# 6

# The Software Passes the Test When the User Fails It

## Constructing Digital Models of Analog Signal Processors

*Jonathan Sterne*

In a prescient 1987 essay, John Mowitt noted a changing tendency in attitudes toward signal processing that accompanied the growing adoption of digital technologies for making, storing, and transmitting music. He wrote that "the fetish of noise reduction has gone hand in hand with the aggressive marketing of distortion boosters and other less obvious instrumental sources of noise."[1] Digital audio recording had been sold in terms of its promises of perfect fidelity to an original source and low signal distortion, but it turned out that these marketing and engineering points were not universally desired by users.[2] Recording engineers spent the 1970s and earlier fighting against the specific noise signatures of electroacoustic devices such as compressors, tape machines, and amplifiers. They devised clever workarounds and strategies to minimize the noise and signal transformation inherent in the technologies they used.[3] Now, in 1987, Mowitt noted, the drive was not just for perfect definition, but for noise and distortion as sonic effects, from the crackle of an LP record grafted onto a hip-hop recording made with a digital sampler, to simulations of tape distortion, to simulations of complex amalgams of musical instrument gear—amplifiers, compressors, echoes and delays, distortion pedals.

With thirty years' retrospection, it is clear that what Mowitt presented as an irony of digital signal processing for music was actually its historical condition. Digital signal processing could only make sense by marking itself out as distinct from analog signal processing. Once there was more than one way for everyday musicians and engineers to process and reproduce sound, aspects of the older way came into relief as aesthetic choices. If, as the martial metaphor goes, engineers of the 1960s and 1970s fought a battle against noise and distortion, we would have to ultimately characterize it as a losing battle. The result is that to a trained ear, music recorded in a particular era or a particular

Jonathan Sterne, *The Software Passes the Test When the User Fails It* In: *Testing Hearing.* Edited by: Viktoria Tkaczyk, Mara Mills, and Alexandra Hui, Oxford University Press (2020). © Oxford University Press. DOI: 10.1093/oso/9780197511121.003.0007.

studio environment such as Motown, Abbey Road, or Muscle Shoals had a distinctive sonic signature. That sonic signature was, of course, a combination of many things—the engineers, the equipment, the musicians, the music—and yet it is often plainly audible to listeners. For most, these aspects of sound simply fold into "the music itself," as an aspect of timbre. And timbre, it turns out, is central to the reception and meaning of much recorded music in the twentieth and twenty-first centuries, especially popular music. Even a single note (or sometimes even less) from a well-known popular song is recognizable to listeners who know it, a feat made possible by the variability of instrumentation and production techniques across studios.[4] Musicians and engineers are well aware of this phenomenon—and often (though not always) more sensitive to it than nonmusician audiences. In the context of music production and performance, they may use equipment as a shorthand for whole histories of production, performance, style, and reception, for instance naming the Roland Space Echo to reference a whole history of dub music. This is not unique to popular music or to the twentieth century, as Emily Dolan's piece in this collection demonstrates. The cult of the Stradivarius violin is not that distant from the cult of the Stratocaster guitar.

From stand-alone digital recorders to integrated computer and software systems, early commercial digital technologies presented an unexpected problem for their users because they were at first not legible within this timbral and technical history. "Early digital" now has its own retrospective sound and production style that is a result of a combination of equipment and engineering trends in the 1980s—but it would not necessarily have been apparent to listeners at the time: "It can be compared with somebody who moves into a new house. The first time he looks through the window he only sees the beautiful view. After a few days he detects a small flaw in the glass and from that moment he cannot look through the window without seeing that flaw."[5] If the "perfect sound forever" marketing campaigns of early digital audio promised musicians liberation from the sonic limits of previous generations of recording and signal processing technology, the lived reality is that many people loved their limitations, because they shaped the sound of the music.[6]

In the intervening decades, scholars have noted an "analog revival" in response to digital audio,[7] but an equally remarkable phenomenon has occurred in the world of digital audio itself. Through an extensive labor of translation, hardware and software designers have found ways to model—that is, to emulate or, if full emulation is not possible, at least approximate—analog signal processing technologies in the digital domain. Modeling has become one of the staple offerings of the music technology industry, where one device imitates some aspects of another. "Analog modeling" actually refers to a host of approaches, from attempts to model the behavior of complex circuits down to the component level, to taking actual recordings of analog signal processing devices or even physical spaces and convolving them with other digitally recorded sounds (a process I explain below), to attempts to replicate a vague "feel" or "vibe" using analog devices as inspiration, to just adding depictions of "wood panels" to the window for a software plug-in.

In this chapter, I analyze the process by which engineers in the commercial music technology industry model analog signal processing in the digital domain.[8] Based on participant observation as well as research into the history of the technologies I am writing about, I describe the ways in which digital models at once test the hearing of machinery—how a given technology transduces, registers, and represents sound to itself and to human auditors—and use the machinery to test the hearing of users. The chapter loosely follows the development of two different scenarios. First, I consider a model of a spring reverb device, the AKG BX20, at Universal Audio in Scotts Valley, California. I became involved with Universal Audio as part of a larger ethnographic project on signal processing technologies and the people who produce them. Since the early 2010s, I have visited dozens of companies, laboratories, and workshops. I focus on Universal Audio because they embody a particularly strong example of a set of ideas about modeling and testing, and because of the access they granted me to their modeling process. Following current ethnographic practice, I will sometimes refer to myself (and my hearing) within the narrative so that I do not create the illusion that by dint of my authorship of this chapter I am also a universal or perfect auditor.[9] I then consider modeled amplifiers for electric guitars developed by Line 6 and other modeling companies. For the BX20, I examine the production of a model because I was able to participate in it; for the guitar amplifiers, I consider the reception of models, especially as they move across different user communities.

Most discourse around modeling is producer discourse, which is to say musicians and engineers are primarily talking with one another when they talk about modeling, though it is useful to distinguish people who produce music from people who produce models, even if many of the people who make models also use them. Even if they are themselves musicians, audiences and fans mostly spend little time worrying about the qualities of this or that signal processing technique (if they are even aware of them at all). They will instead simply talk about the music, and *maybe* the sound of the music. When modeling changes the *look* of musicians' equipment, this is sometimes a cause for commentary, a scenario to which we will return later.

## Digital Audio Models of Analog Devices: Tests and Definitions

While every company, and every engineer, has to some extent a unique approach, there is a lexicon of common practice and terminology that exists above and beyond particular implementations. Like testing hearing for other purposes (see Mara Mills and Viktoria Tkaczyk in this volume), the testing of digital models is characterized by a quest for precision and for the quantification of previously unquantified dimensions of sound as it is transduced out of the sonic domain and back again. But unlike in medical or telephonic contexts, a deliberately aesthetic dimension enters into the engineering process when the goal is making music. Telecommunications has its aesthetics as well, to be sure, but the aesthetic aspect is highlighted in the engineering of sound recording and signal processing technology for music. Because "sounding good"—a recurrent trope in engineers' talk—is a foremost concern, we gain a particularly useful insight into the politics of transduction: how cultural, historical, and economic relations are rendered in the sonic realm, and how dimensions of sound and sound technologies come to have value. To "sound good" is to invoke a history of sounds and practices, which implicitly values some aspects of those practices over others. Whether it is a test of how a signal processor receives and renders a signal or a test of how a listener responds to a digital model of an analog device, the moment of testing hearing is where these relations are negotiated, refracted, and brought to life. As part of the emergence of a digital device or piece of software, listening tests—whether strict A/B/X tests or more performative comparisons—help engineers define what is essential to a sonic technology and what is superfluous.

