Millennial matters*

Arnold M. Zwicky, Stanford University

April 2007

For Bernard Comrie on his 60th birthday

Now that we are safely into the third millennium (and twenty-first century) of the Common Era, I'm going back to thinking about what linguists might have to say about the meaning of the word "millennium". I'll start with the notion of "age" (in years) and then creep up on the calendar issues.

Expressions of age are assignments of a number *n* – a number of some unit *y*, like years – to something (a person *p*, in particular) at some moment *t*. The point of the number is to (approximately) measure the span from some fixed initial point *b* (for people, the time of their birth) to *t*. There are many ways to do this systematically. None of them is more "correct" or "mathematically justified" than any of the others; they are all essentially equally well-defined and exact.

One convention is to assign *n* to *p* at time *t* if *t* is in the *n*th unit after *b*. Equivalently, if *n* units will have passed on *p*'s next birthday after *t*, *p*'s age at *t* is *n*. This measure is "prospective to the next birthday". On this measure, the age in years of a newborn is one, and remains so until the next birthday, at which point the age becomes two.

As a measure of the span from birth, this numbers given by this convention are somewhat misleading early in the unit (it's uncomfortable to say of a newborn, "It's one year old"), but more accurate towards the end; they are, however, easily interpretable as ordinals: "of age *n* ys" is equivalent to "in their *n*th *y*".

This is the system used in our calendar for assigning numbers (in years) to events in the (modern) world. The age of this world "at birth" is one year.

This is NOT the system used in western cultures for assigning numbers (in years) to people, however. On this alternative convention, if *n* units passed on *p*'s last birthday before *t*, *p*'s age at *t* is *n*. This measure is "retrospective to the previous birthday". It follows that in this system the age of a newborn in years is zero, and remains so until the next birthday, at which point the age becomes one.

(The system of number words in a natural language usually doesn't have a zero, so that for babies the unit type is shifted from years to months, weeks, days, or even hours – whichever is the largest type of unit that will allow for retrospective age assignment.)

As a measure of the span from birth, the numbers given by this convention are very accurate early in a unit, but rather misleading late in the unit (it's uncomfortable to say of a child that will in two days have been on the earth for two years, "It's one year old").

An alternative is to split the difference, and use a system that "looks to the nearest birthday", rounding the number down or up as the case may be. That is, on this convention, if *n* units (will) have passed on *p*'s closest birthday to *t*, *p*'s age at *t* is *n*. In this system the age of a newborn in years is zero, but becomes one in only six months.

The numbers given by this convention are accurate both early and late in a unit, and only a bit misleading towards the middle of the unit. This system is the one to choose if accuracy in the estimate of span length is important. Either of the other systems will be your choice if the actual moment when one unit of size *y* ends and the next one begins is important. Both of these situations arise.

When relative youth or age is important, without particular concern for the endpoints of units, then it makes sense to speak in terms of age at the nearest birthday. If medical care, for instance, might be somewhat different for two patients whose span from birth differs by a year, then it would make sense to record a patient a week from the end of their thirty-second year on earth, not as thirty-one (as in the usual, retrospective, usage), but as thirty-two. And at least one hospital in my experience did in fact record "age at nearest birthday" and use that in summaries for medical personnel. (This hospital also recorded date of birth, so that retrospective age assignments could also be made, though these were not normally needed for diagnosis and care.)

On the other hand, there are occasions when the endpoints of units are important, namely when these endpoints can be understood as culturally significant points of passage from one stage in life to another, AND when "natural", "culturally based", or "subjective" definitions of these points are believed to be insufficiently "precise" and "objective". The passage from youth to adulthood might well be based on achievement of a certain physical size, on evidences of sexual maturity, on performance on tests of skill, knowledge, strength, etc., on the judgment of certain people in the community, and so on, but these tests are often hard to apply in practice, and as societies grow large, diverse, and complex, their legal systems look for "objective" measures that are both reasonably even-handed and easy to apply. Measures of age that rely on precise (actually, rather over-precise) endpoints then replace "natural" tests for maturity, for a wide variety of purposes: responsibility for one's actions, freedom to move about unsupervised in public at any hour of the day, employment, sexual relations and marriage, property ownership, drinking alcoholic beverages, driving motorized vehicles, voting, and on and on.

