A corpus of rock harmonies and melodies

David Temperley
Eastman School of Music

(collaborative work with Trevor de Clercq and Adam Waller)

EASTMAN
SCHOOL OF MUSIC
UNIVERSITY of ROCHESTER

Most research on harmony and melody has focused on Western classical music. But recently, researchers have started to pay more attention to popular music, particularly rock.

It is clear that rock is *tonal*: some kind of concept of key is operative. It is also basically *triadic:* rock songs are built on progressions of triadic harmonies.

But basic questions arise:

It is clear that rock is *tonal*: some kind of concept of key is operative. It is also basically *triadic:* rock songs are built on progressions of triadic harmonies.

But basic questions arise:

- How is a tonality established in rock? What are the cues that tell listeners what key a piece is in?

It is clear that rock is *tonal*: some kind of concept of key is operative. It is also basically *triadic:* rock songs are built on progressions of triadic harmonies.

But basic questions arise:

- How is a tonality established in rock? What are the cues that tell listeners what key a piece is in?
- Is the distinction between "major" and "minor" keys valid for rock? Or are there quite different kinds of modes and scales?

It is clear that rock is *tonal*: some kind of concept of key is operative. It is also basically *triadic:* rock songs are built on progressions of triadic harmonies.

But basic questions arise:

- How is a tonality established in rock? What are the cues that tell listeners what key a piece is in?
- Is the distinction between "major" and "minor" keys valid for rock? Or are there quite different kinds of modes and scales?
- What are the principles of harmonic organization in rock? Do the classical rules of progression apply (ii ➔V, V➔I, etc.) or are there quite different principles?

It is clear that rock is *tonal*: some kind of concept of key is operative. It is also basically *triadic:* rock songs are built on progressions of triadic harmonies.

But basic questions arise:

- How is a tonality established in rock? What are the cues that tell listeners what key a piece is in?
- Is the distinction between "major" and "minor" keys valid for rock? Or are there quite different kinds of modes and scales?
- What are the principles of harmonic organization in rock? Do the classical rules of progression apply (ii ➔V, V➔I, etc.) or are there quite different principles?

These are controversial questions; there is little consensus on them in the popular-music-theory literature (Moore, 2001; Stephenson, 2002; Everett, 2008).

These questions are not just of academic interest. Many researchers in Music Information Retrieval (MIR) are interested in intelligent processing of musical audio, for various practical purposes:

- Genre classification
- Transcription
- Query systems (e.g. “query-by-humming”)
- Emotion tagging
- Music recommendation

Information about harmony and melody can be very helpful for these tasks.

Data that is hand-annotated with melodic and harmonic information could be be very useful to MIR:

- It can help to resolve basic issues about popular music that arise in the design of MIR systems. (E.g., is it valid to classify songs as "major" and "minor"? What strategies are most effective for "key-finding" in rock?)

- More specifically, it provides "ground truth" data that could be used to train automatic information-retrieval systems. (Data about the frequency of different harmonies, harmonic progressions, melodic scale-degrees and transitions, etc.)

In this talk, I'll report on a long-term project to create a corpus of rock melodies and chord progressions, and some explorations that we've done with this data.

After overviewing the corpus itself, I'll talk about four topics:

1. Scale-degree distributions and the "major/minor" distinction
2. Key-finding
3. Harmonic progression
4. Changes over time

This will be a very "high-level" overview. Read more about all this in our publications (de Clercq & Temperley 2011, Temperley & de Clercq 2013).

The Corpus

Our corpus is based on *Rolling Stone* magazine's list of the "500 Greatest Songs of All Time" (2004). The top 10 songs from the list:

1 Bob Dylan, "Like a Rolling Stone" (1965)
2 The Rolling Stones, "Satisfaction" (1965)
3 John Lennon, "Imagine" (1971)
4 Marvin Gaye, "What's Going On" (1971)
5 Aretha Franklin, "Respect" (1967)
6 The Beach Boys, "Good Vibrations" (1966)
7 Chuck Berry, "Johnny B. Goode" (1958)
8 The Beatles, "Hey Jude" (1968)
9 Nirvana, "Smells Like Teen Spirit" (1991)
10 Ray Charles, "What'd I Say" (1959)

Our corpus includes 200 songs from this list. The list basically spans the period from 1955 to 2000. (We're now working on adding some more recent songs – more about that later.)

In the first stage of the project, Trevor de Clercq and I analyzed the harmony of all 200 songs, in Roman numeral notation.

We created a system for efficiently encoding repeated patterns and sections in a song. So my analysis of the Beatles' "Hey Jude" looks like this:

```
Vr: I | V | | I | IV | I | V | I |
BrP: V7/IV | IV I6 | ii vi6 | V6 V | I |
Br: $BrP $BrP [2/4] I | [4/4] V | |
CP: I | bVII | IV | I |
Fadeout: $CP*18
S: [F] R | $Vr $Vr $Br $Vr $Br $Vr I | $Fadeout
```

"S" expands to a series of symbols (Vr, Br, etc.), each of which expands to other symbols and eventually to chords.

This can then be recursively expanded out to a single long chord progression: The analysis of "Hey Jude":

```
[F] I | V | | I | IV | I | V | I | (etc.)
```

We then transcribed the melodies of all 200 songs as well. (We transcribed just the main melody, not instrumental melodies or backing vocals.) The first verse of "Hey Jude": 

```
[F] [OCT=4] ...5 | 3....356 | v2.....23 | 4.^1..175 |
6.543.........5. | 6.6...6.21.7.16. | 5...v1236 | 5..543.7 |
1.....^5. (etc.)
```

Trevor's and my judgments were in agreement 93.3% of the time on the harmonic analyses and 89.3% of the time on the melodic transcriptions.

We also provide timing data for the harmonies and melodies, allowing them to be aligned with the original audio files.

All this is publicly available at
www.theory.esm.rochester.edu/rock_corpus

Overall scale-degree distribution

First, it's interesting to look at the overall distribution of chromatic scale-degrees (pitch-classes in relation to the tonic) in the corpus.

We can generate scale-degree distributions from our data in two ways:

- from the melodic transcriptions – just count up pitch-classes in relation to the tonic (1, b2, 2, b3, 3, etc....)

- from the harmonic analyses – take each chord to imply one instance of each scale-degree that it contains, e.g. a I chord implies one 1, one 3, and one 5.

