

Generalizing Musical Intervals

Dmitri Tymoczko

Abstract Taking David Lewin's work as a point of departure, this essay uses geometry to reexamine familiar music-theoretical assumptions about intervals and transformations. Section 1 introduces the problem of "transportability," noting that it is sometimes impossible to say whether two different directions—located at two different points in a geometrical space—are "the same" or not. Relevant examples include the surface of the earth and the geometrical spaces representing n -note chords. Section 2 argues that we should not require that every interval be defined at every point in a space, since some musical spaces have natural boundaries. It also notes that there are spaces, including the familiar pitch-class circle, in which there are multiple paths between any two points. This leads to the suggestion that we might sometimes want to replace traditional pitch-class intervals with *paths in pitch-class space*, a more fine-grained alternative that specifies how one pitch class moves to another. Section 3 argues that group theory alone cannot represent the intuition that intervals have quantifiable sizes, proposing an extension to Lewin's formalism that accomplishes this goal. Finally, Section 4 considers the analytical implications of the preceding points, paying particular attention to questions about voice leading.

DAVID LEWIN'S *Generalized Musical Intervals and Transformations* (henceforth *GMIT*) begins by announcing what seems to be a broad project: modeling "directed measurement, distance, or motion" in an unspecified range of musical spaces. Lewin illustrates this general ambition with an equally general graph—the simple arrow or vector shown in Figure 1. This diagram, he writes, "shows two points s and t in a symbolic musical space. The arrow marked i symbolizes a characteristic directed measurement, distance, or motion from s to t . We intuit such situations in many musical spaces, and we are used to calling i 'the interval from s to t ' when the symbolic points are pitches or pitch classes" (xi–xii).¹ As is well known, Lewin goes on to model these general intuitions

Thanks to Jordan Ellenberg, Noam Elkies, Tom Fiore, Ed Gollin, Rachel Hall, Henry Klumpenhouwer, Jon Kochavi, Shaugn O'Donnell, Steve Rings, Ramon Satyendra, Neil Weiner, and especially Jason Yust for helpful conversations. An earlier and more polemical version of this essay was delivered to the 2007 Society for Music Theory Mathematics of Music Analysis Interest Group. My thanks to those who encouraged me to write down my ideas in a less agonistic fashion. For more remarks on Lewin, see Tymoczko 2007, 2008a, and forthcoming.

¹ The phrase "directed measurement, distance, or motion," or a close variant, appears at many points in *GMIT* (e.g., 16, 17, 18, 22, 25, 26, 27). Elsewhere, Lewin describes intervals (group elements of a generalized interval system) as representing "distances" (42, 44, 147). Among Lewin's

mathematically sophisticated readers, Dan Tudor Vuza (1988, 278) and Julian Hook (2007a) take Lewinian intervals to represent quantified distances with specific sizes. See also Tymoczko 2008a.



Figure 1. This graph appears as figure 0.1 in Lewin's *Generalized Musical Intervals and Transformations*. It represents a directed motion from *s* to *t* in a symbolic musical space, with the arrow *i* representing "the interval" from *s* to *t*.

using group theory, embodied in the twin constructions of the "generalized interval system" (GIS) and its slightly more flexible sibling, the "transformational graph."² These innovations have proven to be enormously influential, forming the basis for countless subsequent music-theoretical studies.³ So pervasive is Lewin's approach that newcomers to music theory might well assume that group theory provides the broadest possible framework for investigating "directed measurement, distance, or motion."

It is interesting to note, however, that Lewin's motivating question recapitulates one of the central projects of nineteenth-century mathematics: developing and generalizing the notion of a vector. Intuitively, a vector is simply a quantity with both size and direction, represented graphically by an arrow such as that in Figure 1.⁴ Nineteenth-century mathematicians attempted to formalize this intuition in a variety of ways, including William Hamilton's "quaternions" and Hermann Grassmann's "theory of extended magnitudes," now called the "Grassmann calculus." The modern theory of vectors developed relatively late, primarily at the hands of Josiah Gibbs and Oliver Heaviside, and achieved widespread prominence only in the early decades of the twentieth century.⁵ While this development was occurring, other nineteenth-century mathematicians, such as Karl Friedrich Gauss, Bernhard Riemann, and Tullio Levi-Civita, were generalizing classical geometrical concepts—including "distance," "straight line," and "shortest path"—to curved spaces such as the surface of the earth. It was not until the twentieth century that the Gibbs-Heaviside notion of a "vector," which requires a *flat* space, was defined as an intrinsic object in *curved* spaces. The result is a powerful set of tools for investigating Lewin's question, albeit in a visuospatial rather than musical context.

2 Henry Klumpenhouwer (2006) asserts that Lewin's book opens with a "Cartesian" perspective that is eventually superseded by the later, "transformational" point of view. This interpretive issue is for the most part tangential to this essay, since my points apply equally to Lewinian transformations (understood as semigroups of functions acting on a musical space). In Klumpenhouwer's terms, I am suggesting that even Lewinian transformations cannot capture certain elementary musical intuitions.

3 Indeed, one prominent theorist has declared that David Lewin "created the intellectual world I live in" (Rothstein 2003).

4 "A vector, \mathbf{x} , is first conceived as a directed line segment, or a quantity with both a magnitude and direction" (Byron and Fuller 1992, 1).

5 See Crowe 1994 for a discussion of quaternions, Grassman calculus, and vector analysis. The conflict between advocates of quaternions and vectors is a leitmotif in Pynchon 2006.

The purpose of this essay is to put these two inquiries side by side: to contrast Lewin's "generalized musical intervals" with the generalized vectors of twentieth-century mathematics. My aims are threefold. First, I want to explore the fascinating intellectual-historical resonance between these two very different, and yet in some sense very similar intellectual projects—letting our understanding of each deepen and enrich our understanding of the other. Second, I want to suggest that Lewin's "generalized musical intervals" are in some ways less general than they might appear to be. For by modeling "directed motions" using functions defined over an entire space, Lewin imposes significant and nonobvious limits on the musical situations he can consider.⁶ Third, I want to suggest that there are circumstances in which we need to look beyond group theory and toward other areas of mathematics.

Before proceeding, I should flag a few methodological and terminological issues with the potential to cause confusion. Throughout this essay I use the term "interval" broadly, to refer to vectorlike combinations of magnitude and direction: "two ascending semitones," "two beats later," "south by 150 miles," "five degrees warmer," "three days earlier," and so on. (As this list suggests, intervals need not necessarily be musical.) I am suggesting that Lewin's theoretical apparatus is not always sufficient for modeling "intervals" so construed. Some readers may object, preferring to limit the term "interval" to more specific musical or mathematical contexts. Since I have no wish to argue about terminology, I am perfectly happy to accede to this suggestion; such readers are therefore invited to substitute a more generic term, such as "directed distance," for my term "interval" if they wish. Modulo this substitution, my claim is that directed distances can be quite important in music-theoretical discourse and that geometrical ideas can help us to model them.

I should also say that I am attempting to summarize some rather intricate mathematical ground. The mathematics relevant to this discussion is not that of elementary linear algebra but rather differential geometry, in which a vector space is attached to every point in a curved manifold. Comprehensibility therefore requires that I elide some important distinctions, for instance, between torsion and curvature or between directional derivatives and finite motions along geodesics. I make no apologies for this: My goal here is to communicate twentieth-century geometrical ideas to a reasonably broad music-theoretical audience, a process that necessitates a certain amount of simplification. I write for musical readers reasonably familiar with Lewin's work but perhaps lacking acquaintance with Riemannian geometry, topology, or other branches of mathematics. Consequently, professional mathematicians may need to grit their teeth as I attempt to convey the spirit of contemporary geometry even while glossing over some of its subtleties. It goes without saying that I have done my best to avoid outright falsification.

⁶ In what follows, I use the term *space* broadly, to refer to a set of musical objects. However, many of the musical spaces I discuss are robustly geometrical manifolds or quotients of

manifolds (orbifolds). For more discussion, see Callender 2004, Tymoczko 2006, and Callender, Quinn, and Tymoczko 2008.

This is also the place to emphasize that I am not proposing geometry as an all-purpose, ready-to-hand metalanguage for doing music theory. (Indeed, the perspective I ultimately recommend—in which “intervals” are *equivalence classes of particular motions*—is more or less indigenous to music theory.) My claim here is that it can *sometimes* be useful to borrow ideas and principles from geometry, for example, that a space can have a boundary, or that intervals may not be transportable from point to point. These general principles can often be relevant even in discrete contexts where the technical machinery of geometry is difficult to apply. Mathematically untrained theorists should therefore take heart: I am not demanding that they enroll themselves in graduate classes in geometry, topology, or any other branch of mathematics; instead, I am asking that they consider ways in which common music-theoretical assumptions conflict with basic mathematical principles.

Finally, I want to stress that my motivations are not primarily critical. David Lewin was an enormously creative and intelligent thinker, and his current music-theoretical prominence is beyond dispute. My chief interest is in understanding his ideas clearly and in identifying places where geometry can stimulate further theoretical development. Lewin’s work provides a useful stepping-off point in part because it is so familiar, having been incorporated into the standard tool kit of contemporary theory, and in part because its group-theoretical orientation contrasts fruitfully with the approach I recommend. Thus, although I will be pointing to limitations of Lewin’s work, my ultimate purpose is to suggest that his question—“how do we model directed magnitudes in musical spaces?”—may be richer than even he realized.

I. Intervals as functions

One of the central ideas in *GMIT* is that musical intervals can be modeled as transformations or functions. This means that when we conceive of a single object moving in some musical space, we are conceiving of a transformation that potentially applies to *all* objects in the space. Hence, when a musician moves from G4 to Eb4, she exemplifies a function that can in principle be applied to any note or, indeed, any collection of notes.

