

Article

Symmetry-Breaking Bifurcations of the Information Bottleneck and Related Problems

Albert E. Parker ¹  and Alexander G. Dimitrov ^{2,*} 

¹ Center for Biofilm Engineering, Department of Mathematical Sciences, Montana State University, Bozeman, MT 59717, USA

² Department of Mathematics and Statistics, Washington State University Vancouver, Vancouver, WA 98686, USA

* Correspondence: alex.dimitrov@wsu.edu

Abstract: In this paper, we investigate the bifurcations of solutions to a class of degenerate constrained optimization problems. This study was motivated by the Information Bottleneck and Information Distortion problems, which have been used to successfully cluster data in many different applications. In the problems we discuss in this paper, the distortion function is not a linear function of the quantizer. This leads to a challenging annealing optimization problem, which we recast as a fixed-point dynamics problem of a gradient flow of a related dynamical system. The gradient system possesses an S_N symmetry due to its invariance in relabeling representative classes. Its flow hence passes through a series of bifurcations with specific symmetry breaks. Here, we show that the dynamical system related to the Information Bottleneck problem has an additional spurious symmetry that requires more-challenging analysis of the symmetry-breaking bifurcation. For the Information Bottleneck, we determine that when bifurcations occur, they are only of pitchfork type, and we give conditions that determine the stability of the bifurcating branches. We relate the existence of subcritical bifurcations to the existence of first-order phase transitions in the corresponding distortion function as a function of the annealing parameter, and provide criteria with which to detect such transitions.



Citation: Parker, A.E.; Dimitrov, A.G. Symmetry-Breaking Bifurcations of the Information Bottleneck and Related Problems. *Entropy* **2022**, *24*, 1231. <https://doi.org/10.3390/e24091231>

Academic Editors: Jan Lewandowsky and Gerhard Bauch

Received: 29 June 2022

Accepted: 29 August 2022

Published: 2 September 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Keywords: information bottleneck; optimization; annealing; gradient flow; bifurcations; symmetry

1. Introduction

This paper analyzes bifurcations of solutions to constrained optimization problems of the form

$$\max_{q \in \Delta} F(q, \beta) = \max_{q \in \Delta} \left(\sum_{i=1}^N f(q^i, \beta) \right) \quad (1)$$

as a function of a scalar parameter β and a quantizer or classifier $q = (q^1, \dots, q^N)$ with $q^i \in \mathbb{R}^K$. The real-valued function f is sufficiently smooth, and Δ is the constraint space of valid quantizers, a convex set of discrete probabilities (simplices).

This type of problem arises in Rate Distortion Theory [1,2], Deterministic Annealing [3] and biclustering [4]. The specific motivations for the abstract problem formulation given in (1) are the Information Bottleneck [5] and Information Distortion [6] functions

$$\max_{q \in \Delta} F(q, \beta) = \max_{q \in \Delta} (D(q) - \beta I(Y; T)). \quad (2)$$

These were proposed in [5,7] to analyze the Markov chain $X \rightarrow Y \rightarrow T$ in which $X \rightarrow Y$, characterized by a probability $p(X, Y)$, is the original system of interest, characterized by its mutual information $I(X; Y)$, and T is a simplification (quantized version of) Y . Here we work mainly with discrete versions of Y and T , with cardinalities $|Y| = K$ and $|T| = N$. Typically $N \ll K$. $I(Y; T)$ is the mutual information between the K objects in Y and the

N clusters in T . The goal is to cluster K objects in Y into N clusters in T given inputs X such that the function F is maximized in $[q^i]$; the probability that the j th element of Y is classified as being a member of the cluster with label $i \in T$. We call such a set of conditional probabilities a *stochastic quantizer*, or just a *quantizer*, to relate to the vector quantization literature [8]. The annealing parameter $\beta \in [0, \infty)$.

It has been shown that finding hard-clustering solutions to (2) is NP-complete (combinatorial search) when $D(q)$ is the mutual information $I(X; T)$ [9], as in the Information Bottleneck [5,10,11] and the Information Distortion [7,12,13] methods. Information Bottleneck (IB) approaches are gaining in penetration into multiple scientific and engineering domains [14–18]. As they typically involve the nonlinear optimization problem (2), there is need for optimization methods for such problems that can avoid the rise in complexity implied by the NP-complete hard-clustering solutions [9]. Originally, Tishby et al. [5] approached this problem with an algorithm inspired by the Blahut–Arimoto approach to solving Rate-Distortion types of problems ([2], Chapter 10). The “self-consistent” equations in [5] optimize both the quantizer and the “relevance” distribution $p(x|t)$. However, unlike the classic Blahut–Arimoto algorithm, which can guarantee convergence to a unique solution to its iterative scheme because of the convex geometry of the two state spaces, the “self-consistent” equations have no such guarantee due to the more-complicated geometry of *three* convex sets over which the optimization is performed, as also noted in [5]. Accordingly, in this work, we use the original optimization problem (2) over a single variable: the quantizer (conditional probability) $q(t|y)$. It may be possible that a related Blahut–Arimoto style optimization coupled to the bifurcation structure of its gradient flow discussed here can lead to additional insights into this problem, but we consider this beyond the scope of this particular manuscript.

We have investigated the structure of soft-clustering annealing-type methods that reach the hard-clustering solution in the limit of the annealing parameter [19,20] through a series of bifurcations. A bifurcation in this context is a point that is a solution (q^*, β^*) to (2) such that the number of solutions to (2) changes in a small neighborhood of (q^*, β^*) . Because a bifurcation corresponds to a point at which some of the objects Y have just been classified, in the IB literature, a bifurcation is usually referred to as a phase transition. One of the goals of this and related work is to understand why annealing-type algorithms, such as the original optimization heuristics in [5,10], work as well as they do. This can help with designing further optimization heuristics and can assess how close those can get to the global solutions to IB problems. We believe that this amalgamation of optimization theory and dynamical systems theory, as stated in [19,20], can provide a solid foundation with which to address such optimization challenges.

Because of the form (1) of F , it possesses certain symmetries. That is, the value of $F(q, \beta)$ does not change (is invariant) under arbitrary permutations of the vectors q^i . In other words, F is S_N -invariant. The form (1) further implies that the Hessian $d_q^2 F(q)$ is block diagonal with blocks $\{d_{q^i}^2 f(q^i)\}_{i=1}^N$. These conditions are met by the Information Distortion function [6],

$$F_H(q, \beta) = H(T|Y) + \beta I(X, T), \tag{3}$$

where $H(T|Y)$ is the entropy, and by the cost function used in the original IB method [5],

$$F_{IB}(q, \beta) = -I(Y, T) + \beta I(X, T), \tag{4}$$

which is the focus of this manuscript. Both the Information Distortion and Information Bottleneck problems have the form given in (1) and (2). Importantly, $d^2 F_{IB}(q)$ has a “perpetual kernel” since each block $d^2 f(q^i)$ has the eigenpair $(0, q^i)$ for every q [20]. In other words, the Hessian $d^2 F$ is singular for every q and every value of β . This makes bifurcation detection challenging because bifurcations can usually be detected by identifying isolated singularities of $d^2 F$. This degeneracy is a consequence of the translational symmetry of

F_{IB} : if $\mathbf{k} \in \ker d_q^2 F_{IB}(q^*)$, then $F_{IB}(q^*) = F_{IB}(q^* + t\mathbf{k})$ for all $t \in \mathfrak{R}$ such that $q^* + t\mathbf{k} \in \Delta$. At bifurcations of solutions to (4), the translational symmetry never breaks.

To better understand bifurcations of solutions to problems of the form (1), which includes the problems (3) and (4), we consider the gradient flow

$$\begin{pmatrix} \dot{q} \\ \dot{\lambda} \end{pmatrix} = \nabla \mathcal{L}(q, \lambda, \beta)$$

Equilibria of this flow correspond to critical points of (1), where \mathcal{L} is the Lagrangian with respect to the constraints imposed by Δ , and λ is the vector of Lagrange multipliers.

Previous work showed that when d^2F is generically non-singular, as occurs for the Information Distortion (3), then there are isolated singularities of $d^2\mathcal{L}$ that indicate possible bifurcations of solutions to (1). In this case, an $M > 1$ -dimensional $\ker d^2F$ necessitates an $M - 1$ -dimensional $\ker d^2\mathcal{L}$, which admits a bifurcation of solutions to (1) where symmetry breaks from S_M to $S_m \times S_n$ for every $m, n > 0$ such that $m + n = M$ [21].

Here we allow d^2F and $d^2\mathcal{L}$ to be singular for every $q \in \Delta$, as occurs for the Information Bottleneck (4). That is, the perpetual kernel for d^2F implies that $d^2\mathcal{L}$ also has a perpetual kernel $\ker d^2\mathcal{L} = \mathcal{K}_p(q)$, which means that the eigenvalue crossing condition that must occur at a bifurcation (i.e., $d^2\mathcal{L}$ must have a zero eigenvalue at a bifurcation) [20] is never satisfied in \mathcal{K}_p . There are a few challenges due to the existence of the perpetual kernel (i.e., degeneracy) of the Information Bottleneck that we address in this paper. First, detecting bifurcations may be problematic because one cannot simply monitor the determinant of either d^2F or $d^2\mathcal{L}$. Second, the standard theory that assures the existence of bifurcating branches, the Equivariance Branching Lemma, cannot be applied directly. Lastly, the spaces that contain the bifurcating solutions are always at least two-dimensional, which makes tracking the bifurcating solutions problematic.

Here we address two of these three challenges. We show that at a bifurcation, new eigenvalue(s) of d^2F_{IB} and $d^2\mathcal{L}$ must cross zero, causing $\ker d^2\mathcal{L}$ to expand so that $\ker d^2\mathcal{L}(q^*) = \mathcal{K}_p \cup \mathcal{K}^*$, where \mathcal{K}^* is the span of the eigenvectors with crossing eigenvalues. Instead of detecting bifurcations by the expensive process of monitoring the expansion of $\ker d^2\mathcal{L}$ (from \mathcal{K}_p to $\mathcal{K}_p \cup \mathcal{K}^*$), we give a simple way to check the eigenvalue crossing condition for annealing problems $F = G(q) + \beta D(q)$ as in (2) [20]. We prove the existence of the bifurcating branches by adapting the standard proof for the Equivariant Branching Lemma. This newly developed theory guarantees that bifurcating branches exist in \mathcal{K}^* , are generically pitchforks, and that symmetry breaks from S_M to $S_m \times S_n$. Additionally, we give conditions to check whether the pitchforks are subcritical or supercritical, and how stability of the bifurcating branches relates to optimality in the optimization problem (1).

2. Bifurcation Analysis

2.1. Equivariant Branching Lemma

The Equivariant Branching Lemma relates the subgroup structure of a symmetry group Γ with the existence of symmetry-breaking bifurcating branches of equilibria of $\dot{x} = f(x, \beta)$. Observe that we present a version that does not require absolute irreducibility. For a proof see [22] p. 83.

