The Loudness War: Background, Speculation and Recommendations

Earl Vickers[1]

[1] STMicroelectronics, Santa Clara, CA 95054
earl.vickers@st.com

ABSTRACT

There is growing concern that the quality of commercially distributed music is deteriorating as a result of mixing and mastering practices used in the so-called "loudness war." Due to the belief that "louder is better," dynamics compression is used to squeeze more and more loudness into the recordings. This paper reviews the history of the loudness war and explores some of its possible consequences, including aesthetic concerns and listening fatigue. Next, the loudness war is analyzed in terms of game theory. Evidence is presented to question the assumption that loudness is significantly correlated to listener preference and sales rankings. The paper concludes with practical recommendations for de-escalating the loudness war.

1. INTRODUCTION

"Loudness war" is a term applied to the ongoing increase in the loudness of recorded music, particularly on Compact Discs, as musicians, mastering engineers and record companies apply dynamics compression and limiting in an attempt to make their recordings louder than those of their competitors. [1][2]

Given the incredible technological advances of the last half-century, one might expect that by now we should live in a musical paradise, with a thriving music industry and recordings of amazing depth, texture and dynamic range. Instead, the industry is in decline and "we're making popular music recordings that have no more dynamic range than a 1909 Edison Cylinder!" [1]

In fact, early acoustical recorders had a dynamic range of up to 20 dB [3], which is more than the range of most recent recordings. Figure 1 shows waveform displays derived from a 1909 Edison cylinder recording of "Down Where the Big Bananas Grow" [4] and "My Apocalypse" from Metallica's 2008 *Death Magnetic* CD; note that the former appears to have a much wider range of levels. When analyzed with the "TT Dynamic Range Meter" [5], the earlier recording yields a DR ("dynamic range") value of 15 dB compared to DR = 3 for the Metallica song.

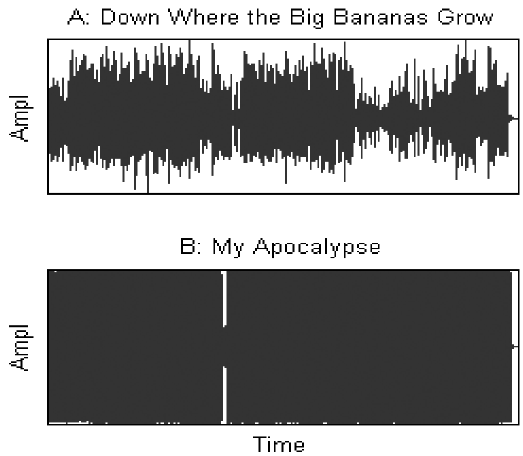

Figure 1. Waveforms, A: "Down Where the Big Bananas Grow" (1909 Edison Cylinder), and B: "My Apocalypse" (Metallica, *Death Magnetic*, 2008), left channel.

In the short term, louder tends to sound better[1], so over-compression (or "hypercompression") is commonly used to squeeze more and more loudness into the recordings. Hypercompression has been facilitated by the development of multiband compressors in the form of hardware boxes and software plug-ins. [1][2] It is widely believed that hypercompression may damage audio quality by removing dynamics, creating musical clutter and reducing the excitement and emotional power of the music. It is also believed that hypercompressed music may cause "listening fatigue," which can discourage continued or repeated listening. [2]

Note that this paper is not concerned with the final playback level, which can always be adjusted to the listener's taste; instead, it relates to the effects of processing designed to make recordings sound louder than others having the same peak level.

Despite a large number of articles in the popular press as well as some recent workshops, there have been few technical papers focused primarily on the loudness war, perhaps because this is a rather slippery subject in between art and science. The topic is awash in speculation, conjecture and unstated assumptions. This paper will attempt to explicitly state the assumptions, categorize the conjecture and provide additional speculation.

The paper is organized as follows. Section 2 provides some background, including a brief history of the loudness war as waged on several fronts such as vinyl records, radio, television and compact disc; it also reviews how multiband compression and the nonlinearity of the ear have helped fuel the loudness war. Section 3 examines the problems most often attributed to the loudness war, with a focus on aesthetic concerns and listening fatigue. Section 4 looks at the loudness war in terms of game theory; it also presents evidence suggesting that, in practice, loudness may be largely irrelevant to listener preference and commercial success. Section 5 recommends some specific de-escalation strategies, and Section 6 gives a summary, including suggestions for further research.

1.1. Terminology

In this paper, "loudness" and "dynamic range"[2] will generally be used in the colloquial sense of the terms. The word "compression" will refer to dynamics compression, not data compression from lossy codecs, unless specifically noted.

2. BACKGROUND

2.1. A Brief History of the Loudness War

This article will focus largely on CDs and downloadable music, but we will begin by looking at some of the other loudness wars that have been waged over the years.

2.1.1. Vinyl Records

Phil Spector was one of the earliest participants in the vinyl-era loudness war [2][8]; his "wall of sound" technique featured large groups of musicians, including guitarists playing in unison and two bassists playing in fifths, processed through an echo chamber. [9] The echo chamber's natural reverberation increased the loudness and density of the sound by boosting the average RMS level for a given peak amplitude (i.e., by lowering the crest factor). Basses playing in fifths may have triggered

[1] This applies primarily when comparing two versions of the same recording. When comparing two different songs, content differences may overwhelm differences in loudness, as discussed in Section 4.3.1.

[2] In general, the author prefers the terms "dynamics" or "dynamic spread" instead of the overly specific "dynamic range" for reasons noted elsewhere [6][7], but the latter term has the advantage of familiarity.

the psychoacoustic illusion of deep bass using a mild version of the “missing fundamental” or “virtual pitch” effect; it is conceivable that Spector liked this technique because it created a big sound on small transistor radios. At any rate, the large ensembles, reverberation and dense sound were all part of his “Wagnerian approach” to creating “little symphonies for the kids.” [9]

In the mid-1960s, the Motown record company adopted a standard called “Loud and Clear,” which used a number of methods to maximize the apparent loudness while maintaining clarity. Song durations rarely exceeded three minutes, primarily to obtain airplay, but also because longer durations required lower recording levels. For songs that increased in level during the course of the tune, loudness envelopes were applied to reduce the level in 0.5 dB steps; these reductions were seldom audible, since people cannot detect decreases in level as readily as they detect increases. In order to combat the poor high frequency response of AM radios and portable record players, and to mask the second harmonic distortion of the midrange (due to “tracing” distortion during mastering), levels were boosted in the 8-10 kHz region. Low bass frequencies were filtered out using a 70 Hz brick wall high-pass filter, and perception of the low bass was restored by boosting the second harmonic of the bass guitar. Finally, half-speed cutting was used to reduce recording amplifier distortion during disc mastering. [10]

Obtaining increased loudness with vinyl records required tradeoffs and difficult choices. Excessive bass levels could cause the needle to jump out of the groove [11]; as a result, louder signals required wider grooves and shorter playing times. As LPs began seeking longer durations, loudness became a lower priority. [2] In addition, since vinyl records could not be played in the car, there was less pressure to squash the dynamic range in order for the quiet parts to be heard over road noise [12].

2.1.2. Radio

Robert Orban used the term “loudness war” in a 1979 article discussing excessive compression and limiting for FM radio broadcast. The article reviewed the battle for ratings and the pressure on station owners to boost loudness at the expense of broadcast quality. Orban stated that “two or three years ago it seemed that many stations were finally realizing that better radio could improve ratings. And the major myth brought over from AM radio – that a louder signal, regardless of quality, attracts more listeners – appeared to be losing its strength.” But within a couple of years, he saw a new escalation of the loudness war, often powered by putting multiple compressors in the processing chain. [13]

Orban and Foti used the term “loudness war” again in a 2001 articled entitled “What happens to my recording when it’s played on the radio?” In the late 1990s, they began seeing recordings that had been pre-distorted by brute-force clipping to increase their loudness. When broadcast on FM radio, clipping does not succeed in increasing the on-air loudness; instead, it just exaggerates the distortion, due to the phase rotator in the typical broadcast processing chain. [14]

2.1.3. Television

The loudness war has not affected television as much it has radio, perhaps because viewers tend to choose TV stations primarily according to which program is playing. While there is a great deal of interest in achieving consistent loudness, level discrepancies are not necessarily based on each station trying to be louder than the next. Consumers often regard louder TV content as disturbing and undesirable, especially when it comes to commercial advertisements. Even though quiet, understated TV ads can be surprisingly effective by drawing viewers in instead of driving them away [15], commercials nevertheless often try to gain a loudness advantage through hypercompression [16]. Standards such as the EBU’s R128 recommendation are aimed at minimizing unwanted changes in loudness [17].

2.1.4. Movies

While there is little apparent advantage to be gained from a loudness war between movie theatres, since the viewers have already paid and constitute a more or less captive audience, trailers in particular are sometimes considered too loud. In general, though, standards such as Dolby’s monitor level calibration recommendation have helped prevent a loudness war by enabling theatres to reproduce the levels selected by the director and sound mixers. [18]

2.1.5. Compact Discs

By the 1980s, Bob Katz began to notice that CDs seemed to be getting louder every year [2]. While vinyl LPs increased in level by perhaps 4 dB over 40 years, the average CD levels went up by almost 20 dB in 20 years. [1][19] In 2001, Speer stated that “much of the music we listen to today is nothing more than distortion with a beat. Great music is suffering because it lacks dynamic range” [20].

According to Robert Dennis, “When one has a crest factor (difference between peak and RMS levels) of 10 dB, we are operating at a level of reasonability in pop sound quality. This was the approximate crest factor that could be obtained with analog tape recording.... This is a way better sound quality than today’s pop CDs released during the loudness wars.... Today mastering engineers often use dynamic processing to reduce the final crest factor to as low as 6 dB.” [10]

Figure 2 shows the increase in RMS levels of pop music CDs from 1980 to 2000, and Figure 3 shows the decline in dynamic range values from 1980 to 2010.

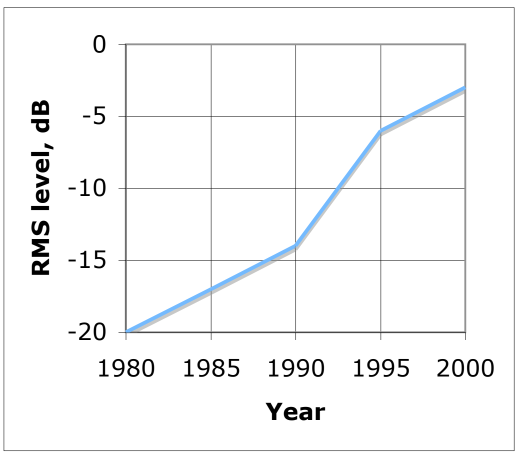

Figure 2. Average RMS levels of ‘hottest’ pop music CDs, 1980-2000, (data from Katz, [1]).

2.2. Nonlinearity and the Loudness War

As a result of its rather Rube Goldbergian design, the ear has a number of nonlinear behaviors, including a reduced sensitivity at low and high frequencies, particularly at low listening levels. Fletcher and Munson of Bell Labs showed that the perceived loudness of tones can be characterized by a family of frequency response curves that become flatter at high levels [22]. A set of equal-loudness contours is shown in Figure 4.

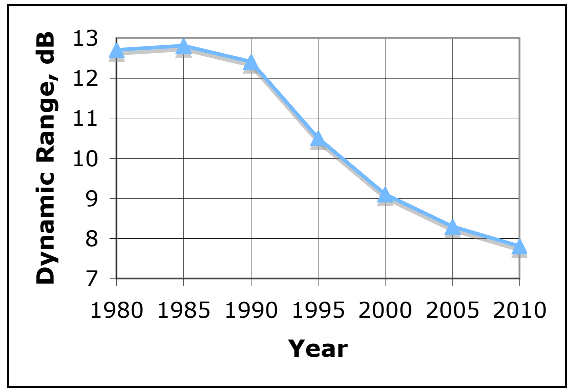

Figure 3. Average “dynamic range” of albums listed on “The Unofficial Dynamic Range Database” [21], 1980-2010, measured with the TT Dynamic Range Meter [5].