This is a brilliant, creative, powerful, and problematic idea. To see why, imagine that someone tells you “the interval $G \rightarrow Eb$ plays an important role in the first movement of Beethoven’s Fifth Symphony.” According to Lewin, the term “the interval $G \rightarrow Eb$ ” refers to a *function*, the “descending major third,” that takes any pitch class as input and returns as output the pitch class a major third below. This would seem to imply that we can paraphrase the statement as “the descending major third plays an important role in Beethoven’s Fifth.” But of course this is mistaken, for what is presumably being said is that the *specific* interval $G \rightarrow Eb$, rather than $F \rightarrow Db$ or $B \rightarrow G$, is particularly important in Beethoven’s piece. As far as I know, there is no standard Lewinian term for specific directed motions, such as $G \rightarrow Eb$, as opposed to generalized functions

such as “the descending major third.” Yet ordinary musical discourse clearly needs both: We cannot do musical analysis if we limit ourselves to general statements about transformations on an entire musical space.

This leads to a second and more interesting problem: Particular directed motions cannot always be converted into generalized functions. A familiar nonmusical example illustrates the difficulty. Suppose Great-Aunt Abigail takes a train from Albany to New York City, moving south by roughly 150 miles. Does this directed motion also define a unique function over every single point on the earth?⁷ A moment’s reflection shows that it does not: It is not possible to “move south by 150 miles” at the south pole, or indeed at any point less than 150 miles away from it. Furthermore, at the north pole *every* direction is south, and hence the instruction “move south 150 miles” does not define a unique action.⁸ Consequently, there is no obvious way to model Great-Aunt Abigail’s train journey using functions defined over the entire earth. And yet it would seem that the phrase “south 150 miles” refers to a paradigmatic example of directed motion in a familiar space.

In fact, the general situation is even worse than this example suggests. We navigate along the earth with respect to special points—the “North Pole” and the “South Pole”—that lie along the earth’s axis of rotation. (In essence, “north” means “along the shortest path to the north pole,” while “south” means “along the shortest path to the south pole.”) On a mathematical sphere, however, there is no principled way to choose a pair of antipodal points. Let’s therefore put aside the terms *north* and *south* and think physically about what happens when we slide a small rigid arrow along the earth’s surface. Figure 2 shows that the result depends on the *path* along which the arrow moves: If we slide the arrow halfway around the world along the equator, it points east, but if we move the vector along the perpendicular circle, it ends up pointing in the *opposite direction!* For this reason, there would seem to be no principled way to decide whether two distant vectors point “in the same direction” or not.⁹ Mathematically, there simply is no fact of the matter about whether two particular directions, located at two different points on the sphere, are the same.

7 Mathematically, a function is a set of ordered pairs $\{(x_1, y_1), \dots, (x_n, y_n)\}$, with no two pairs sharing the same first element. We are looking for pairs (x, y) where y represents the unique point 150 miles south of x , with the x ranging over the earth’s entire surface.

8 In fact, there is a famous theorem of topology stating that there are no nonvanishing vector fields defined over the two-dimensional sphere. It is therefore impossible to define a set of arrows covering the earth, all equally long, and varying smoothly from point to point. See, for example, Eisenberg and Guy 1979.

9 There are some subtle issues here that could be confusing. First, the earth is embedded in three-dimensional space, and

this gives us a means of comparing distant vectors; geometers, however, are typically concerned with the sphere as a space unto itself, and not as embedded in another space. Second, Figure 2 transports arrows by sliding them rigidly, as if they were physical objects. However, we operate with directions like “go 150 miles south” somewhat differently: The direction “south” turns slightly relative to the equator as one moves west, so that it always faces the south pole; consequently, the endpoints of an east-west arrow would get closer together as they move toward the poles. Mathematically, the difference here is between a connection with *curvature* and one with *torsion*. This subtlety is not relevant in what follows.

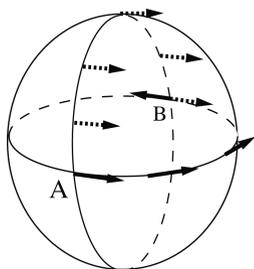


Figure 2. Transporting an arrow along a sphere produces path-dependent results. Here, sliding an east-pointing arrow halfway around the earth along the equator produces an east-pointing arrow. Sliding it halfway around the earth along a perpendicular path produces a west-pointing arrow.

The question is whether theorists would ever want to deal with such spaces in which arrows cannot be uniquely transported from point to point. The answer is yes. Figure 3 shows a (continuous) two-dimensional space in which points represent two-note chords; the space is a Möbius strip whose left edge is “glued” to the right edge with a half twist that links the appropriate chords. Since it would take us too far afield to discuss this space in detail or to tease out its various analytico-theoretical virtues, the following discussion relies on just a few basic facts: Points in the space represent unordered pairs of pitch classes, while directed distances represent *voice leadings*, or “mappings” that specify paths along which the two notes in the dyad move.¹⁰ In particular, the leftmost arrow on the figure corresponds to the voice leading $(B, F) \rightarrow (C, E)$, in which B moves up by semitone to C, while F moves down by semitone to E.¹¹ Now suppose we want to slide the arrow in Figure 3 so that it begins at $\{C, F\}$. If we move the arrow rightward, it points to $\{D, F\}$. However, if we move it leftward, off the edge of the strip, it reappears on the right edge facing downward. (Remember, the right edge is glued to the left with a “twist” that causes arrows to turn upside down.) We can then slide it leftward so that the head of the arrow points to $\{G, B\}$. As in the spherical case, our arrow ends up

¹⁰ Readers who are unfamiliar with this space are referred to Tymoczko 2006, which shows that any two-voice passage of music can be translated into a sequence of directed motions in this space. (See also the online supplement to Callender, Quinn, and Tymoczko 2008, which uses the figure to uncover nonobvious relationships in an intermezzo by Brahms.) Briefly, it can be shown that any point in \mathbb{R}^2 assigns pitches to a two-voice ensemble; any line segment in \mathbb{R}^2 represents a “voice leading” in which the instruments move. (See Tymoczko 2008b for more on voice leading.) The length of the line segment represents the size of the voice leading. When we disregard octave and order, we form the quotient space \mathbb{T}^2/S_2 , which is a Möbius strip whose boundary

is singular. (A useful fundamental domain for the space is bounded by $(0, 0)$, $(-6, 6)$, $(6, 6)$, and $(0, 12)$.) In the quotient space, the images of line segments in \mathbb{R}^2 sometimes appear to “bounce off” the singularities like balls reflecting off the bumpers of a pool table. The reason to use a continuous rather than discrete space is to preserve the association between line segments (now generalized to include those that “bounce off” the singularities) and voice leadings.

¹¹ Thus, the arrow represents motions such as $(B_4, F_5) \rightarrow (C_5, E_5)$ and $(F_3, B_5) \rightarrow (E_3, C_6)$, but not $(B_4, F_5) \rightarrow (C_5, E_6)$, since here a note moves by eleven semitones. See Tymoczko 2008b.

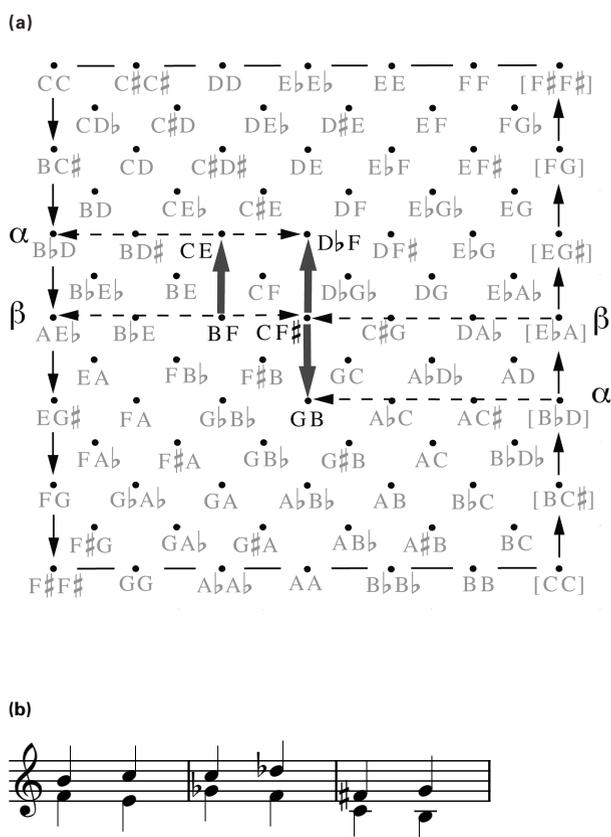


Figure 3. (a) A Möbius strip representing two-note chords. Depending on how we move the arrow $(B, F) \rightarrow (C, E)$ to $(C, F\sharp)$, we will end up with two different results: If we move the arrow rightward, we obtain $(C, F\sharp) - (D\flat, F)$. If we move it leftward off the edge of the strip, it reappears on the right edge upside down; sliding it leftward produces $(C, F\sharp) - (B, G)$. (b) Musically, these two methods of “sliding” correspond to transposition up by semitone and down by perfect fourth, respectively.

pointing in different directions, depending on how we move it between the two points.¹²

We can interpret the geometry as follows: The arrow $(B, F) \rightarrow (C, E)$ represents a voice leading in which B moves up by semitone and F moves down by semitone. When we slide the arrow horizontally we are transposing the voice leading: One of the notes in the new chord will correspond to B and will move up by semitone; the other will correspond to the F and will move down by semitone. But of course, there are two ways to do this for any pair of

¹² Geometrically, there is a significant difference between the Möbius strip and the sphere: The former is *nonorientable*, while the latter has *curvature*. In both cases, however, the result of sliding a vector around the space depends on the path.

tritones; for example, we can transpose B to C and F to F \sharp , which corresponds to sliding the arrow rightward, or we can transpose B to F \sharp and F to C, which corresponds to sliding the arrow leftward off the edge of the Möbius strip, so that it reappears on the right side, upside down (Figure 3). Thus, the geometrical fact that we cannot uniquely transport vectors from point to point reflects the *musical* fact that there are two equally good ways to transpose one tritone onto another.

