Some parentheses around algorithmic composition

WARREN BURT

Melbourne, Australia
E-mail: waburt@melbourne.DIALix.oz.au

Algorithmic composition and gambling, two activities associated (in the public mind) with the use of chance, are contrasted. Gambling is based on the concepts of winning and losing. Algorithmic composition is not, or should not be. Problems of mappings of information from one medium to another are considered, along with problems of reception for artworks made with these methods. In the end, the quality of attention given to an artwork may be more critical to its reception than any methods used to construct it. *Mr. Yasser's Piano*, *Tuning the Furniture of Chaos* and *Pi and the Square Root of Two*, three recent algorithmic compositions by the author, are considered in the light of these thoughts.

Many people over the years have said to me, half jokingly, 'Since you're so involved in randomness in art, you ought to go to the track (play the lottery, go to the casino, etc.). That will be a *true* test of your randomness.'

I always laughed at/with this comment – its very absurdity and incorrect logic always provided an ironic giggle. One reason for the laughter, though, was the fact that I was (and am) almost totally uninterested in gambling. This was underscored on a recent trip to Louisiana. My host, composer and computer graphics artist Catherine Schieve, invited me to visit Casino Rouge, one of two casinos in Baton Rouge that have been built into old Mississippi riverboats. The casino was everything she promised, and more – garish, gaudy, noisy, manic, and with a gloriously chaotic sound environment to die for. We also played the 25 cent and 5 cent slot machines (hey, yeah! high rollers, us!): starting with $5, she won $10, then quit. Starting with $3, I won $4, then lost it all. A bit of mild fun to get into the spirit of the place, but, for me, hardly of any lasting or compelling interest.

All around us, though, were hundreds of people for whom the slots and tables obviously *were* of lasting and compelling interest. The air of frantic anxiety was so thick you could smell it, even through the haze of cigarette smoke. At one point Catherine's friend, art historian Carl Volkmar, observed that most of the people on the boat probably liked their music and art really regular and straight, but were dedicated to the use of chance in their financial lives, while he, Catherine and I were very cautious with our finances, but used chance extensively in our art. I replied that I remembered that John Cage had said that the use of chance should be restricted to areas where you wouldn't get hurt. So you could use chance to pick notes, but you shouldn't use it with mushrooms. Nor, I then added, in a burst of (tediously middle-class?) moralism, should you use it with the grocery money, which it appeared to me that many of the people on the boat were doing.

Later, though, it occurred to me that there was an essential difference between gambling and algorithmic composition, and that a misunderstanding of that difference was what prevented many people from understanding the latter, because the main use of chance they are familiar with is the former.

(This phenomenon is not just limited to algorithmic composition. Since 99 per cent of the photography that people see is in advertising – photography used to sell something – when people see art photography, they may (perhaps falsely) assume that the photographer, like the advertising photographer, is using their photography to sell the subject matter of the photograph. This, of course, may not be the case – the photographer may intend critical comment, ironic comment, no comment, or simply be interested in some formal visual aspect of the subject, regardless of what others perceive as its 'content'.)

In gambling, one is aiming at a desired end – to increase one's money. There are clear criteria for success and failure. Winning and losing, after all, is the name of the game.

(One composer friend buys a weekly $1 lottery ticket. She says that the odds are a lot higher on the lottery than on her next grant application, but that the lottery payoff will be much greater than the penny-ante amounts arts funding currently provides, and that the $52 she spends yearly on lottery tickets is usually far less than the amount she spends on materials for just one grant application.)

A popular misconception of the use of chance in art is that it too should be judged by criteria of winning and losing. For many, winning means, 'sounds like something I'm already familiar with', or, 'makes me happy in ways I know'. Hence all the silly

Organised Sound **1**(3): 167–72 © 1996 Cambridge University Press

https://doi.org/10.1017/S1355771896000234 Published online by Cambridge University Press

writings and speakings about how music made algorithmically doesn't sound like Mozart, jazz, rock, or whatever.

(There are, of course, ways in which chance, randomness, and other (non-chance-based) algorithmic procedures *can* be used to assemble music that sounds like already-existing music. Just constrain the limits of your choices enough, or use it to choose between blocks of material that already sound familiar. The Mozart dice game, and Peter Gannon's algorithmic accompaniment program *Band-in-a-Box* are but two examples of this. And recently, as part of the musical accompaniment to computer graphic artist Robert Randall's *Visual Haiku*, I had the fun of using Russ Kozerski's algorithmic software *Sound Globs* to make some happily clichéd 'Hawaiian ukelele music', complete with crashing ocean waves, twittering birdies, and an in-tune-and-time, but ethnically inappropriate steel drum solo. Hey, it was music from some tropical ocean environment or other! Who's keeping score, anyway?)