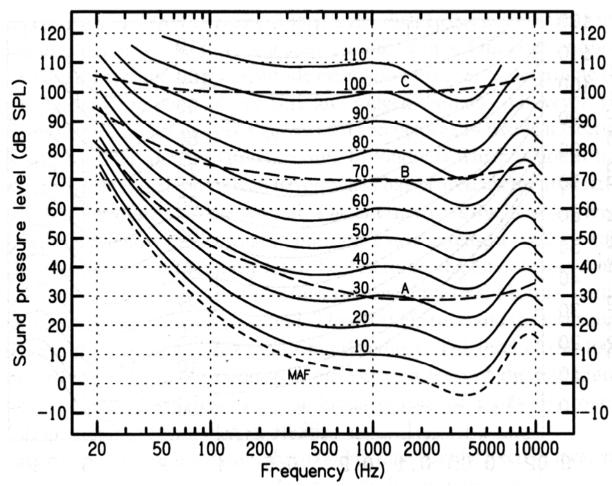

Figure 4. Equal loudness contours, with inverted A-, B- and C-weighting curves superimposed, from [23]. (Note that these contours don’t necessarily apply when multiple frequencies are present simultaneously.)

As a result of this behavior, boosting the playback volume results in a more linear response and makes it easier to hear details at low and high frequencies. Milner wrote: “This quirk of hearing has played an important role in enabling the loudness war. If you play the same piece of music at two different volumes and ask people which sounds better, they will almost always choose the louder, partly because more of the frequencies are audible.” [2]

In an AES workshop, Katz mentioned testing a new set of A/D and D/A converters and finding that the output seemed to have more depth, more space, a wider sound and more inner detail, compared to the original signal. He discovered that the converters were slightly out of calibration, adding 0.2 dB of gain. [19]

Katz also mentioned that if you compress a piece of music, making it 1 dB louder than the more open and dynamic original, even experienced listeners will say that the original material sounds compressed, when in fact the louder recording is the one that is more compressed. [19]

2.3. Multiband Compression

While the ear is nature's original multiband audio compressor, the development of powerful hardware and software compressors has facilitated the loudness war. During the mastering stage, compression is typically used to even out levels, bring out certain details and make the mix more coherent [1]. Compression also makes a recording seem louder for a given peak level. Multiband compression allows more aggressive processing with less "pumping" or intermodulation distortion than with single-band compression. [24]

Compression may also be useful during the playback stage, depending on the listening environment. Lawson wrote, "iPods are emblematic of mobility, and mobility means being surrounded by noise, so these devices need to drown that other noise out. Consequently, the more steadily loud the music is, the better." [25]

Many office workers listen to music throughout the workday to mask unwanted co-worker conversations [26]; compressed music can be more effective for this purpose due to its consistent level. In noisy environments such as cars and planes, compression keeps the background noise from covering up the quiet parts of the music (or, conversely, enables the music to block out the noise). Late night listening also benefits from compression, which allows the quiet parts to be audible while keeping the loud parts from disturbing sleeping housemates.

Providing the option of compression during playback essentially gives listeners the ability to remaster the recording to suit individual preferences and listening situations. Intelligent playback compression can also refrain from re-compressing music that is already hypercompressed [6].

Ideally, music as released would be mastered to retain a relatively wide dynamic range, because it can always be compressed further at the playback stage if needed. However, the belief that louder songs sell better has led to hypercompression of the master recordings.

3. WHAT ARE THE PROBLEMS?

Two main questions of interest are:

1. Can listeners distinguish between relatively uncompressed and hypercompressed audio, and if so, what difference do they hear?

2. Whether or not listeners consciously notice any difference, is hypercompressed audio more fatiguing?

These questions correspond to the main problems attributed to the loudness war: aesthetic concerns and listening fatigue. Other possible problems may include large level differences between recordings from different eras, increased incidence of hearing damage and decline of the music industry.

3.1. Aesthetic Concerns

Hypercompression has been accused of removing dynamics and making music sound "squashed" [1], creating musical clutter and reducing depth and texture [1][11][20], robbing music of its excitement and emotional power [11][20], producing an uncohesive sound due to the multiband compressor's continually changing frequency response [1][24], amplifying mono information and reducing stereo width [1], and reducing the punch of transients [1].

While many people have expressed concern about the damage done to music by hypercompression, there are few if any studies regarding the extent to which people can actually hear the difference. Clearly there is a continuum of effects, from subtle compression that would sound transparent even to the most careful listener, to extreme settings that virtually anyone would notice. The danger is that, as the loudness war escalates, the pressure to use excessive amounts of compression may continue to increase.

Some examples of the effects of hypercompression (created for the purpose of illustration) are shown in Fig. 5. The strong verse-chorus distinctions in a highly dynamic Björk song (Fig. 5A, as released) virtually disappear with hypercompression (Fig. 5B). Comparing a 1994 recording of Ravel's Boléro (Fig. 5C) to a hypercompressed version (Fig. 5D), a listener could hardly fail to notice the difference, especially since the piece was intended as an experiment in dynamics [27]. Likewise, imagine applying fast, extreme compression and limiting (Fig. 5F) to the sudden loud chord at the end of the opening theme of the second movement of Haydn's Symphony No. 94 (Fig. 5E). Finally, consider the loss of dynamics in "Stairway to Heaven" (Fig. 5G) with hypercompression (Fig. 5H).

In these admittedly extreme examples, the effects should be obvious: hypercompression makes the "It's Oh So Quiet" verses as loud as the choruses, takes the crescendo out of the "Boléro", removes the surprise from the "Surprise Symphony," and turns the "Stairway to Heaven" into a sidewalk.

3.1.1. Loss of Excitement and Emotion

One of the main complaints about hypercompression is that it flattens the dramatic and emotional impact of the music. Levitin stated that "The excitement in music comes from variation in rhythm, timbre, pitch and loudness. If you hold one of those constant, it can seem monotonous." [11]

Lawson used Róisín Murphy's song "Overpowered" as an example: "...as cleverly assembled as 'Overpowered' is, the use of dynamic range compression arguably limits its potential for pleasure. The recording is not as aggressively mastered as some, but when the song sounds like it is 'supposed' to hit a peak in volume – 'As science struggles...' – there is no actual increase... In order to be beguiled the listener has to *imagine* that the volume has increased...." [25]

Rodgers speculated that when we remove dynamics, we place a distance between ourselves and the composer or performers, who intended the music to have dynamic changes. This may prevent us from bonding with the music as we did before music was so portable, when people were more likely to engage in active listening. "By now relegating music to background through this capacity to over-compress, we've made it truly a background object and perhaps not as emotionally rewarding as it potentially could be." [28]

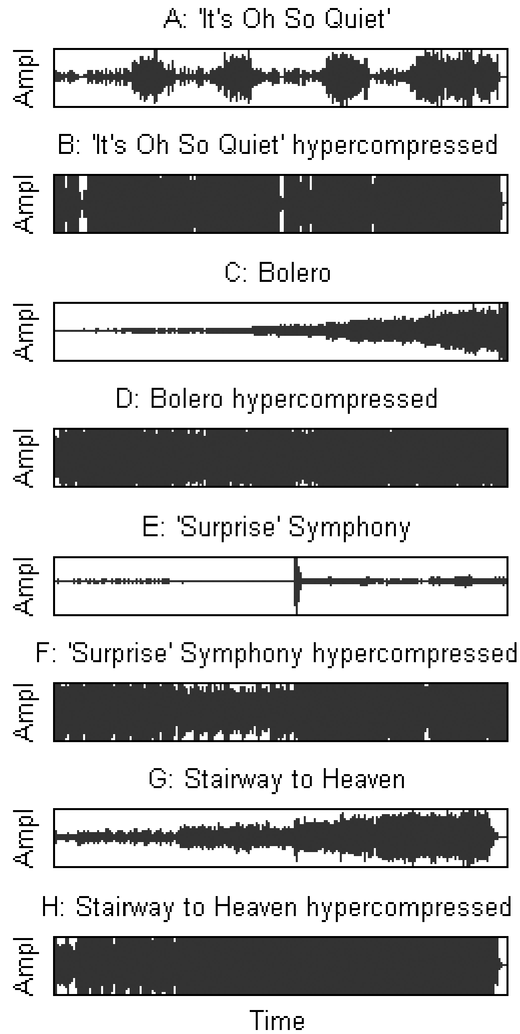

Figure 5. Waveforms of "It's Oh So Quiet," "Bolero," Haydn's Symphony No. 94, and "Stairway to Heaven," with and without hypercompression (applied by the author).

3.1.2. Remastering

Without access to the original recordings, consumers may have no basis for comparison to determine whether audible damage has been done. Recently remastered recordings, however, have frequently been criticized for excessive compression compared to the initial releases; examples cited include remasters of the Rolling Stones albums *Sticky Fingers* and *Some Girls*, the 2003 *Red Hot Chili Peppers Greatest Hits, Beatles I,* and the 2007 Led Zeppelin compilation, *Mothership*. [11][29]

Lawson notes that "record labels will often advertise that the disc has been 'remastered from the original master tapes,' implying that... what makes its value greater from that of previous releases is that it is closer to the source material.... That the master tape itself is a kind of inaccessible fetish object – albeit one that most consumers would be unable to play anyway – is suggested by the series of Abba reissues from 2001: an insert depicts the covers of each album with the catalog numbers beneath, but in the center is a photo of the boxed master tapes captioned 'not for sale.' And no wonder: they were remastered yet again in 2005. As long as Universal Music Group has financial incentive to keep revisiting the Abba catalogue, they need to appeal to the notion that each time they are getting closer and closer to the original, bringing out further unheard details.... The irony is that the remastering of Abba's catalogue has, with respect to dynamic range, probably gotten further from the master tapes with each new release." [25]

A frequently cited example of hypercompression is Metallica's *Death Magnetic* album, which was released simultaneously for the *Guitar Hero* game and as an audio CD. The CD version was much louder and more compressed than the game version, making it easy for fans to compare the two. So far, over 21,000 people have signed an online petition asking Metallica to remaster the CD with less compression. [30][31]

Different listeners do not always agree on the effects of hypercompression. Levine's "The Death of High Fidelity" article in *Rolling Stone* listed Dylan's *Modern Times* as being one of the few modern CDs with a wide dynamic range and sense of spaciousness [11]. However, the article "Tears of Rage: The Great Bob Dylan Audio Scandal" listed the same CD as being "significantly damaged" and "strangely wearisome" compared to the LP, which had "depth and nuance and poignancy and richness and warmth"; comparison displays showed the CD version with flat-topped waveforms suggestive of significant digital limiting [32].[3] The discrepancy suggests that the effects of compression as commonly employed may not always be immediately obvious.

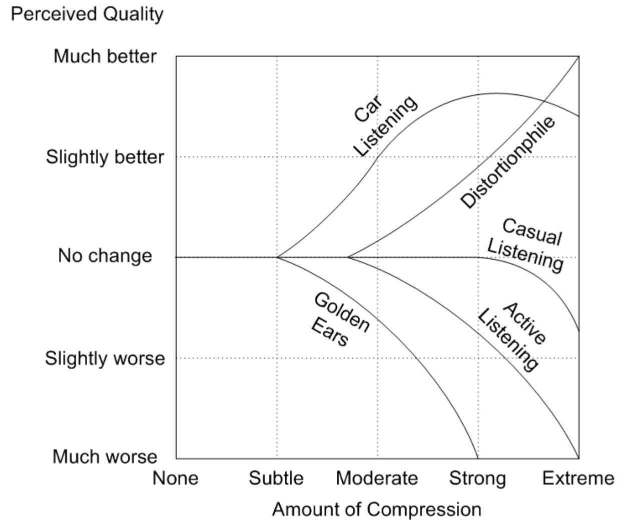

Figure 6. Illustration of how the relationship between amount of compression and perceived quality might vary for different types of listeners, listening styles and listening situations. (The curves are not based on actual data.)

