

Chronological Aging in Continuous Time

Robert L Eberlein
James P Thompson
David B Matchar

Duke-NUS Graduate Medical School
Health Services and Systems Research
8 College Road, Singapore 169857
Singapore
+65 6601 1362

robert.eberlein@duke-nus.edu.sg
james.thompson@duke-nus.edu.sg
david.matchar@duke-nus.edu.sg

Abstract

Aggregate continuous time formulations used in System Dynamics models result in the implicit mixing of individual constituents of levels. Normally, this is acceptable as the heterogeneous nature of model variables implies that some individuals counted in a level's value will pass others and exit earlier. In models with a focus on chronological aging and age-related characteristics this phenomenon, which we call cohort blending, can result in large distortions. Though these distortions can be reduced by using aging chains, they persist in a significant way even when using one-year grouping in the aging chains. As an alternative, we introduce an approach we call "continuous cohorting" in which populations are tracked with cohorts sized the same as the computational interval of the model. This approach eliminates the blending problem with minimal notational and moderate computational burden. The resulting models display quantitative, though limited qualitative, differences from their traditional counterparts and are more easily defensible with demographers and others who study population.

Introduction

System dynamics models generally use a continuous time conceptualization and change levels based on their instantaneous values. For most applications this is appropriate, and higher order time delays can be well represented by a cascaded sequence of levels with proportional outflows – that is, an aging chain. (See Sterman (2000), Ch. 12, for a discussion of aging chains.) In some

cases the inherent blending of constituent parts that results from the use of cascaded levels is not appropriate. The most common example of this is a population in which individuals age chronologically, and that is the focus of this paper. Other situations for which exact temporal order is important, such as the rotation of inventory based on shelf life expiration, are likely to be more mechanical in nature, but the same principles apply.

To begin we lay out the blending problem in its simplest form, using a single stock population model. Because it is so simple and familiar, this model allows a clear explanation of the disaggregate assumptions implicit in continuous feedback formulations. We then extend this model by dividing the population into different age cohorts in cascading levels. The results of this extension are intuitive, but even with highly differentiated age groups the persistence of significant blending effects is not. We present one standard solution to the blending problem, discuss its shortcomings, and provide an alternative approach we call continuous cohorting that alleviates many of these drawbacks. Finally, we present three examples using the continuous cohorting technique that elucidate more general qualitative issues.

Background

In the field of System Dynamics, we are familiar with the aggregate representation of population as shown in Figure 1. Most come across this in the early stages of learning, and it is the structure that appears in *World Dynamics* (Forrester, 1971), though in that model there are more effects on both births and deaths.




Figure 1: Aggregate population model

In an arithmetic sense this model is completely correct; population does increase with each birth and decrease with each death. However, a real population is composed of individuals of different ages, and no such distinction exists in a single-stock population model. Each individual

born immediately becomes indistinguishable from all other people. This is sometimes called perfect mixing, and in this context it means that chances of having offspring, or dying, are not related to the age of the individual. In steady state this has the interesting implication that the average age of the population will equal the average lifetime which is also the average age of people dying.

This symmetry of age between the living and the dying makes it difficult to consistently relate past birth patterns to the current age distribution using simple formulations. For example, a coflow could be formulated to track average age by using population as the inflow to the numerator.¹ This would show the inverse relationship between population growth rate and average age, but the coflow itself assumes perfect mixing, which means the average age at death is still equal to the average age of the population. That is, if the average lifetime is 70 years, then the average age reported through the coflow also would be 70 years. This is consistent with the single population level formulation but not with the way people actually age.

In the single-stock formulation, what is true of the outflow is true of the inflow. Any change in the level will immediately change births in this simple formulation. The single stock masks an implicit assumption of such formulations: as soon as an individual is born, that individual is fertile and fertility lasts to death. This divergence from observed realities of individual development makes the completely aggregate representation difficult to apply for many problems. Still, as demonstrated by *World Dynamics*, it can be a very useful formulation that allows investigation of fundamental feedback relationships. The subject of essential dynamic detail is something we will return to again.