For these purposes it is now exquisitely important just when, say, eighteen years have elapsed since the day you were born. It's not important whether this moment is conventionally labeled as "becoming eighteen" or "becoming nineteen", so long as people in general use the same scheme of labeling.

In fact, there are quite a few other types of endpoint-centered systems that use a point of reference *r* other than *b*. All that's important is that *r* is reasonably close in time to *b*, that *r*-type events recur, and that the period from one *r*-type event to the next is roughly constant (hence usable as the unit *y*) – so that measuring *y*s from *r* is acceptably precise as a measure of time elapsed from *b*. In all the examples that follow, the unit *y* is in fact a year (in the sense of a solar cycle), although other systems are certainly imaginable.

You could, for instance, use someone's saint's day as *r*. Someone's age in years would then be one on their first saint's day after *b*, two on the following one, etc. I don't know of any place where this system is actually used, but it's certainly possible.

The systems that ARE used have the advantage of choosing an *r* that's the same for everybody: the vernal equinox, 1 January (or 15 April or any other fixed date), the king's birthday, the traditional date for the creation of the world, whatever. (If natural events have roughly one-year periods, even things like the beginning of the monsoon season could be used; ditto, Passover or Easter or Chinese New Year.) you just start counting with the first *r* after *b* and then go from there. The marvelous thing about such systems is that once you're one year old (as things are reckoned in this system), your actual birthday is irrelevant; everybody gets one year older on the same day. And the endpoints are sharp. This is the sort of scheme used for reckoning the ages of horses.

At any rate, we now have one big class of circumstances where endpoint-centered systems are good: where a precise endpoint defines a change in legal status.

Endpoint-centered systems have another good point, one that is a little harder to appreciate, but also turns on having a precise moment of passage. Remember that this moment of passage is associated with a change in the number *n* of *y* units, let's say from *N*-1 to *N*. The moment of passage is when *N* "comes into being", so to speak – the "beginning" of *N*. Why should that be important?

It's important because huge numbers of cultures assign special values to certain numbers and their multiples: two, three, five, seven, ten, twelve, twenty, twenty-five, fifty, sixty, and a hundred are some bases for what counts as an "auspicious" number. In western cultures, five, ten, twenty-five, fifty, and a hundred are especially auspicious. Hence, the moment when we pass from *N*-1 to an *N* that's a multiple of one of these numbers is an especially auspicious and important moment, a "good luck" moment, signalling a milestone in life, an achievement. We expect these moments to be celebrations, festivals, signal events. We are inclined to think that someone who dies just short of their one-hundredth birthday (as we reckon these things) has missed something.

(Note that none of this has anything directly to do with the way the numbers are represented as numerals in decimal notation. Multiples of five and ten were auspicious for the Romans and for lots of other Europeans well before decimal notation, involving the place-holding symbol 0, made it to the West. You don't, in fact, have to be able to write ANYTHING to find certain numbers auspicious; it could just be a cultural fact about certain numbers, like twelve. Very often it's a matter of NUMBER WORDS, not numerals. The auspicious numbers are, for the most part, those whose names form the basis of the system of number words in a language; the especially auspicious ones tend to be those with simple, non-composite, names. Like ten, hundred, and thousand for English. It's true that the decimal notation reinforces this sense of specialness for the change from *N*-1 to one of those really auspicious numbers *N* – the "odometer effect" – but the sense of specialness in no way depends on the decimal notation.)

Now this is all a matter of culture (and language), not mathematics; it's not something "in" the numbers themselves.

Here's where it gets complex. We have four sets of facts about culture/language: (a) a scheme for assigning numbers – ages – to things at particular moments in time, (b) in such a way as to measure time in *y*s elapsed from *b*; (c) a scheme for assigning cultural significance to the endpoints of certain spans; and (d) a scheme for assigning "goodness" to certain numbers.

Sometimes these things all work together like a charm. If we're talking about people's ages, and we choose for the scheme in (a) the retrospective assignment, then the number in (a) is identical to the elapsed-*y* measure desired in (b), which is, moreover, a suitably endpoint-centered measure as desired in (c); and indeed the number assigned in (a) is correctly auspicious (or not), in accord with (d). The retrospective assignment says that I passed from age fifty-nine to age sixty on 6 September 2000; on that day sixty years had elapsed since my birth, which means that this was a pretty significant birthday, sixty being a multiple of ten.

