LET ME TAKE IT DOWN
THE MATHEMATICS BEHIND THE MOST FAMOUS EDIT IN ROCK 'N' ROLL[1]

Jason I. Brown, Dalhousie University and Robert Dawson, St. Mary's University

The song *Strawberry Fields Forever* was a landmark in the musical landscape of The Beatles. Lennon[2] wrote the song while vacationing in Spain in 1966, and the fresh air must have done him good. A number of takes of the song were done in the studio. John was in a quandary after listening and relistening to the takes. He liked the first part of Take 7, and the second part of Take 26, and passed the problem of melding the two together onto the record producer, George Martin. Martin was in a bit of a fix, as the keys were different. Take 7 was in the key of A, while Take 26 was up a whole tone, in the key of B. A further problem Martin encountered was that the tempos of the two takes were different. Take 7 was at a tempo of about 85 beats per minute, while Take 26 was recorded at about 107 beats per minute.

Today, digital sound editing software allows anybody to change pitch and tempo independently. But in the 1960's, the state of the art was varispeed audiotape technology, which allowed the playback (or recording) speed to be selected precisely. A faster playback would increase both the pitch and the tempo of Take 7; slowing the tape would decrease both for Take 26. Could some combination of these be used to fix both problems?

To answer this question, we need to represent pitch and tempo in comparable fashions. Tempo is usually measured in beats per minute (BPM); it usually ranges from about 40 (a very slow beat) to about 200. The human ear cannot easily distinguish changes of less than about 7%; experiments have shown still lower sensitivities for increases in tempos below about 100 BPM. Apart from this, the "just noticeable difference" is roughly constant across moderate to fast tempos, with slight differences between experienced musicians and nonmusicians (see figure 1). Larger differences are also roughly proportional; the effect of putting triplets into a 4/4 beat is similar at any speed. We conclude, then, that subjective tempo is logarithmic.

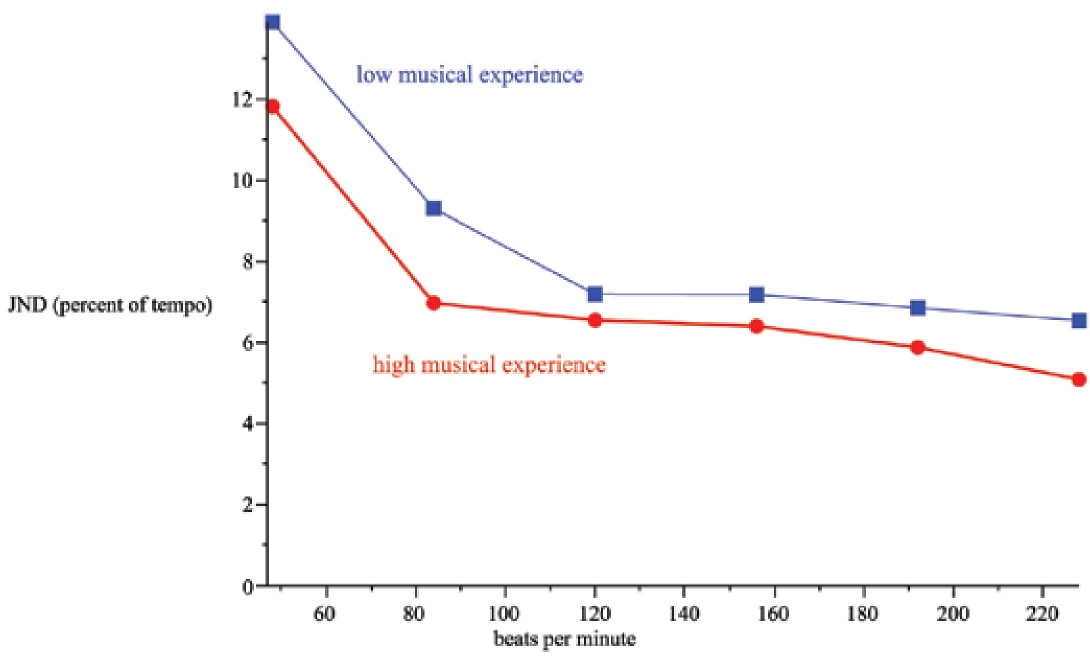

Figure 1: Just Noticeable Differences for Increasing Tempos (based on [1]).

While musicians usually represent pitch using note names and octaves, for scientific purposes it is more usefully represented in "frequency", measured in cycles per second. To a scientist, tempo is also a frequency, though much slower (much as radio, microwaves, light, and X-rays are all electromagnetic radiation). While a musician would not usually need to do it, tempo (in BPM) can be converted to pitch (in cycles per second) by multiplying by 60; dividing converts in the other direction. As with light and radio, the ranges do not really overlap; the lowest note on a piano, A0, corresponds to a dizzy 1650 BPM! (Speed drummers do approach this range.)

Musical intervals that we hear as identical at different pitches are not the same number of cycles per second wide; they represent the same ratios. In particular, an octave always represents a frequency ratio of

(1) We dedicate this article to what would have been John Lennon's 70th birthday this year.
(2) While John didn't excel in school when he was young, one of his favourite authors was Lewis Carroll, which was the pseudonym of Charles Lutwidge Dodgson, a mathematician. John's appreciation for the logical puzzles and word play in Carroll's books pops up in a number of Beatles' songs, most notably in *I Am The Walrus*.