To illustrate a situation where the single stock population model clearly breaks down, consider the prevalence of dementia in a population. Dementia, the most widely recognized form of which is Alzheimer's disease, is a progressive and currently irreversible condition that occurs for the most part in older individuals. Research in this area initiated by one of the authors (Thompson et al. (2011)) inspired the work reported here, though the principles apply to any population for which tracking chronological age is important.

Because dementia most frequently manifests in people over the age of 60, we will look at the behavior of that population group. This perspective simplifies the discussion and is equivalent to assuming a constant prevalence of dementia. Later we will present results for a model with a

more detailed representation of age dependent onset and progression of dementia, but the effects of cohort size and progression are essentially the same.

To compute the number of people over 60 from only a single population stock, which has no memory of when people entered the stock, we assume that a constant fraction (say, 0.25) of the population is in that age group. Using an idealized five-year increase in births from a steady state in which birth and deaths match, we get results shown in Figure 2.




Figure 2: Simulation results for a one stock population model

Given our assumption, the population 60 years old and over mimics the total population and this begins to rise as soon as births increase. This simple approach is not very satisfactory; we are effectively asserting that individuals entering the population are spread out over all ages from the instant they are born. For looking at age dependent conditions this spreading is problematic and so we consider refinements to the age structure.

Aging Chains

Using a series of stocks to represent population, it is possible to get beyond the condition in which the average age of the population is the same as the average lifetime. With an aging chain each stock will have an average age equal to the incoming average age plus the average residence time, so that the average age of the population is the average across the stocks and smaller than the total average lifetime. For example, consider the four level population decomposition such as the one used in the model from *The Limits to Growth* (Meadows et al.

(1971, Meadows et al. (1974)). This, with two adjustmentsⁱⁱ in order to be consistent with later examples, is shown in Figure 3.

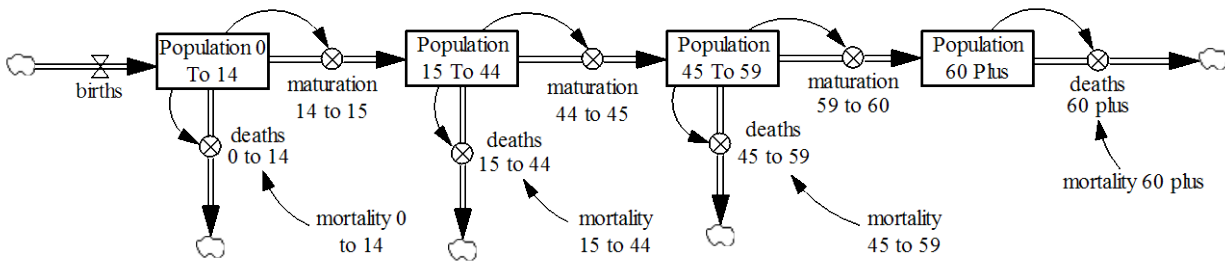


Figure 3: An aging chain with population divided into in four cohorts

With this model it is possible to distinguish different age groups, and we can look directly at the number of people 60 and older. Doing this, with the same test input on births we used for the single stock model we get the results shown in Figure 4.




Figure 4: Simulation results for a four stock population model

The response is more reasonable than it was for the one level model, but it is clear that the timing is not correct. Within ten years of the jump in births, the number of people 60 and older shows a noticeable rise. The timing of the peak, which depends strongly on the specific mortality rates assumed for the different ages, is at year 65 when one would expect the peak to occur between years 70 and 75. The decline, on the other hand, is very slow and the elevated number of people 65 and over persists well past time 100, when it reasonably should be almost back to the original value. It is not shown in this graph, but for this simulation the return takes till about time 180

which implies a significant population of people over 150 years in age. This is an unreasonable result that easily opens models to considerable criticism.