A/B and A/B/X tests are among the most common listener testing regimes for audio technologies for music. Listeners are given two sounds, "A" and "B," and asked whether they can tell the difference. The A/B/X test gives the listeners sounds A and B, and then a third sound, "X," which they are asked to identify as A or B. If the user is right half the time or less, they are guessing and the two sounds are considered indistinguishable. A/B/X tests have a long history in consumer audio. It is not clear when they emerged as the standard for testing digital models of analog processors, but some of the earliest entrants into the digital modeling business, such as Native Instruments, were already doing user testing with A/B/X in the late 1990s. People who build digital models of analog processors will use this kind of testing in constructing the model, but also in final quality assurance testing before releasing a program to the general public.[10]

Listening tests matter to engineers because they design digital models to reproduce the sonic characteristics and processing behavior of other devices (which they then model in code). The digital model may or may not reproduce other aspects of the experience of using the technology. All digital models have software dimensions, and some digital models exist purely as software. A digital model of an analog device such as a guitar amplifier or reverberator might exist in the following forms:

- As a plugin to enhance the functionality of other software, for example, in digital recording and mixing
- As a piece of stand-alone computer software
- As a smartphone app
- As a stand-alone hardware device—a physical box—with digital signal processing built into it
- As a physical box or device that shares the functionality of the analog device and has similar controls and interface elements
- As a physical box or device that shares the functionality of the analog device but looks completely different

To be sure, design matters. It is part of the user experience of the model, and hardware characteristics such as knobs and materials, or graphical interfaces, or the lack of either in a command-line software environment imply a whole set of intended meanings and user scripts. Yet it is important to note that sonically speaking, the model can be completely accurate even if it looks or feels nothing like the device it is said to model. This range of interpretations is most clear if we consider software models. Software may look entirely different from the analog device it models, with interface features that do not match the analog interface at all, or it may have a skeuomorphic appearance, mimicking the device's look. A skeuomorph is, following Katherine Hayles, "a design feature that is no longer functional in itself but refers back to a feature that was functional at an earlier time."[11] Hayles treats skeuomorphism as a visual phenomenon, and graphic user interfaces generally do as well. But as I show in this chapter, it is both a visual and sonic phenomenon. Mowitt observed that when confronted with new, "improved" technologies, musicians and artists immediately began finding ways to get them to behave—and sound—like older technologies. To understand the relation between visual and sonic skeuomorphism, consider the interfaces of three software-based analog-modeling filters. Figure 6.1 shows a skeuomorphic interface for a digital model of the Roland SH-101 synthesizer, a picture of which is set right below it. Although the layout is not identical, the same functions are contained in


**Figure 6.1** Togu Audio Line's Bassline 101 versus a Roland SH101. The "VCF" appears in roughly the same place in both the device interface and the skeuomorphic software interface.

Author's screenshot, from *The Prodigy Live*, http://theprodigy.info/equipment/sh-101.shtml.

the top row of controls for both; the labels and look are maintained. Thus, the layout of controls for the software version will be familiar to anyone who has used the hardware synthesizer, and the filter (the section labeled VCF) is located in the same position as on the control panel of the analog device. This is a classic example of skeuomorphism, like the "desktop" metaphor used in graphical operating systems for computers with its "file folders" and "recycling bins."

Software models may also depart from any specific analog hardware. Figure 6.2 shows a SoundToys plugin called FilterFreak (top) that is supposed to model different behaviors of analog filters. It *looks* like a


**Figure 6.2** Diminishing skeuomorphism: SoundToys FilterFreak versus Ableton Auto Filter.
Author's screenshots.

skeuomorphic representation of a specific analog device, but it represents no device in particular. The screenshot shows knobs meant to represent adjustable parameters, frequency curves set inside frames meant to resemble oscilloscope screens, and on/off toggles represented as switches, and the interface even has gratuitous "wood panels" on its sides, evoking the look of 1970s audio hardware. In contrast, the Auto Filter plugin for Ableton Live (the lower part of Figure 6.2), which performs similar functions to FilterFreak, has no obvious visual similarity to any hardware, real or imagined.[12] There are still circles, squares, and frequency curves, but no attempt is made to have them look like knobs, buttons, or screens. Instead, they are presented to the user in a flat, easy-to-render, Bauhaus-inspired design and paired with standard graphical user interface elements such as pulldown menus (which FilterFreak conceals behind buttons). Yet the Auto Filter plugin does also model analog hardware. The "PRD" setting in the middle indicates that its algorithm is set to emulate a resonant four-pole ladder filter similar to that on a Moog Prodigy synthesizer. Command-line-based audio programs such as Pure Data, which are controlled entirely through text commands and have no graphical user interface at all, also have the means to model the workings of the Moog. Although one *can* have a more or less skeuomorphic interface for digital models of an analog device, it is not *necessary*.

But what exactly *is* a digital model of an analog device? Every term in that sentence raises thorny historical and epistemological questions. Recent scholarship has problematized the concepts of *analog, digital,* and *model*. For the purposes of this chapter, I will keep close to the actors' categories, which are very much *not* precise analytical definitions. In the music technology world, *analog* refers to a retrospective definition that encompasses all audio technologies before or without a digital element. However, it should be noted that not all technologies grouped under this definition are actually analog audio technologies, and not all of them are analog in the same way. This retrospective usage of the term "analog" actually works to legitimate the digital technologies that are said to come after them, by lumping everything outside digital technology into a single, if incoherent, category. *Digital* is generally assumed to refer to any technology with a microprocessor, dating it to the 1970s and integrated circuits—but this too is an actors' category that elides the ideological work of the definition, since in many cases the things other than the microprocessor (including some of the things that are at the core of what a technology does—sending sound out of a speaker, say) do not exist in the digital domain at all, and digital technologies do not necessarily need to have processors. Like the term "analog," the term "digital" cannot be used without an attendant value judgment: critics will refer to digital sound as "cold" or "lifeless," while other users will praise digital technologies for their miniaturization, affordability, control options, flexibility, or verisimilitude.[13]

As for *model*, historians and philosophers of science have endlessly debated what models are and do. Models may "aim to lay bare the essential principles according to which this or that domain of phenomena operate,"[14] but in signal processing contexts, they tend to do so dynamically, so that relations may change in real time. That, indeed, describes how people in the digital modeling industry talk about what they are doing. What counts as essential or superficial is a social and cultural question—and thus the model does epistemological and cultural work at the same time.[15]

The final term in the formulation, "device," has been the subject of much less discussion in recent scholarship. In science and technology studies, it is perhaps most famously associated with Albert Borgmann's "device paradigm," whereby

> those aspects of properties of a device that provide the answer to "What is the device for?" constitute its commodity, and they remain relatively fixed. The other properties are changeable and are changed, normally on the basis of scientific insight and engineering ingenuity, to make the commodity still more available. Hence every device has functional equivalents, and equivalent devices may be physically and structurally very dissimilar from one another.[16]

