Thoughts on “Where we stand” and the age of audiences

Bill Harris, Facilitated Systems

One thing that caught my attention was your claim of aging audiences. In http://www.artsjournal.com/sandow/2006/11/important_data.html, you note that, among other facts, the average audience age went from 45 in 1992 to 49 in 2002. You think about the potential causes and implications of such a development.

One way to think of such a problem is to see if we can "operationalize" it: can we generate an operational description of the events that we conjecture are playing out in the real world, is that operational description similar structurally to the real problem, and is that operational description capable of generating similar behavior? (See the link near the end of http://facilitatedsystems.com/weblog/2007/01/systems-language-for-business.html for more on operational thinking.)

In simpler and more specific words, can I create a computer model that captures your hypotheses, and does that model behave as your data shows? Success doesn't mean I've proven anything, but failure might indicate a need for modified hypotheses. Enough success, combined with a bit of triangulation, can strengthen our belief in those hypotheses.

I like to start with really simple models, adding complexity only when it becomes necessary. Often we can learn the most from those simple models.

Here's a simple model of an "aging chain" that might represent classical music audiences.

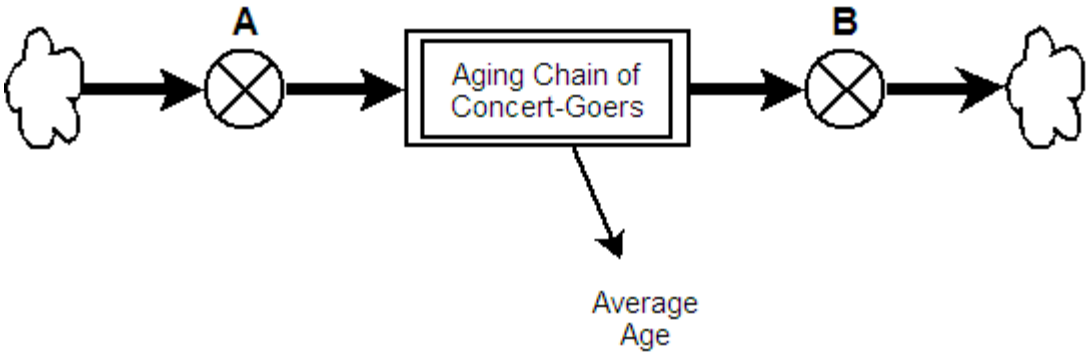

The "aging chain" in the middle is a series of "stocks," one for each decade of age (arbitrarily twenties through sixties; I don't think the dynamics change much if we add or subtract a decade at one end or the other). People move through those stocks, taking ten years to move from one to the other. I've aggregated those stocks into one mega-stock to simplify the graphics; if I had drawn the entire picture, you'd have seen five rectangles, connected in a chain by flows (pipes), instead of that one bigger rectangle.

The "clouds" at the left and right simply mean we don't care where those people come from or where they go, at least for the sake of understanding concert audiences; that's outside the purview of this model. There are two flows, shown here as valves called "A" and "B" on pipes that flow from one stock (or cloud) to another. Flow A represents the number of new concert-goers per year, and flow B represents the number of people leaving the concert-going world each year. In this simple model, 50-year-olds don't all of a sudden decide to become concert-goers, and 30-year-olds don't all of a sudden give up on classical music. Those are constraints we can lift later, if we want to.

Now it's easy to talk about what changes the number of concert-goers: if A is greater than B, you get more concert-goers, while if B is greater than A, you get fewer. If A = B, the aggregate audience size stays static.

I've shown "Average Age" as a statistic we can calculate to describe concert-goers, the same as we can describe their total number.

I created a computer simulation of that model. I set the initial population of concert-goers to 1 million, spread evenly in age across the five decades (stocks) I modeled. If we have 20,000 new concert-goers each year, we'll exactly replace the number of concert-goers who depart each year, and the number of concert-goers will remain constant.

Because you hypothesized that young people had stopped coming to concerts, I tested the model by running it over a 200-year period with 20,000 new concert-goers in the first 10 years and then half that thereafter. Would that replicate the data you were quoting?

Here's the graph of the number of new concert-goers per year:

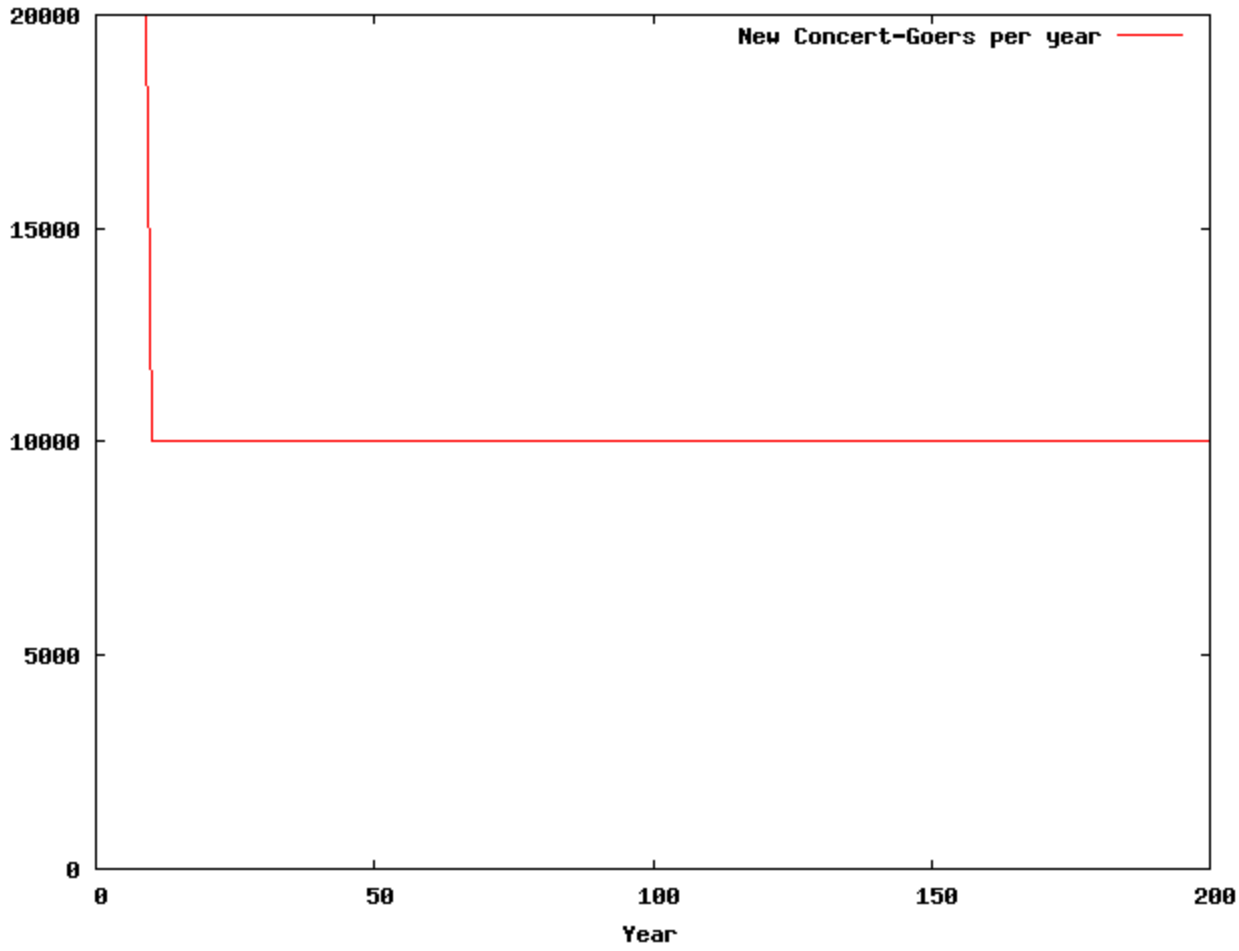
New Concert-Goers per year
20000
15000
10000
5000
0
0
50
100
150
200
Year

It stays constant at 20,000 per year until year 10, whereupon it drops to 10,000 per year throughout the rest of the 200 year time horizon of this simulation.

Here's a graph of the total concert-going population from that model:

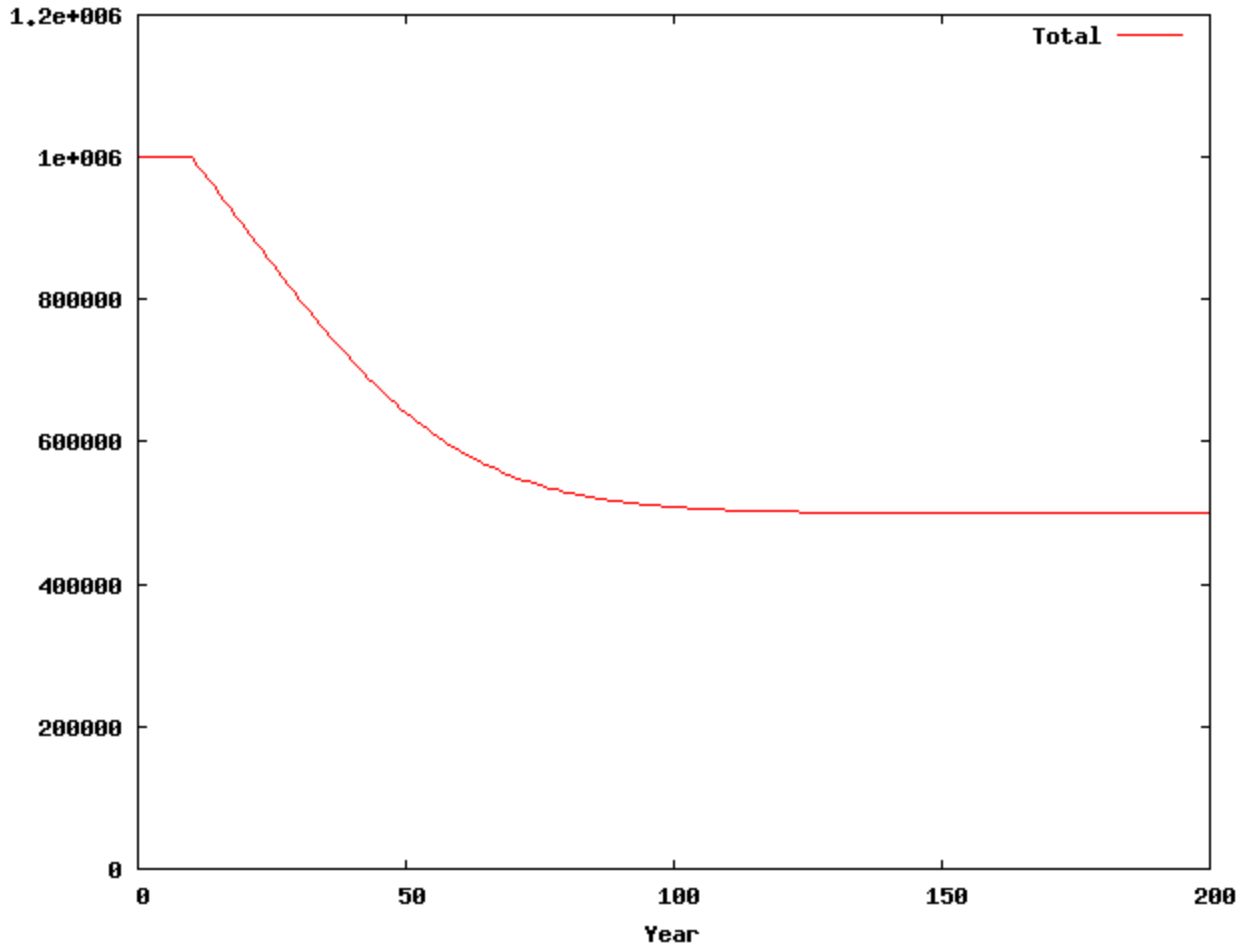

AS you can see, the number of concert-goers stays at 1 million for 10 years and then begins a steady decline, stabilizing at 500,000 at about year 80.

Here's the average age:

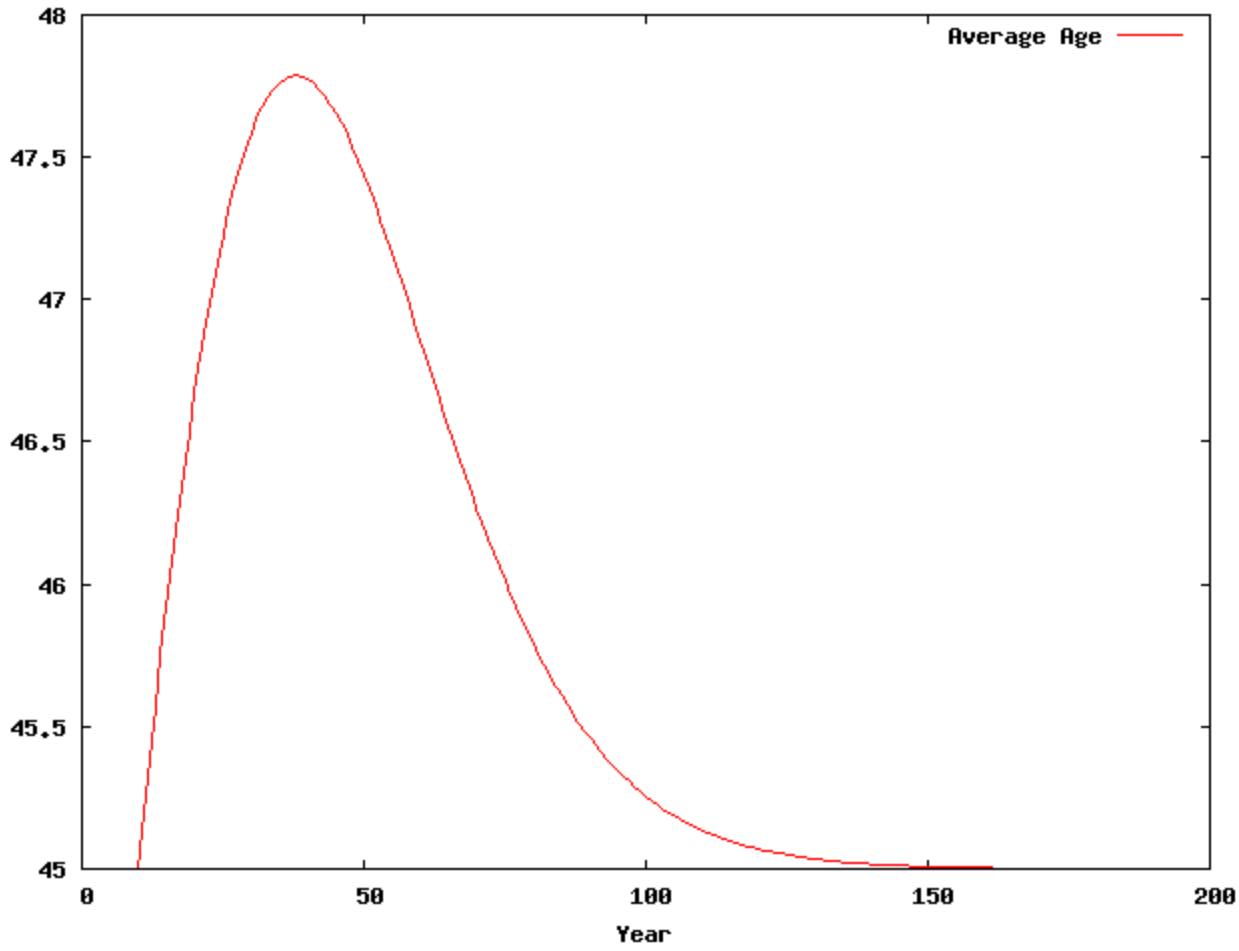

Note that the vertical axis starts at 45, not 0.

At first glance, this seems intriguing. Audience population is declining and average age increases, at least for a bit. Then, under these assumptions, average audience age drops back to the same 45 years old. Could that be? Is it possible that we're just on the front end of a declining audience size, and audience age will correct itself naturally?

There's something else a bit off here, though. In "Where we stand (2)," you provide a graph that indicates the mode of age has gone up about a decade from 1992 to 2002, consistent with your cohort theory, and you note elsewhere that the average age rose by about 4 years. This model only shows a 2.8 year increase in average age, and that's over about 30, not 10, years.

What if the decline in the number of newcomers was a more gradual and continuous decline? Here's an exponentially-declining number of new concert-goers each year, starting at 20,000 and declining to half that in 69 years:

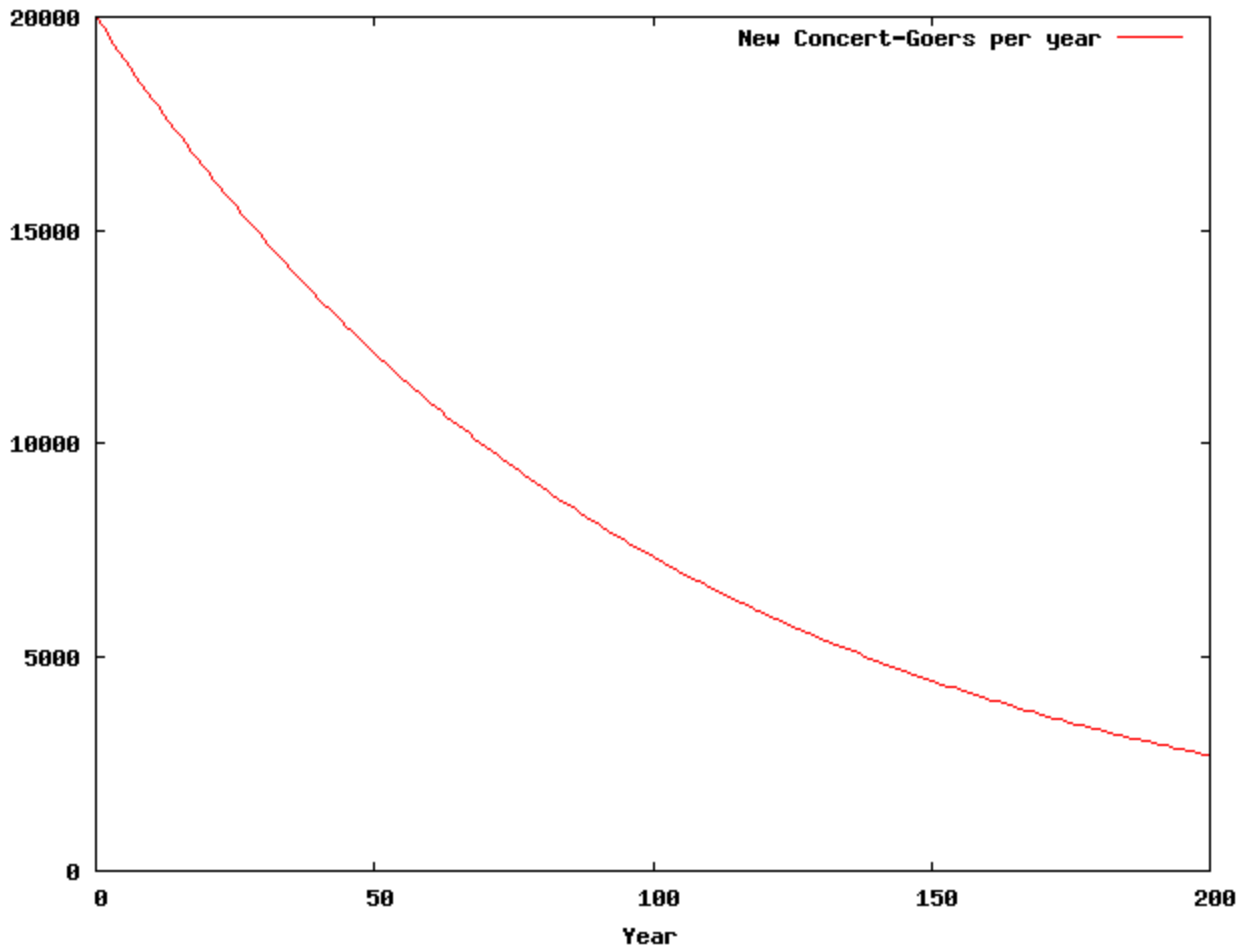

Since the exponential doesn't drop as fast, you might expect the number of total concert-goers to drop more slowly; since the number of new concert-goers continues to drop forever, you might expect the total concert-going population to continue to decline. You'd be right:

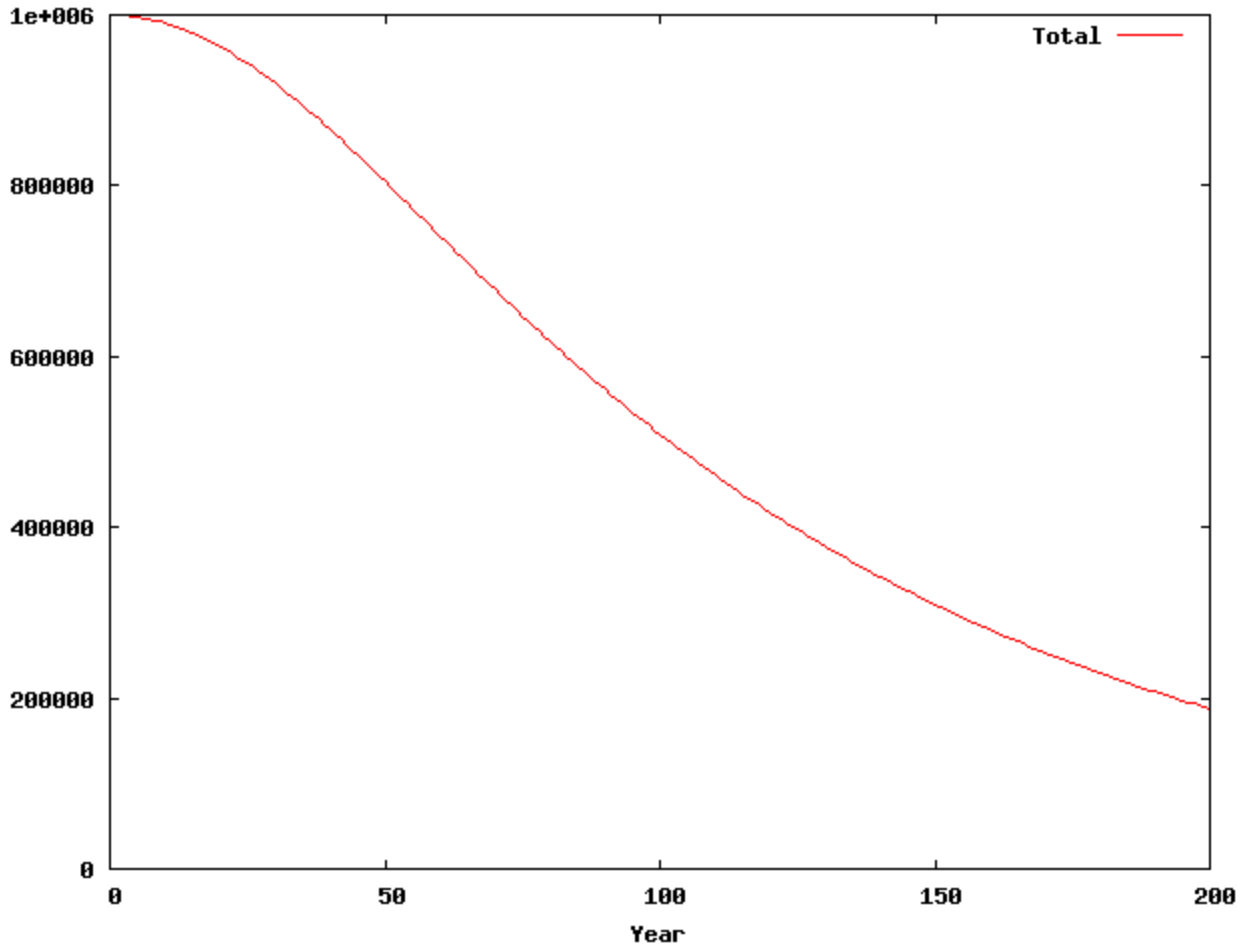

Remember that the drop-off starts in the first year in this experiment, not the tenth year.