C: I IV | V I |

Degree	Count
1	3
b2	0
2	1
b3	0
3	2
4	1
#4	0
5	3
b6	0
6	1
b7	0
7	1

Scale-degree distributions in rock melody and rock harmony

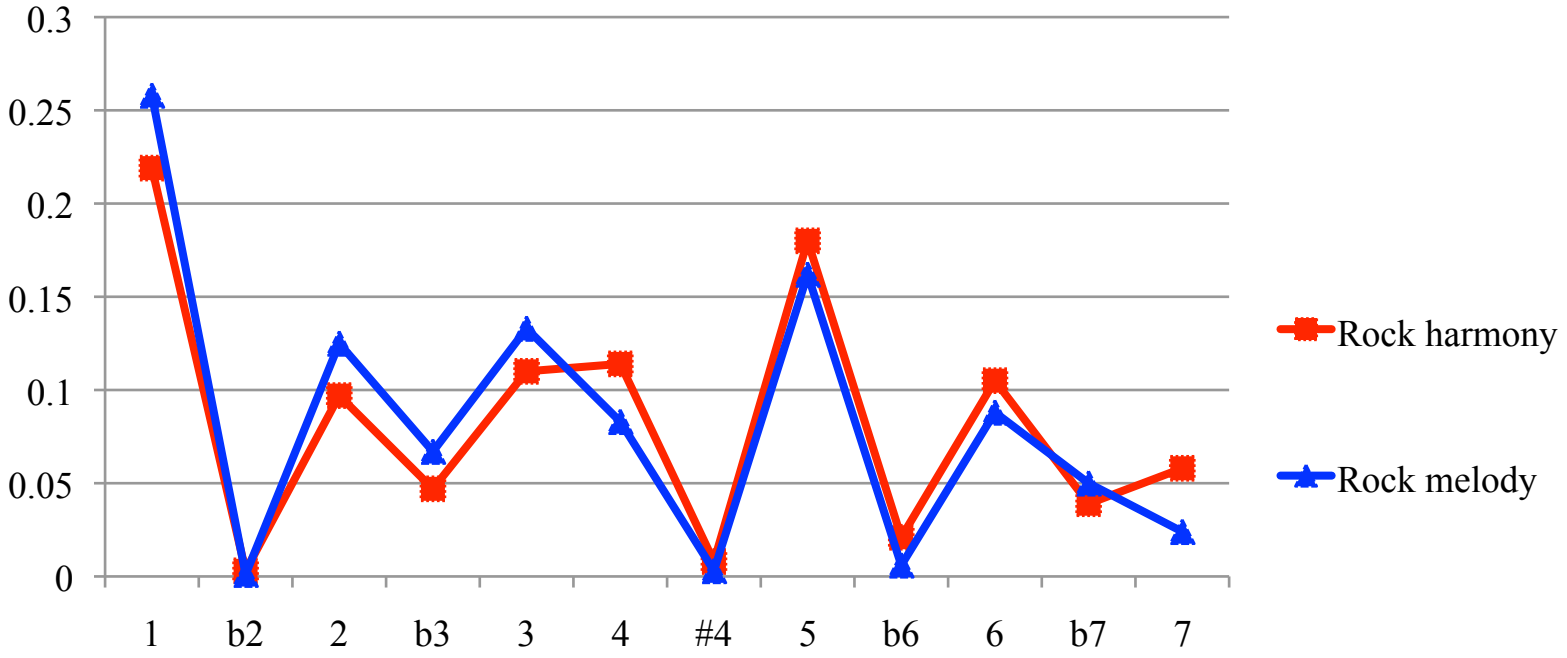

The two distributions are quite similar. In both, the lowest values are for b2 and #4.

(Data for common-practice music – not shown here – is similar!)

Scale-degree distributions in rock melody and rock harmony

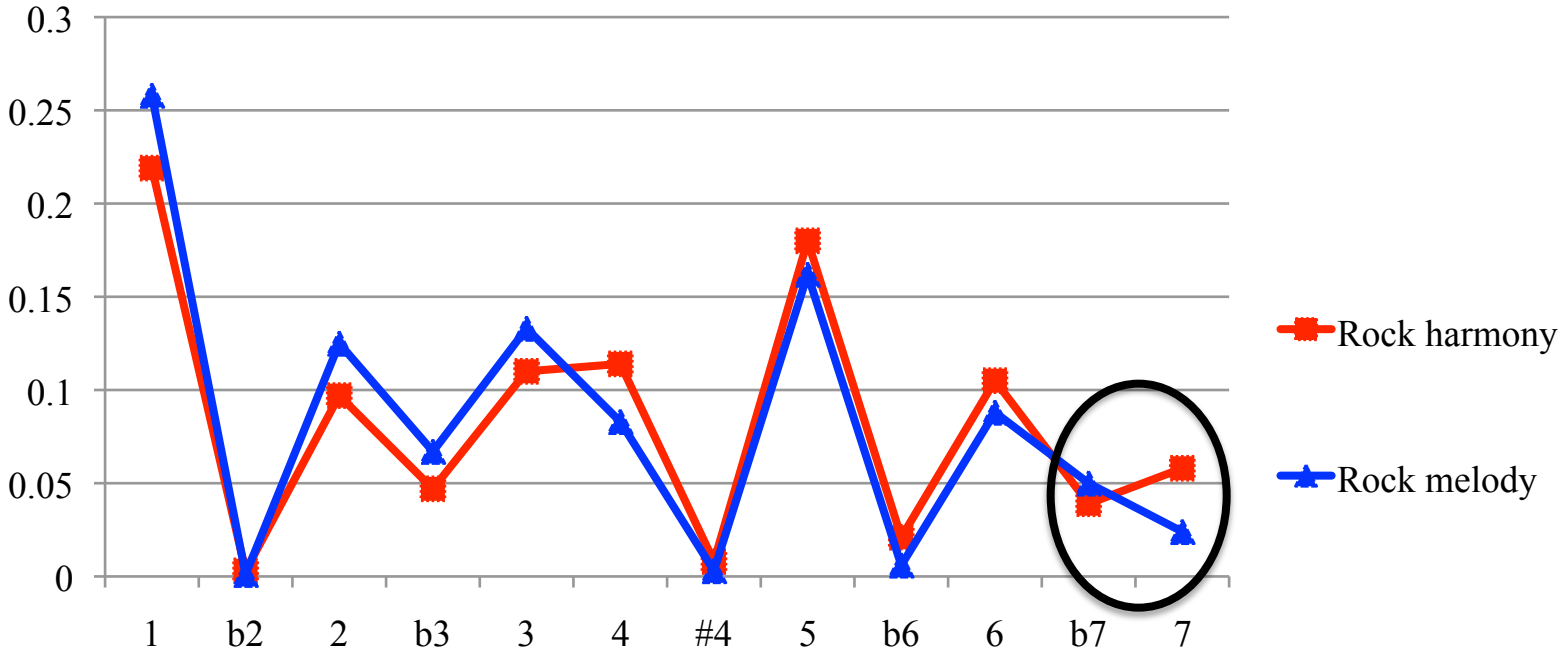

In the harmony distribution, the seven degrees of the major scale have the highest values. This is true in melody distribution too, *except* that b7 is more common than 7.

Clustering

We want to know whether rock songs naturally group into categories in terms of their pitch organization — perhaps major vs. minor, or perhaps some other system of categories.

We explored this question using a statistical "clustering" technique (a variant of the K-means technique), not described in detail here. The basic idea: Using the scale-degree distributions of individual songs, find ways of grouping the songs into categories, so that the average distribution for each category is a good fit to all the songs in the category.

We did this for both the melodic data and the harmonic data; I'll focus on the melodic data.

Distributions for the two categories (C1 and C2) revealed by the clustering analysis of the melodic data

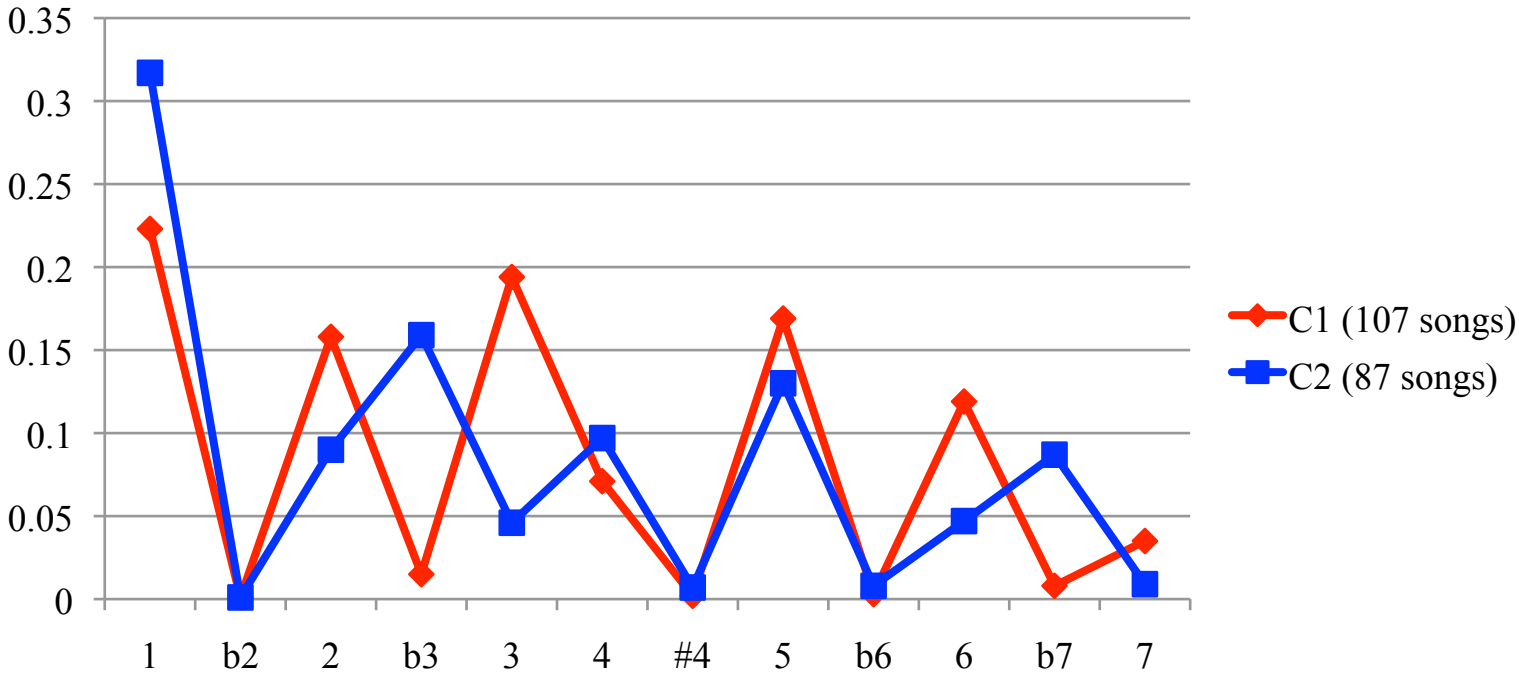

Distributions for the two categories (C1 and C2) revealed by the clustering analysis of the melodic data