Aside from uses such as this, though, it seems to me that one *proper* use of algorithmic composition techniques, chance, randomness, mappings, morphings, etc. (all of which, I acknowledge, are different things), is not duplicatory (assembling lookalikes and soundalikes of previously existing works), but exploratory (finding (making, uncovering, stumbling across, perceiving, learning to comprehend)) that which does not yet exist (or in Herbert Brun's too-often-quoted-to-even-be-bothered-with-footnoting phrase, 'We're interested in the music we *don't* like, yet.').

(Capitalism and other Economist forms of thinking (Marxism, etc.) have instilled in us the idea that winning and losing does exist in the arts. The person who has the most people loving their work, and buying their commodities while they are alive is the winner, the others are the losers. Classical music and other museum-based cultures add the proviso to this – people whose works achieve significant exposure and/or sales after their deaths are also considered winners. In fact, this museum-based view often considers these artists bigger winners than those who achieve big sales while alive. Contrast this with Japanese philosopher Daisetz Suzuki's attitude to philosophy: 'That's why I love philosophy. No one wins.'[1])

That is, I, at any rate, find it much more valuable to use algorithmic methods as a means of finding out what I don't know, rather than making what I do know. For me, it has been a method of expanding my tastes, a tool for personal growth, rather than a method of expressing what I already felt.

(It has been a method of argumentation and refutation for me, as well. For example, annoyed by many people (mostly scientists writing on chaos in the mid-1980s) who said that the human mind would not find interesting the results of randomness where all elements had the same weighting (the structural equivalent of white noise), I wrote my composition *Mr. Yasser's Piano* (1990). In this piece, I used Joseph Yasser's idea[2] of dividing the 19-tone equal-tempered scale up into 12-note chromatic, 7-note diatonic, and 5-note pentatonic subsets. All the pitch choices in each section of the piece were made using this totally random weighting. That is, in a section of the piece that used a 7-note diatonic subset, all the pitches had an equal probability of occurring. I used a rather slow overall tempo for the piece, and a piano timbre with lots of reverberation so that the resulting chords and melodies would form a sensuous 'wash' of sound, allowing the harmonic qualities of the scale to linger and be savoured. In all the times I have played this piece since then, not one person has said that the melodic or harmonic choices were dull or uninteresting, even some of those very scientists who had at first made the comments about equal-random weighting being dull! Of course, they may not have been listening, or have been able to listen, to the moment-to-moment harmonic details of the piece in a musically meaningful way. Perhaps they were seduced by the quality of the reverberation, or the slightly bittersweet sound of the 19-tone harmonies. Clearly, in dealing with the perception of musical compositions (rather than simply sequences of test-tones, for example) there are many more variables to be taken into account than the reductive tests of psychoacoustics usually admit.)

Algorithmic composition, then, is not gambling. (In fact, some algorithmic procedures, while they may produce unforseen outcomes, are totally deterministic.) It's a method of generating different results than we obtained formerly. Using method 'X' (say, $1/f$ noise) to generate output 'Y' (say, a Coltrane tune) is clearly a dumb (though maybe fun) thing to do.

(I always laugh when I hear someone say they've discovered the 'Beethoven' algorithm, or the 'Bach' algorithm. As if those guys used one way to generate their works! And even if a theorist finds one kind of process at work in a section or sections of a composer's works, there are all sorts of questions yet to be asked which are not covered by that process. For example, if a harmony-generating process were discovered in a piece, that still wouldn't tell us why the composer at that moment chose, say, to have those sounds played slowly on the viola. And even with

[1] John Cage, *Silence*, p. 40. London, Calder and Boyars, 1968.

[2] Joseph Yasser, *A Theory of Evolving Tonality*, New York, American Library of Musicology, 1932.

https://doi.org/10.1017/S1355771896000234 Published online by Cambridge University Press

composers who take pains to explain the structural workings of their compositions and compositional methods, there are clearly many aspects of the pieces that remain unexplained by the processes used to generate *some* aspects of those pieces.)