3.1.3. Is Hypercompression Audible?

Following are a number of possibilities (not necessarily mutually exclusive) regarding the audible effects of hypercompression.

Differences between listeners, listening styles and listening environments

It is possible that some people are able to hear relatively subtle amounts of compression while others are oblivious even to extreme hypercompression. In addition, people may choose to listen in different ways at different times: sometimes listening actively to the music, sometimes listening as background. Finally, the amount of compression may need to be adjusted to suit different listening environments.

One could imagine plotting the relationship between amount of compression and perceived quality for different types of listeners, listening styles and listening situations as shown in Figure 6. For example, experienced mastering engineers ("Golden Ears") might readily perceive a loss of quality from moderate amounts of compression, while typical active listeners might tolerate more compression before noticing any damage to the audio quality, and casual listeners might not notice any effect until the settings become extreme.

[3] Tollerton warns of the danger of trusting waveform displays as an indicator of sound quality [33].

Someone listening in the car may perceive an improvement in their ability to hear the quiet parts of the music, and a fan of loud, distorted genres (denoted "distortionphile") might prefer extreme clipping and distortion.

Music differences

Some types of music, such as classical and jazz, may be very sensitive to dynamics compression, while pop music may tolerate more extreme compressor settings, partly because it often relies less on variation in loudness. Nevertheless, according to Lawson, "one can cite numerous examples of popular music whose impact is largely driven by dynamics. For instance, much of the 'alternative' music of the late eighties and early nineties would actually repeat the same riff for the duration of the song; choruses were distinguished by an increase in volume as much as anything else (Nirvana's 'Smells Like Teen Spirit' is characteristic)." [25]

Hypercompression as aesthetic preference

According to this view, hypercompression may simply be a matter of taste. Some people may prefer well-mastered recordings that they can listen to repeatedly without fatigue, while other demographics (e.g., young males) may prefer the music set to "stun," with a louder, distorted, in-your-face type of production. Hypercompression, clipping, aliasing, distortion, "FSU" plug-ins and other means of acoustic mutilation may be viewed as ways of achieving this aesthetic, particularly for genres such as grunge, heavy metal, glitch and shred.

Desire for musical purity

Some of the animosity toward hypercompression could be motivated as much by the idea that it damages the music as by the reality of the situation. Lawson wrote, "These listeners long for access to the purity of the original recording before it was 'squashed,' but the problem is that the original recording does not, in a sense, exist. Producers and mastering engineers assemble the tracks recorded and create a particular sonic product that can later be revisited and 'remastered.'" [25]

Compression as scapegoat for loss of interest in music

People tend to bond closely with the music they heard in their pre-teen and teenage years. As society ages, listeners may blame hypercompression for a loss of musical interest that may result from other factors such as changes in musical styles, age-related hearing loss and various lifestyle changes.

Changes in sensitivity over time

A certain amount of ear-training or detailed listening time may be necessary to become sensitized to some types of aesthetic damage caused by hypercompression. On the other hand, listeners may gradually habituate to bad audio quality, as many have with the sound of low-quality cell phone codecs. [34] [4]

Macrodynamic and microdynamic effects

Slow, macrodynamic compression effects involve reduced dynamic contrast and loss of expressiveness (e.g., flattened crescendos, and choruses becoming no louder than verses). Such effects are not immediately obvious but may become apparent during the course of the song.

Some side effects of fast, microdynamic hypercompression include a "squashed" sound, a loss of "punch," softer transient attacks, and a busy, cluttered sound [1][11]. In a *Rolling Stone* interview, Bob Dylan said, "You listen to these modern records, they're atrocious, they have sound all over them. There's no definition of nothing, no vocal, no nothing, just like – static" [38]. Differences due to microdynamic compression may be non-obvious to many listeners when comparing long recordings, due to deficiencies of acoustic memory, but may be easier to notice when comparing short excerpts.

It is difficult to separate the effects of microdynamic and macrodynamic compression completely, because

[4] There are preliminary suggestions that some listeners may even come to prefer certain types of artifacts. For example, Sheffield reported that consumers in her NPR tests often preferred low bit-rate coded audio to uncompressed audio [35], and Berger stated that an annual study of Stanford undergraduates showed a growing preference for the metallic or "sizzle" sound of data-compressed formats [36][37].

fast compression may also tend to compress the macrodynamic changes between verse and chorus, etc.

Cultural shifts in listening habits

While music has long been used as sonic background, careful repeated foreground listening seems to be on the wane with the commoditization of music, the decline of audiophile culture, and a "societal shift toward convenience and portability" [39]. Mastering engineer Bob Olhsson said, "Going back to the 60s, a record was a luxury; the idea of it being a commodity was absurd to me. You didn't buy a lot of recordings; you bought recordings that were special to you, and you listened to them over and over. And certainly you are less inclined to listen to a distorted record over and over and over than you are to one that just sounds amazing." [12]

It is possible that a vicious cycle has evolved, in which the increased use of music as sonic wallpaper has encouraged hypercompression, which in turn has rendered the music more suitable to background than to foreground listening.

Distortion due to artifacts and abuse

Modern multiband compressors are carefully designed to sound as "transparent" as possible, with a minimum of audible side effects. However, there are many ways to go wrong, and the list of possible artifacts is surprisingly lengthy [24]. It is possible that some of the worst audible damage blamed on multiband compression may result primarily from equipment abuse, design flaws or artifacts such as clipping, inter-sample clipping [40], aliasing [41], excessive fast limiting or other forms of nonlinear distortion.

Robert Dennis stated that "... two releases with an 8 dB crest factor can have remarkably different sound quality, varying between 'not too bad' and 'horrible.' A big factor behind this difference is the frequency spectrum of release. Distortion in CDs is often (or maybe primarily) odd-order harmonics caused by limiting (especially brick-wall limiting). Harmonic distortion of the bass spectrum tends to be masked by the midrange energy in the mix, something that cannot be said for harmonic distortion of the midrange and high frequency components of the spectrum." [10]

Cumulative distortion

In the digital domain, it is difficult to perform clipping without producing aliasing distortion [35]. The combination of clipping and low-bit-rate lossy encoding may be particularly unpleasant. At a recent AES workshop, Katz demonstrated the results of encoding a clipped recording using MP3 at 96 kbps (which is higher than the bit rates currently used on satellite radio). He subtracted the result from the original clipped recording and showed that the difference consisted of harsh high-frequency noise and distortion. [19]

At a 2009 AES session, Marvin Caesar stated that lossy codecs respond badly to heavy compression and limiting. This effect may be especially problematic with stacked codecs in a typical FM signal chain. As a result, he claimed, "We're foisting pretty rough audio on some of our listeners." [42] In particular, clipping and aggressive dynamics processing may cause aliasing and intermodulation distortion when processed through lossy codecs. [35][43]

3.2. Listening Fatigue[5]

Even if people do not consciously notice any problems, it is possible that hypercompressed music may become physically or mentally tiring over time; listeners may gradually lose interest without knowing why.

Marketing claims of "reduced listener fatigue" frequently appear in sales pitches for various audio products. In the audio engineering literature, listening fatigue has been attributed to a wide variety of causes including fast-acting compression and limiting [35], a lack of variation in loudness over the duration of a recording [35], contradictory location cues [44][45] [46], stereo image processing and active matrix upmix steering fluctuations [35][47], phantom image instability [45], phasiness [48], poor equalization [35], low frequency rumble [49], too much boost between 1-4 kHz [50], clipping and intermodulation distortion [35] [51], Doppler speaker distortion [52], low bit-rate lossy encoding [34][42], stacked codecs [42], excessive

[5] This article will use the terms "listener fatigue" and "listening fatigue" interchangeably, though "listening fatigue" is recommended, as it connotes fatigue from the act of listening, as opposed to whatever fatigue the listener might be experiencing. A separate phenomenon, "auditory fatigue," denotes a temporary reduction of the ear's sensitivity in response to sound exposure [23].

hearing aid amplification of soft background sounds [53], listening to monophonic instead of stereo recordings [54], listening to stereo instead of 3-channel recordings [46], and listening with headphones to audio mixed for speakers [55].

It appears, however, that the single most common use of the term “listener fatigue” in the engineering literature may be in regard to the need for frequent breaks during listening tests in order to avoid tiring the subjects and thereby corrupting the results (e.g., [22], [56] and [57]). This raises the not entirely facetious question of whether extended listening tests designed to measure fatigue might in fact corrupt their own results due to listening fatigue.

Given the number of articles, blog postings and conference workshops on the topic, there is surprisingly little experimental evidence regarding hyper-compression and listening fatigue. In one study, Stone et al reported that rapid multiband compression seemed to increase the amount of effort required to perceive independent sound sources within a complex signal, which may contribute to listening fatigue; however, this study was in the context of multiple simultaneous talkers, not music, and it did not measure fatigue as such [58].

3.2.1. Is Hypercompressed Audio More Fatiguing?

Again, there are a number of possibilities, not necessarily mutually exclusive:

Hypercompression causes fatigue

Despite the lack of published studies, it is widely believed that excessive compression results in listening fatigue, which may in turn discourage close, repeated listening [2]. There is a great deal of anecdotal evidence to this effect.

Southall wrote, “Music is about tension and release. With very ‘hot,’ un-dynamic music there is no release because the sensory assault simply doesn’t let up.... you end up feeling like Alex at the end of *A Clockwork Orange* – battered, fatigued by, and disgusted with the music you love.... I very much doubt that this is just me.” [8]

Mastering engineers are among those closest to the situation, with countless hours of detailed listening to varying amounts of compression. While some mastering engineers feel that the loudness war is merely a matter of finding an acceptable compromise between quality and market pressures, others are deeply concerned about listening fatigue and damaged sound quality. [12]

Milner quotes mastering engineer Greg Calbi as saying that the idea that listener fatigue is a by-product of digital compression is “almost universally held.... I never heard the word ‘fatigue’ once when I was cutting vinyl.” [2]

Mastering engineer Bob Weston stated that “highly compressed or limited music with no dynamic range is physically difficult to listen to for any period of time.... This ‘hearing fatigue’ doesn’t present itself as obviously aching muscles, like other forms of physical fatigue, so it’s not obvious to the listener that he or she is being affected. But if you ever wonder why you don’t like modern music as much as older recordings, or why you don’t like to listen to it for long periods of time (much less over the year), this physical and mental hearing fatigue is a big part of the reason.” [29]

According to mastering engineer Joe Lambert, “Some records will work really loud; others wear on you. They sound great the first two times, but then you just stop listening because they fatigue your ear. I know that and the engineers know that, but the consumer doesn’t know why they’re not listening to those records any more.” [12]

At an AES session on listener fatigue and longevity, Marvin Caesar stated that over tens of thousands of repetitions every year, when comparing competitively loud compressed recordings vs. hypercompressed songs, listeners almost always picked the less compressed one, saying “I could listen longer.” [42] Radio stations apparently adjust the amount of compression they use in an attempt to trade off a louder sound, which may initially attract listeners searching for a station, versus the danger of driving listeners away due to loudness fatigue. Stations looking for a female demographic, in particular, may avoid excessive processing because of a belief that women are particularly sensitive to over-compression. [2][14]

“Urban legend”

It is conceivable that the idea of hypercompression causing listening fatigue may be widely believed simply because it sounds believable, allowing this meme to spread across the internet like other "urban legends." Milner quotes psychoacoustics researcher Stephen McAdams as saying "I am unaware of any studies on this phenomenon.... It may be an urban myth" [2]. Certainly the underlying conflict triggers an emotional response that could encourage the idea's spread: greedy record companies trying to outgun each other in loudness at the expense of our cultural heritage, versus the good guys standing up for pristine audio quality.

