

How should corpus studies of harmony in popular music handle the minor tonic?

Trevor de Clercq^{1†}

¹ Department of Recording Industry, Middle Tennessee State University, Murfreesboro, TN, USA

[†] Corresponding author: tdeclercq@mtsu.edu

Published 16 December 2021; <https://doi.org/10.18061/FDMC.2021.0006>

Author video presentation and/or other conference material: <https://doi.org/10.17605/OSF.IO/WXCTV>

Abstract

Corpus studies of harmony in popular music normally assume a singular tonic pitch assigned to scale-degree 1, which highlights similarities in chord organization between parallel keys. Recently, Nobile (2020) posits a “double-tonic complex” in rock, where two tonics—a major chord and its relative minor—are active simultaneously, such that similarities in chord organization manifest between relative keys. Using Kullback-Leibler divergence as a metric, I assess in a corpus of classical music and a corpus of popular music how well chord organization given a minor tonic is modeled by chord organization in the parallel and relative major. I show that chord organization in the classical corpus is modeled well by the parallel key encoding, but chord organization in the popular music corpus shows mixed results. I thus suggest that corpus studies of harmony in popular music should account for the two different ways (parallel and relative) that chord organization given a minor tonic corresponds to a major key. Possible strategies include separate analyses, dual encodings, and six-based minor.

KEYWORDS: *corpus study, harmony, tonality, popular music, minor keys*

Introduction

A decade ago, David Temperley and I published a corpus study of harmony in rock music (de Clercq & Temperley, 2011). One motivation for this work was that musicologists had proposed different theories about the harmonic organization of rock (Everett, 2004; Moore, 1992; Stephenson, 2002), and we hoped our research could offer some empirical evidence in that regard. Our original corpus, called the “RS 5x20,” included Roman numeral encodings for the top twenty songs from five consecutive decades (the 1950s through the 1990s) on the 2004 *Rolling Stone* magazine list of the “500 greatest songs of all time.”

The results from our study showed various distinguishing aspects of harmony in rock as compared to common-practice music. In the overall distribution of chord roots (Table 1), for example, scale-degree $\flat 7$ was the most common root after 1, 4, and 5—a stark contrast with common-practice norms (Temperley, 2009). We also provided data on how many times one chord root

changed to another (Table 2). Chords with a root of scale-degree $\flat 7$, for example, most often moved to chords with a root of 1, 4, and $\flat 6$ and not often to chords with a root of 2, 5, or 6. That said, the question of whether these findings supported one theory of harmonic organization or another was, we admitted, a matter of interpretation.

Table 1: Distribution of the eight most common chromatic roots in the RS 5x20 corpus.

| Root | n | n/N |
|-----------|-------|------|
| 1 | 3,058 | .328 |
| 4 | 2,104 | .226 |
| 5 | 1,516 | .163 |
| $\flat 7$ | 748 | .081 |
| 6 | 674 | .072 |
| $\flat 6$ | 372 | .040 |
| 2 | 336 | .036 |
| $\flat 3$ | 240 | .026 |

Table 2: Chord transitions in the RS 5x20 for the eight most common chromatic roots, with each cell indicating the number of occurrences from an antecedent chord root (“Ant”) to its consequent (“Cons”).

| Ant | Cons | | | | | | | |
|-----------|------|-----|-----------|------|-----|-----------|-----|-----------|
| | 1 | 2 | $\flat 3$ | 4 | 5 | $\flat 6$ | 6 | $\flat 7$ |
| 1 | | 132 | 94 | 1052 | 710 | 104 | 302 | 470 |
| 2 | 120 | | 2 | 58 | 97 | 0 | 24 | 0 |
| $\flat 3$ | 50 | 6 | | 64 | 2 | 67 | 0 | 41 |
| 4 | 1162 | 30 | 98 | | 514 | 57 | 72 | 90 |
| 5 | 788 | 36 | 6 | 392 | | 6 | 191 | 48 |
| $\flat 6$ | 208 | 1 | 20 | 22 | 22 | | 10 | 78 |
| 6 | 144 | 87 | 0 | 260 | 124 | 21 | | 3 |
| $\flat 7$ | 386 | 0 | 11 | 188 | 26 | 114 | 6 | |