For Borgmann, devices are fundamentally commodities, and though I might quibble as to whether this is sufficient for a *general* model of the device, the referents of most digital models, as well as the models themselves, are certainly commodities. Not only are they bought and sold, but they are subject to commodity fetishism, and, as Louise Meintjes has shown, a great deal of studio practice involves musicians acting on their beliefs in the magic or power contained within equipment.[17]

Borgmann finds that the separation of means and ends in technology raises a serious political problem, because device status is designed to occlude both the inner workings and the social workings of a technology. Borgmann's device paradigm thus partakes of what writers in the science and technology studies tradition have called "black-boxing," what writers in cinema studies have called the "concealment of the apparatus," and what writers in the Marxist tradition have called "reification."[18] All of these terms have different theoretical implications and political resonances, but they all focus on defining which aspects of a technology are to be in the foreground of users' attention and which are to be hidden from users. Thus, in the definition of both the analog *device* and the digital *model*, the circumscription of what counts as "inside" and "outside" the device is what makes the model possible. The listening test plays a crucial role in defining the analog technology retrospectively and the digital model prospectively, since the test performs the divisions of inside and outside, consecrating different aspects of the technology, its sound, and the experience of using it as essential while denigrating other aspects as superficial to the model.[19]

## On Reverbs in Closets

It is February 2012, and I am sitting with my laptop at a makeshift desk in a storage room in Universal Audio's headquarters in Scotts Valley, California. Universal Audio is best known for its work in digital modeling and for reviving old analog equipment. Its reputation was built around the LA-2A leveling amplifier and the 1176 limiting amplifier, both of which were prized by recording engineers for the specific ways they behaved as they affected signals. When it was refounded in 1999, Universal Audio (originally founded in 1958) began manufacturing replicas of these old units, right down to the photocell in the LA-2A's circuitry that used to modulate sound. But the company also built digital models of these and many other devices to be used as software plugins, to enhance the functionality, experience, and sound of music mixing programs such as ProTools. These general-purpose programs, called Digital Audio Workstations (DAWs), record multitrack audio and MIDI data (Musical Instrument Digital Interface—a control protocol for various kinds of devices to communicate with one another) and facilitate mixing, editing, and creative signal processing in real time. Almost every commercial music recording made today passes through a DAW at some point in its production process, and usually the DAW plays a central technical role.

My host is Jonathan Abel, a mathematician and signal processing expert who contributed to Universal Audio's original LA-2A and 1176 digital models.[20] Abel is something of a historian of these devices, having spent a lot of time listening to different models and versions, visiting them in professional recording studios, and discussing their use history with recording engineers. He helped to build Universal Audio's paradigm (along with his colleague Dave Berners and several others) and to set up a course in signal processing for audio at Stanford's Center for Research in Computer Music and Acoustics. Abel and I have talked over his approach to digital modeling at length. He has been generous with his time, explaining the math of analog and digital signal processing to me, and I have sat in on his class. For Abel, there is always more to talk about, another layer to peel away, another pool of questions to dive into. He is fascinated by the smallest mathematical details of what these devices do to sound, and he has a deep sense of aesthetic appreciation for the gear, understanding himself as part of a lineage of signal processing researchers. He frequently spoke about the philosophical questions raised by the mathematical and operational challenges of modeling.

I have hooked up my laptop to an audio interface, which is connected to an AKG BX20 reverberator (AKG is the company and BX20 is the model number). The AKG BX20 is a behemoth of a device: large, heavy, wooden, difficult to move. First released in the late 1960s, it is a classic example of a truly analog device: it uses the behavior of an ensemble of materials—electricity, coils, magnets, springs—as an analog of the behavior of sound in a room containing a microphone and a sound source. Once sounds have been converted to electrical signals, they can be sent through the BX20, where the diffusion behavior of the spring adds a sense of audible ambience to the sound, as if it resounded in a chamber or hall. Controls allow the user to damp the spring, blend the spring's "wet" reverberate signal with the "dry" signal that is not run through the spring, and modulate the volume of the sound.

The BX20 has two channels, so it can work in stereo. Transduced into electricity, sound enters the BX20 through a cable plugged into another device. It is then amplified and transduced into mechanical vibrations, which are sent into a spring, converted back into electricity at the other end of the spring, sent to another amplifier and combined with the input signal at the output. This allows the user to choose how much of the processed signal to combine with the original signal. The feedback of the spring system is also transmitted along with the initial sound as modulated by the moving spring. Once out of the BX20, the combined signal can then be run into a mixing board, amplifier, and speakers and transformed back into audible sound.

BX20s were popular in recording studios from the late 1960s on, adding a signature ambience to various famous recordings (today, this can be heard on albums by artists such as Norah Jones and Jack White). Along with other mechanical reverbs, they were often part of a studio's mystique or signature sound. Lore even developed around some specific, individual devices, such as the Echoplate III (a competitor to the BX20) at Muscle Shoals studio, which supposedly benefited from the particular humidity of Sheffield, Alabama, and contributed to the sound for which the studio was known. When I encountered it, the unit had been relocated to a studio space inside a house in suburban Huntsville, Alabama. I confess that upon hearing it in April 2015, I was not immediately able to distinguish its special characteristics in contrast to other plate reverbs I have heard in person or on recordings before or since. This raises a host of issues around comparing sounds to one another. Did the device need the humidity to have a special sound? Were my own shortcomings as a listener the explanation for why I could not hear the difference? Or was the entire scenario a kind of commodity fetishism, condensing all sorts of history, work processes, musical practice, and recording practice into a single device in the suburban Huntsville closet before me? It is impossible to know precisely why that particular Echoplate III is able to retain its significance to a community of users.

Back in that storage room in California in 2012—also really more of a big closet—Jonathan and I are sending signals through the BX20 and recording them. I am using Ableton Live, a popular recording and playback program that is a cross between a traditional DAW and a sampler, allowing the flexible, fast, and repeated playback and recording of samples of recorded audio. We use three files, set up as samples in my program, that Jonathan has given me: a pulse that makes a "ping" sound; a short sine sweep through the whole audible spectrum that yields a sort of springy “boing”; and a longer sine sweep that sounds like a very broadband siren firing up.

Jonathan changes the settings on the AKG; I run the sound through and record it. We listen back to the recording. After a run through the BX20 on a particularly long setting, he smiles at me and admires the way the sound decays. An ear for the subtle dimensions of timbre and time is a part of the modeling process. If modeling is a producer discourse, the biggest connoisseurs of all are the people making the models. Nearly everyone I met at Universal Audio had an appreciation for the old analog gear they were modeling. Like many digital audio companies I visited—Line 6, Native Instruments, Elektron, Teenage Engineering (to name a few)—Universal Audio has a special space in its building stocked with prized examples of vintage sound equipment, a kind of shrine to the sonic and signal processing histories of recorded music in the second half of the twentieth century. Although Universal Audio specializes in the signal processing part of the recording and mixing process, its studio contains plenty of vintage keyboards, amplifiers, and musical instruments.