Suppose, instead, that we're speaking Shmenglish, which is just like English, except that age numbers are assigned by the prospective scheme. Describing my situation on 6 September 2000 in Schmenglish: On that day I passed from age sixty to age sixty-one. That is, the number assigned in (a) is sixty-one, which is not itself the number of years elapsed since my birth (though this number can be systematically derived from it, by subtraction of one) nor, more important, is it an auspicious number. The auspicious number sixty came up, in the Schmenglish translation, on 6 September 1999.

There's nothing wrong with Schmenglish from the mathematical point of view, and it's not unusable for its cultural purposes, but from the latter point of view it's non-optimal: measuring spans and showing fabulousness don't coincide in the same date.

On to the western calendar. This seems at first entirely parallel to Schmenglish, thanks to a prospective numbering scheme for the years. As a result, the year numbered 2000 has the fabulousness, but the span of two thousand years goes with the year numbered 2001.

The obvious solution would be to replace the prospective numbering scheme by a retrospective one. That's an utterly lost cause. It's not that renumbering is out of the question; the Gregorian calendar reform involved omitting ten calendar days, after all – a crude stipulation that one calendar year had 355 rather than 365 days in it, which does violence to the usual definition of "calendar year", but in the service of an important cultural goal, getting calendar years and solar cycles back in synch. The problem is the gigantic cost of renumbering (or reinterpreting) the vast body of records with dates in them. In some sense the easiest solution would be to mimic the Gregorian calendar reform, but unfortunately what would be required would be not an omission, but the insertion of extra stuff, namely a whole extra year. Having two years numbered 1999, say, or one numbered X in between 1999 and 2000, would not really be a practicable option, given all the record-keeping devices that use the decimal notation for keeping and manipulating records.

So we can't satisfy the span-measuring and fabulousness-showing functions in a single year number. Too bad, but that's scarcely a disaster.

Actually, there's an interesting FAILURE of parallelism between Schmenglish ages and western calendar years. In the case of ages, the span-measuring function is culturally extremely important; it's built into legal systems in vast numbers of ways. In the case of years, there is no culturally significant span-measuring function.

True, the intention was to have a calendar with the birth of Jesus of Nazareth as its *b* moment. But there are no cultural practices that actually depend on how much time has elapsed since this birth, and anyway the calendar we got has it all wrong. It's generally agreed that this birth preceded the conventional initial moment by several years and that it was not during the winter season. Then there are the ten days lost in the Gregorian reform. Most oddly, perhaps, is that along the way the conventional date for the birth of Jesus was set to 25 December rather than 1 January. How all this happened is interesting but beside the point at hand, which is that what we've ended up with is a calendar that has a *b* point that is quite arbitrarily chosen, not to mention a calendar that measures spans that seem to have no cultural significance.

The upshot of my discussion is that there is good reason to celebrate the onset of the year 2000, because two thousand is a wonderfully auspicious number. There is less reason to celebrate the onset of the year 2001, though there's certainly nothing wrong with marking the passage of two thousand calendar years (less ten days) since some arbitrarily chosen day – or with marking the passage of two thousand solar cycles since that day.

There's no issue here of deciding which of these two years REALLY begins the millennium. 2000 has the millennium number, 2001 the span of two thousand years. The word "millennium" simply has two different meanings in ordinary (non-technical) usage.

In fact, it can be argued that 2000 actually was the beginning of the third millennium (and twenty-first century). Let's take a clue from the defective Gregorian year, the one with only 355 days in it. (Which year this was, and which calendar days got lost, differs from country to country.) If a year can be stipulated to be defective by some days, why can't a decade, century, or millennium be stipulated to be defective by a year? Why don't we just say that the first decade in our western calendar was short by a year? Then as a result the first century was similarly short, and the first millennium too. Thereafter nothing's defective. This stipulation performs a cultural function; it aligns the beginning of every decade, century, and millennium (except the very first of each of these) with an auspiciously numbered year.

*This piece is a version of a posting to the Out in Linguistics mailing list from 4 January 2001 (spurred by comments from Ron Butters), offered here as an academic amusement for Bernard Comrie on his 60th birthday.