Note that the same problem can arise even when we consider discrete lattices rather than continuous geometrical spaces. Figure 4, for example, represents single-semitone voice leadings among equal-tempered tritones and perfect fifths.¹³ Were we to use functions to model directed motion in this space, we would encounter the same difficulties described above; for instance, we can move the arrow on Figure 4 three squares to the right, so that it points from {F \sharp , C} to {G, C}, or we can move it three squares to the left, so that it points from {C, F \sharp } to {D \flat , G \flat }. Once again, we cannot make a principled choice from among these two possibilities: The voice leading (E \flat , A) \rightarrow (E, A) is equally related to both (F \sharp , C) \rightarrow (G, C) and (C, F \sharp) \rightarrow (D \flat , G \flat). (In both cases a tritone moves to a perfect fifth by a single-semitone ascent.) Clearly, this is because Figure 4 exhibits the same nontrivial structure as the larger Möbius strip in which it is embedded.

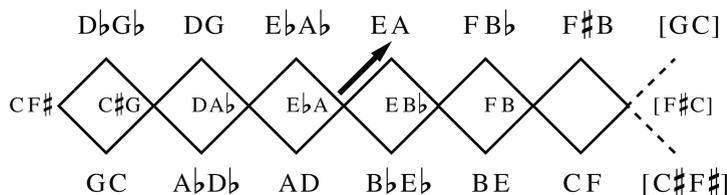


Figure 4. A discrete graph depicting single-semitone voice leadings among equal-tempered tritones and perfect fifths. The figure involves the same topological complexities as those in Figure 3.

It follows that we should not try to “fix” our spaces to remove the “problem”: Our twisted geometry is faithfully reflecting genuine musical relationships. That it takes a sophisticated space such as a Möbius strip to represent elementary musical intuitions is surprising, but no cause for alarm. (In particular, it is no reason to try to force our musical space to conform to theoretical requirements designed for simpler situations.) What we need are theoretical tools that allow us to work with the space as it is. This means rejecting the assumption that directed magnitudes, distances, or motions necessarily

¹³ This diagram is the two-dimensional analogue to Jack Doughett and Peter Steinbach’s “cube dance” (1998). However, I have explicitly shown that the graph is topologically twisted, a feature that can be understood by considering

how it is embedded in Figure 3. Doughett and Steinbach’s discrete graphs, such as “cube dance,” contain similar twists, though this is not evident in the discrete perspective they adopt. See Tymoczko 2006 and 2009 for more.

correspond to functions defined over all the points of the space—for unless we do so, we must declare that the arrow $(B, F) \rightarrow (C, E)$ corresponds to just *one* of the two arrows $(C, F\sharp) \rightarrow (D\flat, F)$ and $(C, F\sharp) \rightarrow (B, G)$, and we cannot make this choice in a nonarbitrary way. Musically, any tritone can resolve semitonally to two different major thirds, a truism that (rather remarkably!) corresponds to the fact that there is no path-independent way to move arrows around the Möbius strip.

Such considerations help explain why contemporary geometers consider vectors to be *local* objects, defined at particular points in a space, rather than as functions over the whole. Formally, they model vectors using what is called the “tangent space” to a point, a space that can be visualized as an ordinary, flat, Euclidean approximation to an infinitely small region of the original space.¹⁴ Figure 5 illustrates this idea in the case of the sphere. The thought here is that a nontrivial space (e.g., the sphere) seems more and more ordinary (and “flat”) the smaller we are; thus, in daily life, the curved surface of the earth is for all intents and purposes equivalent to a flat Euclidean plane. A geometrical space can be understood to have a tangent space attached to each and every one of its points, representing an “infinitely small” region of the space where traditional geometrical intuitions can be safely applied. Vectors, in this picture, live in one and only one tangent space. Since they may not be uniquely transportable from one place to another, we cannot say whether the two vectors in Figure 5 are “the same” or not.¹⁵

We have therefore encountered two reasons to resist the strategy of modeling intervals as functions. Musically, we often need to refer to particular directed motions such as $G \rightarrow E\flat$ rather than $F \rightarrow D\flat$, and this sort of talk is not easily captured within the intervals-as-functions perspective. Mathematically, there are spaces such as the surface of the earth in which it is not possible to transport arrows from one point to another in a unique way; thus, if we insist on modeling “directed distances” as Lewinian functions, we must abjure such spaces even before we know whether they are musically useful. My suggestion here is that these two problems are related: The musical term *interval* is in some ways analogous to the mathematical term *vector*, which is fundamentally a local object existing at a particular point in space. I would therefore recommend that we model musical intervals as particular motions

14 A subtlety: Contemporary mathematicians distinguish vectors-at-a-point from paths-in-the space; for a mathematician the directed distance “go 150 miles south” is not itself a vector. (Technically, vectors provide a way to differentiate functions in the space, whereas directed distances like “go 150 miles south” describe finite motions along geodesics.) However, in sufficiently well-behaved spaces, there is a very close relationship between vectors and paths—in a Lie group, for example, the operation of *exponentiation* relates them. For a good introduction to this and other geometrical ideas in the essay, see Nakahara 2003 or the more accessible Penrose 2004.

15 Mathematicians would say that the agglomeration of tangent spaces has the structure of a *fiber bundle*, with the sphere being the *base space* and the infinite Euclidean plane being the *fiber*. Though the term *fiber bundle* may seem forbiddingly technical, it is simply a new name for what we have already learned: In a fiber bundle, there is no unique way to associate the points of one fiber (e.g., the vectors in the tangent space at one point) with those of a distant fiber (e.g., the tangent space at another point).

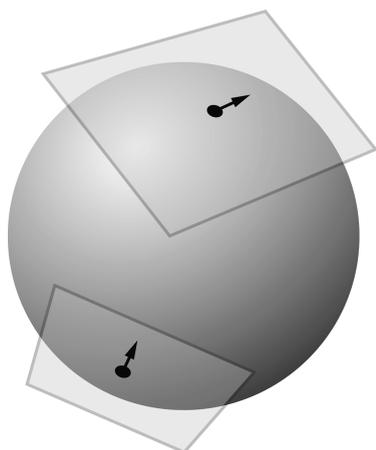


Figure 5. Mathematicians think of vectors as inhabiting the “tangent space,” that is, a Euclidian space that is tangent to a point in the original space. Intuitively, a tangent space is a flat (Euclidean) approximation to a very small region of the curved space—represented here by planes attached to the surface of a sphere. In some spaces, there is no unique way to associate vectors in one tangent space with those in another. Thus, we cannot say whether the two vectors in the two tangent spaces shown here are “the same” or not.

first and foremost: “G4 moves down by four semitones to E \flat 4,” “the tritone {B4, F4} resolves to {C5, E4} by semitonal motion,” and so on. In some cases, we can organize these particular motions into larger categories, perhaps so as to define functions over an entire space. But there is no musical or mathematical reason to demand that we *always* be able to do so. Confronted with spaces like the sphere or Möbius strip—spaces in which *local motions do not define global functions*—we must reach beyond group theory and toward a more geometrical conception of “directed distance.”

II. Boundaries and paths

I now turn to other problems with Lewin’s approach to musical intervals—problems that arise even when we can transport arrows from point to point. These result from Lewin’s requirement that intervals always be defined at *every* point in the space and that they be represented by functions whose identity is entirely determined by their inputs and outputs.¹⁶ As we will see, these requirements further constrain the range of applicability of Lewin’s theory by prohibiting spaces with boundaries and by eliding the various paths that might connect the same points.

¹⁶ In the second half of *GMIT*, Lewin turns from “intervals” to more flexible objects called “transformations.” However, since transformations are functions defined over an entire

musical space (see *GMIT*, definitions 1.2.1, 1.3.2, and 9.3.1), they are still subject to the various difficulties described in this essay.

Let's begin with the space of two-note set-classes, shown in Figure 6.¹⁷ Each point represents an “interval class” or equivalence class of unordered, transpositionally related pairs of pitch classes. (Mathematicians would say that this is the quotient of the Möbius strip in Figure 3 by horizontal motions—or in topological terms, the “leaf space” of the transpositional “foliation” of two-note chord space.) The endpoints are the unison $\{0, 0\}$ and tritone $\{0, 6\}$, the smallest and largest interval classes. The other equal-tempered dyadic set-classes are equally spaced along this line. (Note that for our purposes we can think of this space either as continuous, containing all possible points $\{0, c\}$ where $0 \leq c \leq 6$, or as discrete, with just the seven labeled points.) Distance represents the size of the minimal voice leading between the set-classes' elements, measured using the “taxicab” metric. For example, $\{0, 0\}$ is one unit away from $\{0, 1\}$ since we can transform any unison into a semitone by moving one of its notes by one chromatic step; the unison is six units away from the tritone, since we need to move one note by six chromatic steps to get from one interval to another.¹⁸ Once again, any two-note chord determines a point in this space, and any voice leading between two-note chords determines a line segment.¹⁹

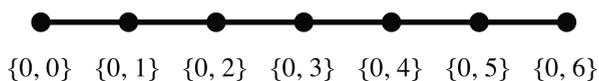


Figure 6. The space of two-note set classes. Distance in the space represents the size of the minimal voice leading between elements of the respective set classes; for instance, $\{0, 2\}$ is two units away from $\{0, 4\}$ because it is possible to transform any major second into a major third by moving one note by two semitones.

Now suppose we want to represent the directed motion in which an interval becomes two semitones larger. (Visually, this corresponds to moving two units rightward, toward the tritone.) Here we do not encounter problems with transportability: At any point on the interior of the line segment, it is obvious which direction we should move. But of course it is not always possible to move two units in this direction: At $\{0, 5\}$ we can move only one unit to the right, and hence the direction “make this interval class two semitones larger” cannot be obeyed. Once again, it would be wrong to “fix” the space so as to remove the “problem,” as the geometry is faithfully reflecting musical facts—in this case that the tritone is the largest interval class.

¹⁷ Such spaces are described in Callender, Quinn, and Tymoczko 2008.

¹⁸ Alternatively, of course, we can move unison to minor second by moving two notes by half a semitone, or one by a third of a semitone and the other by two-thirds of a semitone; in each case, the aggregate distance moved by both voices is 1.

