


A GEOMETRY OF MUSIC

Harmony and Counterpoint in the
Extended Common Practice

DMITRI TYMOCZKO



A Geometry of Music

OXFORD STUDIES IN MUSIC THEORY

Series Editor Richard Cohn

Studies in Music with Text, David Lewin

Music as Discourse: Semiotic Adventures in Romantic Music, Kofi Agawu

Playing with Meter: Metric Manipulations in Haydn and Mozart's Chamber Music for Strings, Danuta Mirka

Songs in Motion: Rhythm and Meter in the German Lied, Yonatan Malin

Tonality and Transformation, Steven Rings

A Geometry of Music: Harmony and Counterpoint in the Extended Common Practice,
Dmitri Tymoczko

A Geometry of Music

*Harmony and Counterpoint in the Extended
Common Practice*

DMITRI TYMOCZKO

OXFORD
UNIVERSITY PRESS

2011

OXFORD
UNIVERSITY PRESS

Oxford University Press, Inc., publishes works that further
Oxford University's objective of excellence
in research, scholarship, and education.

Oxford New York
Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in
Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Copyright © 2011 by Oxford University Press, Inc.

Published by Oxford University Press, Inc.
198 Madison Avenue, New York, New York 10016

www.oup.com

Oxford is a registered trademark of Oxford University Press

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
electronic, mechanical, photocopying, recording, or otherwise,
without the prior permission of Oxford University Press.

Library of Congress Cataloging-in-Publication Data
Tymoczko, Dmitri, 1969–

A geometry of music : harmony and counterpoint
in the extended common practice / Dmitri Tymoczko.

p. cm. — (Oxford studies in music theory)


Includes bibliographical references and index.

ISBN 978-0-19-533667-2

1. Harmony. 2. Counterpoint. 3. Musical analysis. I. Title.

MT50.T98 2010

781.2—dc22 2009046428

Oxford Web Music 

Visit the companion website at:
www.oup.com/us/ageometryofmusic

For more information on Oxford Web Music, visit www.oxfordwebmusic.com

9 8 7 6 5 4 3 2 1

Printed in the United States of America
on acid-free paper

To the memory of my father, who predicted I'd someday
get tired of rock and want to understand other
musical styles as well.

This page intentionally left blank

ACKNOWLEDGMENTS

In college, I met four people who had a profound impact on my life. Milton Babbitt gave me permission to be a composer, showing me that a serious artist could also be a rigorous thinker. Stanley Cavell opened my ears to philosophy, demonstrating that rigorous thinking could begin with scrupulous honesty. Hilary Putnam, who seemed to know everything, reawakened my interest in science and math. And Noam Elkies taught me what true intellectual and musical excellence looked like, setting standards that I am happy to try to uphold, even though I might never actually meet them.

As a graduate student, I learned how to be a professional musician from Edmund Campion, Steve Coleman, Bevan Manson, David Milnes, and David Wessel. Later, my colleagues at Princeton University provided me with a supportive home, encouraging my nonconformist and cross-disciplinary tendencies. (Paul Lansky and Steve Mackey in particular taught me not to worry about the slings and arrows of outraged forefathers.) Rick Cohn, Dan Harrison, Fred Lerdahl, and Joe Straus all nurtured my fledgling theoretical career. Rick's theoretical work was also quite influential, as it pioneered the use of geometry to model chromatic voice leading. Conversations with Clifton Callender, Ian Quinn, and Rachel Hall were crucial to developing a number of the ideas in this book, as can be seen from our various co-authored papers. I still cherish the memory of fall 2004, when Cliff, Ian, and I—and occasionally Noam Elkies—exchanged excited emails about the links between music and geometry.

Rick Cohn, John Halle, Dan Harrison, Christopher Segall, Bill Sethares, and Jason Yust all read very large portions of the manuscript, providing countless substantive and organizational suggestions. Jason's highly professional copyediting, performed on short notice, caught a number of embarrassing errors. Students in Princeton's Music 309, a mixed undergraduate/graduate theory course, read a first draft of the entire manuscript and offered many useful comments. One of these students, Andrew Jones, produced the audio examples on the book's companion website. Another, Jeff Levenberg, proofread the book and produced the index. Still more comments were provided by Kofi Agawu, Fernando Benadon, Poundie Bernstein, Elisabeth Camp, Mark Dancigers, Robert Gjerdingen, Philip Johnson-Laird, Jon Kochavi, Yuhwon Lee, Steve Rings, Dean Rosenthal, Lauren Spencer, and Dan Trueman. Students in Music 306 (a later iteration of 309) caught an infinite number of typos while the book was in page proofs.

Suzanne Ryan at Oxford University Press was a firm and relentless advocate; her calm presence helped guide this unusual book past some formidable obstacles. As series editor, Rick Cohn was a benevolent overseer from first to last. Important

financial support was provided by Princeton University, the Radcliffe Institute for Advanced Study, and the Guggenheim Foundation.

Finally, conversations with my wife, Elisabeth Camp, have provided more than a decade of intellectual stimulation, emotional support, and just plain fun. Since 2008, conversations with our son Lukas have been equally rewarding, though somewhat less relevant to the concerns of this book.

CONTENTS

About the Companion Website xv

Introduction xvii

PART I *THEORY*

CHAPTER 1 *Five Components of Tonality* 3

- 1.1 The five features 5
- 1.2 Perception and the five features 8
- 1.3 Four claims 11
 - 1.3.1 Harmony and counterpoint constrain one another 12
 - 1.3.2 Scale, macroharmony, and centricity are independent 15
 - 1.3.3 Modulation involves voice leading 17
 - 1.3.4 Music can be understood geometrically 19
- 1.4 Music, magic, and language 22
- 1.5 Outline of the book, and a suggestion for impatient readers 26

CHAPTER 2 *Harmony and Voice Leading* 28

- 2.1 Linear pitch space 28
- 2.2 Circular pitch-class space 30
- 2.3 Transposition and inversion as distance-preserving functions 33
- 2.4 Musical objects 35
- 2.5 Voice leadings and chord progressions 41
- 2.6 Comparing voice leadings 45
 - 2.6.1 Individual and uniform transposition 45
 - 2.6.2 Individual and uniform inversion 45
- 2.7 Voice-leading size 49
- 2.8 Near identity 51
- 2.9 Harmony and counterpoint revisited 52
 - 2.9.1 Transposition 53
 - 2.9.2 Inversion 56
 - 2.9.3 Permutation 58
- 2.10 Acoustic consonance and near evenness 61

CHAPTER 3 *A Geometry of Chords* 65

- 3.1 Ordered pitch space 65
- 3.2 The Parable of the Ant 69
- 3.3 Two-note chord space 70
- 3.4 Chord progressions and voice leadings in two-note chord space 73
- 3.5 Geometry in analysis 76
- 3.6 Harmonic consistency and efficient voice leading 79
- 3.7 Pure parallel and pure contrary motion 81
- 3.8 Three-dimensional chord space 85
- 3.9 Higher dimensional chord spaces 93
- 3.10 Triads are from Mars; seventh chords are from Venus 97
- 3.11 Voice-leading lattices 103
- 3.12 Two musical geometries 112
- 3.13 Study guide 114

CHAPTER 4 *Scales* 116

- 4.1 A scale is a ruler 116
- 4.2 Scale degrees, scalar transposition, and scalar inversion 119
- 4.3 Evenness and scalar transposition 122
- 4.4 Constructing common scales 123
- 4.5 Modulation and voice leading 129
- 4.6 Voice leading between common scales 132
- 4.7 Two examples 136
- 4.8 Scalar and interscalar transposition 140
- 4.9 Interscalar transposition and voice leading 144
- 4.10 Combining interscalar and chromatic transpositions 150

CHAPTER 5 *Macroharmony and Centricity* 154

- 5.1 Macroharmony 154
- 5.2 Small-gap macroharmony 156
- 5.3 Pitch-class circulation 158
- 5.4 Modulating the rate of pitch-class circulation 161
- 5.5 Macroharmonic consistency 164
- 5.6 Centricity 169
- 5.7 Where does centricity come from? 177
- 5.8 Beyond “tonal” and “atonal” 181
 - 5.8.1 The chromatic tradition 181
 - 5.8.2 The scalar tradition 186
 - 5.8.3 Tonality space 189

PART II *HISTORY AND ANALYSIS*CHAPTER 6 *The Extended Common Practice* 195

- 6.1 Disclaimers 196
- 6.2 Two-voice medieval counterpoint 197
- 6.3 Triads and the Renaissance 200
 - 6.3.1 Harmonic consistency and the rise of triads 202
 - 6.3.2 “3 + 1” voice leading 204
 - 6.3.3 Fourth progressions and cadences 207
 - 6.3.4 Parallel perfect intervals 210
- 6.4 Functional harmony 212
- 6.5 Schumann’s Chopin 214
- 6.6 Chromaticism 217
- 6.7 Twentieth-century scalar music 220
- 6.8 The extended common practice 224

CHAPTER 7 *Functional Harmony* 226

- 7.1 The thirds-based grammar of elementary tonal harmony 226
- 7.2 Voice leading in functional tonality 231
- 7.3 Sequences 238
- 7.4 Modulation and key distance 246
- 7.5 The two lattices 252
- 7.6 A challenge from Schenker 258
 - 7.6.1 Monism 261
 - 7.6.2 Holism 263
 - 7.6.3 Pluralism 264

