

Typical Chords in Typical Song Sections: How Harmony and Form Interact in a Corpus of Rock Music  
by Trevor de Clercq

Hello everyone. There is no paper handout for this talk, but a PDF of the slides can be downloaded at the address shown at the bottom of the screen.

[NEXT] My talk today will be divided into four parts, and you can track the progress of the talk in the lower left-hand corner of the screen. I should mention that in this talk, I will use the term "rock music" in a broad way, to encompass rock, country, blues, rap, metal, soul, and other commercial music styles. OK, Let's begin.

== Introduction ==

In modern analyses of rock music, a song is partitioned via a limited set of labels, such as intro, verse, chorus, etc. Because these form labels are so ubiquitous, a central task for music theorists has been to explain what these terms mean and what features typically associate with them. For example, consider the form label "chorus". [NEXT] Generally speaking, theorists describe the chorus as the "focal" point of the song, as more "memorable" and having more "energy" than other song sections. [NEXT] A chorus is also characterized as typically having a thicker texture, with an increase in volume and instrumentation, often through the addition of background singers. [NEXT] The lyrics of a chorus are said to typically deliver a "general message," include the title of the song, and repeat upon future iterations of that chorus in the song. [NEXT] In terms of pitch, chorus melodies are said to often be less pentatonic, more coordinated with their supporting harmony, and have slower rhythms. As you can see, there are many aspects that play into our understanding of just this single form label. [NEXT] Many of these aspects are exemplified in the song "Just the Way You Are" by Bruno Mars. Let's listen [NEXT].

Notice that the harmonic content in this song does not really help distinguish the chorus from the verse section, since both share the same chord progression. Many rock songs are structured this way, built off a single repeating chord progression, and so it often appears that harmony plays no role in distinguishing between different section types. Perhaps in part for that reason, existing descriptions of how harmony and form interact in rock have been rather general, if not sometimes conflicting. [NEXT] For instance, Everett states that a chorus typically has "more dramatic" harmonies than a verse. Yet at the same time, the harmonies of a chorus section are said to be relatively stable. Bridge sections, in contrast, are said to often have "complex" chord changes. Note that these are very generic statements; it is not clear exactly what constitutes complex, dramatic, or stable harmonic progressions.

[NEXT] The role of tonic may play an important part. For instance, Everett states that verse and chorus sections typically prolong the tonic harmony. Endrinal and Stephan-Robinson make similar statements, although they seem to imply – somewhat in contradiction to Everett – that tonic plays a more central role in chorus sections than it does in verse sections. [NEXT] With regard to bridge sections, Neal notes that they typically explore non-tonic areas, such as the subdominant or flat-side keys. One aspect that has been discussed in somewhat more detail is the chords that typically begin or end song sections. [NEXT] Stephenson, for example, states that most verses begins on tonic while other song sections typically start off-tonic. [NEXT] Other authors have noted that chorus sections

usually end on tonic. [NEXT] In Summach's dissertation, we find statistical evidence that supports the notion that tonic more often starts verse sections and more often closes chorus sections, although many songs also thwart these trends.

These last two slides have summed up the bulk of current scholarship on how harmony and form interact in rock music. Undoubtedly, much more remains to be discovered. One of the main challenges to theorists in this regard, I believe, has been that fact that section labels reside entirely within the realm of perception, and perception is a messy thing. In essence, section labels can be considered what cognitive scientists refer to as "natural categories." Natural categories are complex, and we understand them not through neat and tidy definitions but rather through a complicated web of various attributes [NEXT], none of which are required for category membership. For example, when we hear a louder section, we are apt to hear it as the chorus because many passages we have previously understood as chorus sections are louder. Yet not all attributes associated with a section label appear in all instances of that label. In the chorus of the Bruno Mars song, for example, there were no background singers, even though the addition of background singers is something we associate with chorus quality.

My goal in this paper is thus to provide more concrete data, using statistical tools, as to what sorts of characteristics we can find in the harmonic domain that associate with the typical formal sections in rock music. In so doing, I hope to shed more light on those harmonic factors that potentially influence our perception of these sections in rock music.

== Methodology ==

[NEXT] Today's paper follows up on corpus work that Davy Temperley and I published in 2011. In that paper, we analyzed rock harmony in only a very general way. Our corpus was drawn from Rolling Stone magazine's list of the "500 greatest songs of all time." [NEXT] Here are a few excerpts from the list; as you can see, the stylistic range is rather broad – from rock to country to rap to metal. That being said, the complete list is weighted somewhat towards the 1960s and 1970s. Our 2011 article reported data from a subset of only 100 of these songs, but we have expanded the corpus to now include 200 songs total.