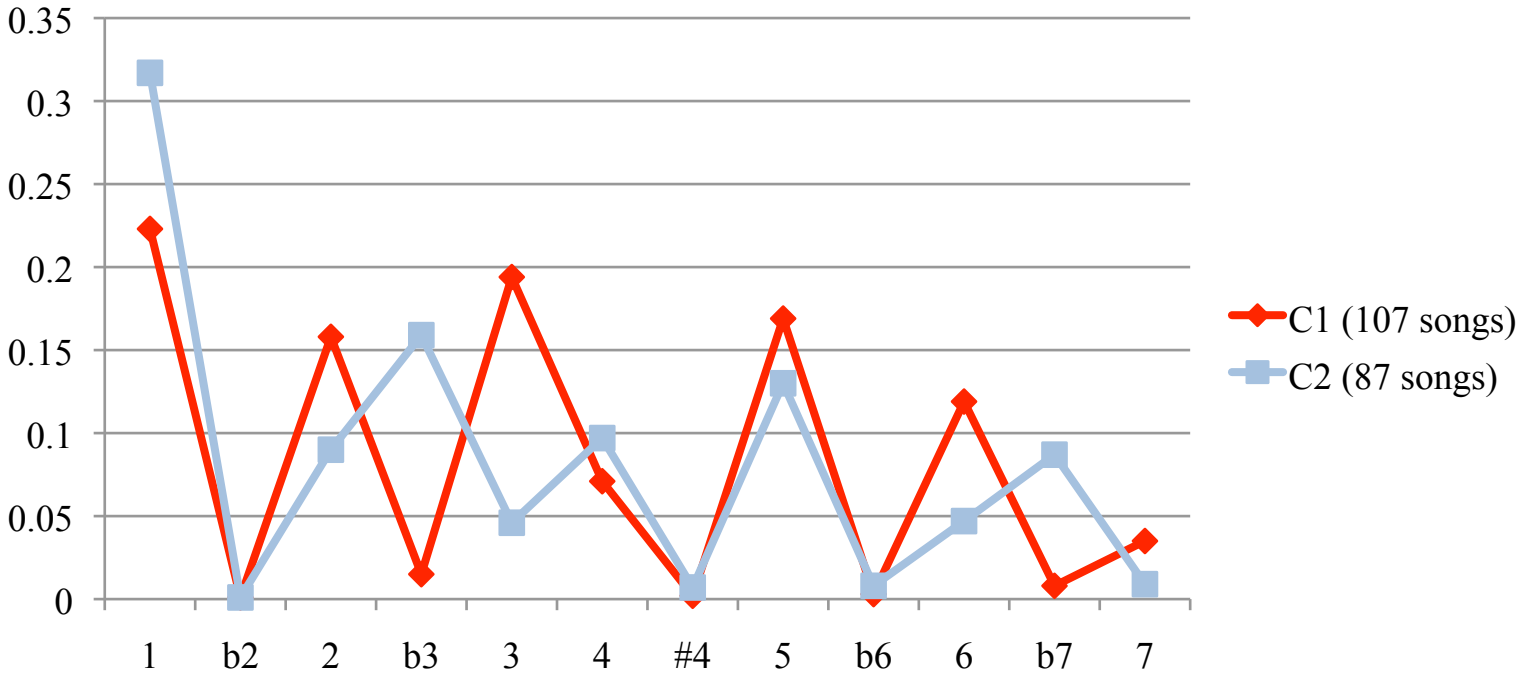

The C1 distribution reflects the major scale; 1, 3, and 5 are most common, then the other two major-pentatonic degrees 2 and 6, then 4 and 7.

Distributions for the two categories (C1 and C2) revealed by the clustering analysis of the melodic data

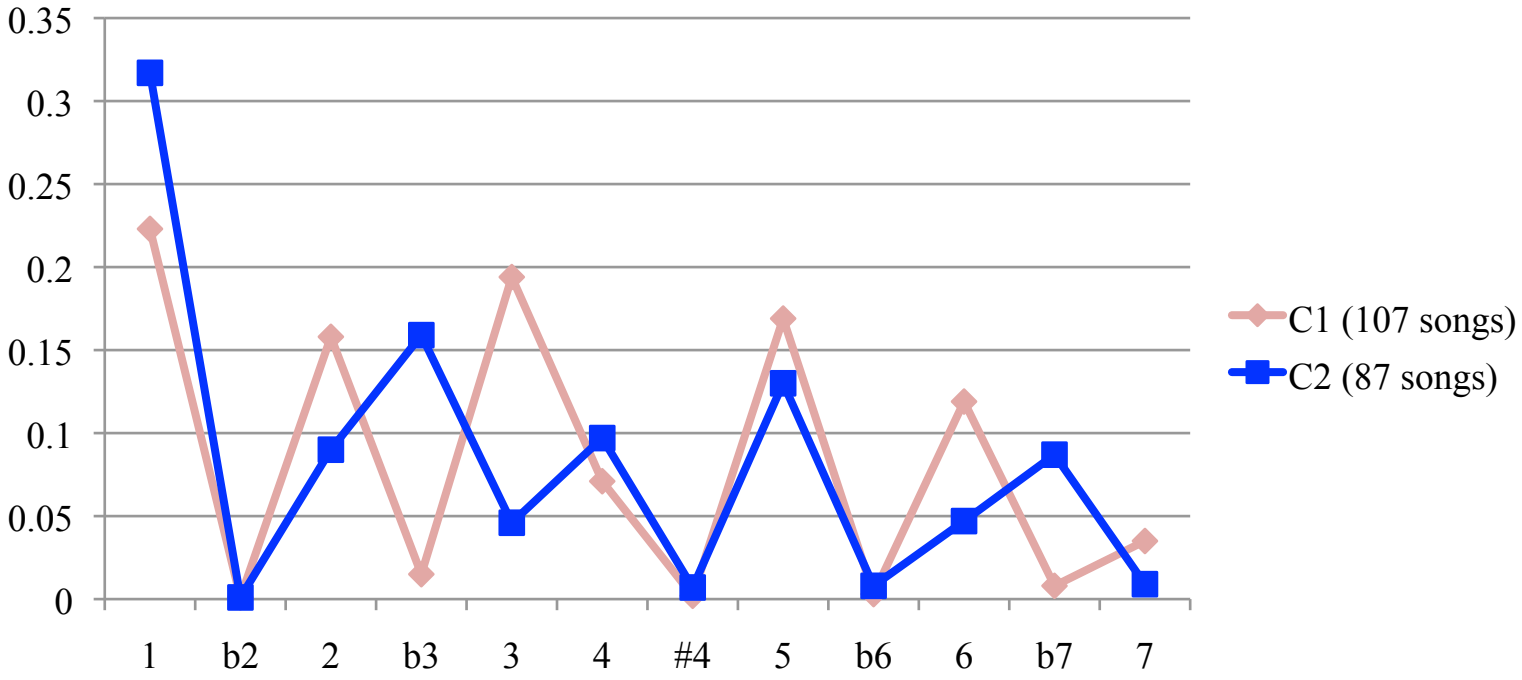

C2 is perhaps a kind of "minor": 1, b3 and 5 are most frequent, then 4 and 2. But b7 > 7 but 6 > b6; this is quite unlike classical minor, where 7 > b7 and b6 > 6. (More like Dorian mode?) Also the value for 3 is quite high.