CHAPTER 8 *Chromaticism* 268

- 8.1 Decorative chromaticism 268
- 8.2 Generalized augmented sixths 272
- 8.3 Brahms and Schoenberg 276
- 8.4 Schubert and the major-third system 280
- 8.5 Chopin’s tesseract 284
- 8.6 The *Tristan* prelude 293
- 8.7 Alternative approaches 302
- 8.8 Conclusion 304

CHAPTER 9 *Scales in Twentieth-Century Music* 307

- 9.1 Three scalar techniques 308
- 9.2 Chord-first composition 314

9.2.1	Grieg's "Drømmesyn" ("Vision"), Op. 62 No. 5 (1895)	314
9.2.2	Debussy's "Fêtes" (1899)	316
9.2.3	Michael Nyman's "The Mood That Passes Through You" (1993)	318
9.3	Scale-first composition	322
9.3.1	Debussy's "Des pas sur la neige" (1910)	322
9.3.2	Janáček's "On an Overgrown Path," Series II, No. 1 (1908)	326
9.3.3	Shostakovich's F \sharp Minor Prelude and Fugue, Op. 87 (1950)	329
9.3.4	Reich's <i>New York Counterpoint</i> (1985)	332
9.3.5	Reich's <i>The Desert Music</i> , movement 1 (1984)	336
9.3.6	The Who's "I Can't Explain" (1965) and Bob Seger's "Turn the Page" (1973)	339
9.4	The subset technique	341
9.4.1	Grieg's "Klokkeklang" ("Bell Ringing"), Op. 54 No. 6 (1891)	341
9.4.2	"Petit airs," from Stravinsky's <i>Histoire du soldat</i> (1918)	343
9.4.3	Reich's <i>City Life</i> (1995)	345
9.4.4	The Beatles' "Help" (1965) and Stravinsky's "Dance of the Adolescents" (1913)	346
9.4.5	The Miles Davis Group's "Freedom Jazz Dance" (1966)	349
9.5	Conclusion: common scales, common techniques	349
CHAPTER 10 <i>Jazz</i> 352		
10.1	Basic jazz voicings	353
10.2	From thirds to fourths	357
10.3	Tritone substitution	360
10.4	Altered chords and scales	365
10.5	Bass and upper voice tritone substitutions	370
10.6	Polytonality, sidestepping, and "playing out"	374
10.7	Bill Evans' "Oleo"	378
10.7.1	Chorus 1	378
10.7.2	Chorus 2	381
10.7.3	Chorus 3	382
10.7.4	Chorus 4	384
10.8	Jazz as modernist synthesis	387
	<i>Conclusion</i>	391

APPENDIX A	<i>Measuring Voice-Leading Size</i>	397
APPENDIX B	<i>Chord Geometry: A More Technical Look</i>	401
APPENDIX C	<i>Discrete Voice-Leading Lattices</i>	412
APPENDIX D	<i>The Interscalar Interval Matrix</i>	418
APPENDIX E	<i>Scale, Macroharmony, and Lerdahl's "Basic Space"</i>	424
APPENDIX F	<i>Some Study Questions, Problems, and Activities</i>	428
	<i>References</i>	435
	<i>Index</i>	445

This page intentionally left blank

ABOUT THE COMPANION WEBSITE

www.oup.com/us/ageometryofmusic

The website contains audio files corresponding to each of the musical examples, as well as a few additional soundfiles and movies that further illustrate points in the text. It is primarily intended to assist readers who have trouble playing or imagining the notated examples.

This page intentionally left blank

INTRODUCTION

When I was about 15 years old, I decided I wanted to be a composer, rather than a physicist or mathematician. I had recently switched from classical piano to electric guitar, and although I exhibited no obvious signs of compositional talent, I was fascinated by the amazing variety of twentieth-century music: the suavely ferocious *Rite of Spring*, which made tubas sound cool; the encyclopedic *Sgt. Pepper*, which contained multitudes; the hypnotic repetitions of Philip Glass, whose spirit seemed also to infuse the music of Brian Eno and Robert Fripp; and the geeky sophistication of art rock and new wave. I was aware of but intimidated by jazz, which seemed to be perpetually beyond reach, like the gold at the end of a rainbow. (Told by my guitar teacher that certain chords or scales were jazzy, I would inevitably find that they sounded flat and lifeless in my hands.) To a kid growing up in a college town in the 1980s, the musical world seemed wide open: you could play *The Rite of Spring* with your rock band, write symphonies for electric guitars, or do anything else you might imagine.

I was somewhat surprised, therefore, to find that my college teachers—famous academics and composers—inhabited an entirely different musical universe. They knew nothing about, and cared little for, the music I had grown up with. Instead, their world revolved around the dissonant, cerebral music of Arnold Schoenberg and his followers. As I quickly learned, in this environment not everything was possible: tonality was considered passé and “unserious”; electric guitars and saxophones were not to be mixed with violins and pianos; and success was judged by criteria I could not immediately fathom. Music, it seemed, was not so much to be composed as constructed—sembled painstakingly, note by note, according to complicated artificial systems. Questions like “does this chord sound good?” or “does this compositional system produce likeable music?” were frowned upon as naive or even incoherent.

I studied many things in college: seventeenth-century masses, eighteenth-century Lutheran chorales, and twentieth-century avant-garde music. I learned about Heinrich Schenker, who purported to reduce all good tonal pieces to a small number of basic templates. I absorbed mathematical tools for constructing and deconstructing atonal compositions. But I did not learn anything whatsoever about jazz, Debussy, Ravel, Shostakovich, Messiaen, or minimalism. In fact, even the music of Wagner and Chopin was treated with a certain embarrassment—acknowledged to be important, but deemed suspiciously illogical in its construction. Looking back, I can see that the music I encountered was the music my teachers knew how to talk about. Unfortunately, this was not the music I had come to college wanting to understand.

Twenty years later, things are different, and a number of the barriers between musical styles have fallen. Many composers have returned to the tonal ideas that my own

teachers deemed irrelevant. Electric guitars now mix freely with violins, and everything is indeed permitted. But despite this new freedom, tonality remains poorly understood. We lack even the most rudimentary sense of the musical ingredients that contribute to the sense of “tonalness.” The chromatic music of the late nineteenth century continues to be shrouded in mystery. We have no systematic vocabulary for discussing Debussy’s early 20th-century music or its relation to subsequent styles. Graduate students in music often know nothing about bebop, or about how this language relates backward to classical music and forward to contemporary concert music. As a result, many young musicians are essentially flying by the seat of their pants, rediscovering for themselves the basic techniques of modern tonal composition.

The goal of this book is to understand tonality afresh—to provide some new theoretical tools for thinking about tonal coherence, and to illuminate some of the hidden roads connecting modern tonality to that of the past. My aim is to retell the history of Western music so that twentieth-century tonality appears not as an aberration, the atavistic remnant of an exhausted tradition, but as a vital continuation of what came before. I hope that this effort will be useful to composers who want new ways to write tonal pieces, as well as to theorists, performers, and analysts who are looking for new ways to think about existing music.

While my primary audience consists of composers and music theorists, I have tried to write in a way that is accessible to students and dedicated amateurs: technical terms are explained along the way, and only a basic familiarity with elementary music theory (including Roman numeral analysis) is presumed. (Specialists will therefore need to endure the occasional review of music-theoretical basics, particularly in the early chapters.) More than anything else, I have attempted to write the sort of book I wish had existed back when I first began to study music. It would make me happy to think that these ideas will be helpful to some young musician, brimming with excitement over the world of musical possibilities, eager to understand how classical music, jazz, and rock all fit together—and raring to make some new contribution to musical culture.

PART I

Theory

This page intentionally left blank

CHAPTER 1

Five Components of Tonality

The word “tonal” is contested territory. Some writers use it restrictively, to describe only the Western art music of the eighteenth and nineteenth centuries. For them, more recent music is “post-tonal”—a catch-all term including everything from Arvo Pärt’s consonances to the organized sonic assaults of Varèse and Xenakis. This way of categorizing music makes it seem as if Pärt, Varèse, and Xenakis are clearly and obviously of a kind, resembling one another more than any of them resembles earlier composers.

“Tonal” can also be used expansively. Here, the term describes not just eighteenth- and nineteenth-century Western art music, but rock, folk, jazz, impressionism, minimalism, medieval and Renaissance music, and a good deal of non-Western music besides. “Tonality” in this sense is almost synonymous with “non-atonality”—a double negative, most naturally understood in contrast to music that was deliberately written to contrast with it. The expansive usage accords with the intuition that Schubert, the Beatles, and Pärt share musical preoccupations that are not shared by composers such as Varèse, Xenakis, and Cage. But it also raises awkward questions. What musical feature or features lead us to consider works to be tonal? Is “tonality” a single property, or does it have several components? And how does tonality manifest itself across the broad spectrum of Western and non-Western styles? Faced with these questions, contemporary music theory stares at its feet in awkward silence.¹

The purpose of this book is to provide general categories for discussing music that is neither classically tonal nor completely atonal. This, in my view, includes some of the most fascinating music of the twentieth century, from impressionism to postminimalism. It also includes some of the most mysterious and alluring music of Chopin, Liszt, and Wagner—music that is as beloved by audiences as it is recalcitrant to analytical scrutiny. My goal is to try to develop a set of theoretical tools that will help us think about these sophisticated tonal styles, which are in some ways freer and less rule-bound than either eighteenth-century classical music or twentieth-century modernism.