What about the average age of concert-goers? Since the decline in new, young concert-goers continues forever, you might expect the age boost to last forever. Since the decline is less drastic, you might expect the age boost to be less drastic. Let's see:

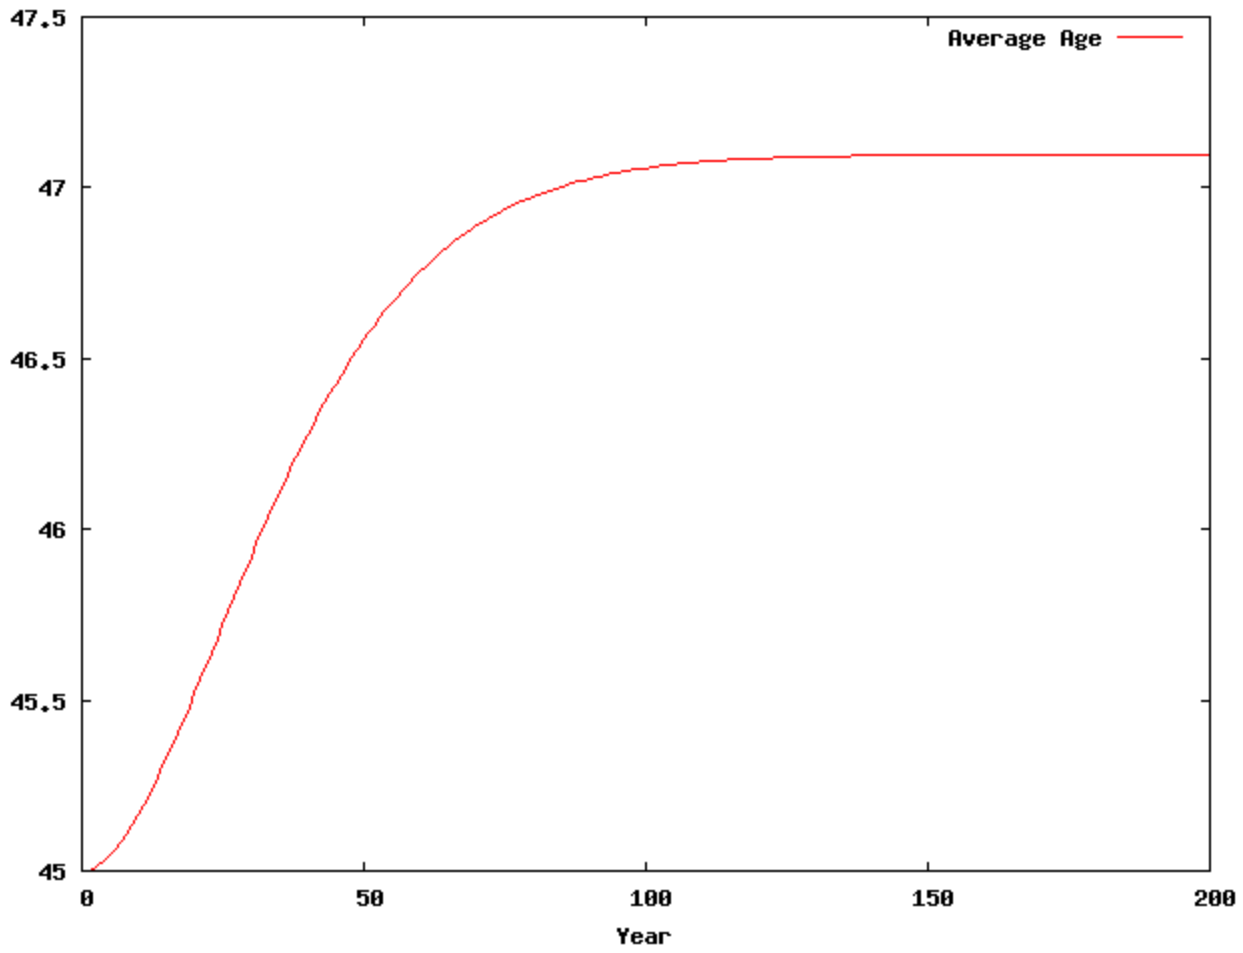

You'd be right in both cases, although the difference in the peak average age is miniscule: 47.0971 vs. 47.7867 years.

So neither explanation seems to account for the drastic aging of the concert-going audience as reported in the data.