Theorem 1 (Equivariant Branching Lemma). *Let f be a smooth function $f : V \times \mathfrak{R} \rightarrow V$ that is Γ -equivariant for a compact Lie group Γ and a Banach space V . Let Σ be an isotropy subgroup of Γ with $\dim \text{Fix}(\Sigma) = 1$. Suppose that $\text{Fix}(\Gamma) = \{0\}$ and the crossing condition $d_{\beta x}^2 f(0, 0)\mathbf{x}_0 \neq 0$ for $x_0 \in \text{Fix}(\Sigma)$. Then there exists a unique smooth solution branch $(tx_0, \beta(t))$ to $f = 0$ with isotropy subgroup Σ .*

For an arbitrary Γ -equivariant system where bifurcation occurs at (x^*, β^*) , the requirement in Theorem 1 that the bifurcation occurs at the origin is accomplished by a translation. Assuring that the Jacobian vanishes, $d_x f(0, 0) = 0$, can be effected by restrict-

ing and projecting the system onto the kernel of the Jacobian. This transform is called the Liapunov–Schmidt reduction (see [23]).

The Equivariant Branching Lemma does not directly apply to yield bifurcating branches for the problem (1) at q for which d^2F is singular for the following reasons:

- \mathcal{K}_p and \mathcal{K}^* have independent bases, which implies that each is invariant to the action of S_N , and so the decomposition $\ker d^2\mathcal{L}(q^*) = \mathcal{K}_p \times \mathcal{K}^*$ shows that S_N does not act absolutely irreducibly on $\ker d^2F(q^*)$, but it does act absolutely irreducibly on each of these disjoint subspaces separately. This is why we present a version of the Equivariant Branching Lemma that does not require absolute irreducibility.
- The Liapunov–Schmidt reduction onto $\ker d^2\mathcal{L}(q^*)$ is clear, but not onto \mathcal{K}^* .
- $\text{Fix}(S_m \times S_n) \cap \ker d^2\mathcal{L}(q^*)$ is two-dimensional with basis

$$\{(n\mathbf{v}, \dots, n\mathbf{v}, -m\mathbf{v}, \dots, -m\mathbf{v}), (n\mathbf{y}, \dots, n\mathbf{y}, -m\mathbf{y}, \dots, -m\mathbf{y})\},$$

where $\mathbf{v}, \mathbf{y} \in \mathbb{R}^K$.

We address these issues in the manuscript and show that a small modification of the Equivariant Branching Lemma allows for similar analysis to be successfully applied to Information Bottleneck-style problems such as (2) with minimal modifications to the original algorithm from [20].

2.2. A Gradient Flow

We now lay the groundwork necessary to determine the bifurcations of local solutions to (1)

$$\max_{q \in \Delta} F(q, \beta),$$

where $F = \sum_{i=1}^N f(q^i, \beta)$, which includes as a special case the Information Distortion (3) and Information Bottleneck (4) problems. The convex set of discrete conditional probabilities is

$$\Delta := \left\{ q \in \mathbb{R}^{NK} \mid \sum_{i=1}^N q_k^i = 1 \quad \forall k : 1 \leq k \leq K \text{ and } q_k^i \geq 0 \quad \forall i, k \right\}.$$

Due to the form of F , it has the following properties:

1. $F(q, \beta)$ is an S_N -invariant, real-valued function of q , where the action of S_N on q permutes the component vectors $q^i, i = 1, \dots, N$, of $q \in \Delta$.
2. The $NK \times NK$ Hessian $d_q^2F(q, \beta)$ is block diagonal, where the i th $K \times K$ block is $d^2f(q^i)$.

The Lagrangian of (1) with respect to the equality constraints from Δ is

$$\mathcal{L}(q, \lambda, \beta) = F(q, \beta) + \sum_{k=1}^K \lambda_k \left(\sum_{i=1}^N q_k^i - 1 \right). \tag{5}$$

The scalar λ_k is the Lagrange multiplier for the constraint $\sum_{i=1}^N q_k^i - 1 = 0$, and $\lambda \in \mathbb{R}^K$ is the vector of Lagrange multipliers $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_K)^T$. The gradient of the Lagrangian in (5) is

$$\nabla \mathcal{L} := \nabla_{q,\lambda} \mathcal{L}(q, \lambda, \beta) = \begin{pmatrix} \nabla_q \mathcal{L} \\ \nabla_\lambda \mathcal{L} \end{pmatrix},$$

where $\nabla_q \mathcal{L} = \nabla F(q, \beta) + \Lambda$ and $\Lambda = (\lambda^T, \lambda^T, \dots, \lambda^T)^T \in \mathbf{R}^{NK}$. The gradient $\nabla_\lambda \mathcal{L}$ is a vector of K constraints

$$\nabla_\lambda \mathcal{L} = \begin{pmatrix} \sum_i q_1^i - 1 \\ \sum_i q_2^i - 1 \\ \vdots \\ \sum_i q_K^i - 1 \end{pmatrix}.$$

Let J be the Jacobian of $d_q \nabla_\lambda \mathcal{L}$

$$J := d_q \nabla_\lambda \mathcal{L} = \underbrace{\begin{pmatrix} I_K & I_K & \dots & I_K \end{pmatrix}}_{N \text{ blocks}}. \tag{6}$$

Observe that J has full row rank. The Hessian of (5) with respect to the vector $\begin{pmatrix} q \\ \lambda \end{pmatrix} \in \mathfrak{R}^{NK+K}$ is

$$d^2 \mathcal{L}(q) := d^2 \mathcal{L}(q, \lambda, \beta) = \begin{pmatrix} d^2 F(q, \beta) & J^T \\ J & \mathbf{0} \end{pmatrix}, \tag{7}$$

where $\mathbf{0}$ is $K \times K$. The $NK \times NK$ matrix $d^2 F(q) := d_q^2 F(q, \beta)$ is the block diagonal Hessian of F with $K \times K$ blocks $\{d^2 f(q^i, \beta)\}_{i=1}^N$.

The dynamical system whose equilibria are stationary points of (1) is the gradient flow of the Lagrangian

$$\begin{pmatrix} \dot{q} \\ \dot{\lambda} \end{pmatrix} = \nabla \mathcal{L}(q, \lambda, \beta) \tag{8}$$

for \mathcal{L} as defined in (5) and $\beta \in [0, \infty)$. The equilibria of (8) are points $\begin{pmatrix} q^* \\ \lambda^* \end{pmatrix} \in \mathbf{R}^{NK+K}$ where

$$\nabla \mathcal{L}(q^*, \lambda^*, \beta) = 0.$$

The Jacobian of this system is the Hessian $d^2 \mathcal{L}(q, \lambda, \beta)$ from (7).

Remark 1. By the theory of constrained optimization [24], the equilibria (q^*, λ^*, β) of (8) where $d^2 F(q^*, \beta)$ is negative definite on $\ker J$ are local solutions of (1). Conversely, if (q^*, β) is a local solution of (1), then there exists a vector of Lagrange multipliers λ^* so that (q^*, λ^*, β) is an equilibrium of (8) (this necessary requirement is called the Karush–Kuhn–Tucker conditions) such that $d^2 F(q^*, \beta)$ is non-positive definite on $\ker J$.

2.3. Equilibria with Symmetry

Next, we categorize the equilibria of (8) according to their symmetries, which allows us to determine when to expect symmetry-breaking bifurcations.

Let $q \in \text{Fix}(S_M)$ for some $1 \leq M \leq N$. Then there exists a partition of $\{1, 2, \dots, N\}$ into the sets \mathcal{U} and \mathcal{R} , where $|\mathcal{U}| = M$, so that $q^i = q^j$ if and only if $i, j \in \mathcal{U}$. Clearly, $d^2 F$ has M identical blocks, $\{d^2 f(q^i)\}_{i \in \mathcal{U}}$.

To ease the notation, and without loss of generality, we set

$$\mathcal{U} := \{1, \dots, M\} \text{ and } \mathcal{R} := \{M + 1, \dots, N\}.$$

To distinguish between the blocks of $d^2 F$, we write

$$B := d^2 f(q^i) \text{ for } 1 \leq i \leq M \text{ and } R_i := d^2 f(q^i) \text{ for } M + 1 \leq i \leq N. \tag{9}$$

As mentioned in the introduction, we assume that for each $q \in \Delta$, each block $d^2f(q^i)$ always has at least a one-dimensional kernel with basis vector(s) which depend on q . Thus, $\dim \ker d^2F \geq N$. At an equilibrium of $(q^*, \lambda^*, \beta^*)$ of (8) where $q \in \text{Fix}(S_M)$, we consider the following three cases:

1. $\dim \ker d^2F(q^*) > N + 1$;
2. $\dim \ker d^2F(q^*) = N + 1$;
3. $\dim \ker d^2F(q^*) = N$.

We will show that the first case necessitates a symmetry-breaking bifurcation (Theorem 3). In the second case, there is no bifurcation (Corollary 1). Finally, in the third case, we expect a saddle node [21], a symmetry-preserving bifurcation.

We are able to distinguish between the three cases above by considering which blocks of $d^2F(q^*)$ have kernels that have more than one dimension. This motivates the following definition.

Definition 1. An equilibrium $(q^*, \lambda^*, \beta^*)$ of (8) is M -singular (or, equivalently, q^* is M -singular) if:

1. $q \in \text{Fix}(S_M)$ so that $q^i = q^j$ for every $1 \leq i, j \leq M$.
2. For B , the M block(s) of the Hessian defined in (9), $\ker B$ has dimension 2 with basis vectors $\mathbf{v}, \mathbf{y} \in \mathbb{R}^K$. \mathbf{v} is associated with the crossing eigenvalues, and \mathbf{y} is associated with the constant zero eigenvalue of B .
3. The $N - M$ block(s) of the Hessian $\{R_i\}_{i \in \mathcal{R}}$, defined in (9), each have a one-dimensional kernel with basis vector $\mathbf{z}(i) \in \mathbb{R}^K$.
4. The vectors \mathbf{v}, \mathbf{y} and $\{\mathbf{z}(i)\}$ are linearly independent.
5. The matrix

$$A := B \sum_{i=M+1}^N R_i^- + MI_K \tag{10}$$

is nonsingular. R_i^- is the Moore–Penrose inverse of R_i . When $M = N$, we define $A := NI_K$.

We wish to emphasize that we showed in [21] that requirements 2–5 in Definition 1 hold generically.

A straightforward calculation shows that every block of the Hessian d^2F of the Information Bottleneck cost function (2) is singular for every (q, β) , and the basis for $\ker d^2f(q^i)$ is $\mathbf{y} = q^i$ for $1 \leq i \leq M$ and $\mathbf{z}(i) = q^i$ for $M + 1 \leq i \leq N$ (Lemma 42 in [25]), which assures that these vectors are linearly independent, as in Definition 1.4. At a bifurcation, the kernels of the identical blocks B expand by \mathbf{v} as in Definition 1.2. Using the notation above, $\mathbf{y} = q^i$ for each $i \in \mathcal{U}$, and $\mathbf{z}(i) = q^i$ for each $i \in \mathcal{R}$.