1 Fétis (1840/1994), who popularized the term “tonality” in the early nineteenth century, was one of the first scholars to try to provide a general account of the phenomenon. (For more on the early history of the term “tonality” see Simms 1975 and Dahlhaus 1990.) More recently, Joseph Yasser (1975), Richard Norton (1984), William Thomson (1999), Brian Hyer (2002), Carol Krumhansl (2004), and Matthew Brown (2005) have offered different perspectives on the subject.

More specifically, I will argue that five features are present in a wide range of genres, Western and non-Western, past and present, and that they jointly contribute to a sense of tonality:

1. *Conjunct melodic motion*. Melodies tend to move by short distances from note to note.
2. *Acoustic consonance*. Consonant harmonies are preferred to dissonant harmonies, and tend to be used at points of musical stability.
3. *Harmonic consistency*. The harmonies in a passage of music, whatever they may be, tend to be structurally similar to one another.
4. *Limited macroharmony*. I use the term “macroharmony” to refer to the total collection of notes heard over moderate spans of musical time. Tonal music tends to use relatively small macroharmonies, often involving five to eight notes.
5. *Centricity*. Over moderate spans of musical time, one note is heard as being more prominent than the others, appearing more frequently and serving as a goal of musical motion.

The aim of this book is to investigate the ways composers can use these five features to produce interesting musical effects. This project has empirical, theoretical, and historical components. Empirically, we might ask how each of the five features contributes to listeners’ perceptions of tonality: which is the most influential, and are there any interesting interactions between them? For instance, is harmonic consistency more important in the context of some scales than others? Theoretically, we might ask how the various features can *in principle* be combined. Is it the case, for example, that diatonic music necessarily involves a tonic? Conversely, is chromatic music necessarily non-centric? Finally, we can ask historical questions about how different Western styles have combined these five tonal ingredients—treating the features as determining a space of possibilities, and investigating the ways composers have explored that space.

This book is primarily concerned with the theoretical and historical questions. I am a musician, not a scientist, and although I will sometimes touch on perceptual issues, I will largely leave empirical psychology to the professionals. Instead, I will ask how composers *have combined* and *might combine* the five features. Part I of the book develops theoretical tools for thinking about the five features. Part II uses these tools to argue for a broader, more continuous conception of the Western musical tradition. Rather than focusing narrowly on the eighteenth and nineteenth centuries (the so-called “common practice period”), I attempt to identify an “extended common practice” stretching from the beginning of Western counterpoint to the music of recent decades. The point is to retell the history of Western music in such a way that the tonal styles of the last century—including jazz, rock, and minimalism—emerge as vibrant and interesting successors to the tonal music of earlier periods.

My central conclusion is that the five features impose much stronger constraints than we would intuitively expect. For example, if we want to combine conjunct melodic motion and harmonic consistency, then we have only a few options, the most important of which involve acoustically consonant sonorities. And if we want to

combine harmonic consistency with limited macroharmony, then we are led to a collection of very familiar scales and modes. Thus the materials of tonal music are, in Richard Cohn’s apt description, “overdetermined,” in the sense that they are special or distinctive for multiple different reasons.² This suggests that when we look closely, we should find important similarities between different tonal styles: since there are only a few ways to combine the five features, different composers—from before Palestrina to after Bill Evans—will necessarily use the same basic techniques. In the second part of this book I make good on this claim, tracing common practices that connect the earliest examples of Western counterpoint to music of the very recent past.

1.1 THE FIVE FEATURES

Let’s consider the five features in more detail, with an eye toward understanding why they might be so widespread throughout Western and non-Western music. A preference for *conjunct melodic motion* likely derives from the features of the auditory system that create a three-dimensional “auditory scene.”³ An eardrum, in effect, is a one-dimensional system that can only move back and forth. From this meager input our brains create a vivid three-dimensional sonic space consisting of individually localized sounds: the phone ringing in front of you, the honk of a car horn outside your window, and the sound of a droning music theorist off to your right. To accomplish this dazzling transfiguration, the brain relies on a number of computational tricks, one of which is to group sonic events that are nearby in pitch.⁴ Thus, a sequence like Figure 1.1.1a tends to be heard as belonging to a single sound source, whereas Figure 1.1.1b creates the impression of multiple sources. In this sense, small melodic steps are intrinsic to the very notion of “melody.”

Acoustic consonance, or intrinsic sonic restfulness, is another very widespread musical feature.⁵ Many styles make heavy use of consonant intervals such as the octave and perfect fifth, assigning them privileged melodic and harmonic roles. Scales containing a large number of consonant intervals are found in seemingly independent musical cultures, and there is evidence from infant psychology that the preference for consonance is innate.⁶ At present, however, we do not know

Figure 1.1.1 Small movements sound melodic (a), while large registral leaps create the impression of multiple melodies (b).



2 See Cohn 1997.

3 See Bregman 1990, Narmour 1990, and Vos and Troost 1989. Huron 2007 contains data about statistical properties of Western melodies, including conjunct melodic motion.

4 Wessel (1979) suggests that the relevant variable might be the “spectral centroid,” which in normal listening circumstances is highly correlated with pitch.

5 Izumi (2000) and Hulse et al. (1995) indicate that nonhuman animals such as monkeys and birds can distinguish consonance from dissonance.

6 Crowder, Reznick, and Rosenkrantz 1991, Zentner and Kagan 1996 and 1998, Trainor and Heinmiller 1998, Trainor, Tsang and Cheung 2002, and McDermott and Hauser 2005.

for certain how universal or innate this preference is. Fortunately, this issue is largely irrelevant in the present context: what matters is just that many listeners, both Western and non-Western, do have a fairly deep-seated preference for consonant sonorities.

Slightly more general than acoustic consonance is *harmonic consistency*, or the use of sonorities that resemble one another. For example, Figure 1.1.2a features a series of major and minor chords, all audibly similar. Their resemblance gives the passage a kind of smoothness, and we experience the chords as belonging together. Likewise, Figure 1.1.2b uses a series of very dissonant chromatic clusters that also seem to belong together. By contrast, Figure 1.1.2c uses very different-sounding harmonies, switching aimlessly between different sonic worlds. This sort of harmonic incongruity is quite unusual in Western music, and often provokes spontaneous laughter—suggesting that the expectation of harmonic consistency is very strong, even though it is rarely discussed.

Figure 1.1.2 Harmonic consistency using consonant sonorities (a) and dissonant sonorities (b). Sequence (c) does not exhibit harmonic consistency.



In most Western and non-Western music, pitches are drawn from a relatively small reservoir of available notes—typically, between five and eight.⁷ As a result, Western music has a two-tiered harmonic consistency: at the local (or instantaneous) level, a passage like Figure 1.1.3 presents a series of major and minor chords, which are audibly related, while over larger time spans it articulates a scale by using only seven different notes. The scale can thus be considered a kind of “large” or *macro* harmony that subsumes the individual chords.⁸ Even though there is no one instant at which this larger harmony is presented, it nevertheless has a significant effect on our listening experience: scale-based melodies are easier to remember than nonscalar melodies, and notes outside the scale sound more pungent than notes in the scale. As we will see, macroharmonies can be relatively consonant, like the diatonic or pentatonic scale, or relatively dissonant, like the chromatic scale. They can also participate in larger-level voice leadings analogous to those connecting individual chords.

Finally, we often hear some pitches or notes as being more important (or “central”) than others. These pitches tend to serve as points of musical arrival, to which others are heard as “leading” or “tending.” Thus, one and the same sequence of notes—such as that in Figure 1.1.4—can be heard in a variety of ways. If we hear it in a musical

⁷ Burns 1999 and Dowling and Harwood 1986. The range “five to eight” recalls familiar facts about the limitations on human short-term memory (see Miller 1956 on “seven plus or minus two”).

⁸ Thanks to Ciro Scotto for suggesting this term.

Figure 1.1.3 Major and minor triads belonging to the same diatonic scale.



Figure 1.1.4 A single melody will sound different in different harmonic contexts.



context where C is the most stable pitch, then it sounds like a beginning that is in need of some sort of continuation, ending with a comma rather than a period. But in a context where F is stable, it sounds more complete, as if it ends with a period or exclamation point. Centricity is again a very widespread feature of human music, appearing in a large number of seemingly unrelated musical cultures. However, different styles can emphasize the tonic note to different extents: as Harold Powers once noted, there is a much stronger feeling of centricity in Indian music than in Renaissance polyphony.⁹

Of these five features, harmonic consistency is clearly the most culturally specific. The idea that music should consist of rapidly changing chords is a deeply Western idea, in a double sense: it is deep, insofar as it characterizes much Western music since before the Renaissance; and it is Western, since there are many cultures in which the notion of a “chord progression” simply plays no role. (Many traditional non-Western styles are purely monophonic, or feature an unchanging “drone” harmony; however, there are now a large number of syncretistic styles that combine Western harmonies with non-Western melodic and rhythmic ideas.) Acoustic consonance is also somewhat culture-specific: although many cultures make some use of consonant intervals, and although some have recognizably Western conceptions of consonance, other non-Western styles sound quite dissonant to Western ears. By contrast, the three other features—conjunct melodic motion, limited macroharmony, and centricity—are common to virtually all human music. This near universality may be attributable, at least in part, to features of our biological inheritance.¹⁰

Now for an important disclaimer. While I think that typical Western listeners *prefer* music that exemplifies the five features, I do not mean to suggest that such music is intrinsically *better* than any other kind of music. “Tonal,” for me, is not synonymous with “good.” (Nor is “popular,” for that matter: there is plenty of unpopular, nontonal music that I happen to like, from Nancarrow to Xenakis to Ligeti.) In particular, I have no interest in arguing that atonal composers are misguided, fighting against biology, or anything of the sort. Instead, the purpose of this book is an affirmative one: to develop new theoretical tools for thinking about tonality, and to provide new insights into the relations between various musical styles. My hope is that this investigation will be useful to composers and theorists of all varieties, as even the advocates of atonality will stand to gain from a deeper understanding of that which they are trying to avoid.