It is also possible that the lack of published studies on this topic could be due to the "file drawer effect," whereby studies failing to confirm the hypothesis remain unpublished. Nevertheless, the anecdotal evidence seems quite compelling.

Subtle, hard to measure

A recurring theme is that listeners may not consciously notice the difference in audio quality, but with hypercompressed music they will change stations or stop listening after a while [2][12][29][31]. If this is true, it should be possible to measure the effect.

Unfortunately, listening fatigue is particularly hard to measure for a number of reasons, including the need for lengthy and fatiguing tests with many subjects. Psychoacoustics researcher Sheffield is quoted as warning of the difficulty of testing the effects of poor sound quality, noting that consumer responses often seem to defy scientific reason [35]. Section 3.2.3 discusses some of the issues involved in measuring listening fatigue.

Differing susceptibility to fatigue

Different individuals may have different degrees of susceptibility to listening fatigue. Mastering engineers, who have learned to tune in to the effects of compression and who spend countless hours listening to processed audio, may be the "canaries in the coal mine." Johnston mentioned that as you learn to recognize various artifacts over time, you get heightened sensitivity and the artifacts become more annoying [59].

Artifacts and abuse

As with aesthetic damage due to hypercompression, it is possible that listening fatigue may be minimal with moderately strong, competitively loud compression, but substantial when high compression ratios, excessive limiting, fast time constants and clipping are used.[6]

The combination of hypercompression and the typical radio broadcast processing may cause increased fatigue due to audible artifacts. Orban and Foti stated that in extreme cases, hypercompression over the radio "sounds overtly distorted and is likely to cause tune-outs by adults, particularly women" [14]. The harsh, high frequency artifacts from low bit-rate lossy encoding of previously clipped audio [19] (mentioned in Section 3.1) might also cause annoyance and/or fatigue.

3.2.2. Types of Listening Fatigue

It is likely that the term "listening fatigue" may refer to two or more different phenomena. Johnston speculated that listener fatigue may include physical fatigue as well as central nervous system fatigue [59].

Physical fatigue

Physical fatigue, in turn, may consist of two types: mechanical cochlear fatigue and neural or biochemical cochlear fatigue. Keeping in mind that loudness (a perceptual phenomenon) and intensity (a physical phenomenon) are not perfectly correlated, Johnston suggested that mechanical cochlear fatigue may be due to high intensity, since outer hair cell damage appears to correlate to intensity. [59]

High loudness, on the other hand, can cause outer hair cells to depolarize, suggesting the possibility of biochemical fatigue. In addition, the inner hair cell firing rate is largely proportional to loudness (not intensity), also suggesting biochemical fatigue. [59]

Johnston mentioned that listening to something that sounds harsh, even if it's not that loud, can make us want to run screaming in a few minutes; it's not clear

[6] To the extent that listening fatigue and other undesired effects of hypercompression may be due to design flaws, artifacts and abuse of the technology, an all-in-one audio quality software tool might prove useful to mastering engineers. Such a tool could compare 'before' and 'after' recordings and quantify things like intermodulation distortion, changes in spectral balance [24], loss of spectral contrast [60], clipping, inter-sample peaks [40], excessive limiting, aliasing from rapid gain changes [41], etc., in order to red-flag possible problem areas before the master is released.

whether this is a physiological or intellectual process [59]. In a comparison of various loudspeakers, Berkow noticed that more distortion components can give speakers a fuller, richer sound, but they also cause fatigue and "wear the ears out sooner" [49].

Cognitive fatigue

Central nervous system fatigue, or "cognitive fatigue," can result from missing, false or contradictory cues. If cues are incorrect or unavailable, the brain has to do more work to extract the information. With speech, in particular, multiband compression may flatten the signal and destroy the articulation. [59]

Marvin Caesar discussed the importance of transient response: fast attack limiting crushes the leading edge of the signal, reducing intelligibility and forcing the listener to work harder to understand the content [42]. Sheffield found that digital artifacts such as echo, swirling and sibilance [61] from low bit-rate coders seemed to cause fatigue over time with speech, but not with music [34]; it is possible that dynamics compression artifacts could have a similar effect.

It is unclear how cognitive fatigue might apply to music, since there may be less need to extract specific information [59]. However, to the extent that processed audio sounds "unnatural," with a mismatch between what we hear and how things normally sound in the natural world, the brain may puzzle over various discrepancies and distractions. Thus, cognitive fatigue might result from such things as contradictory or incomplete location cues, rapidly changing frequency response [24] and various other artifacts.

Slow (macrodynamic) compression causes a lack of tension and release variation, allowing the brain to regard the music as "background" and tune it out. At an AES workshop on the loudness war, Rodgers said, "When you reduce the dynamic changes, the listener can listen for a longer period of time – it's less cognitively taxing. Attention doesn't have to be woken up from these changes – you can tune it out.... Music becomes this background noise – it is the noise in your own personal signal-to-noise ratio. It's not changing; it's not engaging you cognitively." [28] Thus, macrodynamic compression may actually reduce cognitive fatigue so long as the listener is not actively trying to attend to the information, though it may induce a related phenomenon, boredom, due to the lack of contrast and "negative space."

If the listener is trying to pay attention to the content, on the other hand, macrodynamic compression may cause increased fatigue, perhaps because the brain can have trouble efficiently decoding signals delivered at a constant level [62]. In "Over the Limit," Rowan wrote:

"WHY IS THE LOUDER IS BETTER APPROACH THE WRONG APPROACH? BECAUSE WHEN ALL OF THE SIGNAL IS AT THE MAXIMUM LEVEL, THEN THERE IS NO WAY FOR THE SIGNAL TO HAVE ANY PUNCH. THE WHOLE THING COMES SCREAMING AT YOU LIKE A MESSAGE IN ALL CAPITAL LETTERS. AS WE ALL KNOW, WHEN YOU TYPE IN ALL CAPITAL LETTERS THERE ARE NO CUES TO HELP THE BRAIN MAKE SENSE OF THE SIGNAL, AND THE MIND TIRES QUICKLY OF TRYING TO PROCESS WHAT IS, BASICALLY, WHITE NOISE. LIKEWISE, A SIGNAL THAT JUST PEGS THE METERS CAUSES THE BRAIN TO REACT AS THOUGH IT IS BEING FED WHITE NOISE. WE SIMPLY FILTER IT OUT AND QUIT TRYING TO PROCESS IT." [63]

3.2.3. Measuring Listening Fatigue

In 1961, Barlowe suggested that it would be helpful to have a meter that is calibrated in units of "listening fatigue" [51]. Unfortunately, as Johnston reminds us, currently there are no such units. We might attempt to measure the time spent willingly before the subject "changes stations" or stops listening, but it is hard to distinguish fatigue from other factors such as boredom, annoyance, dislike of the content (e.g., genre or lyrics), dislike of the test setup, or normal fatigue [59]. Other testing methods include seeing how Mean Opinion Scores (MOS) change over the course of extended listening sessions in which each participant hears the samples in a different order [34][64], and measuring how processed speech affects memory for text passages [61].

Setting the playback level

One difficulty is the issue of how to set the playback level during testing. Listening fatigue is presumably a function of level; certainly if the level is inaudible, there is no listening fatigue. Since the perceived loudness is different for compressed and uncompressed music, any comparison should be at the relative loudness levels that listeners would typically use.

This raises the question of how listeners typically set the levels. Some possibilities include:

1. Listeners tend to set the levels according to the maximum short-term loudness they're comfortable with.

   In this case, their ears would get less rest time and would be exposed to more long-term sound energy with hypercompressed music, which would consistently approach the maximum short-term comfort level.

2. Listeners set the playback volume according to long-term loudness (which itself can be defined in various ways – for example, by giving more weight to the louder portions [6][65]).

   In this case, there may be less difference in listening fatigue between the compressed and uncompressed cases, since the listener has already essentially normalized the levels for equal listening fatigue (assuming some correlation between long-term loudness and fatigue).

3. Listeners set the playback volume according to annoyance, not loudness.

4. Listeners initially set the playback volume according to one of the above strategies after listening for a few seconds, but they may adjust the volume again later if the loudness goes outside a preferred range.

5. Different listeners have different strategies regarding level setting.[7]

Recommendations for testing listening fatigue

Given the uncertainties around how people typically set their listening levels, it is important for test subjects to adjust playback levels according to personal taste. Unfortunately, many people will find it difficult to set the levels consistently, and the levels they select may vary according to the order in which recordings are presented. At any rate, there should be generous rest periods between tests of compressed and uncompressed music; it might be best to test these on separate days.

Note that if normalization is done according to some "loudness" descriptor such as the ITU BS.1770 standard [67], we won't know for sure whether we are measuring listening fatigue, deficiencies in the loudness metric, or both.

3.3. Other Side Effects of the Loudness War

3.3.1. Inconsistent Levels from Year to Year

Another problem is the lack of consistency between the levels of older vs. recent recordings. CD levels may differ by as much as 20 dB, depending largely on the year of release [19]. Such level changes, which may be noticed when playing songs from multiple albums in "shuffle" mode, can be quite unpleasant and potentially harmful to speakers or ears.

3.3.2. Possibility of Hearing Damage

Aside from potential damage due to unexpected level jumps from one CD to the next, there is some concern about whether hypercompression itself may increase the potential for hearing damage. The harm caused by excessive noise exposure is related to the total energy received [23], and is therefore a function of the average sound level and duration of exposure [68]. Thus there is some reason to wonder whether repeated listening to over-compressed audio, with its higher average sound levels, might contribute to hearing loss over time.

On the other hand, it is possible for downward compression and limiting to protect the ears from dangerously high peak levels. Ultimately, the question of whether hypercompressed music is likely to increase the incidence of hearing damage may depend, as with listening fatigue, on how listeners set the playback volume: according to average or maximum loudness. If they set the volume based on average loudness, they may essentially normalize for total sound exposure.

Another possibility is that hypercompressed music may cause damage over time due to the lack of "rest periods." Just as people are encouraged to take frequent breaks from the computer keyboard in order to avoid carpal tunnel syndrome, our ears may need to take periodic breaks from peak sound levels to give the hair cells time to recover.

[7] Johnston reported that a very small and inconclusive test suggested that some people seemed to set levels based on averages, while others tended toward equalizing peaks. [66]

Friedemann Tischmeyer, mastering engineer and founder of the Pleasurize Music Foundation [69], claimed that exposure to hypercompressed music that never gives the ears time to rest can have detrimental effects on hearing. "We have already a lot of evidence that this is the main reason for the drastic increase of hearing damage in the young generation.... I personally believe that this is not just a matter of good taste any more, it is a matter of responsibility to protect the pleasure of hearing." [29]

In *IEEE Spectrum*, Katz was quoted as saying "You want music that breathes. If the music has stopped breathing, and it's a continuous wall of sound, that will be fatiguing.... If you listen to it loudly as well, it will potentially damage your ears before the older music did because the older music had room to breathe." [70] This statement suggests the possibility that hypercompressed music could be more damaging than uncompressed music, even after normalization for equal RMS levels, BS.1770 loudness, or user playback level preferences, because of the lack of rest periods.[8]

Many recent recordings include digital clipping in addition to hypercompression. Clipping, aliasing and other forms of nonlinear distortion can sound particularly harsh and unpleasant; it is unclear whether such sounds are physically hard on the ears or merely annoying.