In the years since our 2011 publication, Temperley and I expanded the corpus to include harmonic and melodic encodings for 200 songs, called the “RS 200” (2013). Subsequent published work, both by us and others, has used this corpus to study various topics, including rhythm and meter (Biamonte, 2014), the interaction of harmony and form (de Clercq, 2017), and



probabilistic models of harmony (Korzeniowski, Sears, & Widmer, 2018). In addition, other corpora of chord annotations for popular music have become available (Burgoyne et al., 2011; Koops et al., 2019).

Alongside this corpus work, musicologists continued to propose and refine theories about the harmonic organization of popular music (Biamonte, 2017; Doll, 2017). Of particular relevance here, Nobile recently posited that some rock songs may have not one but two active tonics, a situation he calls the “double-tonic complex” (2020). In short, Nobile asserts that the tonic chords of two relative keys—such as C major and A minor—can coexist, without one necessarily taking precedence over the other. A common example of this is the Axis progression (Richards, 2017), which consists of the four-chord sequence Am–F–C–G and its various rotations and transpositions, such as the Bm–G–D–A loop found in the song “Building a Mystery” by Sarah McLachlan (1997). Prior to Nobile’s work, music theorists would assign Roman numerals to this progression in one of two ways: either as (1) vi–IV–I–V, with the third chord as tonic; or as (2) i–♭VI–♭III–♭VII, with the first chord as tonic. See Doll (2017) for examples of both methods. What does it mean for corpus studies of harmony that both Roman numeral analyses may be correct, given that current corpora of chord annotations indicate only a single tonic or key?

Parallel or Relative Key Relationships?

Consider the implications of the double-tonic complex on a corpus study of harmony. The corpus might, for example, include many instances of the Axis progression, which has been found in upwards of 30% of hit songs in the past fifteen years (Richards, 2017). Assuming for the sake of argument that roughly half of the Axis progressions were analyzed in a major key and the other half in a minor key, we would find a variety of chord roots in the overall distribution. In addition to scale-degree 1, we would find many instances of scale-degrees ♭3, 4, 5, ♭6, 6, and ♭7, making popular music appear to contain a rich palette of chord roots. Yet underlying these results is a rather simple set of merely four diatonic chords.

The problem is that Roman numerals, which traditionally take the tonic chord whether major or minor to be “one,” excel at highlighting the similarity of chord progressions between parallel keys. For classical music, this is a useful feature, since theorists posit similar chord function between parallel keys. Compare, for example, the normative harmonic functions of

chords in a major key (Figure 1) and chords in a minor key (Figure 2), as diagrammed in a prominent music theory textbook (Kostka, Payne, & Almén, 2013, p. 105). Notice that, aside from modifications in chord quality to account for the mode of the key, the diagrams are otherwise identical. We would thus expect scale-degree behavior to be the same between major and minor tonalities. In the key of C major, for example, we expect an A minor chord (vi) to move to an F major (IV) or D minor chord (ii), similar to our expectation in the key of C minor for an A♭ major chord (♭VI) to move to an F minor (iv) or D diminished chord (ii°).




Figure 1: Normative harmonic functions in major keys, as diagrammed by Kostka, Payne, and Almén.




Figure 2: Normative harmonic functions in minor keys, adapted from the diagram by Kostka, Payne, and Almén.

The Axis progression, however, does not display similar chord behavior between parallel keys. In an Axis progression in C major, for example, an A minor chord (vi) goes to an F major chord (IV), whereas an Axis progression in C minor, an A♭ major chord (♭VI) goes to an E♭ major chord (♭III). The Axis progression instead displays similar chord behavior between relative keys.