[NEXT] We each independently analyzed the 200 songs by ear, and so each song has two separate analyses. The harmony was encoded using a recursive notation, as shown here for the song "Da Doo Ron Ron". In our notation, a dollar sign indicates a variable that is defined elsewhere in the encoding. So, for example, the entire song, which is denoted by the letter "S", includes sections such as "In" (which stands for "intro") and "Vr" (which stands for verse). We did not standardize these section abbreviations, although their meaning is usually self-evident. So although our main goal was encoding harmony, we also included a great deal of information about the formal structure of each song. Let's listen. [NEXT]

[NEXT] In our 2011 paper, we reported statistics on the number of instances for each chromatic root, as shown here. For example, bVII was found to be the most common chromatic root after I, IV, and V. [NEXT] These findings don't change much when we look at the distribution of chromatic roots for the 200-song corpus. I will not talk much more about chord instances today, however, because I don't think they are the best way to compare harmony between different song sections. [NEXT] For example, consider the

hypothetical verse and chorus sections shown here. (FYI, a dot simply means that the chord from the measure before continues through in that measure.) If we use "instances" as our assessment tool, we would say the verse has one instance of tonic and one instance of dominant, whereas the chorus has four instances of tonic and four instances of dominant. In both cases, half of the instances are tonic and half are dominant. But the verse is obviously more tonic-heavy, even though there are technically more "instances" of tonic in the chorus. A better assessment tool is the proportional duration of chords. So here, the verse would be 75% tonic and 25% dominant, whereas the chorus would be split 50/50. We did not report proportional or durational information in our 2011 paper, but that will be one of the primary domains I investigate in this paper.

[NEXT] Let's return to the the distribution of chromatic root instances for the 200-song subset. Note how tonic is about a third of all chromatic root instances, and that bVII is much more common than II. [NEXT] If we look at chromatic root durations instead, we see that tonic accounts for almost half of the time overall. As well, there is a much smaller difference between bVII and II. So we can see that, globally speaking, chord duration information provides a different perspective than chord instances.

A few more preliminary issues are worth addressing. As shown in this table, some chords are simply not that common. [NEXT] For instance, out of over 20,000 measures in the corpus, the VII chord comprises only 25 measures total. As the corpus is divided into smaller parts, such as verse and chorus sections, the total amount of time spent on a VII chord becomes even smaller. Consequently, it is difficult to achieve statistically significant results with regard to the behavior of uncommon chords such as VII. The results I report today thus examine only half of the possible chromatic roots: [NEXT] These are I, II, IV, V, VI, and bVII, i.e., the six most typical chromatic roots. This is not to say that other chords (such as III) do not potentially play into our perception of form. Rather, this corpus is simply not large enough to provide sufficient evidence in that regard.

Another consideration is the level of subjectivity in our analyses. [NEXT] For the sake of simplicity, our 2011 study pooled results from both of our harmonic encodings. We felt that approach was justified because there was a relatively high level of agreement between our harmonic analyses. For example, we agreed on the chromatic relative root of a chord about 92% of the time and the absolute root about 94% of the time. In terms of section labels, our level of agreement is nowhere near that high. It is hard to give a precise figure, since we were free to use any form label we saw fit. Nonetheless, I went through our analyses and grouped labels as best I could into more general buckets such as verse, chorus, bridge, etc. All told, my best estimate for our level of agreement is around 67%. [NEXT] Obviously, analyzing form has a much greater subjective element than analyzing harmony. I will therefore not report pooled results; instead, I will report results from both of us individually. When those results agree, we have a stronger case for statistical validity.

[NEXT] As a final issue, be aware that not all songs make use of all section labels. ("DT" here stands for David Temperley, "TdC" for me.) For example, while almost every song had verse material, only about two-thirds had chorus material, and only about one-third had bridge material. As a result, I will focus on how harmonies interact with only these three most typical formal areas. Other section labels are simply too rare in our corpus to provide any meaningful data. [NEXT] For instance, only about 6%

of the songs we analyzed had a prechorus. This should not actually be too surprising, since – as Summach has described in his 2011 article – prechorus sections were not very prevalent until the late 1980s.

== Results ==

Now for some results. [NEXT] This table shows the proportion of time spent on the six most typical chromatic roots overall as compared to the proportion of time spent on these chromatic roots in verse material. I should mention that the p-values used here and throughout my talk were calculated using paired sample t-tests, with each individual song acting as a matched sample pair. As you can see, there is some evidence that saying a musical passage is a verse implies we may expect less II chord than other passages in the same song. But otherwise, verse sections do not seem to differ significantly from the rest of the musical material.