¹⁹ Mathematical note: Since the boundaries of the space are singular, these “line segments” may appear to “bounce off” them. This is because the space is a quotient space or orbifold, as discussed in Tymoczko 2006 and in Callender, Quinn, and Tymoczko 2008.

Interestingly, Lewin addresses this issue in the early pages of *GMIT*. There, he notes that his formal definition of a “generalized interval system” requires that pitch space be infinite: We cannot construct a GIS to model the intervals among a finite space of pitches (e.g., the keys on an ordinary piano keyboard), for precisely the reasons explained above. Lewin responds by writing:

If we can conceive of an element s and if we can conceive of a characteristic measurement, distance, or motion i , then we can conceive of an element t which lies at distance i from s . In certain specific cases, application of this idea may require enlarging practical families of musical elements, to become larger formal spaces that are theoretically conceivable while musically impractical. For instance, we shall need to conceive supersonic and subsonic “pitches” in order to accommodate the idea of being able to go up or down one scale degree from *any* note. (27)

But let us reconsider this in light of Figure 6. We can certainly conceive of a “minor second,” an equivalence class containing pairs such as $\{C, C\#\}$ and $\{E, F\}$. And a phrase like “two semitones larger” certainly seems to define a “directed distance” between interval classes: $\{0, 3\}$ is two semitones larger than $\{0, 1\}$, just as $\{0, 4\}$ is two semitones larger than $\{0, 2\}$. But, contra Lewin, we *cannot* conceive of a pair of points on a circle being more than half a circumference apart.²⁰ Consequently, it is impossible to imagine an unordered pair of pitch classes whose notes are two semitones farther apart than those of a perfect fourth. (The pitch classes $\{0, 7\}$ are *five* semitones apart on the pitch-class circle, not seven, and only one unit farther apart than those of the major third.) Unlike the piano keyboard, the musical space of Figure 6 is *bounded in principle*: Any conceivable dyadic set class must lie on the continuous line segment shown in the figure. For this reason, Lewin’s initial sentence (“if we can conceive . . . then we can conceive”) would seem to be false: It is literally impossible to imagine an unordered dyad of pitch classes lying more than half an octave apart, or less than zero semitones apart.²¹

Related problems can arise when there is *more than one interval* between the same pairs of points. For example, on a circle there are multiple ways to move between any two locations: one can move from 2 to 3 on the ordinary clock face by moving $1/12$ of a circumference clockwise, $11/12$ of a circle counterclockwise, $13/12$ of a circumference clockwise, and so on. (Here again there is no problem of transportability: At any point on the circumference, it is obvious how to move clockwise.) Similarly, on the surface of the earth, one

²⁰ We can imagine *moving* from one point on a circle to another by more than half a circumference, as in Figure 8. But since the distance between two points corresponds to the length of the *shortest* path between them, it is impossible for that distance to be more than half of the total length of the circle. And while it is true that we could consider the infinite space of dyadic *pitch* set-classes, this new space is not a simple extension of our original space: Pitch-class set-classes are different from pitch set-classes, and the

question here is whether we can use Lewinian techniques to model the former.

²¹ It is interesting that Lewin sometimes acknowledges that particular musical spaces do have boundaries and thus fail to meet the definition of a GIS (see *GMIT*, 24–25 and 29–30, in connection with durational space 2.2.5). This admission would seem to conflict with the assertion quoted in the main text. Thanks here to Steven Rings.

can move between two points in two different directions along a great circle; for antipodes, there are infinitely many “shortest paths” between them. Lewin defines intervals as functions—which is to say, “machines” that take one point as input and deliver another as an output. As a result, there is no way to define distinct intervals between the same sets of points; thus, we cannot distinguish the process of moving $1/12$ of a circumference clockwise from moving $11/12$ of a circumference counterclockwise.²²

Musically, this creates difficulties even in mundane analytical contexts. Suppose we want to formalize the statement “G, in any octave, moves downward by four semitones to E \flat .” (We might, for example, claim that this is the germinal intervallic motive in the first movement of Beethoven’s Fifth Symphony.) We would like to distinguish this particular motion both from its transpositions, such as F4 \rightarrow D \flat 4, and from those in which G moves to E \flat along some other path, for instance, by eight ascending semitones. In other words, we want to group together the motions in Figure 7a while distinguishing them from those in Figure 7b and 7c. How can we model this straightforward musical idea?



Figure 7. The directed motions in (a) are similar, since G moves to E \flat by four descending semitones. Those in (b) move G to E \flat by eight ascending semitones, while those in (c) involve notes other than G and E \flat . The similarities between the progressions in (a), and their difference from (b) and (c), are not captured by traditional pitch or pitch-class intervals: As pitch intervals, the progressions in (a) are no more similar than are those in (c), while as pitch-class intervals (a) and (b) are the same.

Clearly, Lewinian intervals will not do the job: As functions over all of pitch space or pitch-class space, they will not distinguish $G \rightarrow E\flat$ from $F \rightarrow D\flat$, nor do they distinguish “G moves down by four semitones to E \flat ” (which we can write $G \xrightarrow{-4} E\flat$) from “G moves up by eight semitones to E \flat ” (or $G \xrightarrow{+8} E\flat$). However, we can model these progressions as vectors in the *tangent space* of a particular point on the circle.²³ As shown in Figure 8, the tangent space is a

22 A subtle point: Lewinian functions can act similarly upon *some* but not *all* points in a space; for instance, Lewin could distinguish the operation T_7 (which transposes every pitch class by seven ascending semitones) from I_7 (which, like T_7 , sends C to G and F \sharp to C \sharp , but acts differently upon the remaining pitch classes). However, the paths “clockwise by $1/12$ of a circumference” and “counterclockwise by $11/12$ of a circumference” link exactly the same sets of points: Given any input point p , they return precisely the same output q .

Here, the only distinction lies in *how they move* the input to the output.

23 This conception of pitch-class intervals as objects in the tangent space is reminiscent of Mazzola’s (2002) distinction between “local compositions” and “global compositions.” I will not attempt to address Mazzola’s book in detail, because its mathematical complexities far outstrip those we are dealing with here.

one-dimensional line approximating an infinitely small region of the circle. Vectors in this line can be represented as real numbers, such as -4 , $+8$, and $+36$. Musically, these can be interpreted as directions such as “start at G and go down four semitones,” “start at G and go up eight semitones,” and “start at G and go down sixteen semitones.” (The first two progressions in Figure 7 are represented by the vector “ -4 ,” while the second two are represented by “ $+8$.”) Figure 8 shows that vectors in the tangent space determine unique *paths* in pitch-class space. For example, the vector -4 , located at G, corresponds to the path in which the pitch class G moves down by four semitones to $E\flat$ (a counterclockwise motion by four-twelfths of a circumference), while $+8$ corresponds to the path in which it moves up by eight semitones to $E\flat$ (a counterclockwise motion by eight-twelfths of a circumference). For every vector in the tangent space, there is a unique path on the pitch-class circle and a unique collection of octave-related progressions such as those in Figure 7a; conversely, for every unidirectional path on the circle there is a unique vector in some tangent space.²⁴ Note that the relevant paths can sometimes wrap around the circle one or more times, corresponding to motions that span more than an octave.

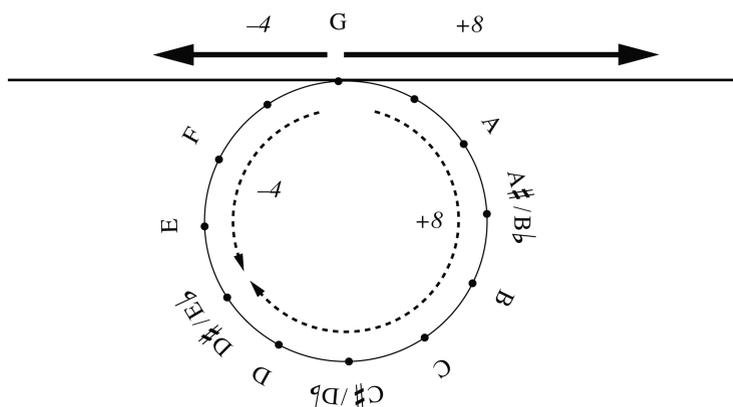


Figure 8. The tangent space to a circle is an infinite line. Two vectors are shown, corresponding to the descending major third $G \xrightarrow{-4} E\flat$ and the ascending minor sixth $G \xrightarrow{+8} E\flat$; the associated paths are shown within the circle. There is a one-to-one correspondence between vectors in the tangent space, unidirectional paths in the circle, and equivalence classes of octave-related pitch intervals such as those in Figure 7a or 7b.

Paths in pitch-class space are analytically quite fruitful: Without them, it is difficult to formalize the thought that the central motive of the first movement of Beethoven’s Fifth Symphony involves G (in some octave) moving

²⁴ Again, the association between vectors in the tangent space and paths in the circle is called “exponentiation” in the theory of Lie groups (Nakahara 2003).

down to $E\flat$ by four semitones, with the specific octave being less important than the particular path.²⁵ But to model this simple intuition we need to reject intervals-as-functions in favor of a conception in which intervals are, first and foremost, *particular motions*. Formally, this can be accomplished by reconceiving pitch-class intervals as *equivalence classes of ordered pairs of pitches*. Particular motions in pitch space can be modeled as ordered pairs; for instance, $(G4, E\flat4)$ represents the motion in which G4 moves downward to $E\flat4$. More general intervals can then be built as equivalence classes of these particular motions; for instance, we can identify “the descending major-third $G \rightarrow E\flat$ ” with the equivalence class containing pairs such as $(G4, E\flat4)$, $(G5, E\flat5)$, and $(G3, E\flat3)$ but not $(G4, E\flat5)$ or $(F4, D\flat4)$. The even more general term *descending major third* can be understood as the equivalence class including pairs such as $(G4, E\flat4)$ and $(F3, D\flat3)$.²⁶ If we follow this procedure, there is no reason to require that intervals be functions, or that they be defined over an entire space, or even that there be just one interval between the same two pitch classes. Instead, we are free to make a wider range of choices about how to group particular motions into more general intervallic categories.²⁷

This moral here is a general one: Intervals, conceived as categories of particular motions, are more flexible than Lewinian intervals-as-functions. As musicians, we need to be careful, lest the functional perspective blind us to interesting possibilities. For an illustration of this phenomenon, consider neo-Riemannian harmonic theory. Lewin modeled “neo-Riemannian transformations” as *functions* that take a single chord as input and return another chord as output. Thus, the neo-Riemannian *Quintschritt* is understood as a function that transforms a major triad into its ascending perfect-fifth transposition while turning a minor triad into its *descending* perfect-fifth transformation.²⁸ Unfortunately, the functional perspective makes it very hard to see how one might apply these concepts to inversionally symmetrical chords such as the augmented triad, the fourth chord, or the diatonic triad: A neo-Riemannian transformation would be a function that moves inversionally symmetrical chords upward *and downward* by the same distance at one and the same time!