2.4. The Kernel at a Bifurcation

The equilibria of (8) change their stability with β , and hence change the solutions to (1). The changes of stability are determined by the kernel of $d^2\mathcal{L}(q^*)$ at a bifurcation point q^* . In this section we show that for any $q \in \text{Fix}(S_M)$ with $M > 1$, $d^2\mathcal{L}(q^*)$ has a perpetual kernel \mathcal{K}_p that is at least $M - 1$ dimensional. The zero eigenvalues associated with the eigenvectors in \mathcal{K}_p remain constant, so that at a bifurcation point $(q^*, \lambda^*, \beta^*)$ of (8) where q^* is M -singular, new eigenvalues of $d^2\mathcal{L}$ must cross zero. Thus, the kernel expands, and the bifurcating directions exist in an “expanded” kernel of $d^2\mathcal{L}(q^*)$, $\ker d^2\mathcal{L}(q^*) = \mathcal{K}^* \times \mathcal{K}_p$.

We determine a basis for $\ker d^2\mathcal{L}$ at an M -singular q^* when $M > 1$. If q is 1-singular with a trivial isotropy group (i.e., no symmetry), then $d^2\mathcal{L}(q^*)$ is non-singular— \mathcal{K}_p disappears. First, we ascertain a basis for $\ker d^2F(q^*)$.

Recall that in the preliminaries, when $\mathbf{x} \in \mathfrak{R}^{NK}$, we defined $\mathbf{x}^j \in \mathbf{R}^K$ to be the j th vector component of \mathbf{x} . We now define the linearly independent vectors $\{\mathbf{v}_i\}_{i=1}^M$, $\{\mathbf{y}_i\}_{i=1}^M$ and $\{\mathbf{z}_k\}_{k=M+1}^N$ in \mathfrak{R}^{NK} by

$$\mathbf{v}_i^j := \begin{cases} \mathbf{v} & \text{if } 1 \leq i = j \leq M \\ \mathbf{0} & \text{otherwise} \end{cases} \quad \mathbf{y}_i^j := \begin{cases} \mathbf{y} & \text{if } 1 \leq i = j \leq M \\ \mathbf{0} & \text{otherwise} \end{cases} \quad , \quad \mathbf{z}_k^j := \begin{cases} \mathbf{z}(i) & \text{if } M+1 \leq j = k \leq N \\ \mathbf{0} & \text{otherwise} \end{cases} \quad (11)$$

where $\mathbf{0} \in \mathfrak{R}^K$, and \mathbf{v} and \mathbf{y} are defined in Definition 1.2. For example, if $M = 2$ and $N = 3$, then $\mathbf{v}_1 := (\mathbf{v}^T, \mathbf{0}, \mathbf{0})^T$ and $\mathbf{v}_2 := (\mathbf{0}, \mathbf{v}^T, \mathbf{0})^T$.

Due to the block diagonal form of $d^2F(q^*)$, it is easy to see that the $N + M$ vectors defined in (11) form a basis for $\ker d^2F(q^*)$.

Now, let

$$V_i = \begin{pmatrix} \mathbf{v}_i \\ \mathbf{0} \end{pmatrix} - \begin{pmatrix} \mathbf{v}_M \\ \mathbf{0} \end{pmatrix}, \quad Y_i = \begin{pmatrix} \mathbf{y}_i \\ \mathbf{0} \end{pmatrix} - \begin{pmatrix} \mathbf{y}_M \\ \mathbf{0} \end{pmatrix}, \quad Z_k = \begin{pmatrix} \mathbf{z}_k \\ \mathbf{0} \end{pmatrix} - \begin{pmatrix} \mathbf{z}_N \\ \mathbf{0} \end{pmatrix} \quad (12)$$

for $i = 1, \dots, M - 1$ and $M + 1 \leq k \leq N - 1$ where $\mathbf{0} \in \mathfrak{R}^K$. From (7), it is easy to see that these three sets of vectors are in $\ker d^2\mathcal{L}(q^*)$. The next theorem shows that $\{V_i\}_{i=1}^{M-1} \cup \{Y_i\}_{i=1}^{M-1}$ are a basis for $\ker d^2\mathcal{L}(q^*)$. This natural partition of the basis vectors shows that $\ker d^2\mathcal{L}(q^*)$ can be written as $\ker d^2\mathcal{L}(q^*) = \mathcal{K}_p \times \mathcal{K}^*$. According to Definition 1, the ‘‘perpetual kernel’’ corresponding to constant zero eigenvalues of $d^2\mathcal{L}(q^*)$ is generated by

$$\mathcal{K}_p = \langle \{Y_i\}_{i=1}^{M-1} \rangle .$$

The part of the kernel that arises at a bifurcation corresponding to eigenvalues crossing zero is

$$\mathcal{K}^* = \langle \{V_i\}_{i=1}^{M-1} \rangle .$$

The vectors $\{Z_k\}$ do not contribute to $\ker d^2\mathcal{L}(q^*)$.

Theorem 2. *If q^* is M -singular for $1 < M \leq N$, then $\{V_i\} \cup \{Y_i\}$ from (12) are a basis for $\ker d^2\mathcal{L}(q^*)$.*

Proof. To show that $\{V_i\}_{i=1}^{M-1} \cup \{Y_i\}_{i=1}^{M-1}$ span $\ker d^2\mathcal{L}(q^*)$, let $\mathbf{k} \in \ker d^2\mathcal{L}(q^*)$ and decompose it as

$$\mathbf{k} = \begin{pmatrix} \mathbf{k}_F \\ \mathbf{k}_J \end{pmatrix} \quad (13)$$

where \mathbf{k}_F is $NK \times 1$, and \mathbf{k}_J is $K \times 1$. Hence,

$$\begin{aligned} d^2\mathcal{L}(q^*, \lambda^*, \beta)\mathbf{k} &= \begin{pmatrix} d^2F(q^*, \beta^*) & J^T \\ J & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{k}_F \\ \mathbf{k}_J \end{pmatrix} = \mathbf{0} \\ \implies d^2F(q^*, \beta)\mathbf{k}_F &= -J^T\mathbf{k}_J \\ J\mathbf{k}_F &= \mathbf{0}. \end{aligned} \quad (14)$$

Now, from (6) and the fact that d^2F is block diagonal, we have

$$\begin{pmatrix} d^2f(q^1) & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & d^2f(q^2) & \dots & \mathbf{0} \\ \vdots & \vdots & \dots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & d^2f(q^N) \end{pmatrix} \mathbf{k}_F = - \begin{pmatrix} \mathbf{k}_J \\ \mathbf{k}_J \\ \vdots \\ \mathbf{k}_J \end{pmatrix}. \quad (15)$$

We set

$$\mathbf{k}_F := (\mathbf{x}_1^T \ \mathbf{x}_2^T \ \dots \ \mathbf{x}_N^T)^T, \quad (16)$$

and using the notation from (9), then (15) implies

$$\begin{aligned} B\mathbf{x}_i &= -\mathbf{k}_j \text{ for } 1 \leq i \leq M \\ R_i\mathbf{x}_i &= -\mathbf{k}_j \text{ for } M + 1 \leq i \leq N. \end{aligned} \tag{17}$$

It follows that $\mathbf{x}_i = R_i^- B\mathbf{x}_1$ for every $M + 1 \leq i \leq N$. By (14), we have that $\sum_{i=1}^N \mathbf{x}_i = \mathbf{0}$, and so

$$\begin{aligned} \sum_{i=1}^M \mathbf{x}_i + \sum_{i=M+1}^N \mathbf{x}_i &= \mathbf{0} \\ \implies \sum_{i=1}^M \mathbf{x}_i + \sum_{i=M+1}^N R_i^- B\mathbf{x}_1 &= \mathbf{0}. \end{aligned}$$

By (17), for every $1 \leq i \leq M$, \mathbf{x}_i can be written as $\mathbf{x}_i = \mathbf{x}_p + d_i\mathbf{v} + e_i\mathbf{y}$, where $\mathbf{x}_p \in \text{range}(B)$, $d_i, e_i \in \mathfrak{R}$, and \mathbf{v} and \mathbf{y} are the basis vectors of $\ker B$ from Definition 1.2. Thus,

$$\begin{aligned} B \sum_{i=1}^M (\mathbf{x}_p + d_i\mathbf{v} + e_i\mathbf{y}) + B \sum_{i=M+1}^N R_i^- B(\mathbf{x}_p + d_1\mathbf{v} + e_1\mathbf{y}) &= \mathbf{0} \\ \Leftrightarrow (B \sum_{i=M+1}^N R_i^- + MI_K)B\mathbf{x}_p &= \mathbf{0} \\ \Leftrightarrow B\mathbf{x}_p &= \mathbf{0} \end{aligned}$$

since $A = B \sum_{i=M+1}^N R_i^- + MI_K$ is nonsingular. This shows that $\mathbf{x}_p = \mathbf{0}$. Therefore, $\mathbf{x}_i = d_i\mathbf{v} + e_i\mathbf{y}$ for every $1 \leq i \leq M$. Now (17) shows that $\mathbf{k}_j = \mathbf{0}$, and so $\mathbf{x}_i \in \ker R_i$ for $M + 1 \leq i \leq N$, which implies that

$$\mathbf{x}_i = c_i\mathbf{z}(i) \text{ for } M + 1 \leq i \leq N.$$

Hence, $\mathbf{k} = \begin{pmatrix} \mathbf{k}_F \\ \mathbf{0} \end{pmatrix}$, where $\mathbf{k}_F^i = \begin{cases} d_i\mathbf{v} + e_i\mathbf{y} & \text{if } 1 \leq i \leq M \\ c_i\mathbf{z}(i) & \text{if } M + 1 \leq i \leq N \end{cases}$, from which it follows that

$$J\mathbf{k}_F = \sum_{i=1}^N \mathbf{x}_i = \sum_{i=1}^M d_i\mathbf{v} + \sum_{i=1}^M e_i\mathbf{y} + \sum_{i=M+1}^N c_i\mathbf{z}(i) = \mathbf{0}. \tag{18}$$

Linear independence (Definition 1.4) implies that $\sum d_i = \sum e_i = d_i = 0$. Thus, $\mathbf{k}_F = \sum_{i=1}^{M-1} d_i(\mathbf{v}_i - \mathbf{v}_M) + \sum_{i=1}^{M-1} e_i(\mathbf{y}_i - \mathbf{y}_M)$. Therefore, the linearly independent vectors $\{V_i\} = \left\{ \begin{pmatrix} \mathbf{v}_i - \mathbf{v}_M \\ \mathbf{0} \end{pmatrix} \right\}$ and $\{Y_i\} = \left\{ \begin{pmatrix} \mathbf{y}_i - \mathbf{y}_M \\ \mathbf{0} \end{pmatrix} \right\}$ span $\ker d^2\mathcal{L}(q^*)$. \square

Corollary 1. *If q^* is 1-singular and has isotropy group equal to the identity, then $d^2\mathcal{L}(q^*)$ is nonsingular.*

Proof. If q is 1-singular, then $d^2F(q^*)$ has a single block B with a two-dimensional kernel. The other $N - 1$ blocks $\{R_i\}$ are distinct with one-dimensional kernels. By constructing the vectors as in (11), we see that $\dim \ker d^2F(q^*) = N + 1$ with basis vectors $\mathbf{v}_1, \mathbf{y}_1, \{\mathbf{z}_i\}_{i=2}^N$. Now, following the proof of Theorem 2, we take an arbitrary $\mathbf{k} \in \ker d^2\mathcal{L}(q^*, \lambda, \beta)$, and then decompose \mathbf{k} as in (13) and (16). The proof to Theorem 2 holds for the present case up until, and including (18). Linear independence now shows that $d_i = e_i = c_i = 0$, which implies that $\mathbf{k} = \mathbf{0}$. \square

Remark 2. *The independent bases given for \mathcal{K}_p and \mathcal{K}^* in Theorem 2 imply that each is invariant to the action of S_N , and so the decomposition $\ker d^2\mathcal{L}(q^*) = \mathcal{K}_p \times \mathcal{K}^*$ shows that S_N does not act absolutely irreducibly on $\ker d^2F(q^*)$. That is, by definition,*

$$d_{\mathbf{x}}r(\mathbf{0}, \beta) \neq c(\beta)I_{2M-2}.$$

The explicit bases show that $\mathcal{K}_p, \mathcal{K} \cong \{\mathbf{x} \in \mathbf{R}^M : \sum [\mathbf{x}]_i = 0\}$, which implies that S_M acts absolutely irreducibly on \mathcal{K}_p and \mathcal{K}^* [26]. Thus, \mathcal{K}_p and \mathcal{K}^* are each S_M -irreducible.