[NEXT] We find bigger differences when looking at material categorized as chorus or bridge. As this table shows, a chorus label implies a significantly smaller percentage of time, on average – about 40% instead of about 50% – will be spent on tonic than in passages that are not chorus material. Not surprisingly, as time shifts away from tonic, time spent on other chords rises in proportion to compensate. The dominant chord seems to be the strongest recipient of this weight shift, although evidence for this effect is not terribly strong, and all other typical chords seem to occupy a greater proportion of chorus material.

[NEXT] The greatest difference is found with the bridge label. With bridges, an even smaller proportion of time appears to be spent on average on tonic – at just over 30%. Again, other chords occupy a greater proportion of bridge material to compensate for the decrease in time spent on tonic, with the most significant relative amount of change occurring for time spent on II chords.

Of course, the real question is whether these changes in proportion hold up when comparing one section label to another. [NEXT] In this table, we see the results of that comparison, which I have simplified into a generic “less or more” effect. Indeed, there is strong evidence that bridge and chorus sections spend a significantly smaller proportion of time on tonic than do verse sections. We also find some evidence that bridge sections spend a significantly larger proportion of time on II than do verse sections. Some other evidence is scattered throughout this chart, but it is not very strong and holds true for only one analyst. The most notable finding, perhaps, is the lack of evidence for any difference between chorus and bridge sections. I'll say more about that in a moment.

At this point, it might help to listen to some examples, if only to make these findings a little more tangible. [NEXT] As shown here, the song “Every Breath You Take” has a verse made of 50% tonic and a bridge made of only 25% tonic. Consequently, the verse sounds more like a place of rest and the bridge like a place of motion. Of course, the hypermetric locations of tonic, subdominant, and dominant play into these feeling as well; the proportion of tonic is simply one aspect of that. Let's listen [NEXT]. Consider also the song “I Still Haven't Found What I'm Looking For”. The verse is made of 75% tonic while the chorus is only 50% tonic. Because so much time is spent on tonic in the verse, it does not demand our attention as strongly as does the chorus that follows. Let's listen [NEXT]. We could also say the chorus draws our attention because the harmonic rhythm increases here, too. Certainly, that is another factor,

and this leads to another way of looking at chord durations.

[NEXT] For example, consider the hypothetical verse and chorus sections shown here. Both eight-bar sections have four bars of tonic total, so each is 50% tonic. But in the verse, there is less harmonic activity, in that the first four bars are just one long span of tonic, whereas in the chorus, there is a chord change every two bars. The difference here is the average duration for each chord, which is calculated by taking the total number of measures of a chord and dividing it by the total number of instances of that chord. In the verse, the average duration of tonic is four bars; in the chorus, it's two bars. So looking at average durations gets more at the issue of harmonic rhythm.

[NEXT] In this next table, we see data on average chord durations for the corpus overall. (This is pooled data, since section labels are not a factor.) I have separated out the tonic chord as its own category, because its average durations are significantly different than all other chords. Once we remove outliers (that's the trimmed average), we find the average duration for a tonic chord to be about twice that of other chords. Nonetheless, the most frequently occurring duration for any chord is a single bar, as shown by the mode statistic. In the next few slides, we will look at if and how these average chord durations change significantly given the section label.

[NEXT] Let us first compare verse and chorus sections, as shown in this table. In both our analyses, the average durations for chords overall are shorter in chorus sections than in verse sections. But this effect turns out to be primarily due to differences in average duration for the tonic chord. We find no significant difference between verse and chorus sections with regard to average durations for non-tonic chords.

[NEXT] The opposite situation is found when comparing verse and bridge sections. Although the average chord durations for bridges are longer in all categories, there is not any significant evidence that this is true for tonic chords. Rather, we find evidence that non-tonic chords typically last longer in bridge sections than verse sections.

Given these results, we may expect significant differences in chord durations between chorus and bridge sections. [NEXT] But while we find that the average chord durations in bridges are indeed longer than in choruses, the t-tests fail to recognize that any of these effects are significant. I did some post-hoc analysis as to why this might be, and it appears that there are about an equal number of cases in which chord durations in the chorus are longer than the bridge and vice versa, even though the average durations themselves are somewhat skewed.

Here again, it is worth pausing to listen to some examples so as to make these findings more tangible. [NEXT] In the song "Just What I Needed", we can see that the verse and chorus sections have the same proportions of tonic. Specifically, there are two bars total of tonic in the verse and four half-bars of tonic in the chorus, which gives both sections a 25% tonic load. But in the chorus, where each tonic lasts only about half a bar, the average duration of this tonic is much shorter than in the verse. Of course, all the chords in the chorus are shorter, so there are certainly other factors at work in this particular song as well. Let's listen [NEXT].

[NEXT] "I Saw Her Standing There" is a good example of a song in which the

non-tonic chords in the bridge have longer average durations than the non-tonic chords in the verse. Here we find that the bridge avoids tonic entirely, which is surely one factor that makes the bridge seem unlike a verse. But note as well that the durations of the chords in the bridge, each of which lasts two bars or more, are much longer than those in the verse, where the non-tonic chords only last a bar. Let's listen [NEXT].