A recent study in the *Journal of the American Medical Association* showed that the prevalence of hearing loss among U.S. adolescents has increased significantly from 14.9% in 1988-1994 to 19.5% in 2005-2006 [71]. While the losses are often slight, they may also be permanent and may get worse with continued exposure. Many experts suspect that the primary cause is the use of earbuds and headphones for listening to portable music. [72] Portable media players can produce maximum levels from 80 to 115 dB(A). A study by the European Union warned that those who listen at high volume for five hours a week receive more noise exposure than permitted in the noisiest factories and workplaces, and that 5% to 10% of listeners are at high risk due to their levels and duration of exposure. [68][73] A recent study of Australian children showed that personal stereo usage was associated with a 70% increase in the risk of hearing loss [74].

[8] This possibility is reminiscent of the finding by Moore et al that hypercompressed speech is louder, even after normalization for equal RMS levels [16].

The extent to which hypercompression contributes to hearing loss is somewhat speculative and difficult to prove; on the other hand, the stakes are high. Noise-induced hearing loss can be thought of as a kind of repetitive stress disorder, and as with other repetitive stress disorders such as carpal tunnel syndrome, most people don't give the issue much thought until the damage is done. Many teens, in particular, consider themselves invulnerable in this regard [72]. At the very least, people should think twice about listening at high levels for hours each week, especially over earbuds or headphones, and especially if the music does not include rest periods.

3.3.3. Decline of the Music Industry

The music industry has declined dramatically in recent years, a development that has been attributed largely to the ease of digital piracy via the internet. Lawson stated that "although for much of the nineties the CD was a major boon for the industry, particularly as consumers replaced their vinyl albums with the new format, it ultimately precipitated the industry's current malaise and presumed future downfall." [25] As of the first half of 2009, CDs still comprised 65% of all music sold in the United States, but digital music sales (from iTunes, Amazon and other online stores) make up a rapidly growing share of the market. [75]

Meanwhile, Nielsen SoundScan reported that 2.5 million vinyl albums were sold in 2009, a 33% increase from the previous year and almost three times the 858,000 units sold in 2007. This was the highest level of LP sales since they began tracking the data in 1991. [76][77][78] While vinyl is still a niche format (only 1% of total album sales [79])**,** if CD sales were to continue falling at a double-digit annual percentage rate while LP sales continue to surge, the two could conceivably break even by 2020.

Part of the increase in LP sales has been attributed to the perception that vinyl has a "warmer, more nuanced sound than CDs and digital downloads" [78], perhaps because of the necessity of using finesse to work around vinyl's physical limitations regarding signal levels [2].

Mastering engineer Bob Ludwig stated, "People talk about downloads hurting record sales. I and some other people would submit that another thing that is hurting record sales these days is the fact that they are so compressed that the ear just gets tired of it. When you're through listening to a whole album of this highly compressed music, your ear is fatigued. You may have enjoyed the music but you don't really feel like going back and listening to it again." [31]

In 2001, mastering engineer Bob Speer wrote, "The record labels blame digital downloads, MP3s, CD burners, and others for the lack of CD sales. While there is some truth to their constant whining, they only have themselves to blame for the steady decline in CD sales. Much of the music being produced today isn't music at all.... It's anti-music because the life is being squashed out of it through over-compression during the tracking, mixing, and mastering stages.... It's no wonder that consumers don't want to pay for the CDs being produced today. They're over-priced and they sound bad." [20]

Southall wrote, "Compression will continue to be abused in the pursuit of loudness for as long as the recording industry believes that louder shifts units.... Global album sales are falling year-on-year, far less mega-million-selling records are occurring... and I think this is because the clamour to make music louder has made it less loveable, and in the long run lovable sells more" [62].

It may be difficult or impossible to prove a causal relationship between the loudness war and the decline of the music industry, but the "Evergreen Project" (see Section 4.3.1) suggested that many of the most influential and best-selling albums in the history of the music industry had a wide dynamic range.

4. FIXING THE PROBLEM

4.1. A Game Theory View of the Loudness War

The terms "loudness race" and "loudness war" are implicit reminders of the similarity to an arms race, in which each party feels compelled to keep up with the others in order to avoid being at a competitive disadvantage. In the loudness war, even parties motivated primarily by audio quality feel forced to use hypercompression to make their recordings competitively loud.

Just as actual war is not fundamentally a chemical or nuclear engineering problem, the loudness war is not primarily an audio engineering problem. Developments in signal processing have certainly facilitated the ongoing increase in loudness, but the underlying problem lies in the domain of game theory. I will review some of the basics of game theory and explore implications for the loudness war, including de-escalation strategies such as educating people about the actual payoffs and changing the payoffs to discourage hypercompression.

4.1.1. Brief Review of Game Theory

Game theory is a branch of mathematics dealing with the analysis of conflict, often using computer simulations. It applies the study of games of strategy to arenas such as war or business. Some of the most interesting topics involve non-zero-sum games, which are partly competitive, partly cooperative. For example, during the cold war, the U.S. and U.S.S.R. had a shared interest in avoiding mutual destruction. [80][81]

The Prisoner's Dilemma

The Prisoner's Dilemma is a classic game theory problem that has been applied to fields as diverse as political science, business, evolutionary biology and sociology. In this scenario, two suspects are accused of committing a crime together; if each pursues his own rational, individual self-interest, each will receive a worse outcome than was possible, hence the dilemma. [80][81]

In 1984, Robert Axelrod published an influential book entitled *The Evolution of Cooperation*, which used computer tournaments of the iterated Prisoner's Dilemma to show how cooperation could emerge from competition between self-seeking individuals, even in the absence of a central authority to police their actions. [80]

Social Dilemmas

The two-player Prisoner's Dilemma does not effectively model a situation like the loudness war, in which there may be dozens, hundreds or thousands of players, each of whom have limited interaction with (and limited leverage over) the others. This type of situation is a "social dilemma," in which individually reasonable behavior leads to a result in which everyone is worse off. As Kollock writes, "a group of people facing a social dilemma may completely understand the situation, may appreciate how each of their actions contribute to a disastrous outcome, and still be unable to do anything about it" [82].

This description applies well to the loudness war. Many mastering engineers are reluctant to apply excessive compression, but they often feel pressured by musicians and record company executives who are motivated to stay competitive with others, even at the possible expense of audio quality. [2]

One common social dilemma is the Tragedy of the Commons, which models situations such as pollution and over-fishing, where benefits are privatized but costs are borne collectively [82]. For the loudness war, possible costs may include not only the damage to our cultural heritage but also potentially the damage to the music industry itself as each ultra-loud recording pressures others to over-compress as well.

4.2. Some Lessons from Game Theory

4.2.1. Avoid Loudness Envy

Katz wrote that “the practice of overcompression is part of a vicious circle of loudness envy” [83]. He stated that “over 90 percent of my clients fall into positions of trying to compare their CD against another CD, and I have to educate them very hard that they have their own volume control, and why the more dynamic lower average level CD sounds much better, and performs better on the radio as well.” [12]

In Axelrod’s Prisoner’s Dilemma computer tournament, the highest-scoring strategies were those that focused on their own success, rather than on doing better than their competitors. The simplest strategy, “Tit For Tat,” never out-scored a single player, but it won the overall tournament by eliciting behavior that allowed both to do well. [80]

Kollock wrote: “This seems to be one of the hardest lessons for individuals to learn, perhaps because of the competitive game as a model in many cultures – if the only metaphor you have is the zero-sum game, you tend to treat everything as if it were a war.” [82]

4.2.2. Don’t Try to Compete with Every Genre

Axelrod’s computer simulation found that, while a “nice” strategy such as Tit For Tat could not thrive when surrounded by players who always defect (“meanies”), a cluster of nice players could thrive so long as some percentage of their interactions were with each other [80].

A somewhat parallel situation in the loudness war relates to clusters such as musical genres and demographics. Even if many hit songs continue to rely on over-compression, there will still be a market for well-recorded songs that aren’t necessarily competing for the same listeners. It is unlikely that jazz, folk and classical recordings, for example, will feel as much pressure to over-compress as hard rock, hip hop and dance music.

The Norah Jones CD *Come Away with Me* sold over 10 million units despite (or perhaps partially due to) having less compression than the other discs released the same year. Its relative lack of processing also helped it win a Grammy for Best Engineered Album, Non-Classical. [20]

4.2.3. Communication About the Issue

To the extent that some of the participants in the loudness war can be characterized as reluctant combatants, increased awareness and discussion of the underlying issues may help persuade them to push back against the pressure to over-compress.

One of the more consistent results from game theory research is the finding that communication between participants can help in a number of ways: it provides an opportunity to appeal to the “right” or “proper” thing to do; it gives group members a chance to make explicit commitments and promises about what they will do; and it reinforces a sense of group identity. [82]

This suggests that group discussions at forums such as AES conventions can help solidify resistance to the loudness war. It also suggests that public pledges, such as the one organized by pleasurizemusic.com, may have some value. [69]

4.3. Educating People About the Actual Payoffs

In game theory, if players are misinformed about the costs and benefits of various strategies, they can build the wrong mental model and end up playing the wrong game [84]. For example, they may play the game as a dilemma when no dilemma exists, or they may play it as a zero-sum game like football or chess when the game is actually non-zero-sum (with the possibility of a win-win outcome).

If louder songs have a commercial advantage, leading to a loudness war in which everyone hypercompresses (effectively negating the advantage), and if the resulting lack of musical dynamics ends up hurting the overall music industry, this would be a classic social dilemma. But if players use hypercompression because of a mistaken belief about the commercial advantage of louder songs, this is no longer a dilemma, merely a tragedy.

In cases where individuals are misinformed about the actual payoffs, education can help encourage changes in behavior. Providing accurate information in the following areas may help de-escalate the loudness war.

4.3.1. Louder, Hypercompressed Music May Not Sell Better

The assumption for many years has been that louder is better in terms of sales, but this is only an assumption, perhaps driven by the zero-sum mindset of "beating the other guy" instead of focusing on making one's own recordings sound as good as possible.

It is difficult to do a true A/B test, where you release two different versions of the same recording (with and without hypercompression) to the same market at the same time and see which one sells better. Do louder recordings get more attention initially, but then suffer a rapid decline in sales due to lack of listener enthusiasm? Would *Come Away with Me* have sold even more if it had been 3 dB louder, or, conversely, would sales have benefited from additional dynamic range?

The Evergreen Project

In the Evergreen project, Johnson visually analyzed spectrograms from a number of the most commercially important albums of the last few decades and found that "the more strongly they sell, the more likely it is that they will have High Contrast characteristics," i.e., a wide dynamic range. Speaking of the album *The Eagles Greatest Hits 1971-1975*, he wrote that it's "gratifying, but unsurprising, to discover that the single most commercially important album in RIAA history contains some of the most striking dynamic contrasts pop music's ever seen.... people want dynamic contrasts." [2][85]

A limitation of this analysis is that cumulative sales figures effectively punish more recent (and likely more compressed) recordings [2]. An analysis limited to albums released during the same era and using sales rankings instead of absolute sales figures would minimize the effect of confounding variables such as the number of years the recording has been available and the decline of the overall music industry. In addition, the use of numerical dynamics metrics would allow statistical analysis of the data.

Commercial success vs. loudness

A dissertation by Dave Viney looked at 30 recent CD singles randomly selected from the UK charts and found no significant correlation between the measured loudness and sales chart position or weeks in chart. Ratings made by a listening panel of 36 producers and engineers showed a negative correlation between assessed "level of overall processing" and top radio and sales chart positions. Recordings perceived as having low amounts of processing were rated as sounding "more pleasant" and "above average quality" and were more commercially successful. However, the category of "overall processing" included all types of processing including EQ, not just compression. [86]

A study of sales rankings vs. dynamic range

Figure 7 shows an x-y scatter plot of scores based on sales rankings vs. dynamic range. The y-axis scores were derived from the *Billboard 200* year end charts for 2002 through 2009 (the years readily available on billboard.biz); these rankings track annual sales of music albums in the United States [87]. The rankings were then modified by subtracting the values from 201 and accumulating the scores of albums that appeared on the charts in multiple years. In this analysis, larger numbers imply more sales. For example, an album that ranked #200 in one year only would receive a score of 1; an album that ranked #1 three years in a row would receive a score of 600.