When one is exploring new materials one often notices certain structures inherent in them. Joseph Yasser, in deriving the 19-tone equal-tempered scale from a series of (flattened) fifths, noticed how the scale could be divided up into the chromatic, diatonic and pentatonic scales mentioned above. This exploration of materials inherent in a structure is, for me, one of the joys of composition. Algorithmic methods provide one way of exploring, and finding, these inherencies.

(Gleefully ignoring those inherencies is also another of the joys of composition, as in the 'Hawaiian music' mentioned above where the qualities of 'Hawaiian-ness' were far more important than the randomness used to generate the details of the composition.) Mostly, when we talk about algorithmic processes, we are talking about mappings of one kind or another – a certain kind of randomness applied to a set of pitches, for example, or a structure from mathematics applied to sound, or language, or a visual image. But these mappings are only the beginning. Even in the most simple realisation of a mapping such as this, how we perceive the mapping, in what contexts, how we observe, how we have *learned* to observe, are often the crucial factors in any judgements we might make as to how 'interesting' or 'dull' the results of the mapping are.

(Exploring inherencies, as stated above, is fun. Whenever I start working with a new microtonal scale, I usually just listen to an octave of it (if it's an octave-based scale!), played by fast eighth-notes with one timbre with equal weighting of all the notes, for quite a while just to get the 'sound' of the scale into my ears. Conversely, my compositional program *randie*, which I wrote in 1991 using Jim Binkley's *Ravel* language, allows me to apply the output of eight different 'random number' generating programs (equal weight, bell curve, $1/f$, the henon attractor, a one-dimensional cellular automaton, etc.) to a constant set of musical parameters to hear the musical results of the different weightings of numbers produced by the different equations. In practice, however, nearly all the pieces I've made with *randie* have used different equations controlling different sets of musical parameters and using different timbres from section to section. An example of this is *Tuning the Furniture of Chaos*, a 1994 interactive live computer music composition where each *c*.90 second section of the piece used one of three different equations, the henon attractor, the cellular automaton, and a set of counters feeding back into each other of my own design, and applied to different limits and timbres. The cellular automaton was always played slowly with piano timbres, the henon attractor used a fast 'Le Marteau'-type ensemble, and the 'feedback counters' were played with moderato bell sounds. The tuning was a 60-note scale derived (in a droll attempt at being 'ethnically Minnesotan' (I was living in St. Paul at the time, as McKnight Visiting Composer Fellow with the American Composers' Forum)) from the last two digits of the phone numbers of sixty furniture stores listed in the 1993 St. Paul, MN Yellow Pages.)

In our collaborative 1979 super-8 film *Der Yiddisher Cowboy*, composer and cultural historian Ronald Al Robboy, discussing our collaborative and film-making methods, and alluding to the early film-making techniques of such cinema pioneers as Alan Dwan, says that I am the ideal collaborator on such a project because of my ability to make a piece out of 'any thin thread of trashy material'. This scene, in which Robboy is talking directly to the camera in extreme close-up, is immediately followed by a short flickering burst of hand-painted film. On closer examination, or showing in slow motion, the hand-painted film is seen to be a portion of the Fibonacci series, written in my handwriting, one number per frame.

(Since the film dealt with many aspects of Eastern European culture and its influence on our lives, this was my loving and irreverent homage to my maternal countryman and teen idol Bartok.[3] With enough fudging, and enough ignoring of inherencies in one's basic material, anything can be transformed into anything else. And though some theorists or critics may want to set boundaries as to which transformations or mappings of material from one medium to another are 'proper' or 'improper', such efforts, when faced with an exuberant, free-spirited creativity, are bound to fail. As Kenneth Gaburo writes, 'As a composer, the statement that a 'given system is the only tenable one' constitutes the only challenge necessary to disprove that system.'[4]

Even within the professional sphere of those using algorithmic and process-oriented methods, I notice a certain level of basic misunderstanding. For example, I have seen many requests in various Internet forums on algorithmic composition for people to talk about the details of the algorithms they use. If these requests are motivated by a curiosity as to the minutiae of some composers' structural procedures, I have no problem with them. But I often get the impression that the requesters are engaged in some sort of search for magic formulae, some set of musical philosopher's stones that will transmute dross into gold.

[3] See Erno Lendvai, *Bartok, an Analysis of his Music*, London, Kahn and Averill, 1979, for a (highly controversial) explication of some aspects of Bartok's use of the Fibonacci series.