⁹ Powers 1958.

¹⁰ Dowling and Harwood 1986, Narmour 1990.

1.2 PERCEPTION AND THE FIVE FEATURES

This book is primarily concerned with what composers do, rather than what listeners hear: the goal is to describe various ways in which the five features have been or might be combined. But it is not possible to avoid perceptual issues altogether. After all, readers have a right to wonder whether my five features do indeed contribute to the experience of tonality, and if so, whether they are the only factors that contribute.

The first question can be easily answered: all one needs to do is constrain randomly generated notes according to each of the five features. Insofar as the constraints cause random music to sound increasingly tonal, or at least ordered, then I am right about their psychological importance. Furthermore, the same experiment can be used not just to show *that* these features have important psychological effects, but also *what particular psychological effects they have*. This is useful because familiar styles tend either to combine many of the features, and hence be fully tonal in a traditional sense, or else to abandon most of them à la radical atonality. Consequently, existing musical works do not always help us to understand the specific contributions made by each of our five components individually.

I strongly urge you to try this experiment for yourself: the results are not subtle, and they demonstrate the powerful psychological effects that can be obtained with simple musical means. (In particular, I encourage you to use the book's companion website, which contains a number of illustrative examples.¹¹) Unfortunately, this is a case where a musical experience is worth a thousand words: no amount of merely verbal description on my part will substitute for your own investigation. Nevertheless, it is worth trying to describe these effects, if only to persuade you to actually perform the experiment. Figure 1.2.1a presents a series of completely random three-note chords, with pitches chosen arbitrarily from the range C2 to C7. It provides a baseline against which subsequent examples can be judged. (Some people, myself included, find this sort of random texture to be oddly appealing.) Figure 1.2.1b constrains the randomness by requiring that the notes move just a few semitones from chord to chord. What results is very rudimentary sort of counterpoint, consisting of three independent melodic strands. Although it is considerably less random-sounding than the pointillistic texture of (a), the harmonic structure of the sequence still sounds somewhat indistinct, providing the ear with relatively little to grab onto. Figure 1.2.1c combines conjunct melodic motion with harmonic consistency, requiring that each chord be a "stack of fourths." The melodic lines, rather than wandering aimlessly, now seem to create chords with a distinctive harmonic identity, which in turn gives the passage a feeling of increased consistency.¹² (This example might seem reminiscent of some Stravinsky or Hindemith.) Finally,

¹¹ www.oup.com/us/ageometryofmusic

¹² With a little practice it is easy to distinguish harmonically consistent sequences from those involving random, unrelated chords. However, it should be said that some chords (such as the major triad or stack of fourths) have a very distinctive sound, while others (such as C-C#-D#) are more generic. It takes more compositional work to create a palpable sense of harmonic consistency using these generic chords.

Figure 1.2.1d restricts the chords to the same diatonic scale. To complete the transition from utter randomness to something recognizably tonal, I have replaced the “fourth chords” of (c) with more consonant diatonic triads. Although the result will not win any composition prizes, it does demonstrate that a kind of rudimentary tonalness is in fact generated by my five features. Indeed, the differences between Figures 1.2.1a–d are striking and unmistakable, even for a layperson with no specialized musical training.

Informal experiments like these suggest that, for typical listeners, the five features play an important role in determining the tonalness (or perhaps “orderedness”) of musical stimuli. Furthermore, I strongly suspect that for many listeners, “tonalness,” “orderedness,” and “pleasantness” are correlated: all else being equal, music displaying many of the five features will be preferred to music that does not.¹³ This constitutes the testable psychological theory lurking at the core of this book. Personally, I think it would be interesting to try to determine not just *that* the five features have important psychological effects, but also their relative importance. What makes a larger difference to listeners’ perceptions, harmonic consistency, acoustic consonance, or conjunct melodic motion? We lack even the most rudimentary data that would allow

Figure 1.2.1 Four randomly generated sequences. Sequence (a) is completely random; (b) exhibits efficient voice leading; (c) exhibits harmonic consistency and efficient voice leading; (d) exhibits harmonic consistency, efficient voice leading, and limited macroharmony.

The figure displays four musical sequences, labeled (a) through (d), in piano notation. Each sequence consists of two staves, a treble clef staff on top and a bass clef staff on the bottom. Sequence (a) shows a completely random sequence of notes and chords. Sequence (b) shows efficient voice leading, with notes moving smoothly between staves. Sequence (c) shows harmonic consistency and efficient voice leading, with notes moving smoothly and chords being consistent. Sequence (d) shows harmonic consistency, efficient voice leading, and limited macroharmony, with notes moving smoothly and chords being consistent, and a limited range of macroharmony.

¹³ Some preliminary studies, conducted by John Muniz, Cynthia Weaver, and Asher Yamplosky (graduate students at Yale University), suggest that conjunct melodic motion may contribute more to “orderedness” than “pleasantness.” Disentangling these issues is a subject for future research.

us to answer this question. One of my hopes is that some psychologist readers will be motivated to undertake the obvious experiments suggested by the ideas I will be discussing.

Figure 1.2.1 shows that the five features *can* contribute fairly dramatically to the sense of tonality. But are they *necessary* for creating tonal effects? And if so, are they the only such features or are there others?

In some ways, I think the question is misguided. The point here is not to police the use of the word “tonality” by setting strict limits on what may or may not be described with the term, but rather to replace the crude opposition “tonal/atonal” with a more nuanced set of distinctions. You can decide for yourself whether to use “tonal” to describe (say) diatonic music without a tonal center, or chromatic music with a strong sense of harmonic consistency; my job is just to show that the term “tonality” typically applies to music that exhibits my five basic features. I should also point out that I am neglecting important issues such as rhythm, motivic variation, timbre, form, performance, and rubato, all of which can contribute to the sense of tonality. (My goal is not to provide a complete theory of all music, but rather to discuss a few general features whose musical importance is for the most part unquestioned.) That said, however, I confess that it is difficult for me to imagine that I would ever want to use the term “tonal” to describe music in which acoustic consonance plays no role, in which there is no conjunct melodic motion or harmonic consistency, in which no tone is heard as central, and which does not limit itself to a relatively small number of pitch classes over short stretches of time. In this sense, it seems that at least some of the five features are necessary for tonality, at least as I personally understand it.

Throughout the twentieth century, composers devised new musical languages—some idiosyncratic, some very widely used—intended to replace traditional tonality. It is instructive to subject these alternative systems to the experimental test described above: that is, to constrain random musical notes according to their basic principles, and to listen for the perceptual differences that result.¹⁴ To that end, Figure 1.2.2 uses “constrained randomness” to investigate one of the most prominent alternatives to tonality, the twelve-tone system.¹⁵ On a first (or even tenth) listening, I do not find the sequences in Figure 1.2.2 to be dramatically different from randomly generated sequences such as Figure 1.2.1b. Indeed, if someone were to present the three passages in a psychology experiment, I doubt I would notice that one was random while the other two were not. This is not to say that twelve-tone structure makes *no* audible difference, or that one cannot learn to hear the coherence in Figure 1.2.2, but rather

14 It is exceedingly difficult to judge the effects of a musical syntax if we only encounter it in the context of complete compositions, since compositional skill may mask the contributions of the syntax itself.

15 Figure 1.2.2a uses three successive rows to generate a sequence of 12 three-note chords. Figure 1.2.2b sets the row in counterpoint against its inversions and retrograde-inversions forms, producing a series of 12 three-note chords. I have borrowed the 12-tone row from Schoenberg’s first consistently twelve-tone piece, his Op. 25 Suite for piano.



Figure 1.2.2
Randomly
generated twelve-
tone music.

that the psychological effects here are relatively subtle—considerably less dramatic than those produced by our five features.

Does this show that twelve-tone music is aesthetically problematic? Or that the twentieth-century quest for alternatives to traditional tonality is fruitless? Not at all.¹⁶ But it does suggest that twelve-tone rows produce less powerful psychological consequences than harmonic consistency, conjunct melodic motion, acoustic consonance, macroharmony, and centrality.¹⁷ And while twelve-tone music is just one twentieth-century musical system, similar comments might be made about other approaches. To my mind, this suggests that the five features are unusually powerful tools for creating musical coherence. To say this is not to deny that alternative tools may in principle exist, but simply to reiterate the basic point that tonality constitutes a fairly unique solution to some elementary compositional problems. If this is right, then the task of providing an alternative to tonality is much more difficult than one might intuitively have imagined.

1.3 FOUR CLAIMS

The argument of this book revolves around four basic claims, each of which concerns ways in which the five features can interact with or constrain one another. In this section I'll briefly outline these claims as a way of foreshadowing some of my central preoccupations.

