therefore $(\mathcal{A}, \partial)$ is stable tame isomorphic to a DGA that satisfies the property stated in Corollary 4.2. We then verify that $(\mathcal{A}, \partial)$ also satisfies the corollary in two steps.

Step 1 The corollary holds for a stabilization $(S(\mathcal{A}), \partial')$ of a DGA $(\mathcal{A}, \partial)$ if and only if it holds for $(\mathcal{A}, \partial)$.

Here, $S(\mathcal{A})$ is obtained from $\mathcal{A}$ by adding two generators x and y in successive degrees, and the differential satisfies $\partial'|_{\mathcal{A}} = \partial$ and $\partial' x = y$. Restricting augmentations of $S(\mathcal{A})$ to $\mathcal{A}$ provides a surjection from the set of augmentations of $S(\mathcal{A})$ to the set of augmentations of $\mathcal{A}$, and this gives a well-defined bijection between homotopy classes of augmentations of $S(\mathcal{A})$ and $\mathcal{A}$. Moreover, for any augmentation $\epsilon\colon S(\mathcal{A}) \to \mathbb{Z}/2$, the linearized homology groups associated to ϵ and $\epsilon|_{\mathcal{A}}$ are isomorphic, so Step 1 follows.

Step 2 If $\varphi\colon (\mathcal{A}_1, \partial_1) \to (\mathcal{A}_2, \partial_2)$ is an isomorphism of DGAs, then the corollary holds for $(\mathcal{A}_1, \partial_1)$ if and only if it holds for $(\mathcal{A}_2, \partial_2)$.

To see this, observe that $\epsilon_2 \mapsto \epsilon_1 \circ \varphi$ gives a bijection from augmentations of $(\mathcal{A}_2, \partial_2)$ to augmentations of $(\mathcal{A}_1, \partial_1)$ that preserves homotopy classes and linearized homology groups. □

Corollary 4.2 provides a means for strengthening the LCH invariant. The set

$$\{\mathrm{LCH}(\epsilon)\}_{\epsilon \in \mathrm{Aug}(D)},$$

along with a count of the number of augmentation homotopy classes associated with each group, is a Legendrian isotopy invariant. The authors are currently unaware of an example where this refinement is able to distinguish knots that are not already distinguished by the LCH invariant taken without regard to multiplicity.

Corollary 4.3 *If ϵ and ϵ' are homotopic, then ϵ and ϵ' are mapped to the same graded normal ruling by the many-to-one map from augmentations to graded normal rulings defined in [16].*

Proof Suppose ϵ and ϵ' are homotopic. By Lemma 3.1, the A–form MCSs $\mathcal{C}$ and $\mathcal{C}'$ corresponding to ϵ and ϵ' are equivalent. By [10, Lemma 3.14], every MCS determines a graded normal ruling. By [10, Proposition 3.15], equivalent MCSs determine the same graded normal ruling. Therefore, $\mathcal{C}$ and $\mathcal{C}'$ determine the same graded normal ruling. In [16], there is an algorithmically defined many-to-one map Ω from $\mathrm{Aug}(D)$ to the set of graded normal rulings of D. In the case of an augmentation ϵ and its corresponding A–form MCS $\mathcal{C}$, $\Omega(\epsilon)$ is the same as the graded normal ruling determined by $\mathcal{C}$ in [10, Lemma 3.14]. Therefore, Ω maps ϵ and ϵ' to the same graded normal ruling. □

References

[1] **Y Chekanov**, *Differential algebra of Legendrian links*, Invent. Math. 150 (2002) 441–483 MR1946550

[2] **Y V Chekanov**, *Invariants of Legendrian knots*, from: "Proceedings of the International Congress of Mathematicians, II", (T Li, editor), Higher Ed. Press, Beijing (2002) 385–394 MR1957049

[3] **Y V Chekanov**, **P E Pushkar'**, *Combinatorics of fronts of Legendrian links, and Arnol'd's 4–conjectures*, Uspekhi Mat. Nauk 60 (2005) 99–154 MR2145660 In Russian; translated in Russian Math. Surveys 60 (2005) 95–149

[4] **T Ekholm**, **K Honda**, **T Kálmán**, *Legendrian knots and exact Lagrangian cobordisms*, preprint `arXiv:1212.1519`

[5] **Y Eliashberg**, **A Givental**, **H Hofer**, *Introduction to symplectic field theory*, Geom. Funct. Anal. (2000) 560–673 MR1826267

[6] **Y Félix**, **S Halperin**, **J-C Thomas**, *Differential graded algebras in topology*, from: "Handbook of algebraic topology", (I M James, editor), North-Holland, Amsterdam (1995) 829–865 MR1361901

[7] **D Fuchs**, *Chekanov–Eliashberg invariant of Legendrian knots: Existence of augmentations*, J. Geom. Phys. 47 (2003) 43–65 MR1985483

[8] **D Fuchs**, **T Ishkhanov**, *Invariants of Legendrian knots and decompositions of front diagrams*, Mosc. Math. J. 4 (2004) 707–717, 783 MR2119145

[9] **D Fuchs**, **D Rutherford**, *Generating families and Legendrian contact homology in the standard contact space*, J. Topol. 4 (2011) 190–226 MR2783382

[10] **M B Henry**, *Connections between Floer-type invariants and Morse-type invariants of Legendrian knots*, Pacific J. Math. 249 (2011) 77–133 MR2764942

[11] **M B Henry**, **D Rutherford**, *A combinatorial DGA for Legendrian knots from generating families*, Commun. Contemp. Math. 15 (2013) MR3037194

[12] **M B Henry**, **D Rutherford**, *Ruling polynomials and augmentations over finite fields*, J. Topol. 8 (2015) 1–37 MR3335247

[13] **J Jordan**, **L Traynor**, *Generating family invariants for Legendrian links of unknots*, Algebr. Geom. Topol. 6 (2006) 895–933 MR2240920

[14] **T Kálmán**, *Contact homology and one parameter families of Legendrian knots*, Geom. Topol. 9 (2005) 2013–2078 MR2209366

[15] **L L Ng**, *Computable Legendrian invariants*, Topology 42 (2003) 55–82 MR1928645

[16] **L L Ng**, **J M Sabloff**, *The correspondence between augmentations and rulings for Legendrian knots*, Pacific J. Math. 224 (2006) 141–150 MR2231655

[17] **J M Sabloff**, *Augmentations and rulings of Legendrian knots*, Int. Math. Res. Not. 2005 (2005) 1157–1180 MR2147057

[18] **J M Sabloff**, **L Traynor**, *Obstructions to Lagrangian cobordisms between Legendrians via generating families*, Algebr. Geom. Topol. 13 (2013) 2733–2797 MR3116302

[19] **L Traynor**, *Legendrian circular helix links*, Math. Proc. Cambridge Philos. Soc. 122 (1997) 301–314 MR1458235

[20] **L Traynor**, *Generating function polynomials for Legendrian links*, Geom. Topol. 5 (2001) 719–760 MR1871403

Department of Mathematics, Siena College
515 Loudon Road, Loudonville, NY 12211, USA

Department of Mathematics, Ball State University
2000 W University Ave, Muncie, IN 47306, USA

mbhenry@siena.edu, rutherford@bsu.edu

Received: 25 July 2014 Revised: 10 April 2015