For all this mythologizing of equipment inside the actual spaces that modeling companies inhabit, treating the equipment as a stand-in or condensation of musical and aesthetic histories, the modeling process is fundamentally a process of demystification. This is crucial: even though finished software products are black-boxed from users—you cannot know the algorithm that governs the behavior of Universal Audio’s BX20 unless you sign a nondisclosure agreement and have a “need to know”—and fascination with old audio gear depends on a layer of mystification, the work of building digital models of analog devices requires every “mystery” of an old analog device such as the spring reverb to be explained or classified.[21] The engineers at Universal Audio want to know how and why the BX20 produces the sounds that it does, and the recordings we took are a path into that process, if not exactly the first step. Taken together, the sounds I sent through the BX20 tested its behavior—how it hears sound and how it plays sound back. By measuring the difference between the files we used to test the device and the recording, by sending the original signal through the device at different settings, Jonathan and the engineers at Universal Audio can construct a model of what the BX20 is doing to the sound and then build a working model able to imitate it.

The file that is the result of the difference between the sounds we sent into the BX20 and the sounds we recorded out of it provides a baseline for evaluating what the unit does to sound. The difference between the two files is treated as an “impulse response” that can be applied to other signals. If the BX20’s impulse response is applied to a dry recording (say, of a voice in a dampened room or studio) through a process called convolution, it will sound as if that voice went through the BX20. In digital signal processing, convolution multiplies the spectra of two audio signals to combine them, expressing their overlap: for instance, if you convolved a recording of my voice with a filter that removes all the low frequencies of my voice, the result would be a sound of my voice with the low frequencies attenuated. But convolution is time invariant; it does not change over time. Since it is based on a recording of a particular sound at a particular time, it cannot give anything other than a snapshot of the BX20's behavior.[22] Rather than representing the chaotic behavior of the spring inside the BX20, it represents one instance of the behavior of the circuitry inside the BX20, a single sonic performance.

Thus, convolution begins the modeling process but does not finish it. In the months following my visit, working with his former graduate student Sean Coffin, Abel first tested the impulse responses taken from the BX20 against the same sounds being run through the BX20. They were seeking a "perceptual match," where the listener cannot tell the difference between a convolved signal and a signal actually sent through a BX20. The listeners here are the people working on the algorithm—Abel, Coffin, friends of both, employees of Universal Audio, and others who may be brought in for listening tests.[23] Like other kinds of listening tests, these test both the user and the technology at the same time. The goal of the listening test is for the technology to pass it, not the user. If the impulse response is good enough, the user will fail the test. If the user passes the test and can tell the two technologies apart, then the impulse response has failed. As Dave Berners, Universal Audio's chief scientist, explained:

> To me the function of listening is to find bugs, and that's my own opinion. We have a really great . . . person at UA, Will Shanks, who does . . . qualification that's separate from our quality assurance team. He'll do A/B tests; sometimes he'll bring in other people to do subjective listening tests. But he's really the person who makes sure that things are going the way we want them to. And he provides us a lot of feedback. He'll tell us qualitatively what's going on when he listens to something. But in my mind, the way it fits into the design process is, if there's something that he can tell that's different, it means that there's a bug. I don't want to have him say, well, it should be a little bit brighter, and then I'll say, okay, I'll just put a little EQ and make it a little brighter. It's not an iterative thing where we try to, based on listening, converge towards something. When the design gets to that stage, it should already be converged. It should be a model of the real process and be identical. And if it isn't, then if something sounds different, the way that it sounds different can tell us what is likely to be wrong. So it's really useful to have the listening feedback [to tell us] what way the sound is not right, it tells us where in the algorithm to look

> for a mistake. But nevertheless I still feel like the listening's there to find problems more than it is to nudge things or tweak them. There's very few cases where we do any design based on the perception. I would feel really vulnerable if that were the case.[24]

Trevor Pinch notes that technological tests are usually understood as performances to be witnessed by others: Will Shanks comes through in Dave Berners's explanation as a kind of super listener, someone who stands in for future audiences. In other cases, a computer program can just as easily serve as a witness for this kind of test as a person can: the goal is simply to know whether an imagined future listener can tell the difference or not. Witnessing the test can be delegated, but there must be a human or technological witness.[25] The goal here is to establish a logic of sonic equivalence between two sounds: a digital recording fed into an algorithm on one side, and a physical device on the other. According to the logic of the test, the equivalence only happens when the two sounds are indistinguishable—but in fact the testing scenario itself, with its A/B/X structure and "this or that" choice, *already* establishes a logic of equivalence. As Berners puts it, "it should be a model of the real process and be identical." It does not matter that one contender is a giant box with springs inside and the other is a piece of code that takes up an infinitesimal space on a hard drive. Those differences are excluded from the test before the fact.

Here, sound is the basis of commensurability between two operationally distinct technologies, and the test performs that moment of commensuration. In the test, the sound—really the hearing of the sound—is what creates the relationship between the two other elements.[26] That is why, for Berners, iterative design or aesthetic judgment would make the test and the engineer "vulnerable": they would undermine the possibility of equivalence. The devices modeled by Universal Audio, such as the 1176 or BX20, are already widely respected in the world of professional audio engineering, and are tied to famous recordings and the sounds achieved by famous musicians. Iterative listening and design has already occurred. The model simply consecrates the tangled musical and social relations that resonate inside the sound. Put simply, the modeling process decrees that "good sound" has already been decided by a history of practice.

## Contested Models: Guitar Amplifiers

The condition of commensurability necessary for digital models of analog devices is also one of the most contested dimensions of the whole enterprise. What constitutes a satisfactory model, for whom, and under what circumstances? When can a musician, artist, or recordist substitute a digital model for the device that it models? My visit to Line 6, a company that made its name with digital models of famous guitar amplifiers, helped me to understand what can happen when different user populations have different ideas about commensurability. An analog guitar amplifier sends sound as an electrical signal through a series of processing stages, where electrical operations have an *analogous* relationship to sonic processes: electrical clipping leading to harmonic distortion, the roar or grind of a distorted electric guitar. As with the studio equipment discussed in the previous section, specific brands and models of guitar amplifiers are famous for the characteristics they impart to sounds.

While the BX20 is targeted for a single use—mixdown of prerecorded material—the guitar amplifier model here is targeted to at least three very different user bases that crisscross recording studios, live performance, and home practice: studio engineers who are not guitarists and are seeking to change the sound of a guitar performance; guitarists who are recording or practicing at home in situations where using a guitar amplifier at full volume would be problematic; and performing guitarists who want the flexibility of having many more types of amplifiers, with significant savings in cost and space (some amplifiers are not only expensive but also larger than refrigerators). For an engineer, an amplifier is another flavor of signal processing that goes into the recording. For an electric guitarist, the amplifier is literally part of the instrument. Whereas the sound generation and amplification on an acoustic guitar are already separate—the strings and fretboard make the sound, the body amplifies it—on an electric guitar sound generation results from a considerably more complex system. Electrical pickups (usually magnets) turn the vibrations of the strings into electrical signals, which are sent down a cable to an amplifier, which processes the sound and then sends it out a set of speakers (indeed, when guitarists speak of the "amp," they are often speaking synecdochically of the amplifier-speaker system). Without the amplifier, the guitar makes a sound, but not very much of one. With the amplifier, the guitar can become something like a controller, as single notes ring out and have impossible sustain, harmonics above and below sounded notes are synthesized and emphasized, and modulations such as built-in spring reverbs (smaller than the AKG) change the sound.