[4] Kenneth Gaburo, 'The beauty of irrelevant music', in *Collaboration One*, La Jolla, Lingua Press, 1973 (now distributed by Frog Peak Music, PO Box 1052, Lebanon, NH 03766, USA).

https://doi.org/10.1017/S1355771896000234 Published online by Cambridge University Press

There are no magic formulae. There is no philosopher's stone. Or, better stated: there is no dross, there is no gold. There is only how you choose to listen, there is only your attention, and the quality of that attention at any given moment.

(There is care-full listening, and there is care-less listening. And these levels of attention can be given to any sound. And any sound can be (to make an excessively Roman Catholic analogy) either an 'occasion of sin', or an 'occasion of redemption'. Or, to give a less moralistic example from another environment heavily influenced by Roman Catholicism: On the same trip to Louisiana, Catherine took me to a Cajun music club. I was fascinated sociologically, but I wasn't enjoying myself musically at all. Then, just before we left, I had a chance to stand on the side of the band, out of the direct range of their loudspeakers, and away from the swirling whirl of the dancers. Sonically, everything changed. The memory of those stunningly beautiful cross rhythms on the triangle, washboard and drums, and the thrilling way they played off the phrasing of the singer, accordionist and guitarist, will be with me until I die.)

Often, when a new idea emerges from one field, it spreads to others. It's almost as if this new idea is a kind of 'intellectual template' that we put up against everything, to see how it fits. A recent example of this is the history of chaos and fractal theory and its uses. Coming out of research into such nonlinear systems as the weather, the structural analogies it suggested began to be applied to many different fields. Soon, people were suggesting that this new kind of mathematics might explain everything from economic factors to psychological behaviour and to the structure of music. Self-similar structures had been around in music for a long time, of course. Schenkerian analysis could unearth many examples of limited self-similarity in classical tonal works, and closer to our own era (arbitrarily to choose just one example among many), the rhythmic structures of Cage's percussion works of the 1940s are overtly self-similar. But along with this looking at pre-existing structures came a hope that somehow this new theory would provide us with a 'golden key' to find finally a truly beautiful 'natural music' lying hidden behind the veil of mathematics.

(We've seen this sort of naivety again and again. From ancient Pythagorean number mysticism to the Fibonacci idolatry of earlier this century, history is littered with the products of this kind of thinking. That some of these products are astonishingly beautiful does not negate the essential naivety of the thinking. More succinctly, I can see great beauty in some object which embodies the properties of the golden section. But I can also see great beauty in a scrap of paper on the sidewalk. That I can see beauty in both means that neither the golden section, nor incidental randomness can claim to be 'the' key to beauty.)

Of course, when people finally sat down and applied the results of Henon equations, Mandelbrot sets, Verhulst equations, etc., to both waveform synthesis and pitch and rhythm selection, they found they were up against the same old eternal problems of mappings. Which pitch set will this numerical output be applied to? Which rhythms? With what timbres? In the end, chaos equations provided fascinating structures, and analogies for other kinds of structuring, but (i) they were structures with their own unique beauties, which had to be learned in order to be perceived (and perceived in order to be learned), and (ii) to make a piece, one needed to make many decisions (or ask many questions) about purely musical materials which related more to the taste of the mapper than to the nature of the equations used. In short, like Pythagorean harmonic methods, or the Fibonacci series, chaos was simply another tool, with no more (and no less!) magic in it than any other.

('Beauty' is also determined by how you look, how you listen. Looking at pages and pages of the Rand Corporation's *A Million Random Digits*. . .[5] is not, for me, a moving experience. But everytime I look closely at any one section (say 100 or so digits) of this numerical transcription of white noise, I am astonished at the beautiful structures that exist on the local level, never again to be repeated, in this supposedly 'unstructured' mass of numbers. Similarly, when, several years ago, I downloaded a table of the first million digits of pi, and I looked at the close detail of some of the numerical patterns found within it, I honestly thought, for a brief moment (and what moments are more precious than those that are briefest?) that I had never seen anything quite so beautiful. Repeat, again and again: A beautiful numeric structure may be different from a beautiful visual structure may be different from a beautiful sonic structure may be different from a beautiful verbal structure may be different from . . .)