LET ME TAKE IT DOWN
THE MATHEMATICS BEHIND THE MOST FAMOUS EDIT IN ROCK 'N' ROLL *continued*

2:1. A semitone (the difference in pitch between two adjacent keys on a piano keyboard) is always a ratio of $\sqrt[12]{2}^{i}:1 \approx 1.06$. So pitch perception, too, is logarithmic. Conveniently for George Martin, this means that if a tune is played back at a different (but constant) speed, the apparent intervals do not change, so it sounds like the same tune.

Interestingly, we discriminate pitch much better than tempo. The JND for tempo represents the same ratio as the easily-distinguished semitone interval. At some frequencies the human ear can distinguish notes a twentieth of a semitone apart or less, though again this deteriorates at low frequencies.

Moving frequencies up by *i* semitones (with *i* being positive when you move up and negative when you move down) corresponds to multiplying frequencies, and hence tape speed, by $\left(\sqrt[12]{2}\right)^{i}$. So, for example, to change the key of Take 7 from A to B, the tape speed would have to be multiplied by $\left(\sqrt[12]{2}\right)^{2} = \sqrt[6]{2}$, which is about 1.12 (that is, 12% faster). Unfortunately, this doesn't get the tempo right! Speeding an 85 BPM track up by this ratio leaves it at about 95 BPM, still significantly slower than the 107 BPM of Take 26.

Because of the human ear's better pitch acuity, leaving some mismatch in the pitch in order to get a closer tempo match was (as George Martin realized) not an option. So what was there to do? In his book *All You Need Is Ears* [2], Martin wrote "I thought: If I can speed up the one, and slow down the other, I can get the pitches the same. And with any luck, the tempos will be sufficiently close not to be noticeable. I did just that ..." He moved Take 7 gradually up one semitone, starting from the beginning of Take 7 until the big edit point about 1 minute later, arriving at the key of Bb, and spliced in Take 26 at the appropriate point, down one semitone to the same key – a natural way to try to solve the problem by meeting in the middle. What happened to the tempos? The tempo of Take 7 moved up from 85 to $85 \cdot \sqrt[12]{2} \approx 90$ beats per minute, while Take 26 slows down from 107 to $107/\sqrt[12]{2} \approx 101$ beats per minute.

Not a perfect match, but close enough to fool many listeners. The change in tempos at the edit point, 11 beats per minute or about 12%, is over the just noticeable difference, which is approximately 9% at that tempo for most people (see figure 1). For more experienced and musical listeners, the just noticeable difference at that tempo is around 7%, and the change obviously bothered Paul McCartney [3], who has an impeccable sense of rhythm: "We could hardly hear the join, but its one of those edits where the pace changes slightly; it goes a bit manic for the second half of the song." Proportionally, it was just the same change as if Take 7 had been sped up by 12%, or Take 26 slowed down! Any change that makes the pitches match would leave the tempos differing by more than the JND.

Experimental data do suggest one solution — but not an artistically viable one. George Martin could have hidden the tempo difference if he elected to slow down the tapes more significantly, past the range of logarithmic response. If he had slowed one take by about an octave, and the other slightly more, the new tempo would have been around 43 beats per minute. For such slow tempos, the JND is much larger, and even Paul might not have noticed. Of course, then the song would change from a 4 minute song to one of 8 minutes, and indeed it would have felt like Strawberry Fields *Forever!*

So why did George Martin choose to move one take up one semitone, the other down one, when other choices would do? From reading his book, it seems most likely that he was happy to split the difference, key-wise, and the fact that the tempos were close was good enough. Sometimes luck and mathematics are both on the side of brilliant people.

I mean it must be high or low. That is, you can't, you know, tune in but it's all right.

That is I think it's not too bad.

Acknowledgements

This article was partially supported by a grant from the Natural Sciences and Engineering Research Council of Canada.

References

[1] M. Ellis, Thresholds of Detecting Tempo Changes, Psychology of Music 19 (1991), 164–169.

[2] G. Martin and J. Hornsby, *All You Need is Ears*, St. Martins Press, New York, 1979.

[3] K. Ryan and B. Kehew, *Recording the Beatles*, Curvebender, Houston, 2008.

DU BUREAU DU PRÉSIDENT *suite*

Que conclure, en particulier dans la perspective du plan à long terme du CRSNG ? La première chose est que nous avons utilisé les fonds qui nous ont été confiés de façon remarquablement efficace. Notre communauté s'est agrandie, et est désormais un acteur important dans le courant international du sujet; en bref, nous avons livré. Nous aurons besoin de plus de ressources- tous ces étudiants à nourrir, et toutes ces nouvelles carrières à développer. Nos institutions continueront à avoir besoin de soin et de maintien, et ceci n'est pas un acquis; MITACS, par exemple, arrive à échéance dans quelques années, dans la partie de son portefeuille qui subventionne les mathématiques industrielles. Finalement, une communauté de recherche en bonne santé présuppose une distribution assez large de fonds recherche- notre moteur doit avoir l'essence pour tous ses cylindres, pas seulement les cylindres numéro trois, quatre et six. Encore du pain sur la planche.