Played at moderately high volumes, the guitar/player/amplifier/room tetrad forms a kind of cybernetic system, a situation crystallized in the concept of *feedback* that is so central to both guitar playing and cybernetic theorizing. Guitar feedback, where the sound of the amplifier vibrates the strings and is detected by the pickups, which in turn send it to the amplifier for further amplification, is an instantiation of the more general cybernetic concept of feedback.[27] This feedback question goes to the heart of why some guitarists reject digital models of amplifiers that otherwise would pass listening tests as indistinguishable from the analog device. Those who reject the digital models will often complain that they don't *feel* the same, though this aspect of feeling is almost impossible to quantify.

The difference between a software model and a physical device in the room can be important for someone playing an instrument, even if it is not audible to a third party. It is one thing for two recordings to be "perceptually equivalent"—for a listener to be unable to tell the difference between, for instance, a Marshall amplifier recorded live and a dry signal run through a piece of software that algorithmically models the behavior of the Marshall amp. For a mixing engineer who is not a guitarist, this equivalence is sufficient. For an audience member listening to the recording, it is also sufficient. But for a guitarist used to a set of physical interactions with an amplifier as part of their instrument, it may not be enough. Marcus Ryle, Line 6's president and cofounder, explained to me:

> In the end, the only way people can actually decide whether we did a good job or not is to listen to the actual amps we modeled, which may not sound the same as the amps you own. But we do double-blind listening tests here with outside folks. . . . We've done it with everyone from artists to enthusiasts. We did it once at the LA amp show, where these are the real aficionados of boutique tube amps. And an interesting by-product that happens, I think at that show close to half the people we offered to did not want to take the test. The test was simply, here's a guitar, there's two amplifiers behind this sheet. Here's your A/B foot switch, just play. Switch as long as you want, play whatever you want and just identify which is the modeling amp. And it seems that there's people that don't want to do that. That gets into issues far beyond technology, right? And the fact that for many people, it's not possible for us to hear with just our ears.[28]

Ryle's "far beyond technology" shows the definitional work attached to the A/B or A/B/X test. By defining the visceral, felt, and unheard dimensions of guitar playing as "far beyond technology," Ryle is defining them as not relevant to the model. The distinct impression I got from our conversation was that the problem was one of mystification—that guitarists believe there is some magic in the devices, whereas to the engineer, they can be explained and those explanations can be operationalized in a model.

There is another possible interpretation, however. It is entirely possible for the sound of a model to be indistinguishable from the sound of an analog amplifier, for the behaviors around harmonic distortion or tone shaping to be identical, and for there still to be a difference. For one thing, the A/B/X test only tests the person's hearing, not the guitar: the pickups "hear" the string (in some languages, the word for "pickup" is the same as "microphone" even though they work on different mechanical principles), but in a room they also recycle the signal from the amplifier, forming the basis of the feedback loop. Even if the sound of an amplified electric guitar can be recorded and recreated with completely different equipment, the experience may not be the same for the guitarist, and this matters precisely because modeling is so saturated with producer discourse. The embodied experience of playing at a lower volume, without a physical amplifier in the room, will be entirely different from the experience of playing at high volume. In *Sensing Sound*, Nina Sun Eidsheim uses the term "sonic reduction" to describe understandings of sound that treat it as a disembodied phenomenon. In the case of playing an electric guitar, sounding and hearing are multisensory phenomena for both the human being and the equipment. If the embodiment of the signal processor matters, then so too must the body of the listener or the user.[29]

The notion of the embodied guitarist is of course a highly gendered one—not to mention one weighted with racial and sexual overtones.[30] Amplifiers also carry tremendous symbolic freight. A wall of amps for a large-scale concert once synecdochically represented the same kinds of power and mastery over nature that control over the guitar did. Their sheer size and sonic force reinforced the imagery of the individual musician as masterful and in control. For quite a long time, however, large guitar amplifiers have simply been unnecessary for most commercial music performance: they are louder than needed, they are less controllable than the installed mixer and amplifier setup found in performance venues, and, for the touring musician not rich enough to hire a fleet of roadies and ship gear across oceans, they are also expensive and difficult to move. Thus, in recent years, performing musicians in some genres have moved away from amplifiers, but in every instance I have found, their alternative is an intentional or inadvertent commentary on the symbolism of the now-absent wall of amplifiers, whether the replacement is a stack of washing machines, eye-searing LED displays, or giant artworks (Figure 6.3).[31]

Thus, to say that the presence or absence of an amplifier is a purely sonic question is to miss the cultural work that amplifiers do for both musicians and audiences. Line 6 deals with this problem of different needs for different user bases by making software plugins, digital modeling boxes that look nothing like amps, and various kinds of amplifiers. As a business, Line 6 takes a pragmatic approach rather than committing to a single philosophy. But for many users, the "device" meaning of amplifier goes beyond its sonic

**Figure 6.3** Extreme metal band Meshuggah in front of a custom stage backdrop.
Photo courtesy of Carrie Rentschler.

characteristics—to gendered, embodied allusions to power and control that are intimately tied up with musical subcultures, genre identities, and the experience of making music. This is precisely the kind of collapse we often find in technologies: the device itself exists within the web of much larger and more sophisticated social relationships—commercial, financial, artistic, experiential, interpersonal. While there is much more to be said about the gender politics of guitar playing, especially given the increasing prominence of women and genderqueer guitarists in recent years, to treat the sound as the "thing" in amplification is clearly to single out only one part of the process. Whether or not there is an audible sonic difference between a Line 6 model and the analog amplifier it models, for many users there remains an irreducible cultural difference that must be negotiated one way or another.

## Proof Is Not Enough

If listening tests are the moment of proof for a digital model, they do not form the only basis on which digital models are built. Impulse responses are not sufficient either, because of their nature as single samples of an otherwise dynamic process. This is where different schools of analog modeling diverge. Some treat the impulse response as the "truest" representation of an analog device, because it bears a relationship of cause and effect to the specific device—this is the approach of the company Acoustica Audio with its Nebula software, and is that found in the secondary market of sample libraries of impulse responses that can be loaded into digital modeling hardware or software. This approach works fine if the analog device is "linear and time invariant," which is to say that the relationship between the input and output of the system is a linear map (given several inputs to the system with corresponding outputs, the output that corresponds to the sum of these inputs will be the sum of their corresponding outputs) and the system is time invariant (whenever you use the system, you will get the same result). But many analog systems are not linear and time invariant. The relationship between inputs and outputs of the analog system may not be a linear map and may not be time invariant—for instance, if a system behaves differently as it heats up (which is the case for tube amplifiers and for some analog synthesizers). Because of the possibility of these nonlinearities, other companies and users argue, a sampling-and-convolution-based approach is often insufficient on its own because analog devices vary their output over time and over settings. One solution is to add carefully calibrated degrees of randomness to convolution processing, which produces nonlinearity, but not in the way that an analog device would produce it. Universal Audio takes a different line, modeling the behavior of the different elements of the device to produce a digital model that both sounds like the analog device and processes like the analog device. There is a good degree of translation involved—the math for digital signal processing and the math for analog signal processing are not the same. But through years of research and experience, companies making digital models, such as Universal Audio and Line 6, build up sets of processing routines that describe different aspects of the way analog devices operate. The amplifier example shows the limits of the sonic model, and yet engineers go to great lengths to match interface and behavioral elements of a digital model to its analog referent, within the limits of the digital domain and interface.