I love algorithmic and process-oriented composing programs. I collect them like some people collect stamps or butterflies. For me, each program is someone's idea or set of ideas as to how a music might be made, or how sound might be structured. The collection of these programs forms an idea bank, a 'possible structure' bank. And when one envisions a structure that doesn't exist, one writes a program (using one of several languages, each of which implies its creator's idea of a possible way of structuring activity), or makes a physical device (Australian composer Ernie Althoff's self-playing sound sculptures are some of the most gloriously wonderful physical

[5] The RAND Corporation, *A Million Random Digits with 100,000 Normal Deviates*, The Free Press, New York, 1955.

https://doi.org/10.1017/S1355771896000234 Published online by Cambridge University Press

algorithmic compositions I know), or creates a set of rules for people to follow (as in Pauline Oliveros' *Sonic Meditations*) so that it will.

(Polyglots have more fun. Often I've had an idea that migrated between several programs (that old 'fitting the template' behaviour again) before I was happy with the fit of idea and program, before the 'language' of the idea, and the 'language' of the language felt comfortable together.)

I want to describe in detail a very simple algorithmic process I used as part of my composition/installation *Pi and the Square Root of 2* (1993–95), a piece which started off in number mysticism and curiosity, and which in the process of becoming composed, transformed into something else. The piece began when I encountered an article in the *New Yorker* magazine[6] about the work of mathematicians Daniel and Gregory Chudnovsky, who built a supercomputer in their apartment using surplus IBM computers as part of a massive parallel processing machine. To test their machine, they calculated the first 4 billion digits of pi. In the article, the author, Richard Preston, talked at length about certain characteristics of pi, and how this 'transcendental' series of random numbers was considered different from an 'algebraic' series of random numbers, such as the series of digits of the square root of two. I wondered if this difference between these series was in any way hearable, and decided to make a composition that used the exact same algorithm with both series of numbers, so that a comparison could be made. I wanted the algorithm to be simple, and the musical results to be simple, so that the focus would be on listening to the 'overall quality' of line, and harmony, produced by each number series. I decided to have each number series controlling the pitch choice, octave choice, rhythm, note duration, and loudness of four voices, and that these two four-voice textures would be played with the same timbres out of very widely separated loudspeakers, sufficiently far apart that if one were very close to one loudspeaker, the sounds of the other would be mostly overshadowed. (Since the series of both digits were in base ten, I decided to use a 10-note scale for the piece. After considering various 10-note subsets of 12 tone tuning, and the use of the 10-tone equal-tempered scale, and rejecting them, I decided to create my own 10-note scale which used the interval of 31,415/20,000 (about 782 cents) (pi expressed as an interval within an octave) and 1,414/1,000 (600 cents) (the square root of two expressed as an interval within an octave) as its generating intervals. So the ratios of pi and the square root of two would be used to make the scale of the piece as well as the note, duration, etc., choices. For those interested, the scale is shown in the table. Further, not satisfied with using this scale on simple instrumental or electronic timbres, I decided to make my own tones, where the partials would be on frequencies that were in tune with the tones of this scale. This complex of inharmonic tones created some lovely bell-like sounds, which (as predicted by William Sethares[7]) sounded really good when played with the pi/root2 scale given above.)

Table

Scale degree	Cents	Derivation
0	0	0
1	128	5*pi − (pi − sqr2)
2	182	2*pi − (pi − sqr2)
3	364	2*pi
4	546	4*pi − (pi − sqr2)
5	600	sqr2
6	728	4*pi
7	782	pi
8	964	3*pi − (pi − sqr2)
9	1,146	3*pi

The algorithm which used the two series of digits was very simple. First, each series went through an 8-stage bucket-brigade delay, so that with each new note generated, each digit of the series advanced to the next stage of the delay. For the first voice, the output of the first stage of the delay (a random digit between 0 and 9) selected which pitch of the 10-note scale was to be played. The output of the second stage was reduced modulo 4 and chose which of 4 octaves the pitch was played in. The third stage output was again reduced modulo 4 and chose a duration of 1, 2, 3 or 4 pulses for the note. The fourth stage, reduced modulo 8, chose one of 8 sustain lengths for the note, and the output of the fifth stage chose one of 10 loudness levels for the note.

For the next voice, different outputs of the delay were used for each parameter. So for the second voice, pitch was determined by stage 2, octave by stage 3, attack time by stage 4, note sustain length by stage 5, and loudness by stage 6. The third and fourth voices similarly used different outputs of the delay for different parameters. Further, each of the four voices also played at a different tempo. The tempos of the 4 voices were related by the proportions 2:3:4:5.