However, we can overcome this difficulty if we are willing to model neo-Riemannian relationships using *equivalence classes of pairs of chords*: We simply assert that the C augmented chord is in the *Quintschritt* relationship to *both* G augmented and F augmented. The point here is to underscore the fact that

25 For more on paths in pitch-class space, see Tymoczko 2005, 2008b, and 2010, as well as Callender, Quinn, and Tymoczko 2008.

26 One could also construct traditional pitch-class intervals in this way.

27 Returning to the example of two-note set-classes, we can identify the interval “two semitones larger” with motions such as $\{0, 0\} \rightarrow \{0, 2\}$, $\{0, 1\} \rightarrow \{0, 3\}$, and $\{0, 2\} \rightarrow \{0, 4\}$. This equivalence class will contain no pair whose first element is

$\{0, 5\}$ since the interval “two semitones larger” is not defined for every set class. For more on this approach to intervals, see Callender, Quinn, and Tymoczko 2008.

28 The collection of neo-Riemannian transformations—the *Schritts* and *Wechsels*—constitutes a set of Lewinian “intervals” between the twenty-four major and minor triads. See Klumpenhouwer 1994 for more.

neo-Riemannian theory categorizes chord progressions by means of transpositional and inversional equivalence: Two chord progressions $A \rightarrow B$ and $C \rightarrow D$ exemplify the same “neo-Riemannian transformation” if and only if there is some transposition or inversion that sends $A \rightarrow B$ into $C \rightarrow D$ (Figure 9).²⁹ Since we can transpose or invert $A \rightarrow B$ even if one of the chords is inversionally symmetrical, there is no reason to restrict neo-Riemannian concepts to asymmetrical chords. Of course, we will no longer be able to model *Quintschrift* as a single-valued *function* that transforms one chord into another; instead, we need to model it as a *relationship* that may obtain between one chord and several others. That is, we group the particular motions C augmented \rightarrow G augmented and C augmented \rightarrow F augmented into the same category, without worrying about whether this produces a function or not.³⁰

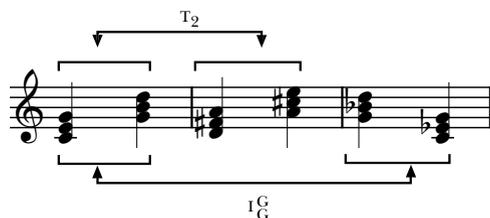


Figure 9. Two progressions exemplify the same “neo-Riemannian transformation” if and only if they are related by transposition or inversion. Here, three instances of *Quintschrift*.

In fact, it is even possible to define augmented-chord analogues to the familiar neo-Riemannian *L*, *P*, and *R* operations. For instance, we could say that two augmented triads $A \rightarrow B$ are in the neo-Riemannian relationship X (where X is some combination of the familiar *L*, *P*, and *R* transformations) if and only if there exist two triads C and D , each either major or minor, such that A and C , and B and D share a major third, with C and D being related by the traditional neo-Riemannian transformation X . From this point of view, C augmented and B augmented are in the *P* relationship, since C augmented is a semitone away from C major, while B augmented is a semitone away from C minor, while C major and C minor are *P* related. The benefit is that we can now imagine music in which neo-Riemannian harmonic ideas are applied to

²⁹ See Tymoczko 2008b. In the language of Callender, Quinn, and Tymoczko (2008), two progressions are *dualistically equivalent* if and only if they are related by *individual* applications of the octave, permutation, and cardinality-changed operations and *uniform* transposition and inversion.

³⁰ Lewin (*GMIT*, 177) asserts that Riemann “did not quite ever realize that he was conceiving a ‘dominant’ . . . as something one *does* to a Klang,” a remark that suggests the functional perspective is internal to Riemann’s thought, though perhaps in a latent and unrealized way. Several

sentences later, however, Lewin seems to criticize Riemann for conceiving of “dominants” in a nontransformational, Cartesian sense—as “*labels for Klänge* in a key, rather than as *labels for transformations* that generate Klänge from a local tonic.” Here I am arguing that Riemannian labels-for-Klänge are actually more powerful than Lewinian transformations, as they can be defined even for symmetrical chords: There is no *function* that is *Quintschrift* of an augmented triad, though we can label two separate augmented triads as the *Quintschrift* of a third.

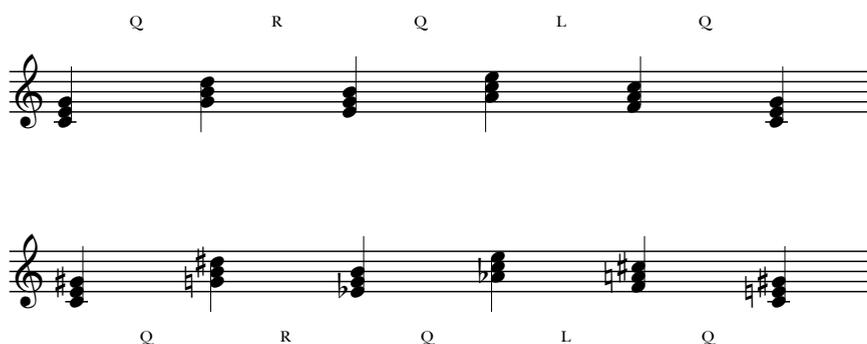


Figure 10. The top staff presents a series of triads exemplifying familiar neo-Riemannian relationships: *Quintschrift*, which relates a major triad to its ascending-fifth transposition (and a minor triad to its descending-fifth transposition); the *Relative* relationship, which relates major and minor triads sharing a major third; and *Leittonwechsel*, which relates major and minor triads sharing a minor third. The bottom staff presents a series of augmented triads linked by analogous relationships; here, two augmented triads are said to stand in a particular neo-Riemannian relationship if and only if they are semitonally related to major and minor triads that are in that same relationship. These relations cannot be modeled by functions.

a broader range of chords; for instance, a composer could write a passage in which the same sequence of neo-Riemannian relationships obtained between an initial sequence of major and minor triads, and a second sequence of augmented triads (Figure 10). Thus, by abandoning the requirement that neo-Riemannian concepts be modeled as *functions*, we gain the ability to imagine new musical possibilities.

III. How large are intervals?

Section I considered spaces in which there was no principled way to transport intervals from one point to another. Section II discussed additional problems that can occur even when the intervals are transportable: Sometimes not every interval is defined at every point, and sometimes there are multiple intervals between the same pairs of points. Section III now narrows the focus even further, considering cases in which these problems are irrelevant. Even here, I suggest, we continue to encounter difficulties with Lewin's formalism.

The basic question I want to ask is whether Lewinian intervals are supposed to have size. This issue arises because Lewin simultaneously asserts that intervals represent “measurements” or “distances”—paradigmatically quantifiable entities—while also proposing a formal group-theoretical model in which magnitudes are not explicitly represented.³¹ Consequently, readers of

31 For suggestions that intervals have size, see note 1.

GMIT need to reconcile the nonmathematical rhetoric of Lewin's book, which often seem to suggest that intervals are quantifiable, with the more formal mathematical portions, which give the opposite impression.³² Depending on which of these we emphasize, we get a very different picture of the work's fundamental goals.

Consider, for specificity, the following three ways of modeling equal-tempered pitch-class intervals:

Interval System 1 has as its space the twelve equal-tempered pitch classes.

The intervals consist of the set $\{0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11\}$, with elements combining under addition mod 12. Each interval i represents the directed distance from a pitch class to the pitch class i semitones above it.

Interval System 2 has as its space the twelve equal-tempered pitch classes.

The intervals consist of the set $\{0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11\}$, with elements combining under addition mod 12. Each interval i represents the directed distance from a pitch class to the pitch class i perfect fourths above it.

Interval System 3 has as its space the twelve equal-tempered pitch classes.

The intervals consist of the set $\{0, 1, 2, 3, 4, 5, 6, 7, 8, 9, t, e\}$. Elements combine under addition mod 12, with the symbol "t" acting like "10" and "e" acting like "11." Each interval i represents the directed distance from a pitch class to the pitch class i semitones above it, with "t" acting like "10" and "e" acting like "11."

Question: Are the three interval systems mere notational variants of one another, or are they importantly different? It would seem that there are two possible answers. The first is that the three systems are all fundamentally equivalent because they differ only in the labels assigned to their intervals. Thus, if we treat interval names such as "1," "2," "10," and "t" simply as uninterpreted symbols, there is no important difference among these GISs. (Mathematicians would say that the three interval groups are "permutation isomorphic" and are in that sense equivalent.)³³ Consequently, we should not say that "interval system 1 asserts that the semitone is smaller than the major second, but interval system 2 asserts the converse."³⁴

³² Hall 2009 also notes inconsistencies between Lewin's formalism and his informal writing.

³³ "Permutation isomorphism" is a more concrete analogue of "isomorphism," applicable when groups act upon particular sets (Robinson 1996). However, the notion is somewhat superfluous in this context: Lewinian interval groups are simply transitive, so abstract isomorphism of interval groups implies permutation isomorphism of the relevant group actions.