I return to my work with Universal Audio to explain how this happens. After the BX20's impulse responses were judged to be perceptually equivalent, Abel and Coffin essentially reverse-engineered it. Using the signal flow diagram underlying the BX20, patents, schematics, and prior research, Abel and Coffin tried to create an algorithm that would reproduce the behavior of the BX20 at equivalent settings. As Abel explained to me: "We made a computational model of the impulse responses that would perceptually reproduce the measurements at the measured knob positions."[32] In other words, using a skeuomorphic interface that looked like a two-dimensional representation of the BX20's control panel (Figure 6.4), their digital model *should work just like* the analog BX20.

A visual representation of a knob turned to 3:00 on Abel and Coffin's model should work just like a knob turned to 3 on the analog device. Of course, "work just like" has to be operationalized here, and the only two options for operationalizing it are listening tests and measured response curves. Again, the testing scenario defines functions for the device and its digital model.

But so does prior work in the field. Abel had previously been involved in modeling another famous device, the Roland RE-201 Space Echo—a signature psychedelic effect used in everything from dub reggae to progressive rock to techno. Part of the Space Echo's "spaceyness" comes from a spring reverb attached to its main function, a tape delay. Abel and Dave Berners had spent considerable time studying spring physics and modeling the physics of the spring inside the Space Echo, developing a theory of how it transformed sound and a mathematics to represent and operationalize that theory in Universal


**Figure 6.4** Universal Audio's skeuomorphic panel.
Universal Audio website, http://www.uaudio.com/uad-plugins/reverbs/akg-bx-20-reverb.html.

Audio's model.[33] Thus, when confronted with the BX20, Abel already had a working model of a spring and could begin by adjusting it for the BX20's much larger spring system. This was also true of the BX20's other components, such as the amplifiers. Universal Audio essentially had a library of building blocks and, crucially, a set of submodels of how those different elements of technologies interact. This process of developing models of components and relations among components converts signal processing into a kind of story that can be told about a device: if it has these kinds of op-amps, and those kinds of potentiometers, then Universal Audio's engineers expect it will behave in a certain way.

To be clear, a narrative is not a universal, ontological condition of signal processing, but it may be a common social precondition for digital models and anything that involves reverse-engineering or repair.[34] Engineers need a way to explain what happens inside a circuit, which has both a temporal dimension (a series of events occur) and a spatial dimension (a topology or shape). Narratives, stories of what happens when, why, and how, become the glue for keeping together understandings of the parts of a circuit in an account of its behavior as a whole. Each submodel—how a spring works, how a resistor works—contributes to engineers' un-black-boxing of the hardware device. The narrative is the moment that identifies its completion. So when Universal Audio's engineers assembled their operational submodels and tuned them, they constructed a working behavioral model of the BX20, which was simultaneously a story of how it worked. At that point, they aimed to get the behavioral model to produce sounds that were perceptual matches to the impulse responses of the original BX20. Thus, the tests that Abel, Berners, and Ryle all spoke of were tests of the device as well as of the story of the device. In testing, a chain of equivalence is set up:

| | | |
|---|---|---|
| Analog BX20 at a given setting | ←listener→ | Impulse responses at a given setting |
| Impulse responses at a given setting | ←listener→ | BX20 operational model |
| Analog BX20 | ← → | BX20 digital model |

As Berners explained, this final round of testing is "quality assurance"—a kind of epistemic sealant. The goal is perceptual equivalence at the level of hearing. The software passes the test when the user fails it.

But another set of equivalences is also being shored up, because both the visual and sonic dimensions of the BX20 are rendered skeuomorphically. Universal Audio's business model is to depict the analog device as much as possible in the digital interface. The knobs on the digital model of the BX20 interface are a perfect example: they are not necessary for the software code to do its work on the transduced sound. Given that most users of Universal Audio software are working with computer mice, a slider would probably be easier to use, but because the interface *looks like* the BX20, it helps to reinforce the idea that the interface *sounds like* the BX20. Employing visual skeuomorphism to represent sonic commensurability is a long-standing practice in the modeling world, and is a longer part of the history of commercial music technologies. Leo Fender used auto-body paints for guitars in part because they were widely available and in part to associate his guitars with other kinds of commodities; Tara Rodgers has shown that wood panels on synthesizers come out of a tradition of wood-paneling other consumer goods to make them seem more "organic" and draw attention away from the social relationship in which they are embedded. The same can be said here: a skeuomorphic interface, as opposed to an interface designed around user practice in software, creates an analogy of practice for users and a logic of equivalence—one may be substituted for the other.[35]

This substitutionalist logic has been in operation for most of the history of commercial use of digital audio software. Even something as basic as the visual representation of an analog mixing board in ProTools, a popular digital mixing program, follows this logic.[36] As Evan Brooks, ProTools's co-inventor, explained to me, "To have a separate mixer was really just an attempt to gain—to give people some familiarity with the process. When you're moving from analog over to a digital way of doing things we didn't completely change their overall view, and so we felt that—people's workflow was still kind of, at least mentally if nothing else, divided into concepts of tracking and then mixing."[37] Brooks here is referring to the mixer window in ProTools, which is meant to represent in two dimensions on the computer screen the physical controls of a hardware mixer.

Visual skeuomorphism is rhetorical; it is a story about the ordering of tasks and operations in the creative process. In preserving the task sequence from analog device to digital model, the modeling industry can be said to be inherently conservative. It uses a rhetoric of democratization—more people have more access to more tools at lower cost—and this is probably true.[38] In that context, the model performs a kind of canonizing process, aestheticizing a set of relations into "workflow"—record, then mix; reverberate this way, not that way. To offer "something familiar" is to begin to set the terms of how musical work is supposed to be done, and the use of models points back to tradition. This conservatism is not an inherently bad thing—tradition is an important part of music education in almost all contexts, and practice and imitation is an easy way for beginners to learn outside of formal educational contexts. At another level, the conservatism actually masks a transformation of relations; behind the avowed conservatism of modeling is industrial competition and change. Brooks's "attempt to give some familiarity" links together two technical practices that have no necessary, preordained relationship and suggests a relationship of descent and morphology in the gesture: "See, listen, it is the same as . . . ." Other software companies have gone in other directions, with other consequences.

Thus, procedurally, to create a digital model of an analog audio device, engineers follow a process of defining, testing, refining, and redefining, with listening tests at every stage of the process. But socially, to create a digital model requires establishing relations of equivalence at each stage. Listening—whether done by people or machines—is a crucial part of this practice, but so is the development of a model of *how* the analog device "does what is does." As it is developed, the model circumscribes essential dimensions of the analog device and brackets off nonessential elements. At different stages, listening tests establish equivalence and consecrate it.