What if we take a drastic approach and cut off new concert-goers totally after 10 years? Under those conditions, here is the total concert-going audience:

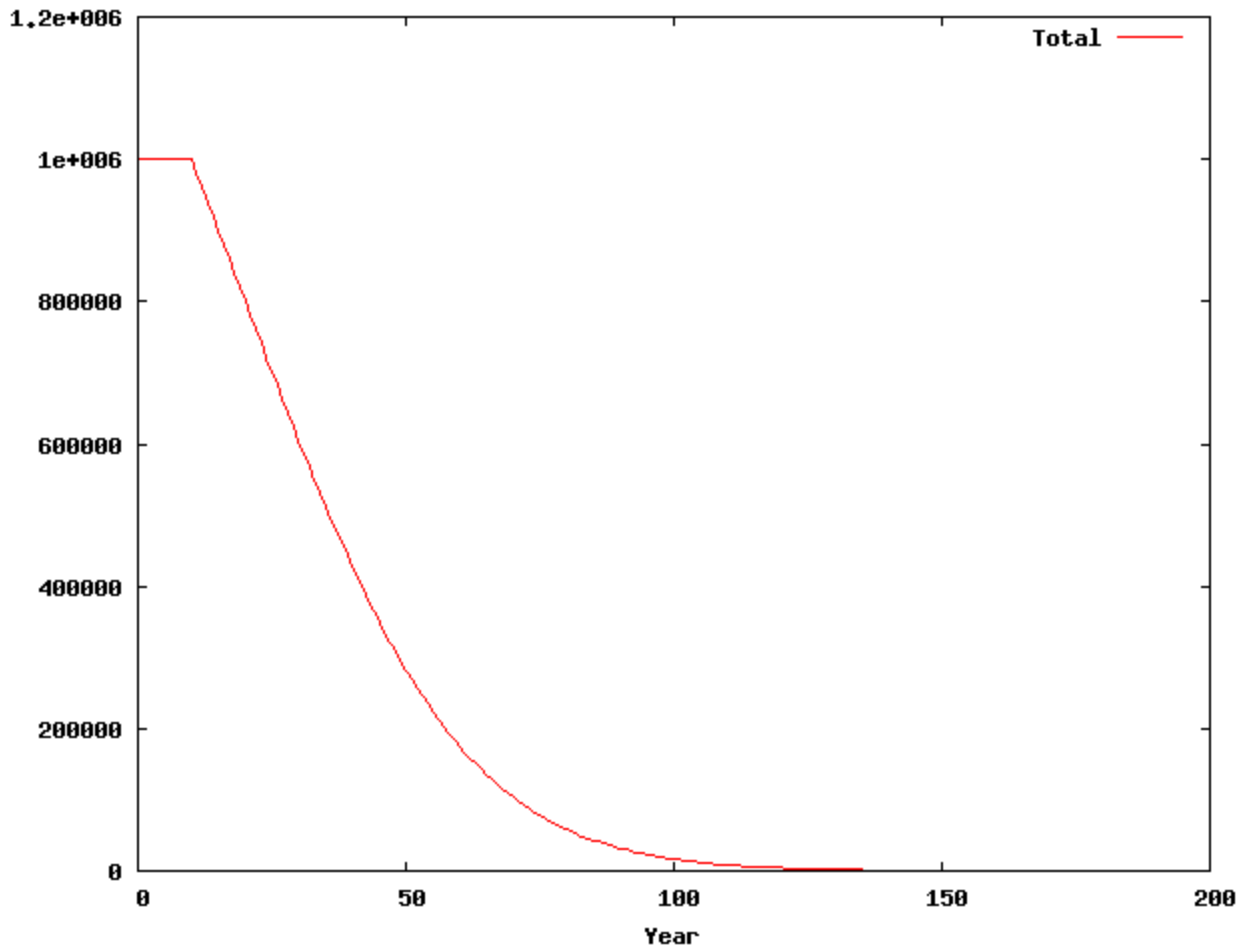

As you can see, the concert-going public drops to nothing (technically, the model shows fewer than 10,000 people by year 108); it is already down to 494,703 by year 36 (26 years after young people stopped becoming concert-goers).

What about the average age?