The x-axis "dynamic range" data were obtained from the "Unofficial Dynamic Range Database" at www.dr.loudness-war.info [21]. This site lets people post statistics analyzed using the TT Dynamic Range Meter, which quantizes dynamic range values to the nearest integer for ease of comprehension [5] [69]. While data from most of the albums in the *Billboard 200* year end charts have not yet been posted on the dynamic range site, 173 albums appeared in both databases.

In Figure 7, the (blue) diamonds represent the data points; for example, Metallica's *Death Magnetic* CD is shown at the left side of the diagram, with a score of 325 and a dynamic range (*DR*) of 3 dB. The black line is the trendline showing the least-squares best fit of the available data points; the trendline equation is given by *score* = 169.8 – 0.4421 x *DR*.

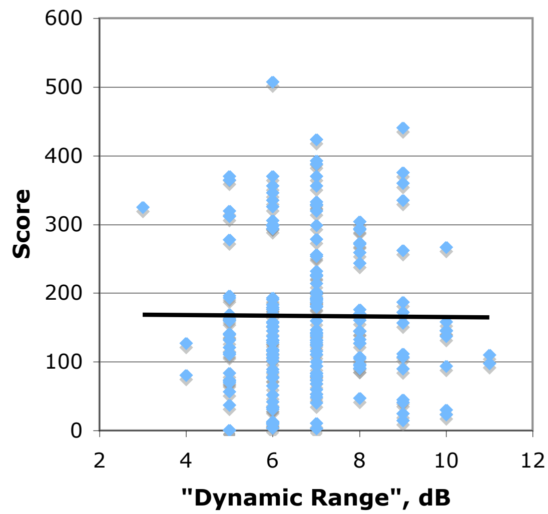

Figure 7. "Sales" scores as a function of quantized dynamic range. (Larger scores imply more sales.) The (blue) diamonds represent data points; the black line is the trendline.

The trendline shows a small negative correlation between dynamic range and sales; however, such correlations are most reliable when the coefficient of determination, r^2, is at or near 1. In this case, r^2 = 0.0000356, suggesting a very weak relationship between dynamic range and sales.

Visually, the data points appear to be widely scattered, which is confirmed by statistical analysis. The t-observed value is given by

$$t_{obs} = r\sqrt{\frac{n-2}{1-r^2}} = -0.0781,$$

where n = 173, the number of data points. Since $|t_{obs}|$ is much less than the α = 5% t-critical value of 1.974, the correlation between dynamic range and sales is not statistically significant at the 5% level (or, for that matter, even at the 90% level). Given an F-observed value of 0.00609, analysis of the two-tailed F distribution shows a 93.8% probability that a $|t_{obs}|$ value this high could have occurred by chance. [88][89]

Unlike Johnson's Evergreen study, which included albums from several decades, this study focuses on the peak hypercompression years of 2002 – 2009; it says little about how well "High Contrast" music might sell today, since the widest dynamic range in this data set was only 11 dB. However, it does suggest the possibility that, at least in the modern era, additional hypercompression may yield little or no sales advantage.

There are many possible flaws in this analysis; for example, the data sets are not perfectly symmetrical or normally distributed, the sales rankings are not a linear function of number of sales, participants in the dynamic range database were self-selected, and certain genres of music (e.g., country and hip-hop) seem to be underrepresented. Furthermore, the analysis uses dynamic range instead of loudness data, though the two are presumably inversely correlated. At any rate, despite a reasonably large number of data points, the available information fails to show a significant correlation between dynamic range and sales for albums on the Billboard 200 year-end charts.

This analysis does not prove that dynamic range is unrelated to commercial success. For one thing, it does not take into account albums that failed to reach one of the top 200 positions. In fact, if we compare all albums listed on the dynamic range database for the year 2009 to those that also appeared in the Billboard charts, we find an average *DR* of 8.3 for all albums vs. a *DR* of 6.7 for those that charted, suggesting that hypercompression might be the price for getting on the charts in the first place.

However, recall that correlation does not imply causation. One possible explanation is that producers may be more likely to hypercompress certain genres and albums that are expected to climb the charts; genres such as folk, jazz and classical may not feel the same pressure to boost loudness. Another possibility is that various music industry gatekeepers, who decide which songs get airplay and other types of promotion, may demand hypercompression due to a firm belief that consumers prefer louder songs. This belief is in no way confirmed by the above data and may simply be a type of groupthink; the lower dynamic range of albums that reach the charts may be due largely to self-fulfilling prophecy on the part of industry insiders. At any rate, the lack of significant correlation between dynamic range and sales for albums that did reach the charts raises questions about the assumption that listeners prefer louder songs.

Effect of loudness on listener's program choices

A recent and important study examined the effects of radio sound processing on listener's program choices. Maempel and Gawlik tested 60 non-expert subjects using various music pieces and one speech recording, processed with a variety of typical radio processing settings that varied in loudness and crest factor. One experiment performed a conventional test whereby the subjects could directly compare different types of processing using the same source material; in this test, as expected, there was a distinct preference for specific processing types, with an apparent advantage for high loudness and high bass. [90]

However, this type of direct sound comparison is never found in realistic broadcast situations. In order to provide some ecological validity, a second experiment allowed the source material and processing to be systematically co-varied. This time, except in the case of the speech recording, the type of sound processing had a marginal and statistically insignificant effect on listener's spontaneous program choices, which were strongly determined by the choice of source material. In this more realistic experiment, the hypothesis that listeners would prefer the louder processing could not be confirmed with respect to a medium effect size ($N = 358$). [90]

Thus, it appears that content is the primary factor, largely swamping any differences due to various types of compression. Furthermore, as the study reminds us, "In any case it should be pondered that such effects of loudness generally don't last very long. They disappear with the first spontaneous loudness correction by the listener, in contrast to the persistent loss of sound quality by the compression and the following risk of medium-term annoyance." [90]

If the perceived sales and listenership advantages of hypercompression turn out to be largely imaginary, this changes the game theory payoff matrix, possibly removing the dilemma entirely. If industry insiders are "playing the wrong game" based on false assumptions [82][84], education about the lack of evidence correlating loudness with listener preference and commercial success could help correct these assumptions and de-escalate the loudness war.

4.3.2. Hypercompression May Not Make Music Sound Louder on the Radio

As mentioned in Section 2.1.2, hypercompressed music typically does not sound louder on the radio, and may in fact sound quieter. Orban and Foti stated, "It sounds more distorted, making the radio sound broken in extreme cases. It sounds small, busy, and flat. It does not feel good to the listener when tuned up, so he or she hears it as background music. Hypercompression, when combined with 'major-market' levels of broadcast processing, sucks the drama and life from music." [14]

4.4. Changing the Payoffs

As mentioned, education about the actual game theory payoffs can help encourage behavioral changes. Another option is to change the game by providing new rewards and punishments. Changes to the payoff matrix can change the type of dilemma or remove the dilemma altogether, reducing or eliminating the incentive toward undesirable behavior. [80][82] For example, the establishment of standards helped remove incentives toward a loudness war in the movie industry.

The loudness war payoffs can be changed in two ways: by making loudness irrelevant so there's no longer an incentive to hypercompress, or by punishing excessive compression and rewarding effective use of dynamics. These strategies can be facilitated by the use of loudness and dynamics metrics, respectively. [9]

Loudness metrics, such as ITU-R BS.1770 [67] and R128's Programme Loudness [17], can be used to enable automated loudness normalization. Loudness normalization (Section 4.4.1) attacks the root cause of the loss of dynamic range: the attempt to achieve unnaturally high peak-normalized loudness levels.

Dynamics metrics, such as R128's "Loudness Range" [17], directly reflect the harm; i.e., the effect on the song's dynamics. These descriptors can help red-flag hypercompressed recordings.

[9] Additional information about loudness and dynamics metrics is available from [7], [91], [92], [93], [94] and references therein.

4.4.1. Loudness Normalization: Making Loudness Irrelevant

The studies in Section 4.3.1 suggest that loudness may already be largely irrelevant to listener preference and commercial success; unfortunately, this idea is contrary to widespread and firmly held assumptions in the industry. Automatically normalizing music to a defined loudness during playback would further remove any motivation for a continual escalation of loudness. If each broadcast or playback device imposes a specified loudness, there will be no competitive advantage to over-compressing the music in the first place. This would minimize the perceived benefit of hypercompression without any destructive effect on the music, since normalization is just a fixed gain. Loudness normalization would essentially make the recording's loudness, or at least its average long-term loudness, irrelevant.

Unfortunately, loudness normalization may not necessarily solve the problem unless its use becomes widespread. If loudness normalization is merely a selectable option (particular if the default is "off"), or if it is unavailable on a significant number of devices, there may still be some pressure to obtain a perceived competitive advantage by increasing the loudness during mastering. Furthermore, legacy devices such as existing CD players would not support loudness normalization. Even if loudness normalization does not become universal, however, its widespread adoption may still reduce the pressure to escalate the loudness war.

Katz wrote, "Your targets [for reversing the loudness war] should be the manufacturers of DVD players, CD players, iTunes, music servers (such as the Squeezebox) and so on. That's a more direct target [than record companies, etc.]. Make loudness normalization in these boxes a standard, not an option. Once the loudness has been normalized, the impetus to push loudness by overcompressing will go away." [95] Katz recommends the "K-System," an integrated metering and monitoring system, as a tool for transitioning from peak normalization toward loudness normalization [1].

Analyzed loudness or playback gain information can be saved with the digital track using a tag format such as ID3v2 [96] or stored on an online database. A recent article by Wolters et al proposed a non-destructive method of controlling playback loudness and dynamic range on portable media players [97]. Their proposal is based on the ITU-R BS.1770 standard [67] and is compatible with existing software and content following the Replay Gain proposal [98]. Products supporting this method have already been released by Dolby Laboratories. [97]

4.4.2. Punishing Hypercompression

Recently, some consumers have started making a point of their displeasure with excessively processed recordings by posting negative reviews on online sites such as the iTunes store and Amazon.com (e.g., feedback about the *Death Magnetic* CD [30] and consumer reviews of remastered recordings by the Rolling Stones and others). Since consumers commonly look at reviews before making an online purchase, music companies would need to consider whether the perceived loudness advantage of hypercompression would be worth the storm of criticism that would likely follow, especially if such reviews become more widespread.

4.4.3. Rewarding Engineering Excellence

The Grammy's Best Engineered Album awards help recognize outstanding achievement, but these are only given to one classical and one non-classical album per year. The AES or other organizations could begin offering a number of awards each year for best mastering, etc., taking into account dynamic range and overall sound quality. Such awards would recognize individual contributors, provide valuable publicity to the award-winning recordings, highlight the importance of mastering and audio engineering, and possibly help de-escalate the loudness war by providing an additional incentive to "do the right thing."

4.4.4. Standards

For cooperation to emerge in the iterated Prisoner's Dilemma, it is necessary for individuals to have information about how other parties have behaved [80] [82]. In the case of the loudness war, appropriate standards, tools and databases can arm consumers with impartial data and empower them to make informed decisions and provide useful feedback.

Defining standards

As mentioned, Dolby's de facto standards have contributed to loudness uniformity in the motion picture industry, and the EBU's R128 recommendation [17] aims to do the same for the broadcast industry. A professional organization or industry group could define loudness targets and dynamics standards for the music industry. Target loudness values could possibly decrease gradually over a number of years, allowing the loudness war to slowly de-escalate without causing a large immediate mismatch in volume between CDs from one year to the next.