By adding additional levels and decreasing the delays between the stocks, we move closer to a result that matches what we would expect for chronological aging. A surprisingly large number of stocks, however, are required to achieve a reasonably close approximation to the way people actually age. To see this we build up an aging chain for ages 0 to 100 broken down into 10, 5, 2 and 1 year age cohorts. Each of these has the same representation as shown in Figure 2, but with different numbers of stocks (10, 20, 50 and 100). The results of simulating the different formulations are shown in Figure 5. ⁱⁱⁱ




Figure 5: Simulation results for different cohort granularity (number of stocks)

It is clear that as the number of levels increases, so does the sharpness of the response (Richardson and Pugh (1981), Ch. 3). We get less spreading and more precision, though changes in the accuracy of the response are not so obvious. The peak in the over 60 population should occur at time 70, precisely 60 years after the number of births returns to its base value. With ten-year cohorts this almost happens, but as we decrease the size of the cohorts to five years and then

two years the peak actually moves to the right toward 78. With the one-year cohort the peak does move back toward the left, but it still occurs six or seven years later than it should. In addition, the increase starts some fifteen years early. Using a one-year cohort clearly does not solve the blending problem.

One technical comment on the simulations shown in Figure 4 and Figure 5 is in order. For an aging chain with no intermediate outflows, increasing the number of stocks in the chain divides the total quantity among the constituent stocks. With a constant input this means that in steady state the total quantity summed across all stocks will not change as the number of stocks changes. When there is an intermediate outflow, however, this is no longer true. Because of the drain in each stock (deaths in this case) the outflow from each level to the next level in the chain will be less than the inflow. If the drainage rate (mortality) for each level is the same, the total population will be the same but the number of people in later age groups will be smaller for smaller cohort sizes.^{iv} When mortality is different for different ages, both the total population and the population for different age groups will change as cohort size changes. To compensate for this in our experiments, we adjusted mortality so that the equilibrium population aged 60 and above would be the same (at one million) for each cohort size. The total populations for the different experiments done here do not, however, match.

To summarize, aging chains allow us to distinguish different age groups and, by making them more granular, a model will simulate chronological aging more precisely. However, the problem of cohort blending persists with one-year cohorts. For demography, where very long times are involved in determining the population distributions, the magnitude of the blending is troubling.^v

Distinct Cohorts

The significance of cohort blending depends on model purpose, and we will return to the question of how blending impacts model results below. Independent of purpose, however, cohort blending is a face credibility issue when the number of people in a population group responds to a change in births before sufficient time has passed or when people linger too long in a past cohort

To overcome the blending problem, a standard solution used by the authors, and that has long been promoted by one of the authors (Eberlein) for models such as Threshold 21 (Barney et al.

(1995)), is the technique of pushing elements from cohort to cohort only at discrete time intervals. Using a standard rate and level formulation this would require something of the form:

$$\text{aging from 1 to 2} = \begin{cases} \frac{\text{population age 1}}{dt}, & \text{when time is a multiple of 1} \\ 0, & \text{otherwise} \end{cases}$$

This discrete formulation will work as intended only when using Euler integration because division by the computational time step for the model is necessary for the intended behavior.

Applying this solution, apparently, solves the blending problem. No changes occur in the number of people over 60 till time 65: the stock value ramps up with the incoming baby boom, slowly decreases as the extra people pass on and more quickly decreases as the age and mortality of the additional elderly increases (Figure 6). By Time 110 in this run the population has returned to its original value. This result uses a computational interval of $1/16^{\text{th}}$ of a year, with results recorded on the year-end.




Figure 6: Simulation results for discrete annual cohorts¹

While this approach apparently takes care of the blending problem, it introduces two other issues that are problematic.

The first is the difficulty of initializing the different cohorts so that their values at the beginning of the simulation are consistent with other reported year-end values. In this model form, at the

¹ Note that the scale for population used in the graph is larger in order to accommodate the higher peak value that results when cohort blending does not occur.

end of a year the surviving members of each cohort, except the last, are moved in their entirety to the next cohort. For the first cohort, however, there is no previous cohort to set values and as a consequence the cohort starts at 0 after the move and then is increased continuously by births. The first cohort is 0 each year-end, and for consistency needs to be initialized to 0. Since no observed population cohort does this, it is necessary to create a special cohort (call it “Age_{<0}”) with this property. In this case births, instead of directly entering the population less than 1 year old, first go into the Age_{<0} group then enter the age under 1 population at the end of the year. This means the age under one population is actually somewhere between one and two years old toward the end of each year which is somewhat confusing. An alternative to adding an Age_{<0} stock is to shift the stocks at the first computation after the end (or start) of the year rather than at the end of the year.