2.5. Liapunov–Schmidt Reduction

To show the existence of bifurcating branches from a bifurcation point $(q^*, \lambda^*, \beta^*)$ of equilibria of (8), the Equivariant Branching Lemma requires that the bifurcation is translated to $(\mathbf{0}, \mathbf{0}, 0)$ and that the Jacobian vanishes at bifurcation. To accomplish the former, consider

$$\mathcal{F}(q, \lambda, \beta) := \nabla \mathcal{L}(q + q^*, \lambda + \lambda^*, \beta + \beta^*).$$

To assure that the Jacobian vanishes, we restrict and project \mathcal{F} onto $\ker d^2\mathcal{L}(q^*)$ in a neighborhood of $(\mathbf{0}, \mathbf{0}, 0)$. This is the Liapunov–Schmidt reduction of \mathcal{F} [23],

$$\begin{aligned} r &: \mathbf{R}^{M-1} \times \mathbf{R} \rightarrow \mathbf{R}^{M-1} \\ r(\mathbf{x}, \beta) &= W^T(I - E)\mathcal{F}(W\mathbf{x} + U(W\mathbf{x}, \beta), \beta) \end{aligned} \tag{19}$$

where $W\mathbf{x} + U(W\mathbf{x}, \beta) = \begin{pmatrix} q \\ \lambda \end{pmatrix}$. The $(NK + K) \times (NK + K)$ matrix $I - E$ is the projection matrix onto $\ker \mathcal{F}(\mathbf{0}, 0) = \ker d^2\mathcal{L}(q^*)$ with $\ker(I - E) = \text{range } d^2\mathcal{L}(q^*)$. W is the $(NK + K) \times (2M - 2)$ matrix whose columns are the basis vectors $\{V_i\} \cup \{Y_i\}$ of $\ker d^2\mathcal{L}(q^*)$ from (12) so that $W\mathbf{x}$ is a vector in $\ker d^2\mathcal{L}(q^*)$. The vector function $U(W\mathbf{x}, \beta)$ is the component of (q, λ) that is in $\text{range } d^2\mathcal{L}(q^*)$ such that $E\mathcal{F}(W\mathbf{x} + U(\mathbf{x}, \beta), \beta) = \mathbf{0}$, $U(\mathbf{0}, 0) = \mathbf{0}$, and

$$d_{\mathbf{x}}U(\mathbf{0}, 0) = \mathbf{0}. \tag{20}$$

The system defined by the Liapunov–Schmidt reduction, $\dot{\mathbf{x}} = r(\mathbf{x}, \beta)$, has a bifurcation of equilibria at $(\mathbf{x} = \mathbf{0}, \beta = 0)$, which are in 1 – 1 correspondence with equilibria of (8). However, the stability of these associated equilibria is not necessarily the same.

It is straightforward to verify the following derivatives ([23] p. 32), which we will require in the sequel. The $(2M - 2) \times (2M - 2)$ Jacobian of (19) is

$$d_{\mathbf{x}}r(\mathbf{x}, \beta) = W^T(I - E)d_{q,\lambda}^2\mathcal{L}(q + q^*, \lambda + \lambda^*, \beta + \beta^*)(W + d_{\mathbf{x}}U(W\mathbf{x}, \beta)), \tag{21}$$

which shows that

$$d_{\mathbf{x}}r(\mathbf{0}, 0) = \mathbf{0} \tag{22}$$

since $\ker(I - E) = \text{range } d^2\mathcal{L}(q^*)$.

Our crossing condition at a bifurcation depends on the matrix of derivatives

$$\frac{\partial^2 r_i}{\partial \beta \partial x_j}(\mathbf{0}, 0) = d_{\beta}d^2\mathcal{L}[\mathbf{w}_i, \mathbf{w}_j] - d^3\mathcal{L}[\mathbf{w}_i, \mathbf{w}_j, L^- d_{\beta}\nabla\mathcal{L}] \tag{23}$$

where the derivatives of \mathcal{L} are evaluated at $(q^*, \lambda^*, \beta^*)$, and L^- is the Moore–Penrose-generalized inverse [27] of $d^2\mathcal{L}(q^*)$. The vectors $\{\mathbf{w}_i\}_{i=1}^{2M-2}$ are the basis vectors of $\ker d^2\mathcal{L}(q^*)$ from Theorem 2.

The $(2M - 2) \times (2M - 2) \times (2M - 2)$ three-dimensional array of second derivatives is

$$\frac{\partial^2 r_i}{\partial x_j \partial x_k}(\mathbf{0}, 0) = d^3\mathcal{L}(q^*, \lambda^*, \beta^*)[\mathbf{w}_i, \mathbf{w}_j, \mathbf{w}_k].$$

In [21], we showed that $\frac{\partial^2 r_i}{\partial x_j \partial x_k}(\mathbf{0}, 0) = 0$ whenever $i = j = k \leq M - 1$. In the present case, there are more zero entries since now the basis vectors $\{\mathbf{w}_i\}$ are of two types: $\mathbf{w}_i = V_i$ for $1 \leq i \leq M - 1$ (basis vectors of \mathcal{K}^*); or $\mathbf{w}_i = Y_{i-M+1}$ for $M \leq i \leq 2M - 2$ (basis vectors of

\mathcal{K}_p , see (12)). We now consider the case when $i, j \leq M - 1$ and $k > M - 1$. All other cases are dealt with using a similar argument. Substituting in for \mathbf{w}_i we have

$$\begin{aligned} \frac{\partial^2 r_i}{\partial x_j \partial x_k}(\mathbf{0}, 0) &= \sum_{\nu, \delta, \eta=1}^N \sum_{l, m, n=1}^K \frac{\partial^3 F(q^*, \beta^*)}{\partial q_l^\nu \partial q_m^\delta \partial q_n^\eta} [\mathbf{v}_i - \mathbf{v}_M]_l^\nu [\mathbf{v}_j - \mathbf{v}_M]_m^\delta [\mathbf{y}_{k-M+1} - \mathbf{y}_M]_n^\eta \\ &= \sum_{l, m, n=1}^K \frac{\partial^3 f(q^{\nu*}, \beta^*)}{\partial q_l^\nu \partial q_m^\nu \partial q_n^\nu} \left(\delta_{ij(k-M+1)} [\mathbf{v}]_l [\mathbf{v}]_m [\mathbf{y}]_n - [\mathbf{v}]_l [\mathbf{v}]_m [\mathbf{y}]_n \right). \end{aligned} \tag{24}$$

The vectors \mathbf{v} and \mathbf{y} are defined in (2). An immediate consequence of this calculation is that $\frac{\partial^2 r_i}{\partial x_j \partial x_k}(\mathbf{0}, 0) = 0$ whenever $i = j = k - M + 1$. Thus, similar arguments show that $\frac{\partial^2 r_i}{\partial x_j \partial x_k}(\mathbf{0}, 0) = 0$ whenever:

- $i = j = k$;
- $i - M + 1 = j = k, \quad i = j - M + 1 = k, \quad i = j = k - M + 1$;
- $i - M + 1 = j - M + 1 = k, \quad i - M + 1 = j = k - M + 1, \quad i = j - M + 1 = k - M + 1$.

Further, we get four different ‘‘cubes’’ of identical entries in the 3-D array. They are:

- For $i, j, k \leq M - 1$, not all equal, the value of the cube is

$$- \sum_{l, m, n=1}^K \frac{\partial^3 f(q^{\nu*}, \beta^*)}{\partial q_l^\nu \partial q_m^\nu \partial q_n^\nu} [\mathbf{v}]_l [\mathbf{v}]_m [\mathbf{v}]_n;$$

- For $i, j \leq M - 1$, not both equal, and $j > M - 1$, the value of the cube is

$$- \sum_{l, m, n=1}^K \frac{\partial^3 f(q^{\nu*}, \beta^*)}{\partial q_l^\nu \partial q_m^\nu \partial q_n^\nu} [\mathbf{v}]_l [\mathbf{v}]_m [\mathbf{y}]_n;$$

- For $i \leq M - 1$ and $j, k > M - 1$, not both equal, the value of the cube is

$$- \sum_{l, m, n=1}^K \frac{\partial^3 f(q^{\nu*}, \beta^*)}{\partial q_l^\nu \partial q_m^\nu \partial q_n^\nu} [\mathbf{v}]_l [\mathbf{y}]_m [\mathbf{y}]_n;$$

- For $i, j, k > M - 1$, not all equal, the value of the cube is

$$- \sum_{l, m, n=1}^K \frac{\partial^3 f(q^{\nu*}, \beta^*)}{\partial q_l^\nu \partial q_m^\nu \partial q_n^\nu} [\mathbf{y}]_l [\mathbf{y}]_m [\mathbf{y}]_n.$$

The points above will prove useful when proving that $d^2 r(\mathbf{0}, 0) = \mathbf{0}$.

The four-dimensional array of third derivatives of r is

$$\begin{aligned} \frac{\partial^3 r_i}{\partial x_j \partial x_k \partial x_l}(\mathbf{0}, 0) &= d^4 \mathcal{L}[\mathbf{w}_i, \mathbf{w}_j, \mathbf{w}_k, \mathbf{w}_l] - d^3 \mathcal{L}[\mathbf{w}_i, \mathbf{w}_j, L^- d^3 \mathcal{L}[\mathbf{w}_k, \mathbf{w}_l]] \\ &\quad - d^3 \mathcal{L}[\mathbf{w}_i, \mathbf{w}_k, L^- d^3 \mathcal{L}[\mathbf{w}_j, \mathbf{w}_l]] \\ &\quad - d^3 \mathcal{L}[\mathbf{w}_i, \mathbf{w}_l, L^- d^3 \mathcal{L}[\mathbf{w}_j, \mathbf{w}_k]] \end{aligned} \tag{25}$$

where the derivatives of \mathcal{L} are evaluated at $(q^*, \lambda^*, \beta^*)$, and L^- is the Moore–Penrose-generalized inverse [27] of $d^2 \mathcal{L}(q^*)$.