¹⁶ First, it is possible that there are gifted listeners who respond strongly and immediately to the non-random features of the sequences in Figure 1.2.2. Second, it is possible that with extensive training ordinary listeners can sensitize themselves to the sequences' structure—as when one gradually starts to discern the details in an all-gray painting. Third, it is possible that there are specific compositional techniques that can make twelve-tone structure psychologically transparent. Fourth, it is possible that some listeners simply enjoy random pitch structures. And fifth, it is possible that twelve-tone music can be attractive in ways that make up for any potential absence of interesting pitch structure.


¹⁷ Schoenberg (1975, p. 215) observed that “composing [using twelve-tone techniques] does not become easier, but rather ten times more difficult” (see also Dubiel 1997).

1.3.1 Harmony and Counterpoint Constrain One Another

Imagine a composer, Lyrico, who would like to combine conjunct melodic motion with harmonic consistency; that is, he would like to write melodies that move by short distances while using harmonies that are structurally similar to one another. Intuitively, it might seem that there are innumerable ways to satisfy these two constraints—that there is an entire universe of syntaxes consistent with these fundamental principles. But in fact, there are just a few ways in which they can be combined.

Consider the simplest possible situation. Suppose Lyrico decides to combine an unchanging “drone” harmony with a moving melodic voice. Harmonic consistency is thus obtained trivially: the chords in the passage will all be similar because *there is only one chord*. Let us further imagine that Lyrico chooses to use a C major chord as his drone. If he were to confine the melody to the notes of this chord, the result would be a series of unmelodic leaps (Figure 1.3.1a). This is because its notes are all reasonably far apart. To obtain conjunct melodic motion, Lyrico can therefore introduce “passing tones” that connect the chord tones by short melodic steps. The result, shown in Figure 1.3.1b, is a *scale* covering an entire octave, in which successive notes are connected by relatively small distances, and in which chord tones alternate with nonchordal “passing tones.” Lyrico can now write melodies that move freely along this scale, alternating between stable and unstable notes.

Figure 1.3.1 (a) Confining a melody to the notes of the C major chord produces large leaps, so it is necessary to add “passing tones” (b).



But suppose Lyrico wakes up one day in a more ornery mood and decides to use the dissonant chromatic cluster {B, C, D \flat } as his drone. Here the compositional situation is reversed: where the major chord is consonant, and has its notes spread relatively far apart, this chord is very dissonant and has all its notes close together. Consequently, by confining himself to the

notes of the chord, Lyrico can obtain conjunct melodies, as in Figure 1.3.2. However, it takes a large number of passing tones to connect the D \flat in one octave to the B in the

Figure 1.3.2 (a) Confining a melody to the notes of the cluster {B, C, D \flat } produces conjunct melodic motion, but changing octaves requires a large number of passing tones (b).



next. Since the resulting scale does not exhibit a regular alternation of stable “chord tones” and unstable “passing tones,” it is difficult to hear the nonharmonic tones as connective devices that simply decorate an underlying harmony.

These two examples suggest a general moral: *harmony and melody constrain one another*. Different types of chords suggest different musical uses. In particular there is a fundamental difference between chords like {C, E, G}, whose notes are all far away from each other, and chords like {B, C, D \flat }, whose notes are clustered close together. When a chord’s notes are clustered close together, it is possible to create conjunct melodies that use only the chord’s notes, but it is not possible to create scales that have a regular alternation between chord tones and passing tones. When a chord’s notes are relatively spread out, it is not possible to create conjunct melodies by using only the notes of the chord, but it *is* possible to create scales with a nice arrangement of chord and non-chord tones.

Now consider a more sophisticated problem. Suppose Lyrico decides to write a C major chord followed by an F major chord, its transposition by ascending perfect fourth. As shown in Figure 1.3.3, *every* note in the C major chord is near some note in an F major chord: C is common to both, E is one semitone from F, and G is two semitones from both F and A. This means that Lyrico can write a sequence of C and F major chords that articulates *three separate melodies*, each moving by small distances. This is *counterpoint*—a group of simultaneous melodies, or *voices*, articulating mappings, or *voice leadings*, between successive chords. For example, the first voice leading in Figure 1.3.3b maps G to A, E to F, and C to C. This voice leading is *efficient* because all the voices move by short distances. Clearly, efficient voice leading is simply conjunct melodic motion in all parts of a contrapuntal texture.

As it happens, the major chord is particularly well suited for contrapuntal music. Figure 1.3.4 shows that any two major chords can be connected by stepwise voice leading in which no voice moves by more than two semitones. This means that Lyrico can write a harmonic progression *without worrying about melody*; that is, for any sequence

Figure 1.3.3 (a) Every note of the C major triad is near some note of the F major triad. (b) It is possible to use a series of C and F major triads to construct three simultaneous melodies.




Figure 1.3.4 Any two major triads can be linked by stepwise voice leading; in the case of the tritone, this requires four voices.

of major triads, there is always some way to connect the notes so as to form step-wise melodies. Conversely, Lyrico can write any melody whatsoever without worrying about harmony, as there will always be some way to harmonize it with a sequence of efficient voice leadings between major chords.


But what if Lyrico writes the chromatic cluster $\{B, C, D\flat\}$ followed by $\{E, F, G\flat\}$, its transposition by ascending fourth? Here, none of the notes of the first chord are within two semitones of any note in the second, and hence there is no way to combine a sequence of these chords so as to produce conjunct melodies (Figure 1.3.5). At the same time, however, the chromatic cluster can do things that the C major chord can't: Figure 1.3.6 shows that it is possible to write contrapuntal music in which individual melodic lines move by short distances *within* a single, unchanging harmony. Clearly, this is possible only because the chord's notes are all clustered together, ensuring that there is always a short path between any two of them. The resulting music produces a feeling of burbling within stasis, a kind of melodic activity within overall harmonic restfulness.¹⁸

Once again, we see that different kinds of chords are useful for different purposes. Chords that divide the octave very unevenly, such as $\{B, C, D\flat\}$, are ideally suited for static music in which harmonies do not change. By contrast, chords that divide the octave relatively evenly, like C major, can be connected to their transpositions

Figure 1.3.5
Chromatic clusters cannot always be linked by efficient voice leading.



Figure 1.3.6
Chromatic clusters allow conjunct melodic motion to be combined with harmonic stasis.



by efficient voice leading, and are therefore suited for contrapuntal music in which harmonies change quickly. Remarkably, the most nearly even chords in twelve-tone equal temperament are the familiar chords of Western music: perfect fifths, triads, seventh chords, ninth chords, and familiar scales. Besides being well suited for contrapuntal music, these chords are all *acoustically consonant*, or restful and stable-sounding. This, then, is an example of Cohn's "overdetermination," a situation in which a familiar musical object is remarkable for multiple reasons: nearly even chords are interesting not just because they permit the combination of harmonic consistency and conjunct melodic motion, but also because they can be acoustically consonant.¹⁹

More generally, it would seem that composers who wish to combine harmonic consistency and conjunct melodic motion have relatively few options: they can either use familiar sonorities in more or less familiar ways, or they can use chromatic chords whose notes are clustered

¹⁸ This sort of texture was popularized during the 1960s by composers such as Ligeti and Lutosławski; precursors to the technique include Bartók's *Out of Doors* ("Night Music") and Ruth Crawford Seeger's String Quartet.

¹⁹ This observation has been made by Agmon (1991) and Cohn (1996), in the special case of the triad. Both writers explained the specialness of the triad in terms of the way it is embedded in a larger collection. In the following chapters, I provide an account that applies to consonant chords more generally and that does not presuppose embedding into a larger scale.

together. The next few chapters will elaborate on this point, using the notion of *near symmetry* to explain exactly how chord structure constrains contrapuntal function. Later chapters will trace the practical consequences of this interdependence, considering a range of superficially different styles, from Renaissance polyphony to contemporary jazz, and showing that they all utilize fundamentally similar procedures. Our theoretical work will demonstrate that these similarities are not simply the byproduct of historical influence, but also of the more fundamental ways in which the five features constrain one another. In other words, they testify to the fact that there are only a few ways to combine harmonic consistency and stepwise melodies.

1.3.2 Scale, Macroharmony, and Centricity are Independent

My second claim is that we need to distinguish the closely related phenomena of scale, macroharmony, and centricity. A *scale*, as I use the term, is a means of measuring musical distance—a kind of musical ruler whose unit is the “scale step.” Relative to the C diatonic scale, the notes E and G are two scale steps apart, since there is precisely one white note between them (Figure 1.3.7).²⁰ These same notes are one step apart relative to the pentatonic scale C-D-E-G-A and three steps apart relative to the chromatic scale. Similarly, C and E are two steps apart relative to the diatonic and pentatonic scales, and four apart relative to the chromatic scale. The three scales therefore give us three different ways of measuring musical distance, and three different estimates of the relative sizes of the intervals C-E and E-G. In principle, we should not ask whether the intervals C-E and E-G are “the same size” unless we also specify a particular musical “ruler.” As we will see, the richness of tonal music lies partly in the way it exploits these various conceptions of musical distance.

If a scale is a musical ruler, then a *macroharmony* is the total collection of notes used over small stretches of musical time. Typically, macroharmonies are also scales: a composer might (for example) use only the white notes on the piano keyboard, while also exploiting the unit of distance defined by adjacent notes (the “scale step”). But in




Figure 1.3.7 Scale steps provide a means of measuring musical distance. The intervals C-E and E-G are two steps large relative to the diatonic scale (a), two and one steps large relative to the pentatonic scale (b), and four and three steps large relative to the chromatic scale (c).