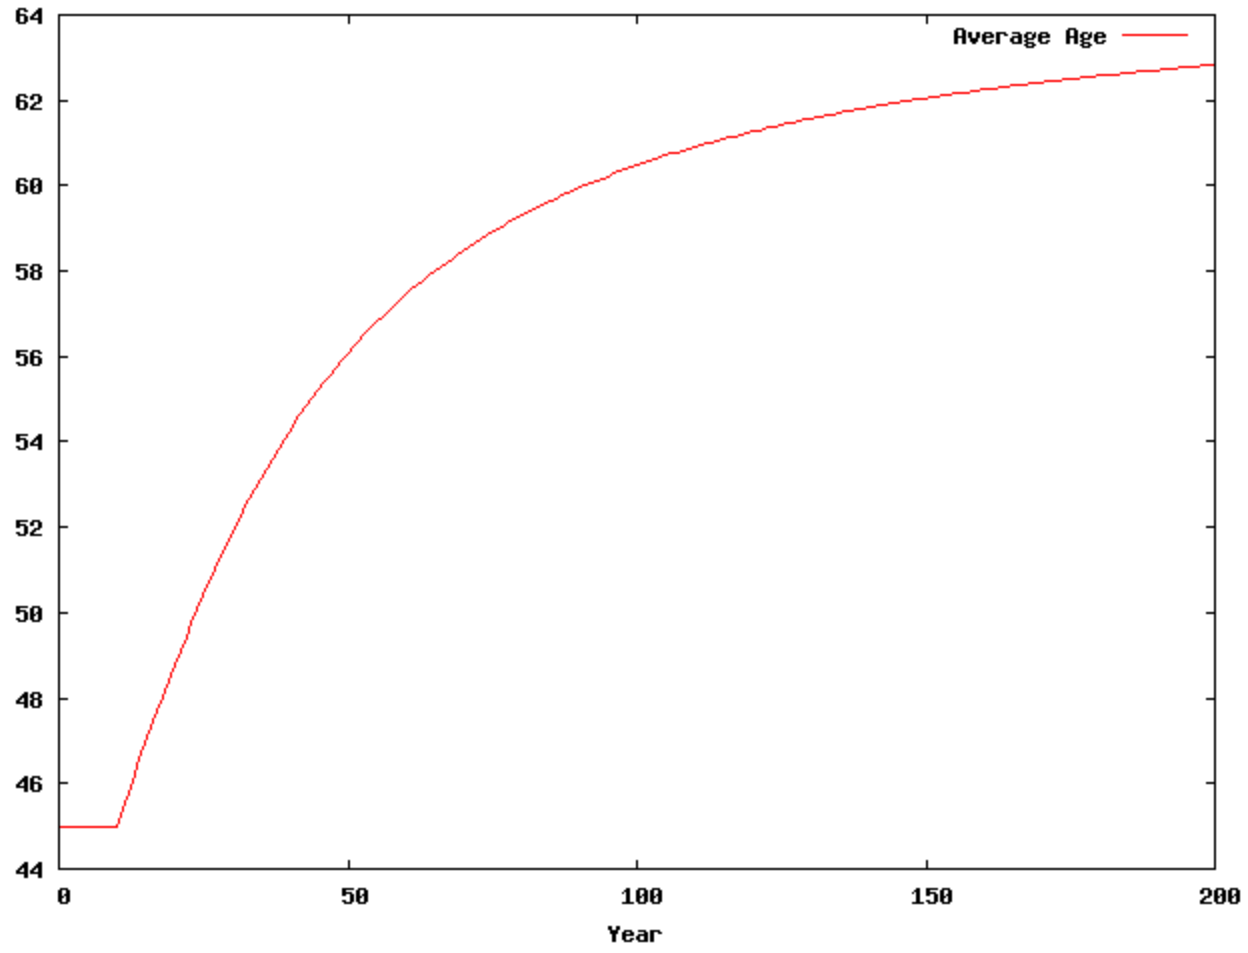

Finally, we're getting drastic changes in average ages. According to http://www.artsjournal.com/sandow/2006/11/important_data.html, the age went from 45 in 1992 to 49 in 2002. In this model, it went from 45 in year 10 to 48.75 in year 20. That's not a bad match, and the model structure seems reasonable.

What's scary is that's a model of /no/ new concert-goers at all! In other words, the data you're showing /could/ be consistent with a sudden change to essentially no new audience members forever. You came close to this same conclusion in today's "The short version."

Now this model doesn't prove there are no new concert-goers. There may be other ways to get similar results. For example, perhaps it's not true that "once a concert-goer, always a concert-goer." Perhaps younger people are starting to attend concerts and then giving up in droves. Perhaps orchestra marketing is drawing in baby boomers who have never attended concerts. Perhaps multiple causes are at work. Perhaps you have other conjectures. Any of these hypotheses could be tested in such a model to see if they are consistent with the reality you've been observing.

What I think is interesting is that a relatively simple model can help shed light on the mental models we create to explain the problems we face. In this case, the first, simple approach suggests things may be as you suggest, with the caution that they /may/ be even more serious than you indicate. I'm curious in your thoughts on all this. I do apologize for the length of this email; I don't yet know how to walk someone through a model such as this without taking a little bit of time.

Drew, did I miss anything fundamental?

I'll be expanding on related ideas using a different model in a column I'm doing for Drew shortly. You can see some of the blog postings I've made about music at http://preview.tinyurl.com/2consf. In particular, http://facilitatedsystems.com/weblog/2006/11/making-sense-with-numbers.html was a very popular posting about the recent Knight report. Drew and I have exchanged other thoughts sparked by your columns, but this note is long enough as it is.