Since $\ker d^2 \mathcal{L}(q^*)$ is not absolutely irreducible, but \mathcal{K}^* is, one might try to define a Liapunov–Schmidt reduction by restricting and projecting $\nabla \mathcal{L}$ onto \mathcal{K}^* . One issue with projecting the reduction onto \mathcal{K}^* is how to define the projection matrix E so that

$$E\mathcal{F} = 0 \text{ and } (I - E)\mathcal{F} = 0 \text{ if and only if } \mathcal{F} = 0$$

holds and $E d_{\mathbf{x}r}(\mathbf{0}, 0)$ is non-singular in $\text{range}(E)$ so that the Implicit Function Theorem assures the restriction $(q, \lambda) = W\mathbf{x} + U(W\mathbf{x}, \beta)$, where $U(W\mathbf{x}) \in \text{range}(d^2\mathcal{L}(q^*))$, and $W\mathbf{x} \in \mathcal{K}^*$ instead of $W\mathbf{x} \in \ker d^2\mathcal{L}(q^*)$ as in (19) [23]. Simply ignoring the space \mathcal{K}_p by considering $U \in \text{range}(d^2\mathcal{L}(q^*))$ and $W\mathbf{x} \in \mathcal{K}^*$ amounts to setting $W\mathbf{x} = k^* + k_p$ and $k_p = 0$. Since $W\mathbf{x} + U$ is still embedded in the larger \mathfrak{R}^{NK+K} , which contains \mathcal{K}_p , then derivatives are affected by the implicit $k_p = 0$ constraint. This constraint $P_{\mathcal{K}_p}(q, \lambda) = k^* + U$ is nonlinear (and may not even be tractable) since \mathcal{K}_p depends on q , where $P_{\mathcal{K}_p}$ is a projection matrix that depends on q (see Theorem 7).

2.6. Isotropy Subgroups $S_m \times S_n$ of S_N

The decomposition $\ker d^2\mathcal{L}(q^*) = \mathcal{K}_p \times \mathcal{K}^*$ shows that $\text{Fix}(S_m \times S_n) \cap \ker d^2\mathcal{L}(q^*)$ is two-dimensional with basis vectors

$$\{(n\mathbf{y}^T, \dots, n\mathbf{y}^T, -m\mathbf{y}^T, \dots, -m\mathbf{y}^T)^T, (n\mathbf{v}^T, \dots, n\mathbf{v}^T, -m\mathbf{v}^T, \dots, -m\mathbf{v}^T)^T\}.$$

Restricted to \mathcal{K}^* , these isotropy subgroups $S_m \times S_n$ of S_M have one-dimensional fixed point spaces. This assures that we can use Theorem 1. We have the following Lemma.

Lemma 1. Let $M = m + n$ such that $M > 1$ and $m, n > 0$. Let \mathcal{U}_m be a set of m classes, and let \mathcal{U}_n be a set of n classes such that $\mathcal{U}_m \cap \mathcal{U}_n = \emptyset$ and $\mathcal{U}_m \cup \mathcal{U}_n = \{1, \dots, M\}$. Now define $\hat{\mathbf{u}}_{(m,n)} \in \mathfrak{R}^{NK}$ such that

$$\hat{\mathbf{u}}_{(m,n)}^i = \begin{cases} n\mathbf{v} & \text{if } i \in \mathcal{U}_m \\ -m\mathbf{v} & \text{if } i \in \mathcal{U}_n \\ \mathbf{0} & \text{otherwise} \end{cases}$$

where \mathbf{v} is defined as in Definition 1.2, and let

$$\mathbf{u}_{(m,n)} = \begin{pmatrix} \hat{\mathbf{u}}_{(m,n)} \\ \mathbf{0} \end{pmatrix} \tag{26}$$

where $\mathbf{0} \in \mathbf{R}^K$. Then the isotropy subgroup of $\mathbf{u}_{(m,n)}$ is $\Sigma_{(m,n)} \subset \Gamma_{\mathcal{U}}$ such that $\Sigma_{(m,n)} \cong S_m \times S_n$, where S_m permutes \mathbf{u}^i when $i \in \mathcal{U}_m$, and S_n permutes \mathbf{u}^i when $i \in \mathcal{U}_n$. The fixed point space of $\Sigma_{(m,n)}$ restricted to $\mathcal{K}^* \subset d^2\mathcal{L}(q^*)$ is one dimensional.

2.7. Bifurcating Branches

Theorem 3. Let $(q^*, \lambda^*, \beta^*)$ be an equilibrium of (8) such that q^* is M -singular for $1 < M \leq N$, and the crossing condition

$$d_{\beta}d^2\mathcal{L}[\mathbf{u}, \mathbf{u}] - d^3\mathcal{L}[\mathbf{u}, \mathbf{u}, L^-d_{\beta}\nabla\mathcal{L}] \neq 0$$

is satisfied. Then there exists bifurcating solutions, $\begin{pmatrix} q^* \\ \lambda^* \\ \beta^* \end{pmatrix} + \begin{pmatrix} t\mathbf{u}_{(m,n)} \\ \beta(t) \end{pmatrix}$, where $\mathbf{u}_{(m,n)} \in \mathcal{K}^*$ is defined in (26), for every pair (m, n) such that $M = m + n$, each with an isotropy group isomorphic to $S_m \times S_n$.

Proof. We mimic the proof of the Equivariant Branching Lemma. Let $\mathbf{u} := \mathbf{u}_{(m,n)} \in \text{Fix}(S_m \times S_n) \cap \mathcal{K}^*$ and let V be a matrix with columns composed of the $M - 1$ vectors $\{V_i\}$. Thus, there exists $\mathbf{x}_0 \in \mathfrak{R}^{M-1}$ so that $\mathbf{u} = V\mathbf{x}_0$. Since $r(\text{Fix}(S_m \times S_n) \cap \mathcal{K}^*) \subseteq \text{Fix}(S_m \times S_n) \cap \mathcal{K}^*$ (for every $\sigma \in S_m \times S_n$, $r(V\mathbf{x}) = r(\sigma V\mathbf{x})$ ($\mathbf{u} \in \text{Fix}(S_m \times S_n)$ that equals $\sigma r(V\mathbf{x})$ (by equivariance)), then $r(t\mathbf{x}_0, \beta) = h(t, \beta)\mathbf{x}_0$, where r is the Liapunov–Schmidt reduction (19), and h is a polynomial in t .

Since \mathcal{K}^* is S_M -irreducible, then $\text{Fix}(S_M) \cap \mathcal{K}^* = \{\mathbf{0}\}$ (otherwise, $\sigma\mathbf{x} = \mathbf{x}$ for some $\mathbf{x} \in \mathcal{K}^*$ for every $\sigma \in S_M$, which implies that $\text{span}(\mathbf{x})$ is an invariant subspace of \mathcal{K}^*).

Now [22] p. 75 shows that $r(\mathbf{0}, \beta) = \mathbf{0}$, and so $h(0, \beta) = 0$, from which it follows that $h(t, \beta) = tk(t, \beta)$. Thus,

$$r(t\mathbf{x}_0, \beta) = tk(t, \beta)\mathbf{x}_0. \tag{27}$$

Differentiating with respect to t yields

$$d_x r(t\mathbf{x}_0, \beta)\mathbf{x}_0 = (k(t, \beta) + td_t k(t, \beta))\mathbf{x}_0, \tag{28}$$

from which it follows that

$$k(t, \beta)\mathbf{x}_0 = d_x r(t\mathbf{x}_0, \beta)\mathbf{x}_0 - td_t k(t, \beta)\mathbf{x}_0,$$

and so $k(\mathbf{0}, 0) = 0$. Furthermore, we see that $d_\beta k(0, 0)\mathbf{x}_0 = d_{\mathbf{x}, \beta}^2 r(0, 0)\mathbf{x}_0 \neq \mathbf{0}$ by assumption (see (23)). This shows that $d_\beta k(0, 0)$ is a non-zero eigenvalue of $d_x r(t\mathbf{x}_0, \beta)$ with associated eigenvector \mathbf{x}_0 . By the Implicit Function Theorem, $k(t, \beta) = 0$ has a non-zero unique solution for $\beta = \beta(t)$. \square

2.8. The Crossing Condition for Annealing Problems

We next determine how to check the crossing condition in Theorem 3 when F is an annealing problem, as in (2)

$$F(q, \beta) = H(q) + \beta D(q).$$

First, we show that the crossing condition can be checked in terms of the Hessian of the function D . Furthermore, when G is strictly concave on $\text{span}(\{\mathbf{v}_i\})$, then the crossing condition is always satisfied, and every singularity is a bifurcation.

Theorem 4. The crossing condition

$$d_\beta d^2 \mathcal{L}[\mathbf{u}, \mathbf{u}] - d^3 \mathcal{L}[\mathbf{u}, \mathbf{u}, L^- d_\beta \nabla \mathcal{L}] \neq 0$$

given in Theorem 3 is satisfied for M -singular q for $M > 1$ if $d^2 D(q)$ is either positive or negative definite on $\text{span}(\{\mathbf{v}_i\})$.

Proof. Let $\mathbf{x}_0 \in \mathbb{R}^{2M-2}$ so that $\mathbf{u} = W\mathbf{x}_0 \in \text{Fix}(S_m \times S_n) \cap \mathcal{K}^*$. Multiplying Equation (21) on the left by \mathbf{x}_0^T and on the right by \mathbf{x}_0 yields

$$\mathbf{x}_0^T d_x r(\mathbf{0}, \beta)\mathbf{x}_0 = \mathbf{u}^T d_{q, \lambda}^2 \mathcal{L}(q^*, \lambda^*, \beta + \beta^*)(I_{NK+K} + d_w U(\mathbf{0}, \beta))\mathbf{u}. \tag{29}$$

By Theorem 2, an arbitrary $\mathbf{u} \in \mathcal{K}^*$ can be written as $\mathbf{u} = \begin{pmatrix} \hat{\mathbf{u}} \\ \mathbf{0} \end{pmatrix}$, where $\hat{\mathbf{u}} \in \text{span}(\{\mathbf{v}_i\}) \subset \ker d^2 F(q^*, \beta^*)$. Substituting this into (29) and observing that $d^2 F(q^*, \beta + \beta^*) = d^2 G(q^*) + (\beta + \beta^*)d^2 D(q^*) = d^2 F(q^*, \beta^*) + \beta d^2 D(q^*)$ yields

$$\mathbf{x}_0^T d_x r(\mathbf{0}, \beta)\mathbf{x}_0 = \beta \begin{pmatrix} \hat{\mathbf{u}}^T d^2 D(q^*) & \mathbf{0}^T \end{pmatrix} (I_{NK+K} + \partial_w U(\mathbf{0}, \beta)) \begin{pmatrix} \hat{\mathbf{u}} \\ \mathbf{0} \end{pmatrix}.$$

Differentiating with respect to β , evaluating at $\beta = 0$, and using (20) yields

$$\mathbf{x}_0^T d_{\mathbf{x}, \beta}^2 r(\mathbf{0}, 0)\mathbf{x}_0 = \hat{\mathbf{u}}^T d^2 D(q^*)\hat{\mathbf{u}}, \tag{30}$$

which must be non-zero since we assume that $d^2 D(q)$ is either positive or negative definite on $\text{span}(\{\mathbf{v}_i\})$. \square

From (30), we can get an expression for ζ , the eigenvalue of $d_{\mathbf{x},\beta}^2 r(0,0)$ with eigenvector \mathbf{x}_0 . Substituting $d_{\mathbf{x},\beta}^2 r(\mathbf{0},0)\mathbf{x}_0 = \zeta\mathbf{x}_0$ and observing that $\mathbf{x}_0^T \mathbf{x}_0 = \mathbf{x}_0^T W^T W \mathbf{x}_0 = \hat{\mathbf{u}}^T \hat{\mathbf{u}}$ yields

$$\zeta = \frac{\hat{\mathbf{u}}^T d^2 D(q^*) \hat{\mathbf{u}}}{\|\hat{\mathbf{u}}\|^2}. \tag{31}$$

The requirement that $d^2 D(q)$ is either positive or negative definite on $\text{span}(\{\mathbf{v}_i\})$ holds when $d^2 G(q^*)$ is either negative or positive definite, respectively, on $\text{span}(\{\mathbf{v}_i\})$.