Monitoring compliance

The cost of monitoring could be minimized by distributing the work among volunteers using a publicly editable collective database. Open source applications or plug-ins could automatically compute dynamics descriptors and upload them to the database whenever a user downloads music or rips a CD. The "TT Dynamic Range Meter" [5] and "Unofficial Dynamic Range Database" [21] can help illustrate how such a system might work, though preferably the descriptors should be based on accepted standards.

Enforcing standards

The loudness war could have been avoided given a central authority empowered to enforce loudness standards. Alternatively, in a perfect world, the major music companies would agree to abide by such standards, recognizing that an end to the loudness war would be in the best interest of the entire industry.

In the absence of either enforcement or voluntary agreements, music listeners can provide appropriate carrot and stick incentives in the form of consumer feedback to online music sites, as mentioned in Section 4.4.2. Given the apparent weakness of any correlation between loudness and commercial success, a relatively small amount of consumer pressure might outweigh any actual sales advantage achieved by over-compression.

5. RECOMMENDATIONS

5.1. Recommendations for Trade Organizations and Standards Groups

- Define recommended target loudness standards for CDs and other music releases.

- Define a dynamic range standard for recordings to receive a "High Dynamic Range" quality rating.

5.2. Recommendations for Manufacturers

- Implement loudness normalization as the default setting in audio playback devices such as CD and DVD players, portable media players, music servers, iTunes software, etc.

- Provide playback compression options, especially for devices like car radios that operate in noisy environments, to reduce the need to compress the masters so that quiet portions will always be audible.

5.3. Recommendations for Consumers

- Encourage manufacturers to include loudness normalization as the default setting.

- Encourage large online retailers such as iTunes (Apple) or Amazon.com to post dynamic range data with each recording, so interested consumers can use this as part of their purchasing decisions.

- Post audio quality reviews on music retailers' web sites.

5.4. Recommendations for Musicians, Producers and Mastering Engineers

- Consider allowing a little more dynamic range than is present in the average recording. This shouldn't make a significant sales difference, and if others do likewise, it could help gradually de-escalate the loudness war.

- Consider mastering a recording the way you would like to listen to it.

5.5. Recommendations for Record Companies

- Consider entering into an agreement with other major companies to follow recommended target loudness guidelines.

- Consider releasing premium audiophile editions with a higher dynamic range (such as the 2009 Beatles box set) for selected recordings.

- Small, independent labels could help grow niche markets of discerning listeners by displaying standardized seals certifying their “High Dynamic Range.”

6. SUMMARY AND CONCLUSIONS

6.1. Research Suggestions

A great deal of additional research is needed on virtually all of the topics discussed in this paper. In particular, formal listening tests are needed to determine the extent to which people can notice artifacts or aesthetic differences due to macrodynamic and microdynamic hypercompression as commonly used. It would be especially interesting to use real data to populate a graph of perceived quality vs. amount of compression for different types of listeners (such as the prototype shown in Figure 6).

Tests are also needed to verify whether hypercompression causes listening fatigue. In addition, study is needed on the effects of broadcast processing and lossy codecs on music that has been hypercompressed or clipped. Finally, a good deal more research is needed into the relationship between loudness, listener preference and commercial success.

6.2. Conclusions

In the music industry, the commonly believed and rarely questioned notion that “louder is better” has helped drive the loudness war from the beginning. There is certainly some truth to this maxim, but it appears to be a vast oversimplification.

A more accurate version of this saying might be: “All other things being equal, louder is better,” keeping in mind that, in general, other things are not equal. The industry may have over-extrapolated from studies showing a preference for the louder of two otherwise identical recordings. While listeners do prefer a louder version of the same recording, loudness does not seem to play a significant role when comparing different songs, nor does it appear to be significantly correlated to sales ranking.

When a listener is deciding what radio station to choose, it is rare that two stations will be playing the same song (conveniently time-aligned for easy A-B comparison), but at different levels. Instead, the choice will be between different songs with different types of processing. The decision about which to listen to is typically made on the basis of genre, melody, beat, vocal style, instrumentation, etc. Differences in loudness appear to play a relatively minor role, unless one of the songs is extremely quiet or excessively loud – note that either of the latter cases could cause the listener to change channels or reach for the volume control.

The research discussed above strongly questions the notion that “loudness trumps everything” [19]. If this idea were true, then Wagner would be considered a greater composer than Bach, Mozart and Beethoven, and a 747 jet would be a greater composer than Wagner. In general, *content* trumps loudness. The ear is more sensitive to things like pitch and rhythm than it is to relatively small differences in loudness. Likewise, the brain is generally more interested in melody, harmony, instrumentation, vocal quality, style, genre, lyrical content, spaciousness, texture, emotion, and other factors than it is in small loudness changes, especially since loudness can easily be adjusted using the volume control. We do not yet have knobs for most of these other factors.

Many of the speculations discussed in this article are difficult to prove. For now, intuition suggests that culture is important, music is important, and how we produce, distribute and preserve our music is important. I encourage well-intentioned people at each stage of the chain – musicians, producers, engineers, manufacturers, music executives and consumers – to take an active role in maximizing audio quality and preserving our cultural heritage.

(Additional material on this topic, including “Metrics for Quantifying Loudness and Dynamics” [7], originally intended as part of this paper, will be posted at http://sfxmachine.com/docs/loudnesswar/ .)

7. ACKNOWLEDGMENTS

The author wishes to thank Martin Walsh, Rob Toulson, Jon Boley and Bob Dennis for their many helpful comments and suggestions.

8. REFERENCES

[1] Bob Katz, *Mastering Audio: The Art and the Science*, Oxford, UK: Focal Press, 2007.

[2] Greg Milner, *Perfecting Sound Forever: An Aural History of Recorded Music*, New York: Faber and Faber, Inc., 2009.

[3] Tom Own and John Fesler, "Electrical Reproduction of Acoustically Recorded Cylinders and Disks," presented at the 70th AES Convention, New York, NY, 1981 October 30 – November 2.

[4] Arthur Collins and Byron Harlan, "Down Where the Big Bananas Grow," (audio recording), Edison Amberol 308, 1909. Archived at the Cylinder Preservation and Digitization Project, U.C. Santa Barbara, http://cylinders.library.ucsb.edu/search.php?queryType=@attr%201=1016&query=308&num=1&start=1&sortBy=&sortOrder=id, accessed 2010 August 14.

[5] "TT Dynamic Range Meter," Tischmeyer Technology, http://www.dynamicrange.de/ar/ar/download#menu1, accessed 2010 September 4.

[6] Earl Vickers, "Automatic Long-term Loudness and Dynamics Matching," presented at the AES 111th Convention, New York, 2001.

[7] Earl Vickers, "Metrics for Quantifying Loudness and Dynamics," http://sfxmachine.com/docs/loudnesswar/loudnessmetrics.html.

[8] Nick Southall, "Imperfect Sound Forever," *Stylus* (online), 2006, May 1. http://www.stylusmagazine.com/articles/weekly_article/imperfect-sound-forever.htm, accessed 2010 June 12.

[9] Richard Williams, *Phil Spector: Out Of His Head*, New York, NY: Outerbridge and Lazard Inc., 1972.

[10] Robert Dennis, personal communication, 2010 August 20.

[11] Robert Levine, "The Death of High Fidelity: In the age of MP3s, sound quality is worse than ever," *Rolling Stone*, 2007 December 27.

[12] Sarah Jones, "The Big Squeeze – Mastering Engineers Debate Music's Loudness Wars," http://mixonline.com/mag/audio_big_squeeze/ , accessed June 25, 2010.

[13] Robert Orban, "FM Broadcast Quality," *Stereo Review*, pp. 60-63, 1979 November.

[14] Robert Orban and Frank Foti, "What Happens to My Recording When It's Played on the Radio?", presented at the AES 111th Convention, New York, NY, 2001.

[15] Seth Stevenson, "Soft Sell: Why quiet, understated TV ads are so effective," Ad Report Card, Slate.com, http://www.slate.com/id/2242966, accessed 2010, July 18.

[16] Brian C. J. Moore, Brian R. Glasberg, and Michael A. Stone, "Why Are Commercials So Loud? – Perception and Modeling of the Loudness of Amplitude-Compressed Speech," *J. Audio Eng. Soc.*, Vol. 51, No. 12, 2003 December.

[17] "Loudness normalisation and permitted maximum level of audio signals," EBU Technical Recommendation R128, European Broadcasting Union, http://tech.ebu.ch/loudness, 2010.

[18] Ioan Allen, "Are Movies Too Loud?", presented at SMPTE Film Conference, March 22, 1997, http://www.dolby.com/uploadedFiles/zz-_Shared_Assets/English_PDFs/Professional/54_Moviestooloud.pdf, accessed June 26, 2010.

[19] Bob Katz, "Turn It Down! Consequences of the Ever-Escalating Loudness Wars," Workshop W20, AES 127th Convention, New York, NY, 2009 October 12.

[20] Bob Speer, "What Happened to Dynamic Range?", CD Mastering Services, 2001, http://www.cdmasteringservices.com/dynamicrange.htm, accessed 2010 June 18.

[21] "Unofficial Dynamic Range Database," http://www.dr.loudness-war.info/, accessed 2010 July 31 – August 6.

[22] Harvey Fletcher and W. A. Munson, "Loudness, its definition, measurement and calculation," *J.*

Acoustical Society of America, Vol. 5, pp. 82-108, 1933 October.

[23] Brian C. J. Moore, *An Introduction to the Psychology of Hearing*, Fourth Ed., London: Academic Press, 1997.

[24] Earl Vickers, "The Non-flat and Continually Changing Frequency Response of Multiband Compressors," presented at the AES 129th Convention, 2010 November 4-7.

[25] Jesse Lawson, "The Compression and Expansion of Musical Experience in the Digital Age," master's thesis, University of Vermont, 2008 October. http://library.uvm.edu/jspui/bitstream/123456789/164/1/Jesse%20Lawson%20MA%20thesis.pdf, accessed 2010 August 21.

[26] Earl Vickers, "Coping with Speech Noise in the Modern Workplace," white paper, The Sound Guy, Inc., 2006. http://chatterblocker.com/whitepapers/conversational_distraction.html, accessed 2010 June 12.

[27] "Maurice Ravel's Boléro: Masterpiece of experimental study?", Vienna Symphonic Library, http://vsl.co.at/en/65/71/1879/1531.vsl, accessed 2010 July 10.

[28] Susan Rodgers, "Turn It Down! Consequences of the Ever-Escalating Loudness Wars," Workshop W20, AES 127th Convention, New York, NY, 2009 October 12.

[29] Chris Neal, "Into the Red," *Music & Musicians*, 2010 January/February.

[30] Ian, "Re-Mix or Remaster Death Magnetic!", GoPetition, http://www.gopetition.com/petitions/re-mix-or-remaster-death-magnetic.html, accessed June 25, 2010.

[31] "The Loudness Wars: Why Music Sounds Worse," NPR, http://www.npr.org/templates/story/story.php?storyId=122114058, accessed June 25, 2010.

[32] Sean Curnyn, "Tears of Rage: The Great Bob Dylan Audio Scandal," *The Cinch Review*, http://www.cinchreview.com/bob-dylan-audio-scandal/570/ , accessed June 25, 2010.

[33] Richard Tollerton, "Waveform Plots Considered Harmful," *Audiamorous: An Audio Engineering Blog*, 2008 September 17, http://audiamorous.blogspot.com/2008/09/waveform-plots-considered-harmful.html, accessed June 26, 2010.

[34] Ellyn Sheffield, "Listener Fatigue and Longevity," Broadcast and Media Streaming Session B9, AES 127th Convention, New York, NY, 2009 October 11.

[35] Francis Rumsey, "If it's loud, does that mean it's bad?", *J. Audio Engineering Society*, Vol. 56, No. 6, 2008 June.

[36] Murad Ahmed and Kaya Burgess, "Young music fans deaf to iPod's limitation," *The London Times*, 2009 March 5. http://technology.timesonline.co.uk/tol/news/tech_and_web/personal_tech/article5847674.ece, accessed 2010 June 5.