This type of relative-key similarity between chord progressions is not limited to the Axis progression. Consider, for example, the song “Rockin’ in the Free World” by Neil Young (1989). The opening chords for the verse and chorus sections are shown in Figure 3. My encoding in the RS 200 (TdC) takes G major to be the tonic of the chorus and E minor to be the tonic of the verse. As a result, the chorus includes the chords I, IV, V, and vi, whereas the verse includes the chords i, ♭VI, and ♭VII. Yet most of the chord symbols in both sections are identical: C, D, and Em. Moreover, the behavior of

these chords is the same between sections, the D chord moving to C and C to Em. My Roman numeral analysis, therefore, obfuscates the very clear similarity in chord palette and chord behavior between the two sections.

| | | | | | | |
|---------------|------|------|-----|-----|------|-----|
| Verse | Em | D | C | Em | D | C |
| (TdC) E = 1 | i | ♭VII | ♭VI | i | ♭VII | ♭VI |
| G = 1 | vi | V | IV | vi | V | IV |
| Chorus | G | D | C | C | Em | |
| E = 1 | ♭III | ♭VII | ♭VI | ♭VI | i | |
| (TdC) G = 1 | I | V | IV | IV | vi | |

Figure 3: The first four bars of harmonic content for the verse and chorus of “Rockin’ in the Free World” by Neil Young (1989), showing chord symbols and Roman numerals with different pitches as scale-degree 1.

How much, then, is this a problem for corpus studies of harmony? The central question, I would argue, is whether chord syntax in popular music given a minor tonic is better modeled by the parallel major, the relative major, or neither. In what follows, I present a statistical analysis of the RS 200 as an initial attempt towards an answer. I then close by suggesting some workarounds and alternative encoding schemes.

Some Statistical Scrutiny

To assess whether harmonic patterns in popular music given a minor tonic correspond more with harmonic patterns in the parallel or relative major, I will use a metric known as *Kullback-Leibler divergence* (or “KL divergence”), also known as relative entropy. In essence, KL divergence shows how surprised we will be by events given a certain set of expectations (McElreath 2020). If we have certain expectations for chord progressions given a major tonic, for example, are we more surprised by the behavior of chords given a minor tonic if we encode that minor tonic as parallel or as relative to the major tonic? Formally, KL divergence (D_{KL}) is determined by calculating the weighted average quotient in log probabilities (here, the natural log) between a target distribution (p) and a model (q):

$$(1) \quad D_{KL} = \sum_i p_i * \ln \left(\frac{p_i}{q_i} \right)$$

Note that KL divergence is not symmetrical. A situation where we expect some events to be rare but in

reality all events are fairly common (some low q_i values, moderate p_i values) generates a greater KL divergence—i.e., is more surprising—than a situation where we expect all events to be fairly common but in reality some events are fairly rare (moderate q_i values, some low p_i values). In other words, we are more surprised by something rare happening than something common not happening.

Table 3 illustrates the application of KL divergence using data from a corpus study by Temperley (2009) that tallied the distribution of chord roots in the musical examples from the third edition of the Kostka-Payne workbook (the “KP corpus”). Presuming chord behavior in these examples follows the diagrams in Figures 1 and 2, we would expect the distribution of chord roots in a minor key to be less surprising when compared to the distribution of chord roots in the parallel major than the relative major. Notice that the greatest proportions (n/N) of chord roots for examples in a major key are scale-degrees 1 (.36) and 5 (.23). Similarly, the greatest proportions of chord roots for examples in a minor key encoded with tonic equal to 1 (a parallel method) are also scale-degrees 1 (.35) and 5 (.27).

Table 3: Diatonic root distributions in the KP corpus, showing KL divergence for excerpts with a minor tonic encoded as parallel or relative to the major tonic.*

| Major | | Minor as Parallel | | Minor as Relative | | | |
|-------|-----|-------------------|--------------|-------------------|----------|--------------|----------|
| Root | n/N | Root | n/N | D_{KL} | Root | n/N | D_{KL} |
| 1 | .36 | 1 | .35 | -0.004 | 1 | .03 | -0.077 |
| 2 | .14 | 2 | .07 | -0.048 | 2 | .08 | -0.041 |
| 3 | .03 | ♭3 | .03 | 0.004 | 3 | .27 | 0.607 |
| 4 | .07 | 4 | .08 | 0.009 | 4 | .08 | 0.009 |
| 5 | .23 | 5 | .27 | 0.044 | 5 | .03 | -0.054 |
| 6 | .07 | ♭6 | .08 | 0.009 | 6 | .35 | 0.540 |
| 7 | .04 | ♭7 | .03 | -0.014 | 7 | .07 | 0.025 |
| chr | .06 | chr | .09 | 0.044 | chr | .09 | 0.044 |
| | | Σ | 0.043 | | Σ | 1.052 | |