We might therefore be led to ask how much of the sound of a digital model is essential and how much is superficial—can digital models of analog devices be said to be *sonic* skeuomorphs as well as *visual* skeuomorphs? Or to put it another way, does a difference matter if you can't hear it? It may matter in a host of ways that are deliberately set aside in the moment of testing. Perhaps the difference matters aesthetically, to the degree that aesthetics are not reducible to measurable perception. But the differences that matter may also lie further out—bodily, yes, but also technologically, culturally, politically, and economically. *Who* gets to signal process and under what conditions is a central question of media theory, and the very question that is left aside at the moment of the listening test.

Further materials related to this chapter can be found in the database "Sound & Science: Digital Histories": https://acoustics.mpiwg-berlin.mpg.de/sets/clusters/testing-hearing/digital-models-sterne.

## Acknowledgments

In addition to audiences at the Max Planck Institute for the History of Science and the International Society for Intermedial Aesthetics, I wish to offer thanks to Carrie Rentschler, Ky Brooks, Gabriella Coleman, Emily Dolan, Mara Mills, Alix Hui, and Viktoria Tkaczyk for their comments on previous drafts, Jonathan Abel for educating me about signal processing and for the access needed to make this paper happen, and Kate Sturge, who is the most thorough editor I know.

## Notes

1. John Mowitt, “Sound of Music,” 194.
2. Of course, the idea of perfect fidelity was itself a fiction, but that is a topic I have covered at length elsewhere. Jonathan Sterne, *Audible Past*, 215–86.
3. Susan Schmidt Horning, *Chasing Sound*.
4. Recounted in Daniel Levitin, *This Is Your Brain on Music*, 151. See E. Glenn Schellenberg, Paul Iverson, and Margaret C. McKinnon, “Name That Tune.”
5. Christer Grewin and Thomas Rydén, “Subjective Assessment.”
6. Robert Poss, “Distortion Is Truth”; Albin J. Zak, *Poetics of Rock*; Simon Frith and Simon Zagorski-Thomas, *Art of Record Production*.
7. Trevor Pinch and Frank Trocco, *Analog Days*; Tara Rodgers, “Synthesizing Sound”; Ian Dunham, “From Kitschy to Classy.” It should be noted that this is not strictly an analog/digital phenomenon. Musicians have long prized older equipment for its supposed characteristics. See Emily Dolan’s chapter in this volume and H. Stith Bennett, *On Becoming a Rock Musician*.
8. All the models considered here were designed by for-profit companies for sale to end users: musicians, audio engineers, producers, podcasters, and others. There is considerable crossover between commercial and academic contexts in music technology research (probably a bigger distinction than between nonprofit and for-profit research), but that is beyond the scope of this chapter. As I discuss later, the profit motive goes partway to explaining some aspects of analog modeling—branding and marketing most notably—but those alone are not sufficient to explain the cultural significance of the phenomenon.
9. This approach is drawn from the anthropology of sound technology. See Louise Meintjes, *Sound of Africa;* Stefan Helmreich, *Alien Ocean*; and David Novak, *Japanoise*, for examples of ethnographers reflexively positioning themselves with respect to their subjects and objects.
10. Author interview with Stephan Schmitt, October 27, 2011. See also Jonathan Sterne, *MP3*, 150–73.
11. N. Katherine Hayles, *How We Became Posthuman*, 17.
12. The two plugins arguably do sound different, but whether that difference is audible in a full mix is a question that would have to be answered on a case-by-case basis.
13. Georgina Born, “Computer Software”; Wendy Hui Kyong Chun, *Control and Freedom*; Derek Robinson, “Analog”; Jonathan Sterne, “Analog”; Benjamin Peters, “Digital.”
14. Lorraine Daston, *Things That Talk*, 226.
15. Joseph Klett, “Baffled by an Algorithm,” 123.
16. Albert Borgmann, *Technology*, 43.
17. Meintjes, *Sound of Africa*, 73–74.
18. Madeleine Akrich, “De-Scription of Technical Objects”; Theresa Hak Kyung Cha, *Apparatus, Cinematographic Apparatus*; Georg Lukács, *History and Class Consciousness*.
19. The recent contest around the terms other than *device* suggests that we need to revisit its meaning as well, especially given its ambiguity in relation to the French *appareil* and *dispositif*, two terms that have been so central to cultural theory of the past generation. See Michel Foucault, “Questions of Method”; Giorgio Agamben, *What Is an Apparatus?*; Karen Barad, *Meeting the Universe Halfway*.
20. “Original digital models” may seem like an odd locution, but it is correct here. Universal Audio has revised its models since Abel’s work on an earlier version of the software,

building on the paradigm he helped establish but making use of the increased processing power inside computers.

21. Strictly speaking, it would be possible to model *some* aspects of a device, and it would also be possible to model a device based purely on measuring inputs and outputs. But in practice, because it is engineers doing the modeling, they tend to reverse-engineer anything that they plan to model.
22. For more on convolution, see Jonathan Sterne, "Space Within Space."
23. I discuss the concept of "expert listeners" brought in for listening tests in Sterne, *MP3*, 163–73.
24. Author interview with Dave Berners, May 13, 2011. All interviews are lightly edited from spoken language for readability.
25. Trevor Pinch, " 'Testing—One, Two, Three . . .," 26.
26. Dylan Mulvin, "Reference Materials," 12.
27. Norbert Wiener, *Human Use of Human Beings*; Steve Waksman, *Instruments of Desire*; Novak, *Japanoise*, 139–68.
28. Author interview with Marcus Ryle, July 15, 2011.
29. Nina Sun Eidsheim, *Sensing Sound*, 2, 120–29. As someone who has used feedback in my own music, I have experienced the interaction between a vibrating instrument and an amplifier, and the only way I can describe it is that it is like shaping electricity and sound together, as if they were clay on a rotating potter's wheel. But it is certainly possible to achieve this with digital audio technologies, so long as there is also an amplifier.
30. Waksman, *Instruments of Desire*, 188; Mavis Bayton, "Women and the Electric Guitar."
31. Other examples are Rush (see http://www.cygnus-x1.net/links/rush/mobile-productions-monthly-09.2011.php, or search for "Rush tour washing machines") and Sleater-Kinney (https://www.inlander.com/spokane/concert-review-sleater-kinney-goes-big-on-tour-opening-night-at-the-fox/Content?oid=18425517, or search for "Sleater Kinney tour LED").
32. Jonathan Abel, email to author, August 5, 2015.
33. Jonathan Abel and Dave Berners, "Digital Emulation."
34. It is no accident that old repair manuals for analog gear are an important resource for engineers who model analog equipment. Whether to model something or to fix it, one needs an understanding of how the component parts fit together. Steven Jackson, "Rethinking Repair."
35. Tara Rodgers, "Into the Woods."
36. See, for example, https://www.youtube.com/watch?v=qIj9WTTkA5o, or search for "Pro Tools mixer window."
37. Author interview with Evan Brooks, February 15, 2012.
38. Walter Benjamin, "Work of Art."

## References

Abel, Jonathan, and Dave Berners. "Digital Emulation of the Roland RE-201 Space Echo." Lecture at the Centre for Interdisciplinary Research in Music Media and Technology, McGill University, November 2, 2006.

Agamben, Giorgio. *What Is an Apparatus? And Other Essays*. Translated by David Kishik and Stefan Pedatella. Stanford, CA: Stanford University Press, 2009.