Distributions for the two categories (C1 and C2) revealed by the clustering analysis of the melodic data

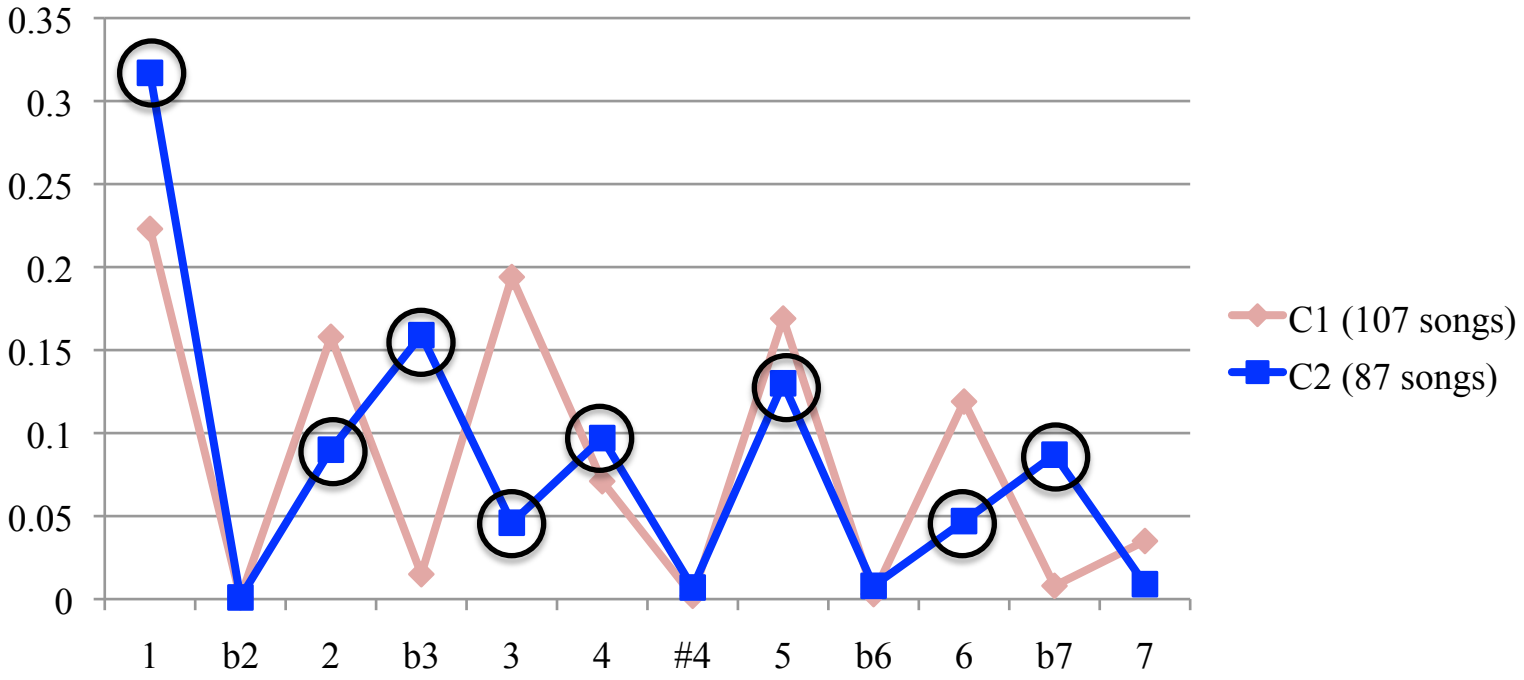

C2 is perhaps a kind of "minor": 1, b3 and 5 are most frequent, then 4 and 2. But b7 > 7 but 6 > b6; this is quite unlike classical minor, where 7 > b7 and b6 > 6. (More like Dorian mode?) Also the value for 3 is quite high.

One could posit an 8-note scale, 1-2-b3-3-4-5-6-b7; this could be seen as a combination of the major and minor pentatonic scales (we call it the "pentatonic union" scale). (It could also be seen as a kind of "blues" scale.)

We did this for the harmonic data too, and found a similar organization of two clusters, a “major” one a “minor-ish” one.

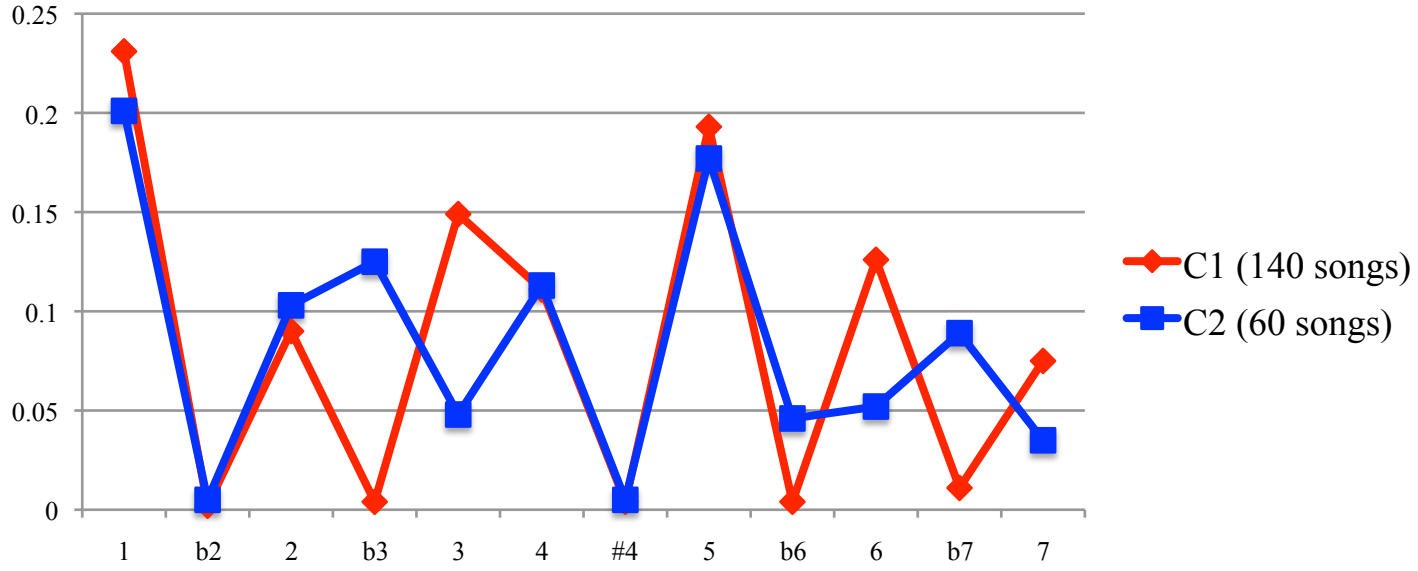

Interestingly, though, songs were not always in the same cluster melodically and harmonically...

The number of songs in the major and minor harmonic and melodic clusters

Melodic → / Harmonic ↓	Major	Minor	Total
Major	98	41	139
Minor	9	46	55
Total	107	87	194

The number of songs in the major and minor harmonic and melodic clusters

Melodic → / Harmonic ↓	Major	Minor	Total
Major	98	41	139
Minor	9	46	55
Total	107	87	194

- About 3/4 of the songs (144 / 194) are either "major/major" or "minor/minor"

The number of songs in the major and minor harmonic and melodic clusters

Melodic → / Harmonic ↓	Major	Minor	Total
Major	98	41	139
Minor	9	46	55
Total	107	87	194

- About 3/4 of the songs (144 / 194) are either "major/major" or "minor/minor"

- The vast majority of the "mixed" songs (41 / 50) are *melodically* minor and *harmonically* major.

Do these categories map on to conventional genres? To some extent.

Heavy metal songs tend to be minor/minor (e.g. Steppenwolf's "Born to be Wild"); soft-rock and pop songs are mostly major/major (Elton John's "Your Song").