Lemma 2. *Let $d^2 F(q^*, \beta^* \neq 0)$ be singular where q^* is M -singular such that $d^2 G(q^*)$ is negative (or positive) definite on $\text{span}(\{\mathbf{v}_i\})$. Then $d^2 D(q^*)$ is positive (or negative) definite on $\text{span}(\{\mathbf{v}_i\})$.*

Proof. If $\mathbf{u} \in \text{span}(\{V_i\}) \subset \ker d^2 F(q^*)$, then $\mathbf{u}^T d^2 G(q^*) \mathbf{u} + \beta^* \mathbf{u}^T d^2 D(q^*) \mathbf{u} = 0$. Since $\mathbf{u}^T d^2 G(q^*) \mathbf{u} < 0$, then $\mathbf{u}^T d^2 D(q^*) \mathbf{u} > 0$. \square

These results are important for the Information Bottleneck problem (2), where $d^2 G(q) = -d^2 I(Y; Z)$ is only non-positive definite on $\ker d^2 F(q^*)$, but is negative definite on $\text{span}(\{\mathbf{v}_i\})$. Thus, every singularity of the Information Bottleneck with $\ker d^2 \mathcal{L}(q^*) = \mathcal{K}^* \times \mathcal{K}_p$ is a bifurcation point. The space \mathcal{K}_p does not contain bifurcating branches since the crossing condition is never satisfied there: for $\mathbf{u} \in \mathcal{K}_p$, $\hat{\mathbf{u}}^T d^2 G(q) \hat{\mathbf{u}} + \beta \hat{\mathbf{u}}^T d^2 D(q) \hat{\mathbf{u}} = 0 + 0$ (by Lemma 42 in [25]), and so (Theorem 109, [25]) $\zeta = \frac{\hat{\mathbf{u}}^T d^2 D(q) \hat{\mathbf{u}}}{\|\hat{\mathbf{u}}\|} = 0$.

2.9. Bifurcation Type

Suppose that a bifurcation occurs at $(q^*, \lambda^*, \beta^*)$, where q^* is M -singular. This section examines the type of bifurcation from which emanate the branches

$$\left(\left(\begin{matrix} q^* \\ \lambda^* \end{matrix} \right) + t\mathbf{u}, \beta^* + \beta(t) \right),$$

whose existence is guaranteed by Theorem 3.

As we showed in [21], the derivative $\beta'(0) \neq 0$ indicates a transcritical bifurcation. If $\beta'(0) = 0$, then the bifurcation is degenerate, and if $\beta''(0) \neq 0$, then we have a pitchfork-like bifurcation. Further, $t\beta'(t) < 0$ for small t indicates a subcritical bifurcating branch, and $t\beta'(t) > 0$ for small t indicates a supercritical bifurcating branch.

Expressions for $\beta'(0)$ and $\beta''(0)$ are derived as follows. Differentiating $k(t, \beta) = 0$ from (27) yields

$$d_t k(t, \beta(t)) + d_\beta k(t, \beta(t))\beta'(t) = 0, \tag{32}$$

so that $\beta'(t) = -\frac{d_t k(t, \beta(t))}{d_\beta k(t, \beta(t))}$. Differentiating (28) with respect to t and then evaluating at $t = 0$ shows that

$$\beta'(0) = \frac{-d_{\mathbf{x}}^2 r(\mathbf{0},0)[\mathbf{x}_0, \mathbf{x}_0, \mathbf{x}_0]}{2\|\mathbf{x}_0\|^2 \zeta} \tag{33}$$

where $d_{\mathbf{x}}^2 r(\mathbf{0},0)[\mathbf{x}_0, \mathbf{x}_0, \mathbf{x}_0] = \sum_{i,j,k} \frac{\partial^2 r}{\partial \mathbf{x}_i \partial \mathbf{x}_j \partial \mathbf{x}_k}(\mathbf{0},0)[\mathbf{x}_0]_i [\mathbf{x}_0]_j [\mathbf{x}_0]_k$ (see (24)). As shown in the proof to Theorem 3, $\zeta = d_\beta k(\mathbf{0},0)$ is the non-zero eigenvalue of $d_{\mathbf{x},\beta}^2 r(0,0)$ with eigenvector \mathbf{x}_0 .

This expression is similar to the one given in [22] p. 90. The numerator can be calculated via (24). In [21], we showed that $\beta'(0) = 0$. We have the same result in the present case.

Theorem 5. *If q^* is M -singular for $1 < M \leq N$, then all of the bifurcating branches guaranteed by Theorem 3 are degenerate, i.e., $\beta'(0) = 0$.*

Proof. To show that the numerator of (33) $d_{\mathbf{x}}^2 r(\mathbf{0}, 0) = \mathbf{0}$, expand r_i , the i th component of r , about $\mathbf{x} = \mathbf{0}$,

$$\begin{aligned} r_i(\mathbf{x}, \beta) &= r_i(\mathbf{0}, \beta) + d_{\mathbf{x}} r_i(\mathbf{0}, \beta)^T \mathbf{x} + \mathbf{x}^T d_{\mathbf{x}}^2 r_i(\mathbf{0}, \beta) \mathbf{x} + \mathcal{O}(\mathbf{x}^3) \\ &= d_{\mathbf{x}} r_i(\mathbf{0}, \beta)^T \mathbf{x} + \mathbf{x}^T d_{\mathbf{x}}^2 r_i(\mathbf{0}, \beta) \mathbf{x} + \mathcal{O}(\mathbf{x}^3), \end{aligned}$$

and so

$$r_i(\mathbf{x}, 0) = \mathbf{x}^T d_{\mathbf{x}}^2 r_i(\mathbf{0}, 0) \mathbf{x} + \mathcal{O}(\mathbf{x}^3).$$

Applying the equivariance relation $Ar(\mathbf{x}, 0) = r(A\mathbf{x}, 0)$, where A is any element of the group isomorphic to S_M that acts on r in \mathbf{R}^{M-1} , and equating the quadratic terms yields

$$A \begin{pmatrix} \mathbf{x}^T d_{\mathbf{x}}^2 r_1 \mathbf{x} \\ \mathbf{x}^T d_{\mathbf{x}}^2 r_2 \mathbf{x} \\ \vdots \\ \mathbf{x}^T d_{\mathbf{x}}^2 r_{M-1} \mathbf{x} \end{pmatrix} = \begin{pmatrix} \mathbf{x}^T A^T d_{\mathbf{x}}^2 r_1 A \mathbf{x} \\ \mathbf{x}^T A^T d_{\mathbf{x}}^2 r_2 A \mathbf{x} \\ \vdots \\ \mathbf{x}^T A^T d_{\mathbf{x}}^2 r_{M-1} A \mathbf{x} \end{pmatrix}.$$

By (24), the diagonal $\frac{\partial^2 r_i}{\partial x_i \partial x_i}(\mathbf{0}, 0) = 0$ for each i as well as for all of the “multi-diagonals”. This shows that $\frac{\partial^2 r_i}{\partial x_j \partial x_k}(\mathbf{0}, 0) = 0$ for every i, j, k (see Theorem 124 in [25]). □

When $\beta'(0) = 0$, we need to compute $\beta''(0)$ to determine whether a branch is subcritical or supercritical. Differentiating (32) and setting $t = 0$ shows that $\beta''(0) = -\frac{d_{\mathbf{x}}^2 k(\mathbf{0}, 0)}{d_{\beta} k(\mathbf{0}, 0)}$. Differentiating (28) twice and solving for $d_{\mathbf{x}}^2 k(\mathbf{0}, 0)$ shows that

$$\beta''(0) = \frac{-d_{\mathbf{x}}^3 r(\mathbf{0}, 0)[\mathbf{x}_0, \mathbf{x}_0, \mathbf{x}_0, \mathbf{x}_0]}{3\|\mathbf{x}_0\|^2 \zeta} \tag{34}$$

where $W\mathbf{x}_0 = \mathbf{u} = \mathbf{u}_{(m,n)}$. Use Equation (25) to calculate the numerator, and $\zeta = d_{\beta} k(\mathbf{0}, 0)$ is the non-zero eigenvalue of $d_{\mathbf{x}, \beta}^2 r(\mathbf{0}, 0)$ with eigenvector \mathbf{x}_0 , for which we give an explicit expression in (31) when F is an annealing problem.

If $\beta''(0) \neq 0$, which we expect to be true generically, then Theorem 5 shows that the bifurcation guaranteed by Theorem 3 is pitchfork-like.

2.10. Stability and Optimality

The next Theorem relates the stability of equilibria (q^*, λ^*, β) in the flow (8) with optimality of q^* in Problem (1). In particular, if a bifurcating branch corresponds to an eigenvalue of $d^2 \mathcal{L}(q^*)$ changing from negative to positive, then the branch consists of stationary points (q^*, β^*) that are not solutions of (1). Positive eigenvalues of $d^2 \mathcal{L}(q^*)$ do not necessarily show that q^* is not a solution of (1) (see Remark 1). For example, see page 668 of [21]. A proof of this theorem is given in [21].

Theorem 6. For each bifurcating branch guaranteed by Theorem 3, \mathbf{u} is an eigenvector of $d^2 \mathcal{L}\left(\begin{pmatrix} q^* \\ \lambda^* \end{pmatrix}\right) + t\mathbf{u}, \beta^* + \beta(t)$ for sufficiently small t . Furthermore, if the corresponding eigenvalue is positive, then the branch consists of unstable stationary points that are not solutions to (1).