* N.B. “chr” refers to non-diatonic (chromatic) roots

The three rightmost columns in Table 3 rotate the data from the “Minor as Parallel” columns so that the minor tonic is taken to be scale-degree 6 in the relative major. Doing so, we can assess how well our expectations from the major key model chord behavior in the relative minor. As the divergence totals in the bottom row show, encoding these excerpts with a minor tonic as parallel to the major key generates less total divergence (0.043) than encoding minor tonics as relative to the major key (1.052). The key of A minor,

for example, thus contains many chords with a root of A and E, which is more like A major than C major.

Table 4 shows the same methodology applied to the RS 200 corpus. This table divides the corpus into songs with a major tonic (I) only (154 songs) and songs with a minor tonic (i) only (32 songs). Songs with both a major and minor tonic were ignored (14 songs). Note that chord roots of scale-degree 1, 4, and 5 are the most common in songs with a major tonic, as in Table 1. But also note that in songs with a major tonic, scale-degree ♭7—which was the next most common root in Table 1—is now subsumed in the chromatic category (“chr”), which in total accounts for a smaller proportion of chord roots than scale-degree 6.

Table 4: Diatonic root distributions in the RS 200, showing KL divergence for excerpts with a minor tonic encoded as parallel or relative to the major tonic.

| Major | | Minor as Parallel | | | Minor as Relative | | |
|-------|-----|-------------------|-----|-----------------|-------------------|-----|-----------------|
| Root | n/N | Root | n/N | D _{KL} | Root | n/N | D _{KL} |
| 1 | .34 | 1 | .33 | -0.009 | 1 | .08 | -0.116 |
| 2 | .04 | 2 | .02 | -0.015 | 2 | .13 | 0.139 |
| 3 | .03 | ♭3 | .08 | 0.088 | 3 | .13 | 0.209 |
| 4 | .26 | 4 | .13 | -0.090 | 4 | .13 | -0.090 |
| 5 | .19 | 5 | .13 | -0.047 | 5 | .16 | -0.029 |
| 6 | .08 | ♭6 | .13 | 0.072 | 6 | .33 | 0.488 |
| 7 | .00 | ♭7 | .16 | 0.610 | 7 | .02 | 0.049 |
| chr | .06 | chr | .02 | -0.020 | chr | .02 | -0.020 |
| | | Σ | | 0.588 | Σ | | 0.630 |

As the KL divergence totals in the bottom row of Table 4 show, encoding the minor key in a parallel way to the major tonic has a roughly similar divergence as encoding it in a relative way (0.588 versus 0.630). To understand why, consider the chord roots in the “Minor as Parallel” column in Table 4 in the key of A minor. The proportions of chord roots on D (4), E (5), F (♭6), and G (♭7) are roughly equivalent. This means that the A minor key can be modeled by A major (the parallel major), in which chords with a root of D (4) and E (5) are common, about as well as it can be modeled by C major (the relative major), in which chords with a root of F (4) and G (5) are common. In other words, the chord distribution for a minor key in popular music exhibits similarities with both the parallel and relative major, unlike the chord distribution for a minor key in classical music (at least as presented in a prominent textbook).