³⁴ Mathematicians often warn against the temptation to overconcretize mathematical objects by treating them as

having additional structure not preserved by the relevant notion of isomorphism (e.g., Lawvere and Schanuel 1997, 89–90). But this is arguably what would be involved in simultaneously asserting (1) that group theory provides the appropriate mathematical framework for modeling intervals, and (2) that the *labels* assigned to particular group elements have quantitative significance. These labels are not preserved by (permutation) isomorphism and hence are not the province of group theory.

The second answer is that the three systems are not the same, because interval labels matter. By assigning the labels “5” and “1” to the interval C–F, interval systems 1 and 2 attribute to it different magnitudes: Interval system 1 says that this interval is five units large, while interval system 2 says that it is one unit large. But here we face a difficulty: If we believe that interval labels matter, then we will certainly want to say that some differences matter more than others. For example, the difference between interval systems 1 and 3 is relatively unimportant, since the symbols “10” and “t[en]” both seem to refer to the same number, as do “11” and “e[leven].” Consequently, if we want to assert that interval names matter, we need to clarify which differences (e.g., that between “1” and “5”) are significant and which (e.g., that between “t” and “10”) are merely orthographical.

It is somewhat remarkable that Lewin himself did not provide a solution to this problem: It is not at all clear whether interval systems 1 and 2, introduced in *GMIT* on pages 17 and 22, are supposed to be the same or different. As a result, readers must choose between two coherent but incompatible readings of his book. The first treats interval labels as insignificant, asserting that Lewin wants us to *abstract away from any intuitions about interval size*.³⁵ The second reading claims that Lewin meant to assign sizes to intervals, using non-group-theoretical structure that is not fully captured by the mathematical machinery he defined.³⁶ However, somewhat surprisingly for a book that aspires to be both rigorous and explicit, Lewin did not discuss (or even mention) this extra structure. On this reading, he simply assumed that his audience would intuitively understand how to quantify the “sizes” of the intervals in the various systems he discussed.

Rather than attempting to solve this vexing interpretive dilemma, let us instead ask how we might augment Lewin’s formalism so as to represent intervals as genuinely quantifiable “distances.” One possibility is to define a *Lewinian interval system* as a combination of three mathematical structures:

- (1) a set (or “space”) of musical objects S ;
- (2) a distance function $D(x, y)$ from ordered pairs of points in the space to nonnegative numbers, representing the “distance” between them; and
- (3) an “interval group” *IVLS*, which acts on S so as to satisfy two conditions:

³⁵ This reading is supported by the general fact that group theory typically abstracts away from issues of “size,” by Lewin’s statement that the transformational attitude “does not ask for some observed measure of ‘extension’” of intervals (*GMIT*, 159) and by Lewin’s remark that GISs with isomorphic interval groups are “essentially the same” (78).

³⁶ This reading is supported by the facts that Lewin announces an explicit intention to model measurements and distances (*GMIT*, 16, 22, 25, 26, 27); asserts that

intervals have “extensions” or sizes according to the traditional/Cartesian conception (159); refers to intervals as “numbers” (19, 21), using such terms as “greater than” (21) and “maximal” (38–40) to compare them; and seems to distinguish between measuring pitch-class intervals “in semitones” and “in fifths” (22). Furthermore, several of Lewin’s mathematically sophisticated readers have concluded that intervals indeed have size (see note 1).

- (a) *Simple transitivity.* For all x, y in S , there is exactly one g in $IVLS$ such that $g(x) = y$.
- (b) *Metric compatibility.* $D(x, g(x)) = D(y, g(y))$, for all x, y in S and all g in $IVLS$.

Conditions 1 and 3a restate Lewin's definition of a "GIS," while conditions 2 and 3b extend the formalism: Unlike Lewin, I require that we explicitly declare how we are measuring distance in the space, with each interval moving all objects by exactly the same distance. (This condition is inspired by geometry, where the concept of a "metric" is central—indeed, the word *metric* is a constituent of *geometric*, and metrics are part of what distinguishes geometrical from topological spaces.)³⁷ These extra conditions allow us to speak of the "size" of the intervals, as we can stipulate that the size of an interval g is simply $D(x, g(x))$.³⁸

This definition provides a straightforward criterion of equivalence for two Lewinian interval systems: Two Lewinian interval systems are *notational variants* if

- (1) they share the same space;
- (2) they share the same distance function; and
- (3) there is a bijective function between their interval groups, $F: IVLS_1 \rightarrow IVLS_2$, such that $g(x) = y$ if and only if $F(g)(x) = y$, for all x and y in the space.³⁹

For example, suppose we convert interval systems 1–3 into Lewinian interval systems as follows: We declare interval "0" to have size 0; intervals "1," "11," and "e" to have size 1; intervals "2," "10," and "t" to have size 2; and so on. (This is the standard notion of "interval class.") Interval systems 1 and 3 are therefore notational variants, while interval system 2 is distinct. This captures the intuitive sense that the first and third systems measure "in semitones," while the second measures "in fourths."

A "Lewinian space" can be defined as a collection of musical objects for which there is some intuitively plausible conception of "interval" that can be modeled using a Lewinian interval system. Lewin's work shows that some of the most familiar spaces in music theory can be understood in this way: linear pitch space, circular pitch-class space, the space of Babbittian time points,

³⁷ Note that for the sake of generality I do not require that a "distance function" satisfy the formal properties of a metric—it simply needs to assign a nonnegative number to every pair of points in the space.

³⁸ "Lewinian interval systems," by departing from Lewin's "GISs," are inherently slightly non-Lewinian. The goal is to explicitly model intuitions that may only be implicit in Lewin's writing—to capture the spirit behind his work, even while departing from the letter of his formalism. In particular, I am trying to find a *minimal* alteration to his ideas that allows us to use a Lewin-like formalism to represent something like directed distances. Here it is relevant that any GIS can

be transformed into a Lewinian interval system simply by assigning a nonnegative number to each of its intervals.

³⁹ Note that this condition is stronger than mere (permutation) isomorphism of interval groups. The octatonic STRANS1 and STRANS2, discussed in appendix B of *GMIT*, share the same musical space and have permutation-isomorphic interval groups, but they are not notational variants since there exists no bijection satisfying condition 3. (The same can be said for the interval systems constructed using the T1 and *neo-Riemannian Schritt/Wechsel* groups acting upon the twenty-four major and minor triads.) Thanks here to Jason Yust.

and so on. The question immediately arises whether we can characterize these spaces more precisely, thereby demarcating a “space of spaces” in which Lewinian techniques are viable. Sections I and II suggest two answers. First, Lewinian spaces are *homogeneous* spaces in which every point looks the same as every other: In particular, there are no special points such as boundaries at which only some intervals are available. Second, in Lewinian spaces, there is a unique way to move intervals throughout the space, so that we can compare an interval at one point to an interval at another. This means that these spaces are very similar to what geometers call *parallelized* spaces. (A parallelized space is one in which there is a recipe for identifying the tangent spaces at different points; some spaces, such as the Möbius strip and the two-dimensional sphere, are inherently unparallelizable.)⁴⁰ Continuous Lewinian spaces are also closely related to Lie groups: homogeneous, unbounded, parallelizable manifolds that also have group structure.

However, there are a number of interesting musical spaces that are *not* Lewinian, including the space of two-note chords (Figure 3) or two-note chord types (Figure 6), as well as their higher-dimensional analogues (see Tymoczko 2006; Callender, Quinn, and Tymoczko 2008). The list could be expanded so as to include any number of more prosaic spaces, such as the space of pitches on an ordinary piano keyboard or the physical space of an actual musical stage. This is because Lewinian spaces, far from being generic or typical, possess the unusual properties of homogeneity and parallelizability. Furthermore, even in a Lewinian space such as the pitch-class circle, it may sometimes be useful to model intervals using *equivalence classes of particular motions* rather than functions. (As we saw in Section II, this will allow us to define “paths in pitch-class space,” which distinguish various ways of moving between pitch classes.) Thus, even though we *can* use Lewinian techniques to model intervals, it may not behoove us to do so.

At this point, I should mention that my ideas intersect with those of Ed Gollin (2000). Much as I have done, Gollin emphasizes that intervals are commonly considered to have size, while group elements are not. Similarly, both Gollin and I try to expand the Lewinian framework so that it explicitly represents distance. However, these expansions proceed in different ways: I use a distance function while Gollin measures distances using the *group structure* of an interval group. Technically, Gollin’s proposal is more restrictive than mine: Any notion of distance that can be modeled using Gollin’s techniques can also be modeled using my “Lewinian interval systems,” but the converse is not true.⁴¹ Furthermore, where I emphasize that Lewinian systems are

⁴⁰ Interestingly, it was not until the late 1950s that mathematicians were able to show that only the one-, three-, and seven-dimensional spheres are parallelizable (Milnor 1958).

⁴¹ Gollin requires that the size of a group element be modeled using its “word length” relative to some privileged set of size-1 generators. However, we might sometimes want to consider interval systems in which the group generators

have different sizes—as in a two-dimensional *Tonnetz* where one axis represents perfect fifths and the other major thirds. Similarly, it may sometimes be that the size of composite elements is not equal to their word length. For instance, we might want to measure elements of the *Schritt/Wechsel* group according to the size of the minimal voice leading between the chords they connect. (That is, we might

themselves merely a subset of the spaces that music-theorists might want to consider, Gollin seems more content with the Lewinian claim that groups suffice for modeling “directed measurements, distances, or motions.” As a result, his proposals do not address the issues discussed in Sections I and II: spaces in which there is no unique way to “transport” an interval from one point to another, and spaces with special points such as boundaries. Despite these differences, however, our two approaches seem to proceed from a similar sense of the limits of *GMIT*’s formalism and of its uneasy relationship to Lewin’s nonmathematical rhetoric.

IV. Analytical repercussions

The preceding sections have discussed some limitations of Lewin’s “generalized interval systems.” But Lewin’s reputation does not rest solely on his work as an abstract system-builder: He is widely and justly admired for the analytical *uses* to which he put his abstractions—specifically, for the dozens of brilliant analyses in which mathematical models provide insight into particular works. It is appropriate, therefore, to end by considering the relation between our theoretical concerns and the everyday business of musical analysis. Why should analysts care about any of the topics we have been discussing?