Akrich, Madeleine. “The De-Scription of Technical Objects.” In *Shaping Technology, Building Society: Studies in Sociotechnical Change*, edited by Wiebe E. Bijker and John Law, 205–24. Cambridge, MA: MIT Press, 1992.

Barad, Karen. *Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning.* Durham, NC: Duke University Press, 2007.

Bayton, Mavis. “Women and the Electric Guitar.” In *Sexing the Groove*, edited by Sheila Whitely, 37–49. New York: Routledge, 1997.

Benjamin, Walter. “The Work of Art in the Age of Mechanical Reproduction.” Translated by Harry Zohn. In *Illuminations*, edited by Hannah Arendt, 217–51. New York: Schocken, 1968.

Bennett, H. Stith. *On Becoming a Rock Musician*. Amherst: University of Massachusetts Press, 1980.

Borgmann, Albert. *Technology and the Character of Contemporary Life*. Chicago: University of Chicago Press, 1984.

Born, Georgina. “Computer Software as a Medium: Textuality, Orality, and Sociality in an Artificial Intelligence Research Culture.” In *Rethinking Visual Anthropology*, edited by Marcus Banks and Howard Morphy, 139–69. New Haven, CT: Yale University Press, 1999.

Cha, Theresa Hak Kyung. *Apparatus, Cinematographic Apparatus: Selected Writings*. New York: Tanam Press, 1980.

Chun, Wendy Hui Kyong. *Control and Freedom: Power and Paranoia in the Age of Fiber Optics*. Cambridge, MA: MIT Press, 2005.

Daston, Lorraine. *Things That Talk: Object Lessons from Art and Science*. New York: Zone Books, 2004.

Dunham, Ian. “From Kitschy to Classy: Reviving the Roland TR-808.” *Sounding Out! The Sound Studies Blog*, June 9, 2014, https://soundstudiesblog.com/2014/06/09/808/.

Eidsheim, Nina Sun. *Sensing Sound: Singing and Listening as Vibrational Practice*. Durham, NC: Duke University Press, 2015.

Foucault, Michel. “Questions of Method.” Translated by Colin Gordon. In *The Foucault Effect: Studies in Governmentality*, edited by Graham Burchell, Colin Gordon, and Peter Miller, 73–86. Chicago: University of Chicago Press, 1991.

Frith, Simon, and Simon Zagorski-Thomas, eds. *The Art of Record Production: An Introductory Reader for a New Academic Field.* Burlington, VT: Ashgate, 2012.

Grewin, Christer, and Thomas Rydén. “Subjective Assessment of Low Bit-Rate Audio Codecs.” Paper presented at the 10th International Conference of the Audio Engineering Society, London, September 1991.

Hayles, N. Katherine. *How We Became Posthuman: Virtual Bodies in Cybernetics, Literature, and Informatics.* Chicago: University of Chicago Press, 1999.

Helmreich, Stefan. *Alien Ocean: Anthropological Voyages in Microbial Seas*. Berkeley: University of California Press, 2009.

Horning, Susan Schmidt. *Chasing Sound: Technology, Culture, and the Art of Studio Recording from Edison to the LP*. Baltimore, MD: Johns Hopkins University Press, 2013.

Jackson, Steven. “Rethinking Repair.” In *Media Technologies: Essays on Communication, Materiality, and Society*, edited by Tarleton Gillespie, Pablo J. Boczkowski, and Kirsten A. Foot, 221–39. Cambridge, MA: MIT Press, 2014.

Klett, Joseph. “Baffled by an Algorithm.” In *Algorithmic Cultures: Essays on Meaning, Performance and New Technologies*, edited by Robert Seyfert and Jonathan Roberge, 111–27. London: Routledge, 2016.

Levitin, Daniel J. *This Is Your Brain on Music: The Science of a Human Obsession.* New York: Plume, 2006.

Lukács, Georg. *History and Class Consciousness: Studies in Marxist Dialectics*. Translated by Rodney Livingstone. Cambridge, MA: MIT Press, 1971.

Meintjes, Louise. *Sound of Africa!: Making Music Zulu in a South African Studio*. Durham, NC: Duke University Press, 2003.

Mowitt, John. "The Sound of Music in the Era of Its Electronic Reproducibility." In *Music and Society: The Politics of Composition, Performance and Reception*, edited by Richard Leppert and Susan McClary, 173–97. New York: Cambridge University Press, 1987.

Mulvin, Dylan. "Reference Materials: The People, Places and Things of Making Measurements." PhD diss., McGill University, 2015.

Novak, David. *Japanoise: Music at the Edge of Circulation*. Durham, NC: Duke University Press, 2013.

Peters, Ben. "Digital." In *Digital Keywords: A Vocabulary of Information Society and Culture*, edited by Ben Peters, 93–108. Princeton, NJ: Princeton University Press, 2016.

Pinch, Trevor. "'Testing—One, Two, Three . . . Testing!': Toward a Sociology of Testing." *Science, Technology, & Human Values* 18, no. 1 (1993): 25–41.

Pinch, Trevor, and Frank Trocco. *Analog Days: The Invention and Impact of the Moog Synthesizer*. Cambridge, MA: Harvard University Press, 2002.

Poss, Robert M. "Distortion Is Truth." *Leonardo Music Journal* 8, no. 1 (1998): 45–48.

Robinson, Derek. "Analog." In *Software Studies: A Lexicon*, edited by Matthew Fuller, 21–31. Cambridge, MA: MIT Press, 2008.

Rodgers, Tara. "Into the Woods: A Brief History of Wood Paneling on Synthesizers." *Sounding Out! The Sound Studies Blog*, May 30, 2011, https://soundstudiesblog.com/2011/05/30/into-the-woods-a-brief-history-of-wood-paneling-on-synthesizers/.

Rodgers, Tara. "Synthesizing Sound: Metaphor in Audio-Technical Discourse and Synthesis History." PhD diss., McGill University, 2011.

Schellenberg, E. Glenn, Paul Iverson, and Margaret C. McKinnon. "Name That Tune: Identifying Familiar Recordings from Brief Excerpts." *Psychonomic Bulletin & Review* 6, no. 4 (1999): 641–46.

Sterne, Jonathan. "Analog." In *Digital Keywords: A Vocabulary of Information Society and Culture*, edited by Benjamin Peters, 31–44. Princeton, NJ: Princeton University Press, 2016.

Sterne, Jonathan. *The Audible Past: Cultural Origins of Sound Reproduction*. Durham, NC: Duke University Press, 2003.

Sterne, Jonathan. *MP3: The Meaning of a Format*. Durham, NC: Duke University Press, 2012.

Sterne, Jonathan. "Space Within Space: Artificial Reverb and the Detachable Echo." *Grey Room* 60 (Summer 2015): 110–31.

Waksman, Steve. *Instruments of Desire: The Electric Guitar and the Shaping of Musical Experience*. Cambridge, MA: Harvard University Press, 1999.

Wiener, Norbert. *The Human Use of Human Beings: Cybernetics and Society*. New York: Doubleday, 1954.

Zak, Albin J. *The Poetics of Rock: Cutting Tracks, Making Records*. Berkeley: University of California Press, 2001.