[37] Jonathan Berger, "How IT Has Changed the Way People Think About Art," https://www.amacad.org/audio/mountain/berger.mp3, accessed 2010 July 12.

[38] Jonathan Lethem, "The Genius of Bob Dylan: An Intimate Conversation," *Rolling Stone*, 2006 September 7.

[39] Blair Jackson, "Issues in Modern Mastering," *Mix*, 2006, December 1. http://mixonline.com/recording/mastering/audio_issues_modern_mastering/, accessed 2010 July 17.

[40] Soren Nielsen and Thomas Lund, "Overload in Signal Conversion," presented at the AES 23rd International Conference: Signal Processing in Audio Recording and Reproduction, 2003 May.

[41] P. Kraght, "Aliasing in Digital Clippers and Compressors, *J. Audio Engineering Society*, Vol. 48, No. 11, pp. 1060-1065, 2000 November.

[42] Marvin Caesar, "Listener Fatigue and Longevity," Broadcast and Media Streaming Session B9, AES 127th Convention, New York, NY, 2009 October 11.

[43] Tim Carroll, "Listener Fatigue and Longevity," Broadcast and Media Streaming Session B9, AES

127th Convention, New York, NY, 2009 October 11.

[44] J. Robert Stuart, "The Psychoacoustics of Multichannel Audio," AES UK 11th Conference: Audio for New Media (ANM), 1996 March.

[45] Michael Gerzon, "Microphone Techniques for 3-Channel Stereo," presented at the AES 93rd Convention, San Francisco, CA, 1992 October 1-4.

[46] Michael Gerzon, "Psychoacoustic Decoders for Multispeaker Stereo and Surround Sound," presented at the AES 93rd Convention, San Francisco, CA, 1992 October 1-4.

[47] Jean-Marc Jot and Carlos Avendano, "Spatial Enhancement of Audio Recordings," AES 23rd International Conference, Copenhagen, Denmark, 2003 May 23-25.

[48] Michael Gerzon, "General Metatheory of Auditory Localisation," presented at the AES 92nd Convention, Vienna, Austria, 1992 March 24-27.

[49] Sam Berkow, "Listener Fatigue and Longevity," Broadcast and Media Streaming Session B9, AES 127th Convention, New York, NY, 2009 October 11.

[50] Frederick Ampel, "When Was the Last Time You Heard a Perfect Room?", AudioControl Technical Paper Number 108, 1995.

[51] Murray Barlowe, "The Relationship Between the Measurements of an Amplifier and Its Listening Quality," presented at the AES 13th Annual Meeting, 1961 October 9-13.

[52] Roy Allison and Edgar Villchur, "The Audibility of Doppler Distortion in Loudspeakers," presented at the AES 70th Convention, New York, NY, 1981 October 30 – November 2.

[53] Robert Ghent, Jr., Michael Nilsson, and Victor Bray, Jr., "Uses of Expansion to Promote Listening Comfort with Hearing Aids," presented at the American Academy of Audiology 12th Annual Convention, Chicago, Illinois, 2000 March 16-19.

[54] Jim Brown, "Systems for Stereophonic Sound Reinforcement: Performance Criteria, Design Techniques, and Practical Examples," presented at the AES 113th Convention, Los Angeles, CA, 2002 October 5-8.

[55] Ole Kirkeby, "A Balanced Stereo Widening Network for Headphones," AES 22nd International Conference on Virtual, Synthetic and Entertainment Audio, 2002 June.

[56] Floyd E. Toole, "Subjective Measurements of Loudspeaker Sound Quality and Listener Performance," *J. Audio Engineering Society*, Vol. 33, No. 1/2, 1985 January/February.

[57] Tobias Neher, Francis Rumsey, and Tim Brookes, "Training of listeners for the evaluation of spatial sound reproduction," presented at the AES 112th Convention, Munich, Germany, 2002 May 10-13.

[58] Michael A. Stone, Brian C. J. Moore, Christian Füllgrabe, and Andrew C. Hinton, "Multichannel Fast-Acting Dynamic Range Compression Hinders Performance by Young, Normal-Hearing Listeners in a Two-Talker Separation Task," *J. of the Audio Engineering Society*, Vol. 57, No. 7/8, pp. 532-546, 2009 July/August.

[59] James D. (jj) Johnston, "Listener Fatigue & Longevity," Broadcast and Media Streaming Session B9, AES 125th Convention, San Francisco, CA, 2008 October 4. (Presentation slides, "Listener Fatigue – Some Speculations," http://www.aes.org/sections/pnw/ppt/jj/fatigue.ppt, accessed June 26, 2010.)

[60] Stephanie Bor, Pamela Souza, and Richard Wright, "Multichannel Compression: Effects of Reduced Spectral Contrast on Vowel Identification, *J. of Speech, Language, and Hearing Research*, 51:1315-1327, 2008.

[61] Ellyn Sheffield, Mary Hinch, Daniel Schwab, and John Kean, "The Effects of Degraded Audio on Memory for Speech Passages," *IEEE Transactions on Broadcasting*, Vol. 55, No. 3, September 2009.

[62] Nick Southall, "Imperfect Sound Forever Revisited," *Stylus* (online), 2006 June 12, http://www.stylusmagazine.com/articles/soulseeking/imperfect-sound-forever-revisited.htm, accessed 2010 July 17.

[63] Rip Rowan, "Over the Limit," ProRec.com, 2002, August 31, archived at http://web.archive.org/web/20070810190708rn_1/www.prorec.com/Articles/tabid/109/EntryID/247/Default.aspx, accessed 2010 July 17.

[64] Ellyn Sheffield, John Kean, Mike Starling, Jan Andrews, Kyle Evans, and Sunny Khemlani, "Results from Subjective Testing of the HD Codec at 16–96 kbps," *IEEE Transactions on Broadcasting*, Vol. 52., No. 2, 2006 June.

[65] B. R. Glasberg and B. C. J. Moore, "A Model of Loudness Applicable to Time-Varying Sounds," *J. of the Audio Engineering Society*, Vol. 50, No. 5, pp. 331-342, 2002.

[66] James D. (jj) Johnston, "The Care and Feeding of Loudness Models," slide presentation, http://www.aes.org/sections/pnw/ppt/jj/loudmodel.ppt, accessed 2010, July 25.

[67] ITU-R, "Rec. ITU-R BS.1770, Algorithms to measure audio programme loudness and true-peak audio level," International Telecommunications Union, 2006.

[68] "Potential health risks of exposure to noise from personal music players and mobile phones including a music playing function," Scientific Committee on Emerging and Newly Identified Health Risks, European Commission, 2008, September 23.

[69] Pleasurize Music, http://www.pleasurizemusic.com, accessed 2010 June 5.

[70] Suhas Sreedhar, "The Future of Music," IEEE Spectrum, August 2007, http://spectrum.ieee.org/computing/software/the-future-of-music/0 , accessed June 25, 2010.

[71] Josef Shargorodsky, Sharon Curhan, Gary Curhan, and Roland Eavey, "Changes in Prevalence of Hearing Loss in US Adolescents," *Journal of the American Medical Association*, 2010 August 18.

[72] Thomas Maugh II, "Can you hear me now? Many teens can't," *Los Angeles Times*, 2010 August 17.

[73] Stephen Castle, "Study Warns of Hearing Loss from Music Players," *New York Times*, 2008, October 13.

[74] Barbara Cone, Melissa Wake, Sherryn Tobin, Zeffie Poulakis, and Field Rickards, "Slight-Mild Sensorineural Hearing Loss in Children: Audiometric, Clinical, and Risk Factor Profiles," *Ear & Hearing*, Vol. 31, No. 2, 2010.

[75] "Digital Music Increases Share of Overall Music Sales Volume in the U.S.," press release, The NPD Group, Inc., 2009 August 18. http://www.npd.com/press/releases/press_090818.html, accessed 2010 August 21.

[76] "A Big Music Year for Jackson, Boyle, Swift, Digital Downloads... and Vinyl?" http://blog.nielsen.com/nielsenwire/consumer/a-big-music-year-for-jackson-boyle-swift-digital-downloads-and-vinyl/, accessed 2010 August 21.

[77] "Billboard / Soundscan: Digital Album Sales Up 32% in 2008," NielsenWire, The Nielsen Company, http://blog.nielsen.com/nielsenwire/consumer/billboard-soundscan-digital-album-sales-up-32-in-2008/, accessed 2010 August 21.

[78] Kristina Dell, "Vinyl Gets Its Groove Back," *Time*, 2008 January 10. http://www.time.com/time/magazine/article/0,9171,1702369,00.html, accessed 2010 August 21.

[79] "FACTBOX: How vinyl record sales stack up," Reuters, 2009 December 8. http://www.reuters.com/article/idINTRE5B74QB20091208, accessed 2010 August 21.

[80] Robert Axelrod, *The Evolution of Cooperation*. New York, NY: Basic Books, 1984.

[81] Earl Vickers, "War and Games," *Creative Computing*, Vol. 10, No. 9, 1984 September.

[82] Peter Kollock, "Social Dilemmas: The Anatomy of Cooperation," *Annual Review of Sociology*, Vol. 24, pp. 183-214, 1998.

[83] Bob Katz, "Compression," http://www.digido.com/compression.html, accessed 2010 July 5.

[84] Giovanna Devetag, "Playing the wrong game: An experimental analysis of relational complexity and strategic misrepresentation," *Games and Economic Behavior*, Vol. 62, No. 2, 2008 March.

[85] Chris Johnson, "Dynamics," Evergreen albums analysis, archived at http://web.archive.org/web/20060820121159/www.airwindows.com/analysis/Evergreens.html, accessed 2010 July 12.

[86] Dave Viney, "The Obsession with Compression," M.A. dissertation, London College of Music, 2008, http://www.dynamicrange.de/sites/default/files/The%20Obsession%20project%20dissertation.pdf , accessed 2010 August 14.

[87] "Top 200 Albums," Year End Charts, *Billboard*, 2002 – 2009, http://www.billboard.biz/bbbiz/charts/yearendcharts/index.jsp, accessed 2010 August 1.

[88] Prem Mann, *Introductory Statistics Using Technology*, Hoboken, NJ: John Wiley & Sons, Inc., 2005.

[89] *Excel* (software), Microsoft Corporation, 2004.

[90] Hans-Joachim Maempel and Fabian Gawlik, "The influence of sound processing on listener's program choice in radio broadcasting," presented at the AES 126th Convention, Munich, Germany, 2009 May 7-10.

[91] Esben Skovenborg and S.H. Nielsen, "Evaluation of Different Loudness Models with Music and Speech Material," presented at the AES 117th Convention, San Francisco, 2004 October 28-31 .

[92] Esben Skovenborg and Thomas Lund, "Loudness Descriptors to Characterize Programs and Music Tracks," presented at the AES 125th Convention, San Francisco, CA, 2008 October 2-5.

[93] Esben Skovenborg and Thomas Lund, "Loudness Descriptors to Characterize Wide Loudness-Range Material," presented at the AES 127th Convention, New York, NY, 2009 October 9-12.

[94] Gilbert Soulodre and Scott Norcross, "Objective Measures of Loudness," presented at the AES 115th Convention, New York, NY, 2003 October 10-13.

[95] Bob Katz, correspondence, quoted in "The life and death of dynamic range: who decides how loud," presented by Rob Toulson and Bill Campbell at Proceedings of The Art of Record Production Conference, Cardiff, Wales, 2009 November.

[96] ID3, http://www.id3.org/, accessed 2010 July 12.

[97] Martin Wolters, Harald Mundt, and Jeffrey Riedmiller, "Loudness Normalization in the Age of Portable Media Players," Presented at the AES 128th Convention, London, UK, 2010 May 22-25.

[98] http://www.replaygain.org, accessed 2010 June 5.