My answer is that these issues are important whenever we would like to think seriously about voice leading. As I have argued elsewhere, it can often be useful to categorize voice leadings on the basis of how their individual voices move, while at the same time disregarding the octave and instrument in which these motions occur.⁴² From this point of view, the first three voice leadings in Figure 11 exhibit the same fundamental structure: Each moves the G major triad to C major by holding the note G constant, moving the note B up by semitone, and moving the note D up by two semitones. The fourth voice leading does *not* exhibit this schema, since it moves its voices very differently, while the fifth diverges from the schema by using completely different chords. My belief is that we can understand voice leading clearly only when we see the first three passages as fundamentally similar, while being importantly different from the others. But it is impossible to categorize voice leadings in this way as long as we confine ourselves to Lewinian pitch or pitch-class intervals. The first three voice leadings in Figure 11 involve different *particular* motions in pitch space and hence are quite different from that perspective. By contrast, the first four voice leadings link the same pairs of pitch classes and are equivalent in that sense. (Note that the equivalence of these voice leadings makes it difficult to assign a “size” to the voice leading that reflects the size of the motions in the actual musical voices.)⁴³ Finally, if we focus only on

want L to be size 1, since L-related triads can be linked by single-semitone motion, with R and R-then-L-then-P being size 2, since they require two semitones of total motion.) This is impossible using Gollin’s techniques, but straightforward using my Lewinian interval systems.

⁴² See Tymoczko 2005, 2008b, and 2010, as well as Callender, Quinn, and Tymoczko 2008.

⁴³ It is important here to distinguish *distance* from *path length*. On an ordinary clock, the number 3 is one unit away from the number 2, but the clockwise path from 3 to 2 is

Figure 11 consists of five musical progressions, labeled (a) through (e), each shown in a grand staff (treble and bass clefs). Progression (a) shows G major (G-B-D) moving to C major (C-E-G) with G fixed, B moving up a semitone to C, and D moving up two semitones to E. Progression (b) shows G major (G-B-D) moving to C major (C-E-G) with G fixed, B moving up a semitone to C, and D moving up two semitones to E. Progression (c) shows G major (G-B-D) moving to C major (C-E-G) with G fixed, B moving up a semitone to C, and D moving up two semitones to E. Progression (d) shows G major (G-B-D) moving to C major (C-E-G) with G fixed, B moving up a semitone to C, and D moving up two semitones to E. Progression (e) shows G major (G-B-D) moving to C major (C-E-G) with G fixed, B moving up a semitone to C, and D moving up two semitones to E.

Figure 11. The first three progressions exemplify the contrapuntal schema “G major moves to C major by keeping G fixed, moving B up by semitone, and moving D up by two semitones.” The fourth progression does not exemplify this schema, since it moves its voices differently, while the fifth uses different chords.

functions in pitch space—such as “up a major second”—then there is no pertinent distinction between the first voice leading and the last: Both hold one voice fixed, move one voice up by semitone, and one voice up by two semitones, even though they apply these functions to completely different notes.

What is needed here is a middle path between pitch and pitch-class intervals, one in which voice leadings are modeled as *collections of paths in pitch-class space*. Far from being an esoteric theoretical contrivance, this approach captures an important aspect of compositional competence, namely, the knowledge of the various contrapuntal routes from chord to chord: “You can move G major to C major by holding G fixed, moving B up by semitone, and moving D up by two semitones.” This is precisely the sort of information that teachers have been passing on to their students for centuries. And though the underlying idea is simple, the analytical payoffs are potentially significant, for once we sensitize ourselves to these basic contrapuntal schemas, we can start to identify complex structural relationships at the heart of sophisticated pieces such as Wagner’s *Tristan* prelude or Debussy’s *Prelude to the Afternoon of a Faun* (see Tymoczko 2008b, 2010). In this sense, the analytical stakes in the present discussion could scarcely be higher: Lewin’s framework, by making it impossible to model paths in pitch-class space, poses a significant obstacle to formalizing voice leading, and hence to understanding a practice central to a huge swathe of Western music.⁴⁴

Here I am echoing my earlier remark that the functional perspective makes it difficult to refer to particular musical motions such as “G, in any octave, descends by major third to E \flat .” I have just observed that the Lewinian approach also makes it difficult to model contrapuntal schemas such as

eleven units long. When we are studying voice leading we are typically interested in *short motions* between pitch classes rather than distances per se. Thus, the fourth voice leading in Figure 11 would typically be considered inefficient, even though the pitch classes in each voice are separated by small distances in pitch-class space. Here each voice takes an unnecessarily long path between pitch classes—which is analogous to moving clockwise from 3 to 2.

⁴⁴ It is interesting that Lewin’s 1998 discussion of voice leading initially models voice leadings as functions, only to reject this perspective in favor of a more flexible formalism. However, he continues to use traditional pitch-class intervals, thereby considering the first four voice leadings in Figure 11 to be the same.

“G stays fixed, B moves up by semitone, and D moves up by two semitones.” In both cases I am talking about the analytical need to situate something that looks like a *pitch-space motion* (“down by four semitones,” “the root is held fixed, the third moves up by semitone, and the fifth moves up by two semitones”) at a particular point in a musical space in which octaves have been discarded (the pitch-class circle, the space of three-note chords). And in both cases, the relevant theoretical concept can be represented, quite literally, as an arrow in a robustly geometrical space (Figure 12; see also Tymoczko 2006). Though superficially similar to the diagram at the very opening of *GMIT*, these arrows cannot be modeled with Lewin’s techniques for all the reasons discussed above.

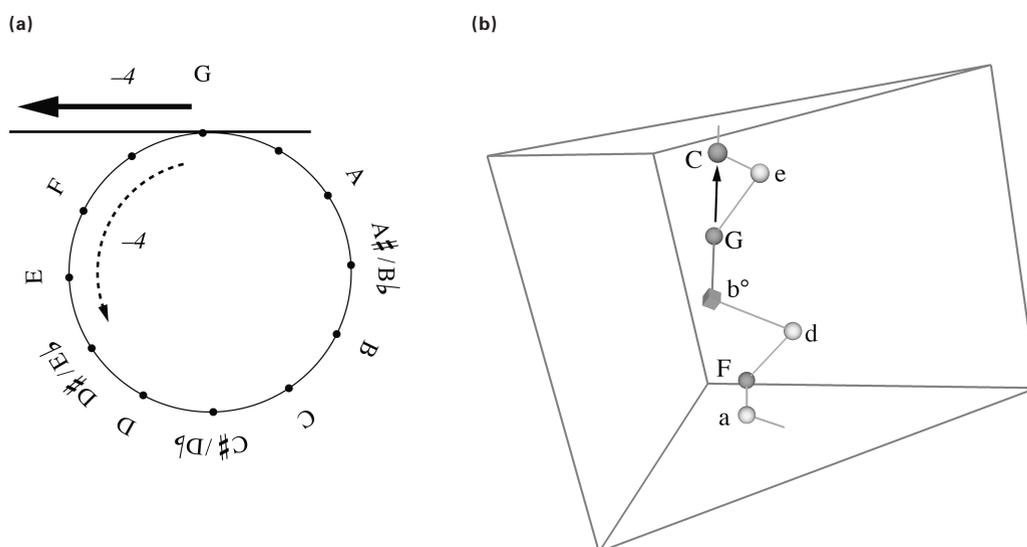


Figure 12. (a) The phrase “G, in any octave, moves down by four semitones to E \flat ” can be represented by an arrow in pitch-class space. (b) The phrase “G major moves to C major by keeping G fixed, moving B up by semitone, and moving D up by two semitones” can be represented by an arrow in the geometrical space representing three-note chords, which is the bounded interior of a twisted triangular 2-torus (or doughnut).

So a concern with voice leading gives us reasons to depart from Lewin’s conception of intervals. But what about the cases where we *can* adopt a Lewinian approach—do we thereby gain access to analytical insights that are otherwise unavailable?

Some tempting responses can immediately be put aside. For instance, one might think that Lewinian interval systems have the advantage of being the most general contexts in which we can apply the tools of traditional set theory—including the notions of transposition, inversion, and set-class. But this is incorrect: The Möbius strip of Figure 3 is not a Lewinian space, yet

transposition and inversion are perfectly well defined; transpositions are represented by horizontal motion, while inversions are represented by reflection around a vertical line.

Alternatively, one might think that Lewinian techniques enable us to appreciate the analogies between structures in different musical domains. For example, chapter 2 of *GMIT* details two different interval systems that are fundamentally similar—one representing pitch-class intervals and the other representing rhythmic intervals between “time points” modulo the measure. The isomorphism between these two interval groups means we can find rhythmic analogues to the traditional techniques of set theory. This, of course, was a central concern of integral serialists such as Milton Babbitt—one of Lewin’s primary teachers and influences. The abstract nature of Lewin’s “generalized interval systems” may seem particularly well tailored to revealing analogies between different musical domains, such as pitch and rhythm.

Again, however, we can often identify such parallelisms even when confronted with non-Lewinian interval systems. Let us return to the sphere, our paradigmatic example of a non-Lewinian space. We can easily imagine an actual, spatial sphere playing a role in electronic music; for example, in an auditorium surrounded on all sides (including ceiling and floor) by speakers, the sphere would represent the directions from which sound can emanate. Moving sound sources could therefore be modeled using paths on the sphere, analogous to the direction “move south 150 miles.” As it happens, the sphere also plays a role in (broadly) set-theoretical contexts as well; for example, it is the space of four-note pitch sequences modulo transposition and positive multiplication.⁴⁵ Hence, it would be possible to write a piece (or analyze one, if such a piece were already written) in which there were precise formal analogies between the motion of a sound source in the auditorium and the various (broadly) set-theoretical manipulations of its pitch structure. This is exactly the sort of cross-modal parallelism central to Babbitt’s work and for which the abstraction of *GMIT* seems well suited. Once again, we can find these parallels even if the underlying interval system is non-Lewinian.