²⁰ Here and elsewhere, I use the term “C diatonic scale” to refer to the notes of the C major scale (i.e. the white notes) without suggesting that the note C is special in any way. I use the term “C major” (or “C ionian”) in contexts where C is a tonal center.

Figure 1.3.8
The music makes use of two scales to create a chromatic macroharmony.



principle it is possible to separate the two phenomena. Figure 1.3.8 shows a passage of “polytonal” music in which the upper staff moves systematically along the C diatonic scale, while the lower staff moves along G \flat pentatonic.

To explain how this music works, we need to postulate two different scales, one for each staff. (Of course it would be possible to combine the notes in both staves into a single chromatic scale, but the resulting collection would be useless for explaining why the individual voices move as they do.) Here, then, is a passage of music that uses pentatonic and diatonic scales to create a chromatic macroharmony. The concept of “scale” allows us to describe the structure within each voice, while the concept “macroharmony” allows us to describe the global harmony they produce.

A second fundamental distinction is between *macroharmony* and *centricity*. “Centricity” refers to the phenomenon whereby a particular pitch is felt as being more stable or important than the others. (In traditional tonality this is the tonic note; in modal theory, it is the “final.”) Chapters 4 and 5 will suggest that macroharmony and centricity are completely independent: it is entirely possible, for example, to write diatonic music in which no note is heard as a tonal center, just as one can write chromatic music with a very clear center. There is, however, a strong historical association between the two, with diatonic music often being centric and chromatic music centerless. The explanation for this connection lies in the fact that some prominent musicians believed that the diatonic scale had a unique “natural” tonic, and hence that a centerless chromaticism was the main alternative to traditional tonality. Other musicians, believing that centricity and macroharmony were independent, felt free to explore a much wider range of scales and modes. Chapter 5 will thus propose that the cleavage between the “scalar” and “chromatic” traditions, exemplified by composers such as Debussy and Schoenberg, was exacerbated by a fundamental disagreement about the relationship between macroharmony and centricity.

Together, scale, macroharmony, and centricity are the three principal components of what I think of as the “general theory of keys”—a set of tools for describing music that is tonal in the broad sense, even though it may not conform to the specific conventions of eighteenth-century tonality. One of this book’s goals is to consider various sorts of “generalized keys” in Western music. Particularly interesting here is the gradual emergence of a musical language that combines a wide variety of macroharmonies and tonal centers. The earliest Western music explored the tonal centers contained within a relatively static and largely diatonic macroharmony (Figure 1.3.9).²¹ Classical music, which reduced the available modes to just major and minor, created large-scale harmonic contrasts by juxtaposing different scalar collections (e.g. G major and C major). It was only in the last decades of the nineteenth century that these two procedures were combined, as composers began to feel free to use *any*

²¹ The presence of unnotated accidentals (*musica ficta*) complicates matters somewhat.

Figure 1.3.9 (a) In earlier music, composers emphasized different centers within a single fixed diatonic collection. (b) In classical music, composers restricted the available modes to two, creating long-term harmonic change by emphasizing *different* major or minor scales. (c) Only in the twentieth-century were these two techniques systematically combined, creating a much wider range of tonal areas to choose from.



of the seven modes of *any* of the twelve diatonic collections. This scalar vocabulary was further extended with the use of additional nondiatonic scales—including the pentatonic, whole tone, melodic minor, harmonic minor, and octatonic—creating hundreds of possible combinations of macroharmony and tonal center. What results is a dazzling proliferation of “generalized keys” providing a wealth of alternatives to traditional major and minor modes.

1.3.3 Modulation Involves Voice Leading

My third claim is that tonal music makes use of the same voice-leading techniques on two different temporal levels: *chord progressions* use efficient voice leading to link structurally similar chords, and *modulations* use efficient voice leading to link structurally similar scales.²² As a result, tonal music is both self-similar and hierarchical, exploiting the same procedures at two different time scales.

Figure 1.3.10 shows the opening of the last movement of Clementi’s D major Piano Sonata, Op. 25 No. 6. The two parallel phrases present a series of three-voice chord progressions: I–V⁶–I followed by I–IV–I and then V⁷/V–V. The bottom staff shows that we can find a higher-level harmonic motion relating two diatonic collections: the first six measures limit themselves to the seven notes of D major, while the rest of the phrase abandons the G \natural in favor of the G \sharp . As we will see in Chapter 4, this modulation, or motion between macroharmonies, can be represented as a voice leading in which the G \natural moves by *semitone* to G \sharp . This means that the music exhibits two sorts of efficient voice leading: on the level of the half measure, there is a sequence of eight efficient voice leadings between triads; while on a larger temporal level there is a

²² This is not all that modulation does, of course, but it is typically part of it.

Figure 1.3.10
Two levels of
voice leading in
Clementi's Op. 25
No. 6.

The figure displays two systems of musical notation for Clementi's Op. 25 No. 6. Each system includes a grand staff (treble and bass clefs) and a bass line with chord symbols. The first system shows a progression from D major (I) to D major sixth (V6), then to D major (I), D major fourth (IV₄), and finally D major (I). The second system shows a progression from D major (I) to D major sixth (V6), then to D major (I), A major fourth (IV), A major seventh (V7), and finally A major (I). A chromatic scale from G to G# is shown at the bottom of the second system.

single efficient voice leading between D major and A major scales, occurring somewhere near the seventh measure of the example.

Figure 1.3.11 shows that similar processes occur in twentieth-century music. The top system depicts the main theme of Debussy's prelude "Le vent dans la plaine," which uses the pitches of E \flat natural minor. When the theme returns, B \flat moves by semitone to B $\flat\flat$, producing a collection that is enharmonically equivalent to F \sharp melodic minor ascending. Debussy's "modulation" is thus analogous to Clementi's, although it involves modes that Clementi himself would never have used. Here, then, we have a familiar tonal technique appearing in the context of a significantly expanded modal vocabulary. As we will see in Chapter 9, this same technique has been used by a number of twentieth-century composers, including Stravinsky, Shostakovich, and the minimalists. There are also a few non-Western styles that use voice leading to link closely related scales.²³

The idea that tonal music is hierarchically self-similar is central to the work of Heinrich Schenker, who claimed that tonal pieces consisted of recursively embedded patterns. The theory I have described is similar to Schenker's insofar as I consider

²³ Morton 1976 and Hall 2009.

B \flat → B $\flat\flat$

Figure 1.3.11
Voice leading
between
macroharmonies
in Debussy’s
“Le vent dans la
plaine.”

tonal music to utilize efficient voice leading at two temporal levels. Unlike Schenker, however, I view macroharmonies and scales (rather than chords or melodic lines) as the primary vehicles of long-range harmonic progression: for me, the long-term voice leading in Figure 1.3.10 connects the D and A major *scales* rather than D and A major tonic triads. By contrast, a Schenkerian would likely interpret the passage—and indeed the entire piece—as involving chord progressions at all hierarchical levels. In Chapter 7 I will return to these issues, suggesting that my approach, though perhaps less unified than Schenker’s, more closely reflects the cognitive processes involved in composition.

1.3.4 Music Can Be Understood Geometrically

My fourth claim is that geometry provides a powerful tool for modeling musical structure. This is because there exists a family of geometrical spaces that depict the voice-leading relationships among virtually any chords we might care to imagine. Some of these (such as the familiar “circle of fifths”) are relatively simple, but others (such as the Möbius strip containing two-note chords) are considerably more complex. One goal of the book is to provide a user-friendly introduction to these musico-geometrical spaces, explaining how they work, and showing how they allow us to visualize a wealth of musical possibilities at a glance.

Suppose, by way of illustration, that our friend Lyrico decides to write music using only the seven triads in the C diatonic collection. After a little exploration, he finds that some of these are closer together than others: for example, he can turn a C major triad into an A minor triad by moving only one note by one diatonic step, whereas he must move each voice to turn C major into D minor (Figure 1.3.12). Pondering this a little further, Lyrico eventually realizes that the diatonic triads can be linked in a “circle of thirds” (Figure 1.3.13), where each chord can be connected to its neighbors by a single-step motion. This circle allows him to define a kind of “distance” according

to which C major and A minor are one step apart (since they are adjacent on the circle) while C major and D minor are three steps apart (since there are two chords between them).

Now suppose that Lyrico's rival Avanta becomes frustrated with all this conservatism. "Why limit yourself to just the triads in that one seven-note scale?" she asks, stamping her foot. "There's a whole world out there beyond the white notes, you know!" Avanta then proceeds to demonstrate some of the musical possibilities not represented in Lyrico's simple circular model: she shows that in the familiar chromatic scale, the C major chord can be linked to *four* separate triads by single-semitone voice leading, and to *seven* triads by a pair of semitone steps (Figure 1.3.14). What is the analogue, in Avanta's expanded musical world, to Lyrico's circular map? How can she depict the voice-leading possibilities between *all* the triads in the chromatic scale?

Figure 1.3.12
Voice leading
between diatonic
triads.