Songs in the "minor/major" category tend to be of several types –

- 1950s rock'n'roll (Elvis Presley's "Hound Dog")

- blues-influenced rock (Beatles' "Can't Buy Me Love," Rolling Stones' "Honky Tonk Woman") 

- soul / R&B (Aretha Franklin's "Respect," the Isley Brothers' "Shout")

Our clustering analysis points to some kind of “major/minor” spectrum as an important dimension of variation in rock. But rock “minor” is quite different from classical minor. Also, a significant proportion of songs are minor melodically but major harmonically.

Key-Finding

We've also experimented with various *key-finding* algorithms for our corpus. For this purpose we split the 200-song corpus into two 100-song sets, one for training (parameter setting) and the other for testing.

Here we just define "key" as a tonic, not categorizing songs as major or minor.

One simple method (similar to methods used for common-practice music—e.g. Krumhansl, 1990):

- Treat the scale-degree distribution of the corpus as a "key-profile."

- Rotate the key-profile values to create specific key-profiles for different keys.

- In identifying a key for a song, find its pitch-class distribution; then choose the key whose profile best matches that distribution. (This can be done using correlation, cross-entropy, or other methods.)

The (harmonic) pitch-class distribution for “Hey Jude” (red), along with the F major key-profile (blue) – a good fit:

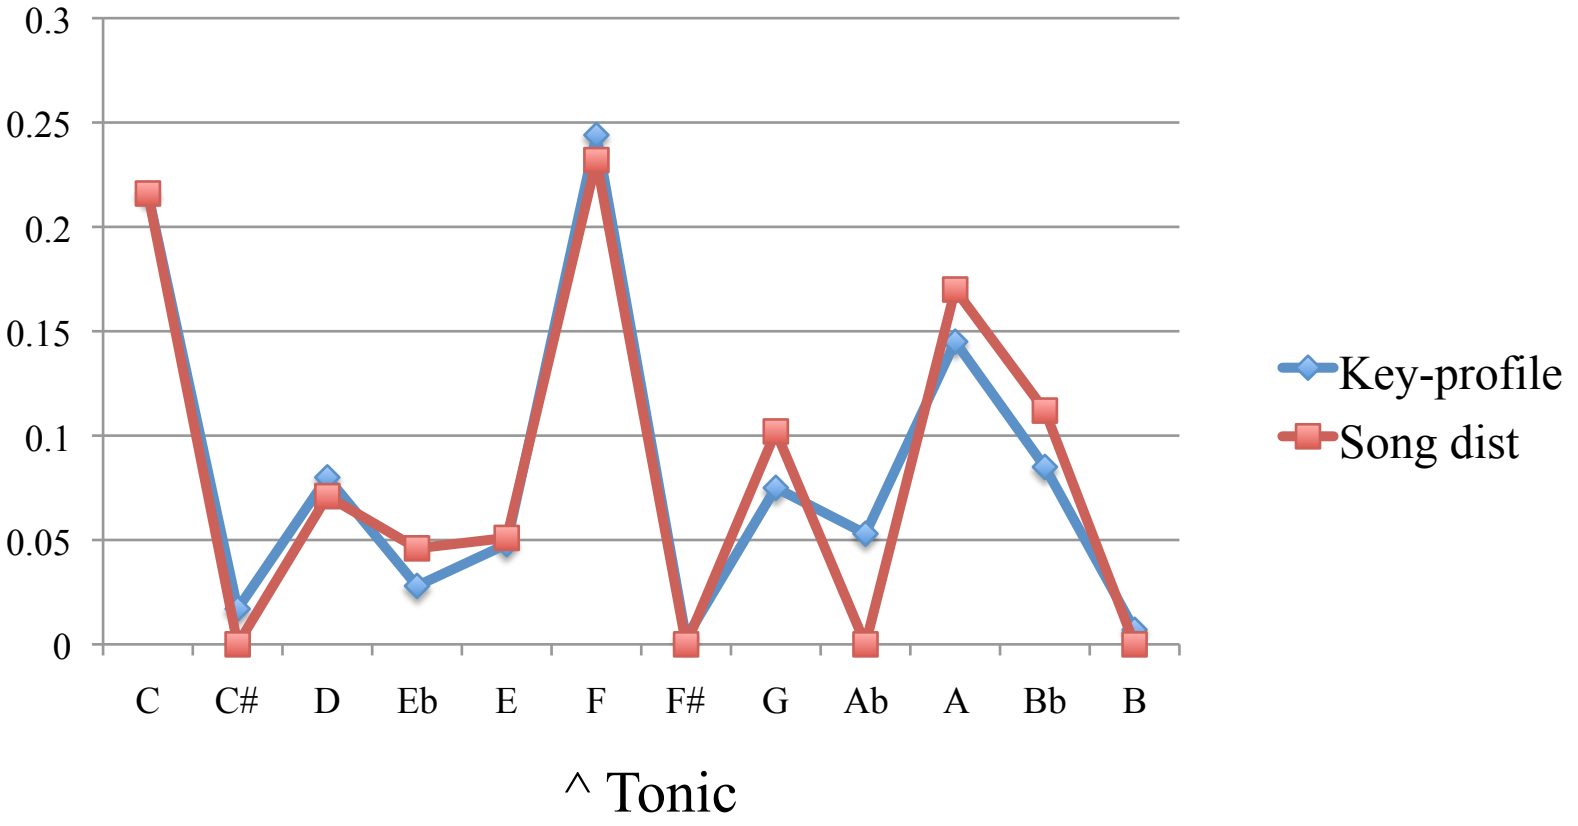

...and now with the F# key-profile – a poor fit:

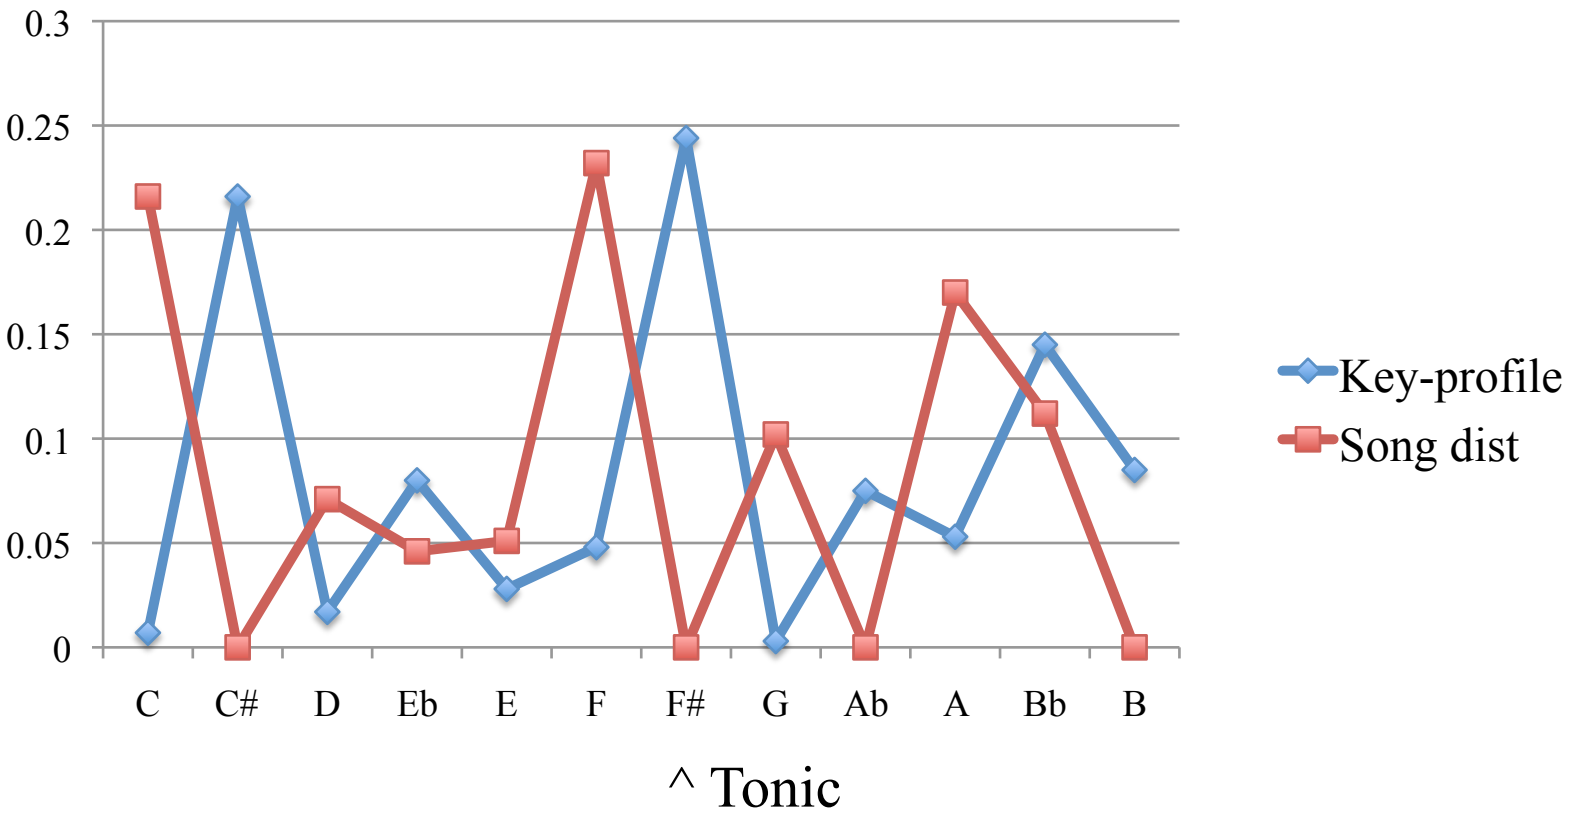

We experimented with this, using both the melodic and harmonic scale-degree distributions, singly and in combination.

The results were OK: at best, we achieved correct results on 86 of the 100 songs (using a combination of the two distributions). But we thought we could do better.

It's been suggested that the most likely tonic for a song is likely to be one that occurs frequently as a chordal root, and in particular, at metrically strong positions (Temperley, 2001; Stephenson, 2002).

C G | C G | C G | (more likely in C)
G C | G C | G C | (more likely in G)

Since our harmonic analyses represent roots and metrical positions, we were able to incorporate this information into a key-finding algorithm – favoring a key whose tonic chord occurs often, and especially on odd-numbered downbeats.

Combining this with the melodic scale-degree profile, we achieved the correct result on 97 of the 100 songs.

The moral: Key-finding in rock can be done most effectively by considering harmonic information directly, and in particular, the *metrical placement* of harmonies.

Harmonic progression

Classical music has strong norms regarding what chord can move to what other chord. ii-V and V-I are "good" progressions; V-ii and ii-I are not. What about rock?

Some theorists have argued that the rock harmony is largely governed by the norms of classical harmony (Everett, 2008).

Others have argued the the opposite norms apply – for example, ascending-5th motions are favored over descending 5ths (Stephenson, 2002).

Which view is correct?

We examined the frequency of different chords and chord progressions in our corpus. The overall frequency of different roots:

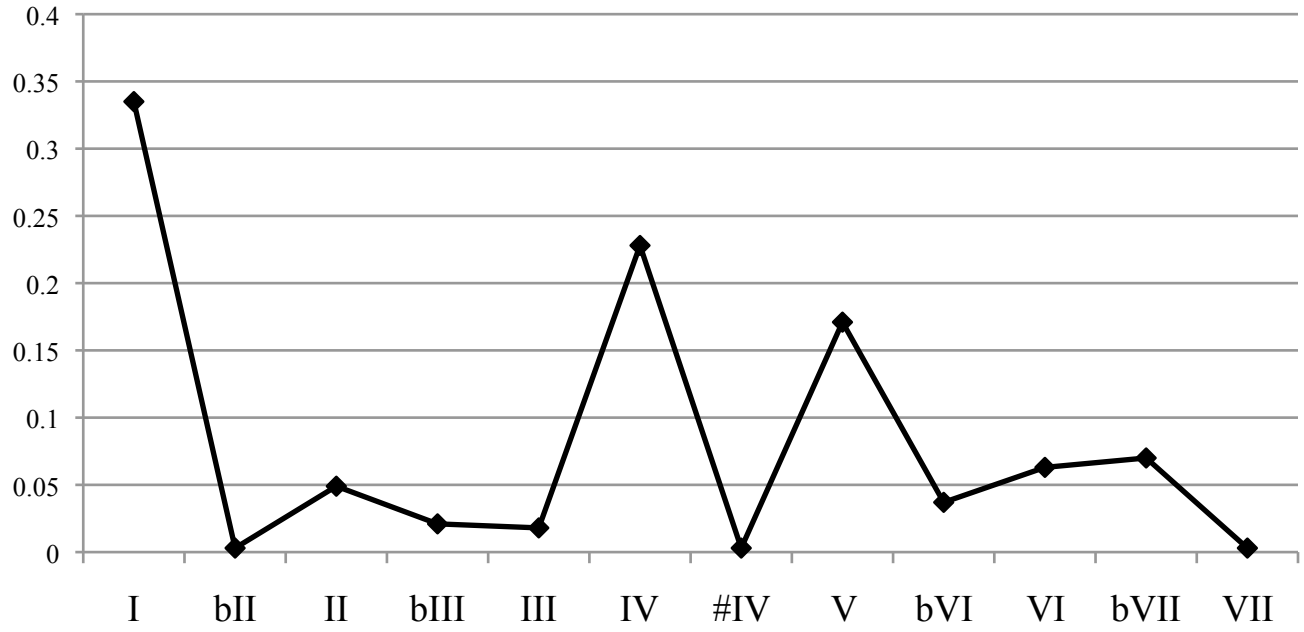

...so IV is the most common chord after I. (This is unlike common-practice music, where V is most common after I.)

The frequency of different root motions:

Interval	−m2	+M3	−m3	+M2	−P4	–	+P4	−M2	+m3	−M3	+m2	TT
Instances	162	326	454	1,384	2,220	–	2,266	1,386	410	412	113	20

The frequency of different root motions:

Interval	−m2	+M3	−m3	+M2	−P4	–	+P4	−M2	+m3	−M3	+m2	TT
Instances	162	326	454	1,384	2,220	–	2,266	1,386	410	412	113	20

Ascending fifth (-P4) and descending fifth (+P4) are about equally common. The same is true for asc/desc thirds, and asc/desc seconds. So, neither classical norms nor "opposite" norms seem to predominate.

The frequency of root motions is inversely related to their "circle-of-fifths" distance.

We also looked at the (proportional) frequency of each chord in "post-tonic" position (immediately following I) and "pre-tonic" position (immediately preceding I).

These frequencies are similar to the overall frequencies of the chords. IV is the most common chord in both pre-tonic and post-tonic position, and overall.