2.11. Structure of the Symmetry Projection

The matrix $P_R(q^*)$ that projects $(q, \lambda) \in \mathfrak{R}^{NK+K}$ onto $\text{range}(d^2 \mathcal{L}(q^*)) \times \mathcal{K}^*$ by annihilating \mathcal{K}_p is important for numerical computations for equilibria of IB, since we may want to take each equilibrium found by Newton’s method and take out any part in \mathcal{K}_p . P_R is written as a function of q since its constitutive vectors \mathbf{y} (from Definition 1) depend on q . The following theorems clarify the structure of this projection.

Theorem 7. $P_R(q) = I - P_{\mathcal{K}_p}(q)$, where $P_{\mathcal{K}_p} = \begin{pmatrix} A & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}$. P_R and $P_{\mathcal{K}_p}$ are $(NK + K) \times (NK + K)$. The matrix A is $NK \times NK$ with N^2 blocks, $\{A_{ij}\}_{i,j=1}^N$, of size $K \times K$, defined by

$$A_{i,j} = \begin{cases} (M-1)\mathbf{y}\mathbf{y}^T & \text{if } 1 \leq i = j \leq M \\ -\mathbf{y}\mathbf{y}^T & \text{if } 1 \leq i \neq j \leq M \\ \mathbf{0} & \text{otherwise} \end{cases}$$

For example, if $M = N = 3$, then

$$P_R = I - \begin{pmatrix} 2\mathbf{y}\mathbf{y}^T & -\mathbf{y}\mathbf{y}^T & -\mathbf{y}\mathbf{y}^T & \mathbf{0} \\ -\mathbf{y}\mathbf{y}^T & 2\mathbf{y}\mathbf{y}^T & -\mathbf{y}\mathbf{y}^T & \mathbf{0} \\ -\mathbf{y}\mathbf{y}^T & -\mathbf{y}\mathbf{y}^T & 2\mathbf{y}\mathbf{y}^T & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix} = I - \begin{pmatrix} (N-1) & -1 & -1 & 0 \\ -1 & (N-1) & -1 & 0 \\ -1 & -1 & (N-1) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \otimes \mathbf{y}\mathbf{y}^T.$$

Proof. Theorem 2 gives the basis of \mathcal{K}_p as $\{Y_i\}_{i=1}^{M-1}$. Let Y be the $(NK + K) \times (M - 1)$ matrix whose columns are the vectors $\{Y_i\}$. For example, if $M = 3$ and $N = 4$, then

$$Y = \begin{pmatrix} \mathbf{y} & \mathbf{0} \\ \mathbf{0} & \mathbf{y} \\ -\mathbf{y} & -\mathbf{y} \\ \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}. \text{ Thus, the matrix that projects onto } \mathcal{K}_p \text{ is } P_{\mathcal{K}_p} = Y(Y^T Y)^{-1} Y^T,$$

and the projection matrix onto $\text{range}(d^2\mathcal{L}(q^*))$ is $P_R = I - P_{\mathcal{K}_p}$. Direct multiplication of $Y(Y^T Y)^{-1} Y^T$, with an appeal to Lemma 34 in [25] to compute the inverse, shows that $P_{\mathcal{K}_p} = \frac{1}{N\mathbf{y}^T \mathbf{y}} \begin{pmatrix} A & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}$. Dropping the constant yields the result. \square

For the Information Bottleneck, the matrix P_R is easy to calculate, since $\mathbf{y} = q^i$ for any $i \in \mathcal{U}$. For example, when $q = q_{\frac{1}{N}}$, then $\mathbf{y}^T \mathbf{y} = \frac{K}{N^2}$ and $\mathbf{y}\mathbf{y}^T = \frac{1}{N^2} \mathbf{1}$, and so

$$P_{\mathcal{K}_p} = \frac{1}{NK} \begin{pmatrix} (N-1) & -1 & \dots & -1 & 0 \\ -1 & (N-1) & \dots & -1 & 0 \\ \vdots & \vdots & & \vdots & \vdots \\ -1 & -1 & \dots & (N-1) & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \otimes \mathbf{1}$$

where $\mathbf{1}$ is a $K \times K$ matrix of 1s. Thus,

$$P_R = I_{NK+K} - \begin{pmatrix} (N-1) & -1 & \dots & -1 & 0 \\ -1 & (N-1) & \dots & -1 & 0 \\ \vdots & \vdots & & \vdots & \vdots \\ -1 & -1 & \dots & (N-1) & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \otimes \mathbf{1}.$$

Theorem 8. The symmetry group S_M commutes with the matrix P_R , which projects onto $\mathfrak{R}^{NK+K} \setminus \mathcal{K}_p$.

Proof. Let $P := P_R$ be the matrix that projects onto $\text{range } d^2\mathcal{L}(q^*) \times \mathcal{K}^* = \mathfrak{R}^{NK+K} \setminus \mathcal{K}_p$. Since $\mathfrak{R}^{NK+K} = \text{range } d^2\mathcal{L}(q^*) \times \mathcal{K}^* \times \mathcal{K}_p$, then any $x \in \mathfrak{R}^{NK+K}$ can be decomposed in the respective subspaces as $x = r + k^* + k_p$. Let σ be an arbitrary permutation matrix in S_M . Then $\sigma P \begin{pmatrix} q \\ \lambda \end{pmatrix} = \sigma P(r + k^* + k_p) = \sigma(r + v)$. Since $\text{range } d^2\mathcal{L}(q^*)$, \mathcal{K}^* and \mathcal{K}_p are all S_M invariant; then $\sigma(r + v) \in \text{range } d^2\mathcal{L}(q^*) \times \mathcal{K}^*$ implies that $\sigma(r + v) = P\sigma(r + v)$, and $\sigma z \in \mathcal{K}_p$ implies that $P\sigma(r + v) = P\sigma(r + v + z)$. Thus, $\sigma P x = P\sigma x$. \square

2.12. Visualizations of Sample Results

We illustrate these structures numerically. In [7], we introduced the toy “Four-blob” probability distribution $p(x, y)$ shown in Figure 1.




Figure 1. The probability distribution $p(x, y)$ for the “Four-blob” toy problem for a system of interest $X \rightarrow Y$. We use this probability to illustrate some results of the bifurcation analysis reported here.

For the Information Distortion problem (3) [7,12,13] and the synthetic dataset composed of a mixture of four Gaussians (Figure 1), we determined the bifurcation structure of solutions to (3) by annealing in β and finding the corresponding stationary points to (1). A typical run of the derived gradient dynamical system tends to follow the main bifurcation branch $S_K \rightarrow S_{K-1}$ from the fully symmetric uniform quantizer $q_{\frac{1}{N}}$ ($N = 4$ here) to the fully resolved deterministic quantizer (hard clustering) seen at the end in Figure 2. The permutation symmetry is also obvious there—the value of the cost function does not change if the classes along the vertical axis in T are permuted/re-labeled. The uniform quantizer $q_{\frac{1}{N}}$ (Item 1 in the figure) plays a special role in the formulation (3), as it is the *unique* solution to the problem for $\beta = 0$ as the maximum entropy solution of $\max_q H(T|Y)$. Its loss of stability at the first bifurcation for increasing β can hence be determined analytically and the first bifurcation structure characterized completely. Because of the “perpetual kernel” of the cost function in (4), the uniform quantizer is just one of a continuous set of “uninformative” quantizers for the IB problem (4): all $\{q(t|y) : q(t|y) = f(t)\}$, having constant probability of assignment of each y to class t , but the assignment weight can be different for different classes. Such a structure does not change the value of the cost function in the IB problem (4) (but does change it for (3), which hence does not have this degeneracy). We address the degeneracy of the IB optimization by projecting onto the subspace that has the correct symmetry (i.e., just the uniform quantizer $q_{\frac{1}{N}}$ in this case), as outlined in Remark 2.

A more-thorough structure of the bifurcation diagram, using the analysis presented above, is shown in Figure 3.

Similar to the results we presented in [28], the close-up of the bifurcation at $\beta \approx 1.038706$ in Figure 3B shows a subcritical bifurcating branch (a first-order phase transition) that consists of stationary points of Problem (1). By projecting the Hessian $\Delta_q(G(q^*) + \beta D(q^*))$ onto each of the kernels referenced in Theorem 6, we determined that the points on this subcritical branch are **not** solutions of (1), and yet they **are** solutions of (2).

Furthermore, observe that Figure 3B indicates that a saddle-node bifurcation occurs at $\beta \approx 1.037479$. That this is indeed the case was proved in [21]. In fact, for any problem of the form (2), these are the only two types of bifurcations to be expected: pitchfork and saddle-node.




Figure 2. The bifurcations of the solutions (q^*, β) to the Information Distortion problem (3). For the mixture of 4 well-separated Gaussians shown in Figure 1, the behavior of $D(q) = I(X; T)$ as a function of β is shown in the top panel, and some of the solutions $q^*(T|Y)$ are shown in the bottom panels. Item 1 shows the uniform quantizer $q_{\frac{1}{N}}$, assigning equal probability of each $y \in Y$ to belong to one of the four clusters in T . Subsequent items 2–5 point to a set of partially resolved quantizations, in which subsets of Y are assigned with high probability to one (2) or more (3–5) classes (dark colors, close to 1), while other subsets are still unresolved (gray levels), albeit as a higher probability than $q_{\frac{1}{N}}$ (darker gray, as some of the classes are excluded after being resolved for another subset). Item 6 shows an almost fully resolved quantizer at sufficiently high β . They become fully resolved (deterministic; $q(t|y) = 1$ or 0) as $\beta \rightarrow \infty$ (not shown).




Figure 3. (A) The bifurcation structure of stationary points of the Information Distortion problem (3), a problem of form (2). We found these points by annealing in β and finding stationary points for Problem (1) using the algorithm presented in [28]. A square indicates where a bifurcation occurs. (B) A close-up of the subcritical bifurcation at $\beta \approx 1.038706$, indicated by a square. Observe the subcritical bifurcating branch, and the subsequent saddle-node bifurcation at $\beta \approx 1.037479$, indicated by another square. We applied Theorem 6 to show that the subcritical bifurcating branch is composed of quantizers that are solutions of (3) but not of (1).

3. Conclusions and Discussion

The main goal of this contribution was to show that information-based distortion-annealing problems such as (2) have an interesting mathematical structure. The most interesting aspects of that mathematical structure are driven by the symmetries present in the cost functions—their invariance to actions of the permutation group S_N , represented as relabeling of the reproduction classes. Such a structure would hold for any biclustering

problem [4] that relies on the intrinsic interaction of a pair of variables for unsupervised clustering. The second mathematical structure that we used successfully was bifurcation theory, which allowed us to identify and study the discrete points at which the character of the cost function changed. The combination of those two tools in [20] allowed us to explicitly compute the value of the annealing parameter β at which the initial maximum of the uniform quantizer $q_{\frac{1}{N}}$ of (1) loses stability. We concluded that for a fixed system $C \rightarrow Y$ characterized by $p(X, Y)$, this value is the same for both problems, that it does not depend on the number of elements of the reproduction variable T , and that it is always greater than 1. We further introduced an eigenvalue problem that links the critical values of β and q for bifurcations, or phase transitions, branching off arbitrary intermediate solutions.