The same process was simultaneously played by four other voices, which used the other number series. As stated above, the two 4-voice textures went to two widely separated loudspeakers. The main audible features of the resulting sound were the regular rhythms that resulted from having only four possible durations, and those related by such regular proportions as 1:2:3:4, the bell-like tones of the inharmonic timbres, and the smooth, consonant harmonies these

[6] Richard Preston, 'The mountains of pi', in *The New Yorker*, March 2, 1992, pp. 36–67.

[7] William Sethares, 'Relating tuning and timbre', in *Experimental Musical Instruments*, Vol. IX, No. 2, 1993, pp. 22–9.

https://doi.org/10.1017/S1355771896000234 Published online by Cambridge University Press

timbres produced when played on the pitches of the scale. Overall, the sound was more like that of a gamelan (an extremely simple gamelan) than anything else. My initial impulse was to make a piece to compare the different qualities, if any, of the two number series. What I arrived at though, was a musical texture where timbre, tuning, and rhythm were the most salient musical aspects. The timbre and tuning were identical for both voices, and the regularity of the rhythm was such that any differences produced by the two number series were obscured, overwhelmed (at least, for me) by the regularity of the rhythm.

At this point, in order to follow my original intention of making a sound world where the differences between the two series of numbers were easily and immediately comprehensible, I could have started over again and made a much simpler musical output. But I decided, instead, to keep the piece as it stood. For I realised that what I had done was to make a sonic object where every aspect of it was suffused with the qualities of my numerical fetishes.

(I had also made a version of the piece that used normal instrumental samples, and I had rejected that. Even though melodic qualities might have been more easily perceived, given the familiar instrumental timbres, the result sounded (to my ears) forced and gratuitous. But the version with the inharmonic timbres and scales based on the proportions of pi and the square root of two sounded tight and coherent, even if it was more difficult to hear the various qualities of the number series with them.)

What had happened was that I had set out to make a demonstration, and along the way, following my instincts, had, instead, made a piece. That is, I began with curiosity. If I had been a scientist, that curiosity would have resulted in a psychoacoustic (or psychoacoustical–mathematical) test. But in my case, the results were not a test, but a complex and hopefully contradictory composition. And I am fully in favour of the idea of contradiction as one essential quality of a work of art.

(I admire Alvin Lucier's work intensely. Somehow, he is able to take the very simplest of sonic manifestations of an idea, and have them be luminous. My habits lead me to make more complex sonic entities than Alvin's, but I feel that his pieces, too, even at their most 'demonstrative', go beyond the simplistic nature of scientific test stimuli, and become sonic entities worthy of the most deep and serious listening.)

(Repeat, again and again: A beautiful numeric structure may be different from a beautiful visual structure may be different from a beautiful sonic structure may be different from a beautiful verbal structure may be different from . . .)

(But they might also be interesting, fascinating, beautiful together, combining in completely unexpected ways that we might be able to learn from.)

My essential reason for using algorithmic methods is to provide me with material that I can learn to listen to. This learning, this listening, will hopefully change me, expand my tastes, expand my abilities to hear exactly what is there in the things I give my attention to. And this exact and careful attention is given with the full knowledge that all decisions made that produced the object could be regarded (by some) as completely and highly arbitrary. But this, to me, is a great pleasure: to float freely in a world of arbitrary decisions (constrained, of course, by my (tediously middle-class? (basic)) moral principles), observing new objects form as a result of those decisions, and learning from those objects. To observe, openly, without judgement, and then to swing capriciously from observation to discernment, and not to be bothered by the difference. To note that many times what first appears 'ugly' later seems 'beautiful', and that what first appears 'beautiful' may later appear 'plain'. And then to learn to perceive the beauty in the 'plain' once more. For me, the importance of algorithmic processes is that they allow me to learn about, and expand, my consciousness and my perceptions.

ACKNOWLEDGEMENTS

Thanks to Catherine Schieve and Carl Volkmar, of the Visual Arts Department at Southeastern Louisiana University, discussions with whom started off this essay. Also thanks to John Fitch, of the Mathematics Department at the University of Bath, England, for supplying me with a file containing the first 32,000 digits of the square root of 2. Work on *Mr. Yasser's Piano* was funded by a grant from the Performing Arts Board of the Australia Council. *Pi and the Square Root of Two* and *Tuning the Furniture of Chaos* were composed with the aid of a McKnight Visiting Composer's Fellowship from the American Composers' Forum, St. Paul, Minnesota.

https://doi.org/10.1017/S1355771896000234 Published online by Cambridge University Press