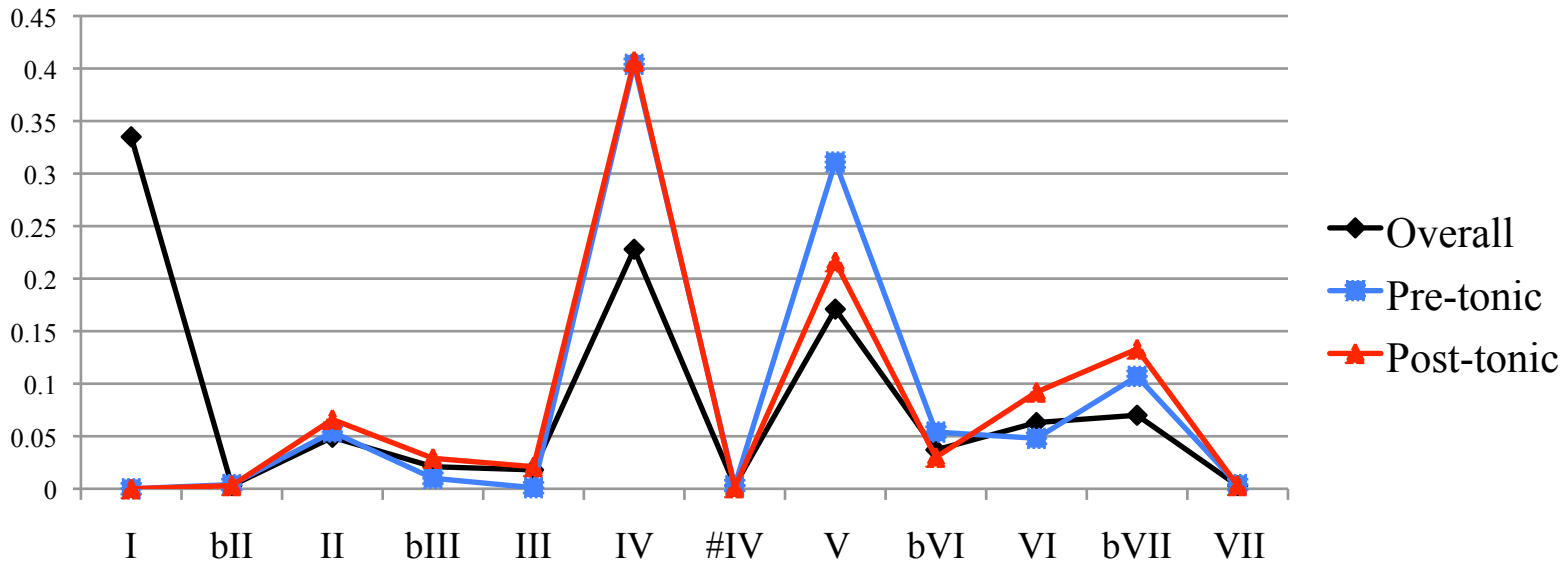

In short: While there are certainly preferences for some chords over others, the temporal *asymmetries* of classical harmony do not seem to be as present in rock. IV-I is as common as I-IV; ascending fifths are as common as descending fifths.

Changes over time

We can also use our data to examine changes over time.

We wanted to extend this investigation to very recent music. The Rolling Stone "500" list basically only goes up to 2000. So recently, Adam Waller and I have added another 40 songs from Rolling Stone's list of the "100 Best Songs of the 2000s". (We've transcribed the melodies; still working on the harmonic analyses.)

With regard to harmony, we examined the distribution of chord roots in each decade, the 1950s through the 1990s.

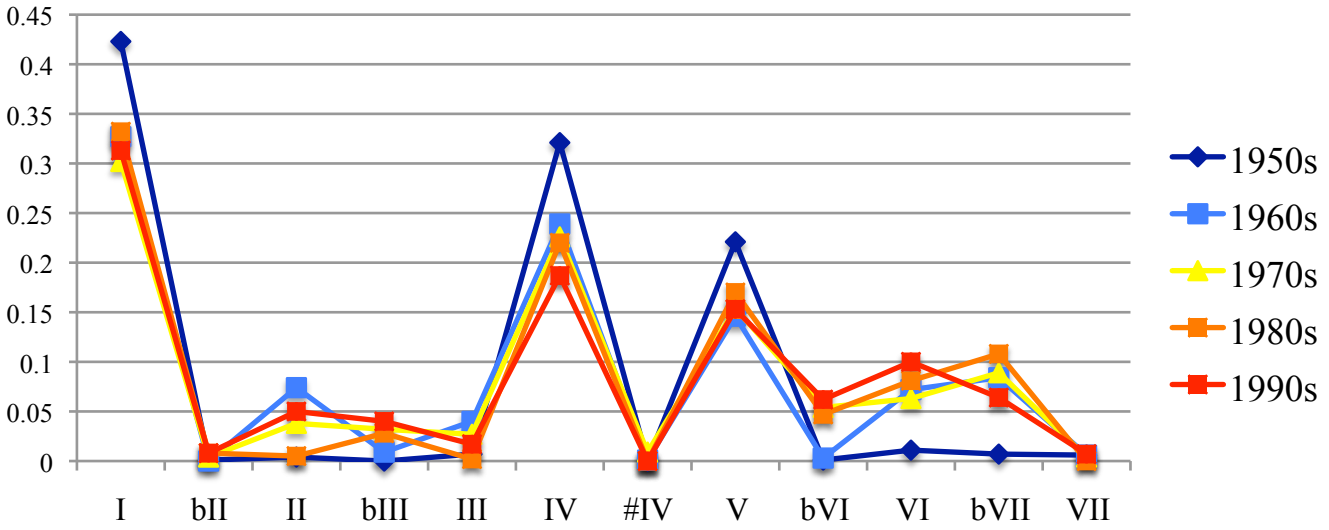

The most noticeable trend is that rock harmony was very simple during the 1950s (just I, IV, and V basically) and then got more complex after that.

In terms of melody, we were curious about the evolution of the "major/minor" distinction. (Here we added in the 40 songs from the Rolling Stone "2000s" list. 10 were rap songs and had no melodic data!)

To make it simple, we simply classified any song in which 3 is more common than b3 as "major"; songs in which b3 > 3 are "minor".

We divided the songs into two historical categories: 1950-1989 and 1990-2009.

The distributions for the "major" category:

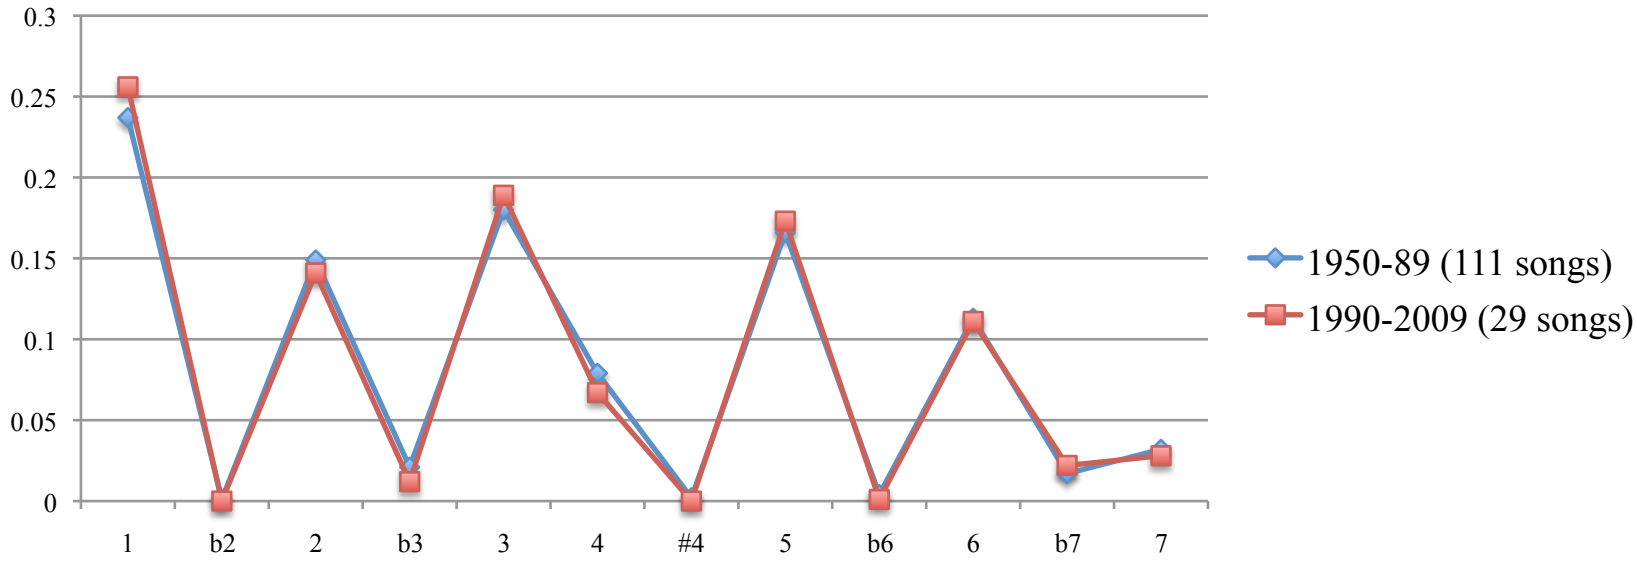

Almost identical between the two periods. (Yes, there are two lines there!)