The answer turns out to be surprisingly complicated: instead of a simple seven-chord circular model, Avanta needs the three-dimensional, 40-chord lattice shown in Figure 1.3.15. This figure provides a map of all the contrapuntal possibilities available to a composer who wants to use traditional triads, but is willing to step outside the confines of a single diatonic scale. This complex-looking construction provides the first hint that ordinary musical questions might sometimes lead to nontrivial geometrical answers. In fact, Chapter 3 shows that Avanta's lattice lives in what mathematicians would call "the interior of a twisted triangular two-torus," otherwise known as a triangular doughnut. This space contains all possible three-note chords in any conceivable scale and any conceivable tuning system. Analogous spaces depict

Figure 1.3.13
Single-step voice-
leading between
diatonic triads
can be modeled
with a circle.







Figure 1.3.14
 (a) One-semitone voice leading and (b) two-semitone voice leading among triads.

Figure 1.3.15 This three-dimensional graph represents single-semitone voice leading between major, minor, augmented, and diminished triads—represented by dark spheres, light spheres, dark cubes, and light cubes respectively. Chords on the same horizontal cross-section are related by major-third transposition; vertical motion corresponds to semitonal transposition, and the top face is glued to the bottom with a 120° twist. We will explore this figure more thoroughly in Chapter 3.



voice-leading relations among four-note chords, five-note chords, and so on. We will use them throughout this book, both for modeling chords and scales and for analyzing specific pieces.

For now, it is enough to note that the world of chromatic voice leading, as represented by Figure 1.3.15, is significantly more complicated than the simple circle representing diatonic triads. Because of this complexity, the fundamental logic animating nineteenth-century music has not always been clearly understood: theorists have sometimes depicted chromaticism as involving whimsical aberrations, departures from compositional good sense, rather than as the systematic exploration of a complex but coherent terrain. This has in turn led historians and composers to depict nineteenth-century chromaticism as pushing tonal logic to its breaking point, such that the step to complete atonality became all but inevitable. I will argue against this point of view, using new geometrical tools to demonstrate that the music of Chopin and Wagner can be just as rigorous as the music that preceded it.

1.4 MUSIC, MAGIC, AND LANGUAGE

Having sketched some major themes, let me now step back to make a few remarks about my approach to music theory. In this book I am primarily interested in the idealized composer's point of view: my goal is to describe conceptual structures that can be used to *create* musical works, rather than those involved in perceiving music. I stress the adjective "idealized," as my goal is not to undertake a historical investigation of the way past composers *actually* conceived of their music, but rather to describe concepts contemporary composers might find useful—in other words, to answer the question "what concepts would be helpful if I wanted to compose music like this?" Music theory, understood in this way, helps composers steal from one another in a sophisticated fashion, allowing us to appropriate general procedures and techniques rather than particular chords or melodies. Of course it can also help performers and analysts understand music "from the inside," by showing how composers make use of the options available to them.

Of course, composer-based music theory cannot ignore listeners entirely: the point is to write music that other people want to listen to, and this can only occur if composers are dealing with musical features that listeners care about. But though composers and listeners need to be synchronized in general, there is room for considerable divergence when it comes to the details. For example, classical composers evidently considered it important to conclude a piece in the tonic key, even though listeners are relatively insensitive to this feature of musical organization.²⁴ This means that theorists should not assume that the cognitive structures involved in *making* music are the same as those involved in *perceiving* it: ideas that are central to the composer's craft, such as the principle that a classical sonata should recapitulate the second theme in the tonic key, may have only a glancing relevance to ordinary listeners.

I find it useful here to consider the analogy with magic. A stage magician uses various tricks to *cause* the audience to have extraordinary experiences—bunnies seem to disappear, beautiful assistants seem to be sawed in half, and so on. Enjoying a magician's performance does not require you to understand how the tricks are done; in fact, understanding may actually diminish your astonishment. Nor is the magician's "ideal audience" composed of professional magicians: the point is to perform the trick for people who will genuinely be *fooled*. In much the same way, I understand composition to be a process of using technical musical tools to ensure that audiences have certain kinds of extraordinary experiences. When composing, I make various choices about chords, scales, rhythm, and instrumentation to create feelings of tension, relaxation, terror, and ecstasy, to recall earlier moments in the piece or anticipate later events. But I do *not* in general expect listeners to be consciously tracking these choices. Listeners who do ("ooh, a dissonant $\sharp 9$ chord in the trombones, in polyrhythm against the flutes and inverting the opening notes of the piece!") are like professional magicians watching each others' routines—at best, engaged in a different

24 Cook 1987, Marvin and Brinkman 1999.

sort of appreciation, and, at worst too intellectually engaged to enjoy the music as deeply as they might.

One might contrast this approach with an alternative, on which the composer's and listener's perspectives are thought to be more closely aligned. Here the relevant analogy is not to magic, but to language: the idea is that composers write music that contains various sorts of patterns, including familiar sequences of chords and keys, thematic recurrences, and so on, while the listener's job is to recover these patterns. Expert listeners recover more patterns than inexpert listeners, and are consequently more qualified to pass aesthetic judgment on particular pieces of music. Insofar as some composers place patterns in their music that cannot be decoded aurally by expert listeners, their compositions are thereby aesthetically flawed.²⁵ On this view, music essentially involves a language-like transmission of syntactic patterns from producer to receiver.

The linguistic model is attractive, and there is no doubt some truth to it. But I think it understates the extraordinary distinctiveness of human language. Linguistic abilities are remarkable both for their accuracy and for their homogeneity: if you say "there is a tiger nearby," I have no trouble repeating the sentence word for word, writing it down, or explaining what it means. Furthermore, almost any English speaker can effortlessly distinguish grammatical utterances such as "there is a tiger nearby" from nongrammatical ones like "nearby tiger is." By contrast, musical abilities are both *heterogeneous* and fairly *inaccurate*. First, there is an enormous spectrum ranging from congenital amusics to gifted listeners with extraordinarily accurate absolute pitch. These individuals have radically different musical capabilities as both producers and consumers of music. Second, while competent language-listeners are typically also competent language-speakers, this is not the case for music: most people enjoy listening to music, while only a few enjoy creating it. Third, even expert listeners lose a large amount of the "signal" in even moderately complicated music. I have had more than three decades of musical training, and yet I—like most other "expert" listeners—would have trouble notating, or recreating at the piano, the notes and rhythms in a ten-second excerpt from an unfamiliar four-voice Baroque fugue. Musically, in other words, I am analogous to a person who cannot reliably understand all the words in the sentence "there is a tiger nearby."²⁶

These differences no doubt reflect the fact that music largely lacks semantic content. A listener who does not understand "be quiet, there's a tiger nearby" is a serious danger to himself and others; consequently, there is tremendous pressure for speakers and listeners to converge on the same interpretation of sentences. By contrast, a listener who does not follow the detailed syntax of a Beethoven symphony, but who nevertheless enjoys it, creates no problems whatsoever. In some important sense *it*

25 See Babbitt 1958, Lerdahl 1988, Raffman 2003, and Temperley 2007.

26 Of course, musical signals are in some respects more complex than spoken sentences: listening to a Bach fugue is like listening to four people speak at the same time. But this should make us wonder why we tolerate higher information content in music.

simply does not matter whether you follow all the details of a piece's syntax; what matters is that you follow the piece well enough to enjoy what you hear. For in the end, the composer's well-being depends on your willingness to listen, not on whether you interpret music in the same way that he or she does. Or to revert to our earlier analogy: what is important is not that you understand the magic trick, but simply that you feel the force of the illusion.

I should pause here to mention that I have a personal perspective on these issues, since I sometimes create music with the aid of algorithms and computers. In these cases I can be fairly removed from the underlying syntactical structure of the music I create, whether that be some complex melodic process, a gradually shifting probability distribution over the twelve chromatic notes, or a harmonic structure derived from a mathematical analysis of Mozart's music. Here my compositional role is that of a gatekeeper or judge, selecting the computer-generated passages that strike me as intuitively compelling, and arranging them, in a collage-like fashion, so as to produce the best musical effect. (Typically, I augment these algorithmic passages with music composed intuitively, producing a final product that blends the human and the inhuman.) The fact that some people seem to *like* this music only serves to highlight the oddness of the linguistic approach. For here the composer is only vaguely aware of the structures that, on the linguistic model, the listener is supposed to be recovering. Given that I am not at all sure that *I* am hearing these structures accurately, it seems presumptuous for me to demand much more from my listeners. Instead, I am hoping that listeners will be *directly affected* by the music in the same way I am—that is, tickled, amused, impressed, or awed.

Faced with this situation, I draw several morals.

First, there is a potential for real divergence between what we might call “composer's grammar” and “listener's grammar.” Listeners may potentially grasp only a fraction of the underlying syntax of a Bach fugue, a Beethoven symphony, or a John Coltrane solo. (Some authors have claimed that this is particularly true when we consider the artificial syntaxes of twentieth-century atonality, and this may be true; but the more important point is that there will be significant gaps whenever we consider music of any complexity.²⁷) In this sense, listening is like trying to catch up to a train that is forever just beyond your reach; indeed, the very fact that we miss so much structure is no doubt part of what leads us to study scores, or listen repeatedly to the same pieces. This means that there are at least two separate projects that music theorists can engage in: modeling what composers actually do, and modeling what listeners actually experience. We should be careful not to conflate these by acting as if listening is “composing in reverse.”

²⁷ See Lerdahl (1988, 2001) and Raffman (2003). Lerdahl asserts that it is desirable for “composer's grammar” and “listener's grammar” to be close together, and suggests this is true for classical music. Furthermore, he postulates lossless musical perception, in which listeners have subconscious but completely accurate access to most of the details in a musical score. He therefore believes the gap between composer and listener is significantly less severe than I do.