Nor do we need Lewinian systems to create recursive structures in which large-scale transformations mirror small-scale intervallic relations. For example, the chord {G, C, E, A, B, F} contains three dyads that can be linked by the pair of voice leadings shown in Figure 13. These voice leadings can in turn be used to transform the entire chord, as shown in Figure 13c. Here, we move each of the component dyads along voice leadings defined by the chord’s internal structure; for instance, we apply the motions linking the first two dyads (“move both notes down by three semitones”) to *each* of the three

⁴⁵ See Callender, Quinn, and Tymoczko 2008. The space of four-note pitch sequences modulo transposition is \mathbb{R}^3 , or ordinary three-dimensional Cartesian space. Identifying sequences related by positive multiplication transforms \mathbb{R}^3 into the two-dimensional sphere.

The figure shows a musical score with six measures, labeled (a) through (f). Each measure contains three staves of music. Measures (a) and (b) show a chord with voice leadings. Measures (c) through (f) show a progression of chords, with (c) being a transformation of (a) and (d)-(f) being a transformation of a different chord.

Figure 13. The chord in (a) can be partitioned into three dyads, linked by the pair of voice leadings in (b). The progression in (c) applies the voice-leading motions in (b) to each dyad of the chord in (a). (In each staff of (c), the first voice leading moves each note down by three semitones, while the second moves the top note down by five semitones and the bottom note down by four semitones.) In this sense, the progression mirrors the internal structure of the initial chord. (d)–(f) show an analogous transformation applied to a different chord.

dyads in the first chord to produce a second six-note chord; to form the third six-note chord, we move the top note of each dyad down by five semitones, while moving the bottom down by four semitones.⁴⁶ This produces a hierarchical structure in which a chord progression echoes relationships within a particular chord—the sort of recursion that plays an important role in Lewin’s later work.⁴⁷ Figure 13d and 13e show another six-note chord that can be partitioned into dyads linked by pairs of voice leadings, with the corresponding chord progression shown in Figure 13f. Musically and mathematically, Figure 13 exhibits the kind of hierarchical relationship Lewin was interested in—even though the underlying space is not Lewinian.

These sorts of examples lead me to suspect that Lewinian systems do not possess striking analytical virtues that non-Lewinian systems lack.⁴⁸ I would therefore suggest that we can best honor the breadth of *GMIT*’s ambitions by relaxing some of the more stringent features of its formalism.⁴⁹ To my mind, the most important lesson of *GMIT* is that we should think carefully about the space of musical objects we would like to consider, about how to measure distance in this space, and about what sort of directed motions we would like to

⁴⁶ Technically, we are parallel-transporting the arrows representing the voice leadings along the path they define in the Möbius strip. This is a well-defined geometrical operation.

⁴⁷ See Lewin 1987 and 1990.

⁴⁸ Of course, this may reflect the limits of my imagination: Perhaps readers will be able to think of further virtues that I have simply overlooked. But even if so, the fact remains that a number of Lewin’s favorite analytical techniques can be transported to more general settings: We can often use the

tools of (generalized) set theory, or find structural analogies between “intervals” in different musical domains, or even uncover recursive patterns whereby the relations between chords mirror the intervals within a single chord. From the standpoint of practical analysis, then, the absence of “GIS structure” does not pose insurmountable, or even significant, obstacles.

⁴⁹ Hook 2007a and 2007b suggest something similar.

investigate. Where I would disagree with Lewin, perhaps, is about the wisdom of constructing overarching mathematical frameworks that promise to satisfy a broad but unspecified range of music-theoretical needs. For the danger is that, in setting up these abstract systems, we will inadvertently make assumptions that limit our ability to imagine new musical possibilities. I believe that traditional pitch-class intervals provide a cautionary example: The unquestioned assumption that intervals are *functions* has led many theorists to be skeptical about the very notion of paths in pitch-class space, and hence of straightforward claims like “the descending major third $G \rightarrow E\flat$ is the central pitch motive in the first movement of Beethoven’s Fifth Symphony.”⁵⁰ And this in turn may have delayed the development of a rigorous approach to voice leading.

Something similar, of course, might eventually be said about the ideas in this essay: It is entirely possible that twenty years from now even modern geometry will seem limited in its own way. So while I am comfortable asserting that geometry can help us develop more flexible alternatives to the Lewinian paradigm, I offer no guarantees that it provides a fully general analytical tool sufficient for all future musical needs. Rather than recommending geometry as a true-for-all-time music-theoretical metalanguage, I would instead like to conclude with a brief for theoretical particularism. Different musical problems require very different theoretical tools. Valuable music-theoretical work can be done, even without overarching frameworks to corral and regiment our thinking. It is possible that even the most comprehensive metalanguage will be limited in ways not obvious to its creator. The real danger, it seems to me, is that abstract theoretical systems will constrain our imaginations in ways that are not apparent until long after the fact.

Works Cited

- Byron, Frederick and Robert Fuller. 1992. *Mathematics of Classical and Quantum Physics*. New York: Dover.
- Callender, Clifton. 2004. “Continuous Transformations.” *Music Theory Online* 10/3. <http://mt0.societymusictheory.org/issues/mt0.04.10.3/mt0.04.10.3.callender.pdf>.
- Callender, Clifton, Ian Quinn, and Dmitri Tymoczko. 2008. “Generalized Voice Leading Spaces.” *Science* 320: 346–48.
- Crowe, Michael J. 1994. *A History of Vector Analysis*. New York: Dover.
- Douthett, Jack and Peter Steinbach. 1998. “Parsimonious Graphs: A Study in Parsimony, Contextual Transformations, and Modes of Limited Transposition.” *Journal of Music Theory* 42: 241–63.
- Eisenberg, Murray and Robert Guy. 1979. “A Proof of the Hairy Ball Theorem.” *American Mathematical Monthly* 86: 571–74.

⁵⁰ My main evidence for this claim is anecdotal: I have been arguing for the importance of paths in pitch-class space for several years now, in both talks and writing (see note 25) and have encountered great resistance—as well as many theorists who remain firmly convinced of the notion’s incoherence. By contrast, I have encountered no analogous

resistance from nontheorists, who lack preconceptions about how we should think about musical intervals. This leads me to suspect that the difficulties surrounding the notion derive from our music-theoretical preconceptions rather than being intrinsic to the idea itself.

- Gollin, Edward. 2000. "Representations of Space and Conceptions of Distance in Transformational Music Theories." Ph.D. thesis, Harvard University.
- Hall, Rachel. 2009. Review of two books by David Lewin. *Journal of the American Musicological Society* 62: 205–22.
- Hook, Julian. 2007a. "Cross-Type Transformations and the Path-Consistency Condition." *Music Theory Spectrum* 29: 1–40.
- . 2007b. "David Lewin and the Complexity of the Beautiful." *Integral* 21: 155–90.
- Klumpenhouwer, Henry. 1994. "Some Remarks on the Use of Riemann Transformations." *Music Theory Online* 0/9.
- . 2006. "In Order to Stay Asleep as Observers: The Nature and Origins of Anti-Cartesianism in David Lewin's *Generalized Musical Intervals and Transformations*." *Music Theory Spectrum* 28: 277–89.
- Lawvere, F. William and Stephen Schanuel. 1997. *Conceptual Mathematics: A First Introduction to Categories*. New York: Cambridge.
- Lewin, David. 1987. *Generalized Musical Intervals and Transformations*. New Haven: Yale University Press.
- . 1990. "Klumpenhouwer Networks and Some Isographies That Involve Them." *Music Theory Spectrum* 12: 83–120.
- . 1998. "Some Ideas about Voice-Leading between PCSets." *Journal of Music Theory* 42: 15–72.
- Mazzola, Guerino. 2002. *The Topos of Music*. Basel: Birkhauser.
- Milnor, John. 1958. "Some Consequences of a Theorem of Bott." *Annals of Mathematics, Second Series* 68: 444–49.
- Nakahara, Mikio. 2003. *Geometry, Topology, and Physics*. New York: Taylor and Francis.
- Penrose, Roger. 2004. *The Road to Reality*. New York: Knopf.
- Pynchon, Thomas. 2006. *Against the Day*. New York: Penguin.
- Robinson, Derek. 1996. *A Course in the Theory of Groups*. New York: Springer.
- Rothstein, Edward. 2003. "Connections: A Seeker of Music's Poetry in the Mathematical Realm." *New York Times*, 28 June.
- Tymoczko, Dmitri. 2005. "Voice Leadings as Generalized Key Signatures." *Music Theory Online* 11/4. <http://mto.societymusictheory.org/issues/mto.05.11.4/mto.05.11.4.tymoczko.pdf>.
- . 2006. "The Geometry of Musical Chords." *Science* 313: 72–74.
- . 2007. "Recasting K-nets." *Music Theory Online* 13/3. <http://mto.societymusictheory.org/issues/mto.07.13.3/mto.07.13.3.tymoczko.html>.
- . 2008a. "Lewin, Intervals, and Transformations: A Response to Hook." *Music Theory Spectrum* 30/1: 164–68.
- . 2008b. "Scale Theory, Serial Theory, and Voice Leading." *Music Analysis* 27: 1–49.
- . 2009. "Three Conceptions of Musical Distance." In *Mathematics and Computation in Music*, ed. Elaine Chew, Adrian Childs, and Ching-Hua Chuan, 258–73. Heidelberg: Springer.
- . 2010. *A Geometry of Music: Harmony and Counterpoint in the Extended Common Practice*. New York: Oxford University Press.
- . Forthcoming. "Dualism and the Beholder's Eye." In *Riemann Perspectives*, ed. Ed Gollin and Alex Rehdig. New York: Oxford University Press.
- Vuza, Dan Tudor. 1988. "Some Mathematical Aspects of David Lewin's Book: *Generalized Musical Intervals and Transformations*." *Perspectives of New Music* 26/1: 258–87.

Dmitri Tymoczko is a composer and music theorist who teaches at Princeton University. His book, *A Geometry of Music: Harmony and Counterpoint in the Extended Common Practice*, will be published shortly. He has just completed an album of pieces that combine jazz, rock, and classical styles. ♪