For the minor category:

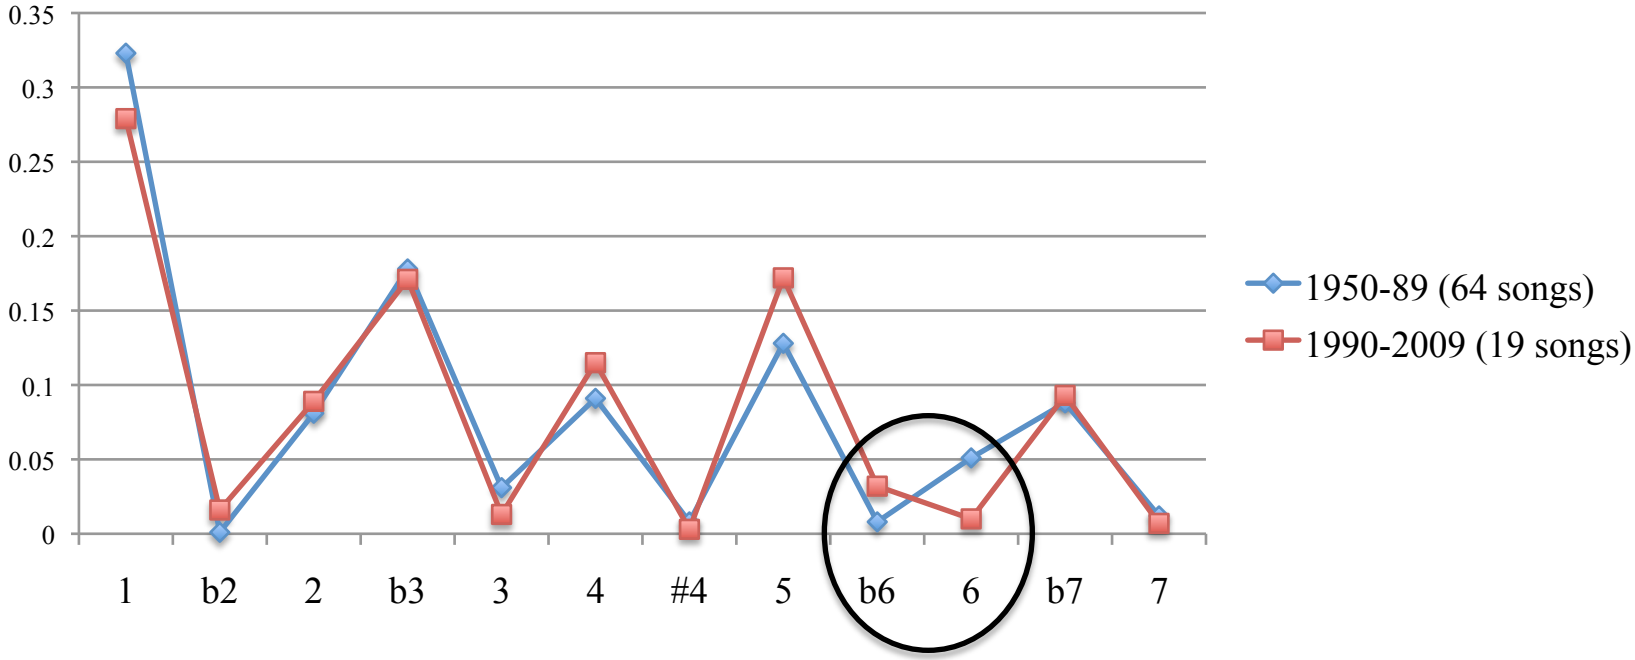

Mostly very similar, with one exception: in the early songs, 6 > b6; in the later ones, it's the reverse.

For the minor category:

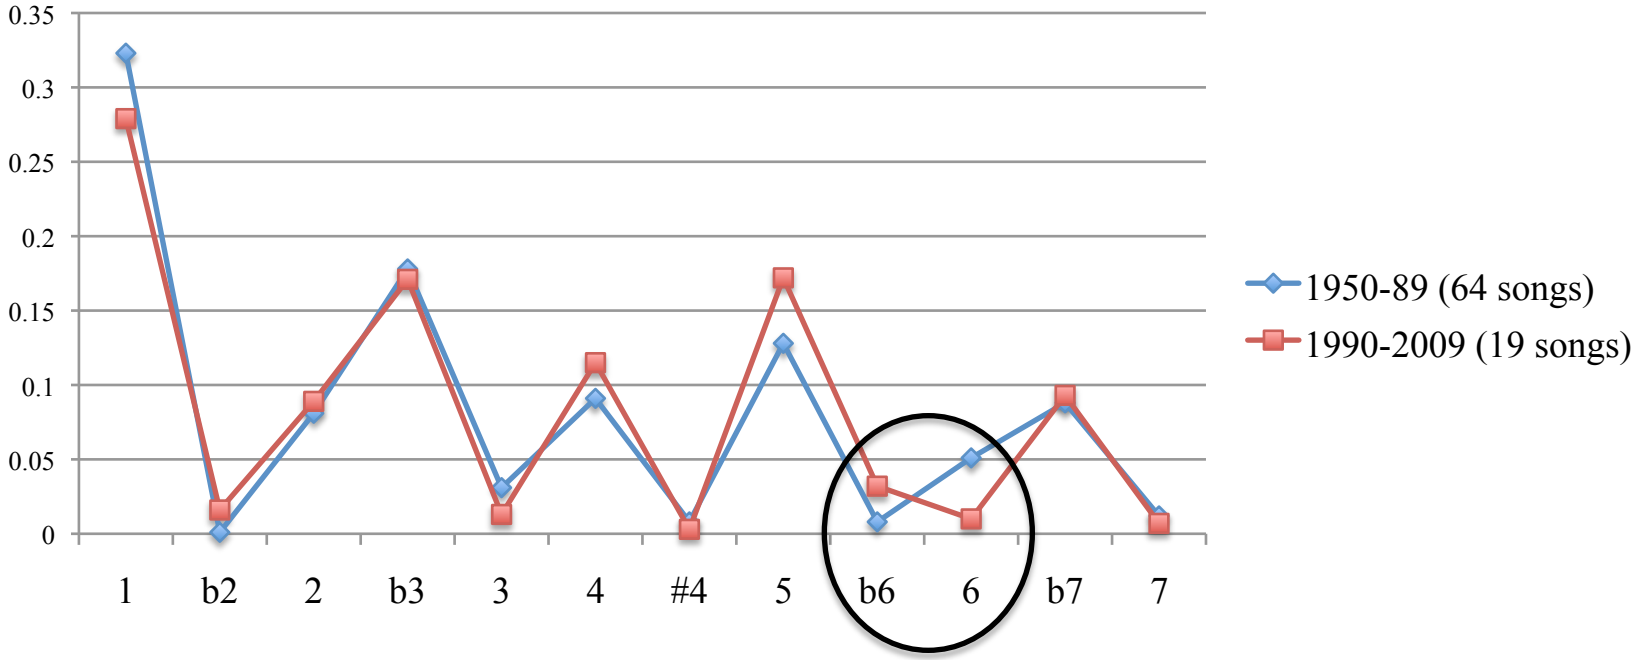

A small detail, but it seems significant. What does it represent?

- The declining influence of the blues?
- The return of "classical" minor?
- A more "negative" emotional tone in more recent music?

.....Requires further exploration!

Stay tuned for further explorations!

Visit www.theory.esm.rochester.edu/rock_corpus to download the corpus, see our publications, and get further information.

Thank you for your attention!

References

de Clercq, T., and D. Temperley. 2011. A corpus analysis of rock harmony. *Popular Music,* 30, 47–70.

Everett, W. 2004. Making sense of rock's tonal systems. *Music Theory Online 10.4.*

Krumhansl, C. 1990. *Cognitive Foundations of Musical Pitch.* New York: Oxford University Press.

Moore, A. 1992. Patterns of harmony. *Popular Music,* 11, 73–106.

Stephenson, K.. 2002. *What to Listen for in Rock: A Stylistic Analysis.* New Haven: Yale University Press.

Temperley, D. 2001. *The Cognition of Basic Musical Structures*. Cambridge, MA: MIT Press.

Temperley, D., & de Clercq, 2013. Statistical analysis of harmony and melody in rock. *Journal of New Music Research*, 42, 187-204.