Second, we should be careful not to assume that there is any one thing that “listening” is. It is possible that some listeners are adept at recovering a large amount of the structure in musical pieces, but many listeners are not. The heterogeneity of musical abilities makes it difficult to abstract away from individual variation in favor of some idealized “competent” listener. Can we describe as “competent” the listener who loves Beethoven, but who performs poorly on standard ear-training tests? What about the listener who hates music but perceives it very accurately? Does the “ideal” musical listener have absolute pitch? A perfect memory for every musical detail? Is the point of listening to music to experience aesthetic enjoyment, or is it to recover a kind of musical “syntax” that the composer placed in his or her music? Personally, I suspect there is no uncontroversial answer to these questions: there simply is no “competent” or “ideal” listener that is analogous to the “idealized speaker” of contemporary linguistics.

Third, listeners’ perceptions may in some respects be more crude and statistical than we would initially think. We do not determine the meaning of sentences by estimating the proportion of nouns to verbs, but we do respond very strongly to relatively crude global features of the musical stimulus. Does the piece use consonant or dissonant harmonies? Does it restrict itself, over moderate spans of musical time, to a small set of notes, and do these notes themselves change over larger time spans? Do melodies in general move by short distances? The answers to these questions tell us an enormous amount about how untrained listeners will respond to a piece. Hence my five features might be compared to a set of basic tools which composers can use to perform their musical magic.

Fourth, musical heterogeneity poses special problems for composers, who confront an audience of widely varying interests and abilities. Traditional tonal composers dealt with this by writing music that was immediately attractive, largely by virtue of exploiting the five features. Many composers also built into their music layers of additional, more complicated structure—complex thematic and formal interrelationships, intricate rhythmic devices, sophisticated contrapuntal tricks, and so on. The result was a kind of music that listeners could engage with in multiple ways: laypersons could simply enjoy a piece for the gross statistical features it shared with many other compositions in the same style, while cognoscenti could become more involved in the subtleties particular to that work.

I find it fascinating that so many twentieth-century musicians chose to abandon this strategy. Composers such as Schoenberg, Webern, Berg, and Varèse—and later Babbitt, Boulez, Cage, Xenakis, and Stockhausen—wrote music that listeners often found quite unpleasant. In many cases, this was because these composers rejected not just acoustic consonance, but also harmonic consistency, conjunct melodic motion, limited macroharmony, and centrality. Some early modernist composers may have hoped that ordinary listeners would adapt to this new musical style, adjusting their ears to the absence of familiar musical structures and coming to enjoy atonality as deeply and directly as traditional music. But after several decades of avant-garde exploration, composers began to realize that this might never occur. The *locus classicus* of this new perspective is Milton Babbitt’s 1958 manifesto “Who Cares if You Listen?”—in which he cheerfully acknowledged that his music was not enjoyed

by laypersons, but only by a specialist musical community analogous to the specialist community of professional mathematicians.²⁸

Rather than criticizing this point of view, let me just say that I am interested in music that does not make this choice: I like music that brings people together, rather than dividing them, and I think the traditional strategy—writing immediately attractive music that also contains deeper levels of structure—is as potent as it ever was. (Indeed, it may be all the more necessary in an economic and cultural environment in which notated music is somewhat marginal.) In this book, therefore, I will be primarily concerned with effects that can be achieved *within* a broadly attractive sound-world. From a technical standpoint, this restriction is relatively unconstraining, as a large number of avant-garde techniques are easily applied within a consonant harmonic context: it is perfectly possible, for example, to write music that is diatonic, or more generally macroharmonically consonant, while also being serial, aleatoric, indeterminate, wildly polyrhythmic, and so on.²⁹ Thus, the choice between superficial accessibility and off-putting dissonance is not forced on us by our interest in particular musical techniques. Instead, it is a relatively independent reflection of our own aesthetic preferences.

1.5 OUTLINE OF THE BOOK, AND A SUGGESTION FOR IMPATIENT READERS

The book is divided into two halves, with the theoretical material front-loaded into the first five chapters. Chapter 2 reviews basic theoretical concepts and introduces simple geometrical models of musical structure, representing pitches as points on a line and pitch classes as points on a circle. These models are then used to investigate the relations between conjunct melodic motion, harmonic consistency, and acoustic consonance. Chapter 3 introduces higher-dimensional “maps” of musical space, providing powerful tools for visualizing the interactions between harmony and counterpoint. Chapter 4 introduces scales, describing them as musical “rulers” that allow musicians to measure the distance between notes; it then identifies a set of familiar scales that are interesting for a number of distinct reasons. Finally, Chapter 5 describes the components of the generalized theory of keys. It proposes various tools for representing macroharmony and centricity, and contrasts two traditions in twentieth-century music: the *chromatic tradition*, which largely abandons scales in favor of highly chromatic textures, and the *scalar tradition*, which makes use of an expanded range of scales and modes. Since Chapters 2–4 are the most technically demanding portion of the book, Appendix F provides a series of study questions to help reinforce the material. These questions make good homework assignments when the book is used in a classroom.

²⁸ Babbitt 1958.

²⁹ Avant-garde techniques have been used in consonant contexts by composers like Conlon Nancarrow, Paul Lansky, Steve Reich (Cohn 1992), and “totalists” such as Mikel Rouse (Gann 2006).

The second part uses these ideas to reinterpret the history of Western music. Chapter 6 proposes that there is an “extended common practice” stretching from the beginning of Western counterpoint to the tonal music of the twentieth century. What links these different styles is the combination of harmonic consistency and conjunct melodic motion: the idea that music should have a two-dimensional coherence, both harmonic (or vertical) and melodic (or horizontal). Chapter 7 uses geometrical models to investigate the functional harmony of the classical period, briefly considering the relation between traditional harmonic theory and the views of Heinrich Schenker. Chapter 8 explores the ways in which nineteenth-century composers exploited efficient voice leading in chromatic space, suggesting that there is more structure to chromatic music than we might expect. Chapter 9 argues that twentieth-century tonal composers used scales to counteract the trend toward a saturated chromaticism, fusing chromatic voice-leading techniques with the limited macroharmony of earlier periods. Finally, Chapter 10 treats what I call the “modern jazz synthesis,” a contemporary common practice that unites impressionist chords and scales, chromatic voice leading, and the functional harmony of the classical era.

The design of the book means that readers will need to absorb a considerable amount of theoretical material before reaching the analytical payoff in the second half. Readers who are less interested in theory for its own sake may therefore want to read the book out of order. Chapter 8 can largely be read directly after Chapters 2 and 3, although the discussion of *Tristan* will be enhanced by familiarity with the material in §§4.8–10. Chapter 9 can be read directly after Chapter 4. This abbreviated path hits the main theoretical highlights (new tools for understanding voice leading and scales), as well as the most important analytical applications: to the chromatic tonality of Schubert, Chopin and Wagner, and to the twentieth-century scale-based tonality of Debussy, Stravinsky and Reich. When using the book in an advanced undergraduate theory class, I have assigned pieces from Chapters 8–9 as homework assignments early in the semester, while students are still reading the introductory chapters. (For instance, I assign Chopin’s E minor prelude and F minor mazurka in the week when students are reading Chapter 3.) That way, students have an opportunity to work with the music on their own before being introduced to my own particular perspective on these remarkable works.

Finally, a word of encouragement: it is possible to understand the gist of later chapters even while remaining somewhat fuzzy about the technical material in Chapters 2–4. So don’t be afraid to forge ahead, returning to earlier sections as the need arises.

CHAPTER 2

Harmony and Voice Leading

The enterprise of “musical set theory” aspires to catalogue all the chords available to contemporary composers. Unfortunately, this project turns out to be more complicated than one might imagine. This is, first, because chords have an intrinsic and sophisticated geometry, with distance being determined by voice-leading size. Second, basic musical concepts such as “transposition,” “inversion,” “triad,” and “chord type” can all be *relativized* to scales; thus two chords may belong to the same category relative to one scale but not another. (To make matters worse, we can also do set theory in an unquantized space that admits a continuous infinity of notes between B and C. This perspective can even teach us useful lessons about the discrete world of ordinary musical experience.) Third, chord tones themselves can be differentiated in terms of their importance, leading to phenomena such as rootedness and centrality. Finally, in the most sophisticated versions of set theory, the objects of comparison may themselves be chord progressions, with musical categorization proceeding flexibly and according to an ever-changing variety of symmetry operations (to be discussed shortly).

The next four chapters will try to address these issues, describing a new approach to chords and rebuilding “musical set theory” from the ground up. This chapter introduces the basic concepts and definitions. We begin with elementary geometrical models of musical structure, representing notes as points on a line (*pitch space*) and on a circle (*pitch-class space*). We then turn to higher-order objects—chord progressions and voice leadings—that describe motion through time. (Because this formalism sets the stage for much of the rest of the book, I encourage you to work through the study questions in Appendix F.) Next we consider the importance of symmetry in music theory, asking under what conditions harmonic consistency and conjunct melodic motion can be combined. This leads to our first significant theoretical result: a general understanding of the interdependence between acoustic consonance, efficient voice leading, and harmonic consistency.

2.1 LINEAR PITCH SPACE

Sound consists of small fluctuations in air pressure, akin to changes in barometric pressure. These fluctuations are heard as having a definite pitch when they repeat themselves (at least approximately) after some period of time t (Figure 2.1.1). The