Admittedly, the divergence of chord distributions in a minor key compared to the parallel major (0.588) is less than the divergence as compared to the relative

major (0.630). We might thus consider that thinking about minor keys in a parallel way is still more appropriate, albeit it perhaps less so than for classical music. This outcome reverses, however, if we consider chord transitions. KL divergence values for root transition probabilities (using Laplace smoothing in cases of zero probability) are shown in Table 5, which compares chord behavior in a minor key with its parallel major, and Table 6, which compares chord behavior in a minor key with its relative major. (For display purposes, the values represent “milli-divergences,” e.g., 371 is actually a divergence of 0.371.) As the overall KL divergences in the bottom righthand corners show, chord transition probabilities in a minor key are closer to those of the relative major (1.353) than the parallel major (1.575). The difference is small, though, and neither shows a close modeling of expectation.

Table 5: KL divergence table (in milli-values) for root transition probabilities in the RS 200, comparing songs with a minor tonic encoded as parallel keys to songs with a major tonic.

| Ant | Cons | | | | | | | |
|------|------|----|-----|-----|-----|------|------|--------------|
| | 1 | 2 | 3/3 | 4 | 5 | 6/♭6 | 7/♭7 | chr |
| 1 | | -6 | 60 | -45 | -22 | -8 | 396 | -11 |
| 2 | -4 | | 0 | 0 | -4 | 0 | 3 | -1 |
| 3/3 | 25 | -2 | | 14 | 4 | 33 | 58 | 1 |
| 4 | -56 | -2 | 45 | | -26 | 13 | 31 | -3 |
| 5 | -6 | 18 | 0 | -13 | | -8 | 8 | -2 |
| 6/♭6 | 104 | 0 | 1 | -6 | 5 | | 204 | 0 |
| 7/♭7 | 371 | 0 | 3 | 94 | 31 | 290 | | 1 |
| chr | -8 | 0 | 0 | -8 | 0 | 0 | 0 | |
| | | | | | | | Σ Σ | 1,575 |

Table 6: KL divergence table (in milli-values) for root transition probabilities in the RS 200, comparing songs with a minor tonic encoded as relative keys to songs with a major tonic.

| Ant | Cons | | | | | | | |
|-----|------|----|-----|-----|-----|-----|-----|--------------|
| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | chr |
| 1 | | 4 | -4 | -46 | -26 | -14 | 0 | -5 |
| 2 | 24 | | 93 | 11 | -6 | 110 | 10 | 0 |
| 3 | 0 | 1 | | -2 | 2 | 257 | 43 | 1 |
| 4 | -22 | -2 | 21 | | -26 | 162 | 0 | 0 |
| 5 | -12 | 11 | 11 | 4 | | 115 | 0 | -2 |
| 6 | 39 | 17 | 124 | -1 | 188 | | 24 | 18 |
| 7 | 2 | 0 | 37 | 0 | 2 | 16 | | 0 |
| chr | 0 | 8 | 0 | -2 | 0 | 14 | 1 | |
| | | | | | | | Σ Σ | 1,353 |

Conclusion

How then should we handle the minor tonic in corpus studies of harmony in popular music?

One recommendation is simply to not merge data from major and minor keys into a single statistical analysis. That is easier said than done, since it is often unclear whether the tonic chord of a song is major or minor. But to the extent that it is possible to separate out tonalities that are clearly major or clearly minor, we should try to do so.

A second recommendation is that when we hope to analyze a corpus that contains encodings of songs in a minor tonality, we must reconcile the possibility that chord behavior in that minor tonality may exhibit similarities to chord behavior in the parallel major, the relative major, or neither. Accommodating for these possibilities requires at least two encodings for a single song. If a song contains the progression Am–F–Dm–F, our encoding (or our analysis of the encoding) should allow for the fact that this could be represented as vi–IV–ii–IV or i–VI–iv–VI, even if we clearly hear A minor to be the tonic. These two approaches need to be strictly distinguished from one another in the final analysis, though, so as not to overstate the complexity of harmonic syntax in popular music.

In this regard, it is worth mentioning that encoding the minor tonic as “one,” which is the default method for classically trained musicians, is not the standard practice among many commercial musicians. Users of the Nashville number system, for example, normally encode the root of minor tonic as scale-degree 6 (de Clercq, 2019), which highlights similar chord patterns between the minor key and its relative major. I am not advocating here that we always encode the root of the minor tonic as scale-degree 6 or as scale-degree 1. Until we have a better understanding of harmonic behavior in popular music, we probably need both methods.