Initializing a model with discrete cohorts to equilibrium is also somewhat more involved than in the case of a pure aging chain. For the aging chain the equilibrium condition requires that:

$$\begin{aligned} \text{population}_{0 \text{ to } 1} &= \frac{\text{births}}{\text{mortality}_{0 \text{ to } 1} + \frac{1}{\text{time in cohort}}} \\ \text{population}_{i \text{ to } i+1} &= \frac{\frac{\text{population}_{i-1 \text{ to } i}}{\text{time in cohort}}}{\text{mortality}_{i \text{ to } i+1} + \frac{1}{\text{time in cohort}}} \\ \text{population}_{n-1 \text{ to } n} &= \frac{\frac{\text{population}_{n-2 \text{ to } n-1}}{\text{time in cohort}}}{\text{mortality}_{n-1 \text{ to } n}} \end{aligned}$$

where n is the total number of cohorts. Assuming that *time in cohort* is 1, the above equations can be simplified somewhat, though maintaining the more verbose form does help reinforce dimensional consistency.

Using the discrete cohort shift, with an Age_{<0} cohort (with no outflow), the equilibrium conditions become:

$$\begin{aligned} \text{population}_{<0} &= 0 \\ \text{population}_{0 \text{ to } 1} &= \text{births} * \text{time in cohort} \\ \text{population}_{i \text{ to } i+1} &= \text{population}_{i-1 \text{ to } i} * (1 - \text{mortality}_{i-1 \text{ to } i} * dt) \frac{\text{time in cohort}}{dt} \end{aligned}$$

$$population_{n-1 to n} = \frac{population_{n-2 to n-1} * (1 - mortality_{n-2 to n-1} * dt)^{\frac{time\ in\ cohort}{dt}}}{1 - (1 - mortality_{n-1 to n} * dt)^{\frac{time\ in\ cohort}{dt}}}$$

These are less intuitive and less numerically accurate than equations for a standard aging chain. If the Age_{<0} cohort is not used, and instead cohorts are shifted after the end of the year, the formulae become even more cumbersome and require an accommodation of the extra step before a shift.

The second issue with discrete cohort shifting is the substantial within-the-year variation associated with each individual cohort. Following each shift, the levels will continue to move based on the rates, which means decreasing by virtue of deaths for all but the Age_{<0} cohort. The graph in Figure 6 was created by simulating with a computational period of 1/16th of a year, but sampling the results only at the end of each year. If we sample at all computational periods, the results look very different (Figure 7).

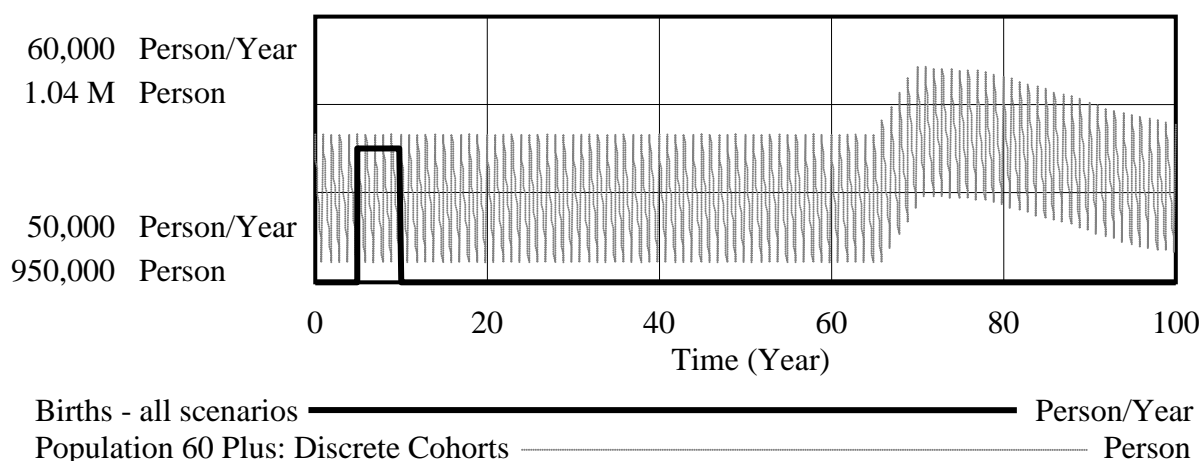


Figure 7: Simulation results for discrete annual cohorts sampled every 1/16th of a year

The enlarged scale in Figure 7 is necessary to capture the within-year movements, and the within-year movements are larger than the changes caused by the boom in births. At first this volatility seems surprising, given that we are looking at slow dynamics in which the shortest effective time constant (from mortality above age 90) is 4 years. Over the course of every year, about 6% of people over the age of 60 die. On reflection this behavior is consistent with our assumptions on mortality. In the observed system more people are aging into that group at the same time, and consequently the change is not nearly so dramatic (and is 0 in equilibrium).