Even though the cost functions F_{IB} (4) and F_H (3) have similar properties, they also differ in some important aspects. We have shown that the function F_{IB} is degenerate since its constitutive functions $I(X; Y)$ and $I(X; T)$ are not strictly convex in q . That introduces additional invariances and singularities that are always preserved, which makes phase transitions more difficult to detect (e.g., the “uninformative quantizers” $q(t|y) = f(t)$ only) and post-transition directions more difficult to determine. In contrast, F_H is strictly convex except at points of phase transitions. The theory we developed here allows us to identify bifurcation directions and determine their stability. Despite the presence of a high-dimensional null space at bifurcations, the symmetries restrict the allowed transition dimensions to multiple co-dimension 1 transitions, all related by group transformations. We achieved that here with three main results. Theorem 8 extended the Equivariant Branching Lemma 1 to the Information Bottleneck case with additional translation invariance. Theorem 4 identified specific conditions at which a bifurcation of the gradient flow (8) occurs. This condition is computable analytically for the initial bifurcation off the uniform quantizer $q_{\frac{1}{N}}$ and with numeric continuation for subsequent bifurcation. Finally, in Section 2.9, we provided checks for the types of bifurcations that occur, giving conditions to detect saddle-node and pitchfork bifurcations and to determine whether pitchforks are supercritical (second-order phase transitions) or subcritical (leading to first-order phase transitions discontinuous in β). The combination of the three results, together with our previous results in [20], completely characterize the local bifurcation structure of Information Bottleneck-type problems with or without the added translation symmetry.

Despite the further development of the bifurcation formalism for IB presented here, there are still open questions that this manuscript did not resolve. In particular, we still cannot confirm or reject the conjecture that the set of S_K symmetric soft-clustering branches connected through symmetry-breaking bifurcations leads to the global hard-clustering optima at $\beta \rightarrow \infty$ (multiple equivalent solutions connected by the permutation symmetry of the problem). We believe this is partially due to a discrepancy between practical observations and theoretical results. In particular, we and other practitioners [29,30] note that the only observed symmetry-breaking bifurcations during optimization are of the kind $S_M \rightarrow S_{M-1}$, while the theory allows for arbitrary $S_M \rightarrow S_m \times S_n$ bifurcations. The latter are known to happen and be stable in other biological systems and circumstances [26,31]. This suggests a research approach of comparing and contrasting the different systems that possess the same S_N symmetry and symmetry-breaking bifurcations to lead to breakthroughs in this application to optimization in the Information Bottleneck problem.

An additional open problem involves the use of continuous variables, already noted in [5] and explored further in [32,33]. This approach, while important for many real-world problems, involves the application of additional mathematical tools, namely Calculus of Variations [34], which further increases the complexity of an otherwise already complex problem. These difficulties are illustrated in a pair of papers [35,36] that use the continuous formulation. They do present some significant results on conditions of learnability, but both papers manage to only get bounds on β under which learnability (optimal solutions beyond the “uninformative” quantizer) can be achieved. This is possibly due to the presence of continuous spectra in covariance operators of continuous quantizers, something that we avoid by focusing on finite spaces. As a consequence, here and in prior work [20], we

show specific values for β for the initial bifurcation from the uniform quantizer, which supports nontrivial clustering. We consider formulation with continuous variables beyond the scope of this manuscript, but look forward to the development of additional techniques to incorporate this important case in the bifurcation framework presented here. Regardless of such developments, any practical problem with numeric optimization will involve discretization of the continuous variables, which effectively converts a continuous problem to the discrete state discussed here.

Author Contributions: Conceptualization, A.G.D.; Formal analysis, A.G.D. and A.E.P.; Investigation, A.G.D. and A.E.P.; Writing—original draft, A.G.D. and A.E.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Gray, R.M. *Entropy and Information Theory*; Springer: Berlin/Heidelberg, Germany, 1990.
2. Cover, T.; Thomas, J. *Elements of Information Theory*; Wiley Series in Communication; Wiley: New York, NY, USA, 1991.
3. Rose, K. Deterministic Annealing for Clustering, Compression, Classification, Regression, and Related Optimization Problems. *Proc. IEEE* **1998**, *86*, 2210–2239. [[CrossRef](#)]
4. Madeira, S.C.; Oliveira, A.L. Biclustering algorithms for biological data analysis: A survey. *IEEE/ACM Trans. Comput. Biol. Bioinform.* **2004**, *1*, 24–45. [[CrossRef](#)] [[PubMed](#)]
5. Tishby, N.; Pereira, F.C.; Bialek, W. The information bottleneck method. In *37th Annual Allerton Conference on Communication, Control, and Computing*; University of Illinois: Champaign, IL, USA, 1999.
6. Dimitrov, A.G.; Miller, J.P.; Aldworth, Z.; Gedeon, T.; Parker, A.E. Analysis of neural coding through quantization with an information-based distortion measure. *Netw. Comput. Neural Syst.* **2003**, *14*, 151–176. [[CrossRef](#)]
7. Dimitrov, A.G.; Miller, J.P. Neural coding and decoding: Communication channels and quantization. *Netw. Comput. Neural Syst.* **2001**, *12*, 441–472. [[CrossRef](#)]
8. Gersho, A.; Gray, R.M. *Vector Quantization and Signal Compression*; Kluwer Academic Publishers: New York, NY, USA, 1992.
9. Mumey, B.; Gedeon, T. Optimal mutual information quantization is NP-complete. In *Proceedings of the Neural Information Coding (NIC) Workshop, Snowbird, UT, USA, 1–4 March 2003*.
10. Slonim, N.; Tishby, N. Agglomerative Information Bottleneck. In *Advances in Neural Information Processing Systems*; Solla, S.A., Leen, T.K., Müller, K.R., Eds.; MIT Press: Cambridge, MA, USA, 2000; Volume 12, pp. 617–623.
11. Slonim, N. The Information Bottleneck: Theory and Applications. Ph.D. Thesis, Hebrew University, Jerusalem, Israel, 2002.
12. Dimitrov, A.G.; Miller, J.P. Analyzing sensory systems with the information distortion function. In *Proceedings of the Pacific Symposium on Biocomputing 2001*; Altman, R.B., Ed.; World Scientific Publishing Co.: Singapore, 2000.
13. Gedeon, T.; Parker, A.E.; Dimitrov, A.G. Information Distortion and Neural Coding. *Can. Appl. Math. Q.* **2003**, *10*, 33–70.
14. Slonim, N.; Somerville, R.; Tishby, N.; Lahav, O. Objective classification of galaxy spectra using the information bottleneck method. *Mon. Not. R. Astron. Soc.* **2001**, *323*, 270–284. [[CrossRef](#)]
15. Bardera, A.; Rigau, J.; Boada, I.; Feixas, M.; Sbert, M. Image segmentation using information bottleneck method. *IEEE Trans. Image Process.* **2009**, *18*, 1601–1612. [[CrossRef](#)]
16. Aldworth, Z.N.; Dimitrov, A.G.; Cummins, G.I.; Gedeon, T.; Miller, J.P. Temporal encoding in a nervous system. *PLoS Comput. Biol.* **2011**, *7*, e1002041. [[CrossRef](#)]
17. Buddha, S.K.; So, K.; Carmena, J.M.; Gastpar, M.C. Function identification in neuron populations via information bottleneck. *Entropy* **2013**, *15*, 1587–1608. [[CrossRef](#)]
18. Lewandowsky, J.; Bauch, G. Information-optimum LDPC decoders based on the information bottleneck method. *IEEE Access* **2018**, *6*, 4054–4071. [[CrossRef](#)]
19. Parker, A.E.; Dimitrov, A.G.; Gedeon, T. Symmetry breaking in soft clustering decoding of neural codes. *IEEE Trans. Inf. Theory* **2010**, *56*, 901–927. [[CrossRef](#)]
20. Gedeon, T.; Parker, A.E.; Dimitrov, A.G. The mathematical structure of information bottleneck methods. *Entropy* **2012**, *14*, 456–479. [[CrossRef](#)]
21. Parker, A.E.; Gedeon, T. Bifurcations of a class of S_N -invariant constrained optimization problems. *J. Dyn. Differ. Equ.* **2004**, *16*, 629–678. [[CrossRef](#)]
22. Golubitsky, M.; Stewart, I.; Schaeffer, D.G. *Singularities and Groups in Bifurcation Theory II*; Springer: New York, NY, USA, 1988.

23. Golubitsky, M.; Schaeffer, D.G. *Singularities and Groups in Bifurcation Theory I*; Springer: New York, NY, USA, 1985.
24. Nocedal, J.; Wright, S.J. *Numerical Optimization*; Springer: New York, NY, USA, 2000.
25. Parker, A.E. Symmetry Breaking Bifurcations of the Information Distortion. Ph.D. Thesis, Montana State University, Bozeman, MT, USA, 2003.
26. Golubitsky, M.; Stewart, I. *The Symmetry Perspective: From Equilibrium to Chaos in Phase Space and Physical Space*; Birkhauser Verlag: Boston, MA, USA, 2002.
27. Schott, J.R. *Matrix Analysis for Statistics*; John Wiley and Sons: New York, NY, USA, 1997.
28. Parker, A.; Gedeon, T.; Dimitrov, A. Annealing and the rate distortion problem. In *Advances in Neural Information Processing Systems 15*; Becker, S.T., Obermayer, K., Eds.; MIT Press: Cambridge, MA, USA, 2003; Volume 15, pp. 969–976.
29. Dimitrov, A.G.; Cummins, G.I.; Baker, A.; Aldworth, Z.N. Characterizing the fine structure of a neural sensory code through information distortion. *J. Comput. Neurosci.* **2011**, *30*, 163–179. [[CrossRef](#)] [[PubMed](#)]
30. Schneidman, E.; Slonim, N.; Tishby, N.; de Ruyter van Steveninck, R.R.; Bialek, W. Analyzing neural codes using the information bottleneck method. In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2003; Volume 15.
31. Stewart, I. Self-Organization in evolution: A mathematical perspective. *Philos. Trans. R. Soc.* **2003**, *361*, 1101–1123. [[CrossRef](#)] [[PubMed](#)]
32. Chechik, G.; Globerson, A.; Tishby, N.; Weiss, Y. Information bottleneck for Gaussian variables. In Proceedings of the Advances in Neural Information Processing Systems 16 (NIPS 2003), Vancouver, BC, Canada, 8–13 December 2003.
33. Chechik, G.; Globerson, A.; Tishby, N.; Weiss, Y. Information Bottleneck for Gaussian Variables. *J. Mach. Learn. Res.* **2005**, *6*, 165–188.
34. Gelfand, I.M.; Fomin, S.V. *Calculus of Variations*; Dover Publications: Mineola, NY, USA 2000.
35. Wu, T.; Fischer, I.; Chuang, I.L.; Tegmark, M. Learnability for the information bottleneck. In Proceedings of the Uncertainty in Artificial Intelligence, PMLR, Virtual, 3–6 August 2020; pp. 1050–1060.
36. Ngampruetikorn, V.; Schwab, D.J. Perturbation theory for the information bottleneck. *Adv. Neural Inf. Process. Syst.* **2021**, *34*, 21008–21018.