In closing, I will admit that I cannot offer a tidy answer to the central question of this paper. Instead, what I have hoped to accomplish is to show how in a corpus of songs drawn from the last century, harmonic behavior in songs with a clear minor tonic has no consistent parallel-key correspondence to harmonic behavior in songs with a clear major tonic, even though it is these parallel-key correspondences that a traditional “one-based minor” Roman numeral analysis highlights. So when I think about the future direction of corpus work on harmony in popular music, I think we need to follow more than the current path we have been on.

References

- Biamonte, N. (2014). Formal functions of metric dissonance in rock music. *Music Theory Online*, 20(2). <https://doi.org/10.30535/mto.20.2.1>
- Biamonte, N. (2017). Pop/rock tonalities. In F. Wörner, U. Scheideler, and P. Rupprecht (Eds.), *Tonality Since 1950* (pp. 89–101). Franz Steiner.
- Burgoyne, J. A., Wild, J., & Fujinaga, I. (2011). An expert ground truth set for audio chord recognition and music analysis. In A. Klapuri & C. Leider (Eds.), *Proceedings of the 12th international society for music information retrieval conference* (pp. 633–638).
- de Clercq, T. (2017). Interactions between harmony and form in a corpus of rock music. *Journal of Music Theory*, 61(2), 143–170. <https://doi.org/10.1215/00222909-4149525>
- de Clercq, T. (2019). The Nashville number system: A framework for teaching harmony in popular music. *Journal of Music Theory Pedagogy*, 33, 3–28.
- de Clercq, T., & Temperley, D. (2011). A corpus analysis of rock harmony. *Popular Music*, 30(1), 143–170. <https://doi.org/10.1017/S026114301000067X>
- Doll, C. (2017). *Hearing harmony: Toward a tonal theory for the rock era*. University of Michigan. <https://doi.org/10.3998/mpub.3079295>
- Everett, W. (2004). Making sense of rock’s tonal systems. *Music Theory Online*, 10(4).
- Koops, H. V., de Haas, W. B., Burgoyne, J. A., Bransen, J., Kent-Muller, A., & Volk, A. (2019). Annotator subjectivity in harmony annotations of popular music. *Journal of New Music Research*, 48(3), 232–252. <https://doi.org/10.1080/09298215.2019.1613436>
- Korzeniowski, F., Sears, D., & Widmer, G. (2018). A large-scale study of language models for chord prediction. In *Proceedings of the 2018 IEEE international conference on acoustics, speech, and signal processing* (pp. 91–95). <https://doi.org/10.1109/ICASSP.2018.8462285>
- Kostka, S., Payne, D., & Almén, B. (2013). *Tonal harmony: With an introduction to twentieth-century music* (7th ed.). McGraw-Hill.
- McElreath, R. (2020). *Statistical rethinking: A Bayesian course with examples in R and STAN* (2nd ed.). CRC. <https://doi.org/10.1201/9780429029608>
- Moore, A. (1992). Patterns of harmony. *Popular Music*, 11(1), 73–106. <https://doi.org/10.1017/S0261143000004852>
- Nobile, D. (2020). Double-tonic complexes in rock music. *Music Theory Spectrum*, 42(2), 207–226. <https://doi.org/10.1093/mts/mtaa003>
- Richards, M. (2017). Tonal ambiguity in popular music’s Axis progressions. *Music Theory Online*, 23(3). <https://doi.org/10.30535/mto.23.3.6>

- Stephenson, K. (2002). *What to listen for in rock: A stylistic analysis*. Yale University.
<https://doi.org/10.12987/yale/9780300092394.001.0001>
- Temperley, D. (2009). A statistical analysis of tonal harmony. <http://davidtemperley.com/kp-stats/>
- Temperley, D., & de Clercq, T. (2013). Statistical analysis of harmony and melody in rock music. *Journal of New Music Research*, 42(3), 187–204.
<https://doi.org/10.1080/09298215.2013.788039>