[NEXT] Earlier, I mentioned that one aspect apparently important to theorists is the opening or closing harmony of a section. Consequently, I show in this table the frequency of the most common chromatic roots that start or end typical song sections. [NEXT] As you can see, Temperley and I analyzed most verse sections as beginning on tonic, with about half ending on tonic. [NEXT] Tonic is somewhat less common as an opening chord in chorus sections, although most chorus sections still begin with tonic. Like verses, about half the choruses end on tonic. [NEXT] The biggest difference is found with bridge sections. In my analyses, the most common chord to begin a bridge is IV, while for Temperley, it was tonic. In both cases, we find that most bridges do not begin on tonic. [NEXT] Bridge sections also seem to be most likely to begin somewhere other than I, IV, or V. [NEXT] In terms of final chords, bridge sections most often end on a dominant.

Many other statistical tests could be run on our corpus with regard to how form and harmony interact. [NEXT] For instance, here is data that looks at the distribution of major, minor, diminished, and augmented chords as well as inverted and root-position chords, for typical song sections. It seems that there are more major chords in chorus sections, for example, but fewer inverted chords. Given limitations of time, though, I will have to "table" such additional findings for another paper.

== Discussion ==

To summarize the main findings I reported today, [NEXT] this corpus shows strong evidence that verse sections spend a greater proportion of time on tonic than do chorus or bridge sections. There is also some evidence that bridge sections spend a greater proportion of time on II chords than do verse sections. There is no evidence, however, of any significant differences in proportion of time spent on any chord when comparing chorus and bridge sections.

[NEXT] In terms of chromatic root durations, evidence exists that, on average, tonic chords have shorter durations in chorus sections than in verse sections. In contrast, we find evidence that non-tonic chords have longer durations in bridge sections than in verse sections. But there is no evidence, once again, for any significant differences in average chord durations when comparing chorus and bridge sections.

[NEXT] Finally, we see that the vast majority of verse sections begin on tonic, and about half end on tonic. Chorus sections are found to less frequently begin on tonic, although most still do; and like verse sections, about half of the chorus sections end on tonic. In contrast, bridge sections most often begin off-tonic, with most ending on dominant.

I should say that one important caveat is important to keep in mind when interpreting these results. In particular, recall that our 200-song corpus is somewhat skewed towards the 1960s and 1970s. A corpus of more modern songs might show different associations between form and harmony. In fact, many authors, such as Covach and Summach, have discussed how songwriting

strategies have changed over time. For example, the AABA form was common during the 1950s and 1960s but became much less common by the late 1970s. So although Temperley and I, for example, used the term "bridge" somewhat indiscriminately in terms of the large-scale song form, it is entirely possible that the harmonic structure of bridge sections in AABA songs differs significantly from bridge sections in verse-chorus songs.

That being said, I believe the findings I presented today possibly shed some light on these historical issues. Foremost, there seems to be no evidence in terms of chord proportions or average durations between bridge and chorus sections. In other words, the general harmonic structures of these two section types may be very similar. This offers one possible explanation as to the mechanisms of the shift from AABA to verse-chorus forms during the early decades of rock. [NEXT] This chart shows how just such a shift could occur. The critical step is in the change from the B section acting like a bridge to the B section acting like a chorus. Given the apparent similarity in general harmonic construction between bridge and chorus sections, this shift may be as simple as changing the final chord in the B section from dominant (typical of bridge sections) to having the final chord in the B section be tonic (more typical of chorus sections).

There are, in fact, many songs in the history of rock that show just how permeable or ambiguous this distinction between bridge and chorus quality can be. [NEXT] Consider, for example, the song "Can't Buy Me Love." The middle section, which is shown at the bottom of the slide, occurs as the B section of an AABA form, and its ending on a dominant chord is similar to traditional bridges from this era. Yet at the same time, this passage sounds very much like a chorus, so much so that George Martin – the famed producer for the Beatles – has referred to this section as the chorus of the song. Let's listen [NEXT].

In conclusion, we should be aware that there are many factors that play into our perception of whether a musical passage in a rock song is a verse, chorus, bridge, or some other section type. Harmony is an important domain, undoubtedly, but other domains may support or contradict the evidence that harmony conveys. For instance, just because one part of a song sits on tonic for an extended period of time does not mean that we must automatically categorize it as a verse. The results I presented today are not meant to act as an analytical system, but rather to help understand the analytical process. Ultimately, the factors that inform our expectations and perceptions about form in rock music are perhaps mostly unconscious and difficult to fully explain, which is perhaps in part why this music continues to fascinate us.