However, the model formulation simulates that aging only at the end of the year, so every cohort decreases by its mortality rate over the course of the year and only at the end of the year does the entire cohort of 59 year-olds suddenly shift to the over-60 population cohort.²

Within-year variation is not always a problem and can be easily hidden by sampling the simulation results at the end of each year. Sparse sampling, however, often hides model shortcomings which can be very difficult to track down.^{vi}

While shifting cohorts is an effective way of preventing cohort blending, it also introduces problems with initializing models and obtaining consistent computations within the year. The approach described in the next section was developed to overcome these shortcomings and provide a more consistent representation of how people age chronologically.

Continuous Cohorting

The same reasons that led us to finer age disaggregation of population and then discrete cohort shifting point us one step further to what we call *continuous cohorting*.³ Using continuous cohorting, the members of each age cohort are updated on every computational interval so that there is no anomalous up and down within a year and no need for an extra indirectly measured cohort to help with initialization.

Creating a model with this level of age detail could be done using an aging chain construct with *time in cohort* set to be the same as the solution interval for the model. In this case the movement between two levels is complete at each step, though care needs to be taken to assure that the mortality associated with each cohort is properly accounted for. A modification to the standard movement equation such as:

$$\text{aging from cohort } i \text{ to } i + 1 = \frac{\text{population in cohort } i - \text{deaths from cohort } i * dt}{\text{time in cohort}}$$

² Demographic models, by computing only mid-year values (effectively assuming time step is 1) sidestep the issue of changes within the year. For example, see Preston, S. H., Heuveline, P. & Guillot, M. 2001. *Demography: Measuring and Modeling Population Processes*, Oxford, Blackwell Publishers. Ch. 1.

³ The term, *continuous cohorting*, describes the simulated outcome. As we describe here, the computations are by necessity discrete and require Euler integration.

will accomplish this. Subtracting deaths in this manner spreads them across all people, including those who would have been aging into the next cohort. This level of detail is normally not included in aging chains because the second term in the numerator is likely to be very small for an appropriately selected solution interval. However, as *time in cohort* approaches the solution interval, the term becomes important and does need to be included. When time in cohort is equal to the solution interval this equation reduces to simply

$$\text{aging from cohort } i \text{ to } i + 1 = \frac{\text{population in cohort } i}{\text{time in cohort}} - \text{deaths from cohort } i$$

While creating continuous cohorts in this manner is conceptually simple, the number of levels involved is very large. With a solution interval of 1/16 of a year to track 100 years of aging requires 1600 levels. Using subscripting and other compacting notational conventions, this can be done reasonably cleanly, but the formulation is still cumbersome. For example, special attention must also be given to breaking up both the initialization process and cohort specific mortality because it is almost certainly undesirable to track either of these in such detail.

By putting these fine time slices into a function outside the model equations, it is possible to get the benefits of near continuous computation while using reasonable age groupings. The specifics of implementation will depend on the application software; we describe the approach we took in Vensim with a user defined external function called COHORT CONTROL.⁴ This function works against a level that itself is broken into more aggregate cohorts of potentially varying length.

The arguments to the cohort management function are the cohort stocks, the length of the age groups (which need not be uniform), the inflow into the first cohort (births) and the number of cohorts in the model. At each computational interval, the COHORT CONTROL function shifts the members of the refined cohorts. Because this shifting is not intermittent there are no surprising moves in variables. For example, using five-year cohorts the results corresponding to those shown in Figures 6 and 7, with sampling at 1/16 of a year are shown in Figure 8.

⁴ A copy of the model and external function file discussed here, along with rudimentary documentation, are available on request from either author.

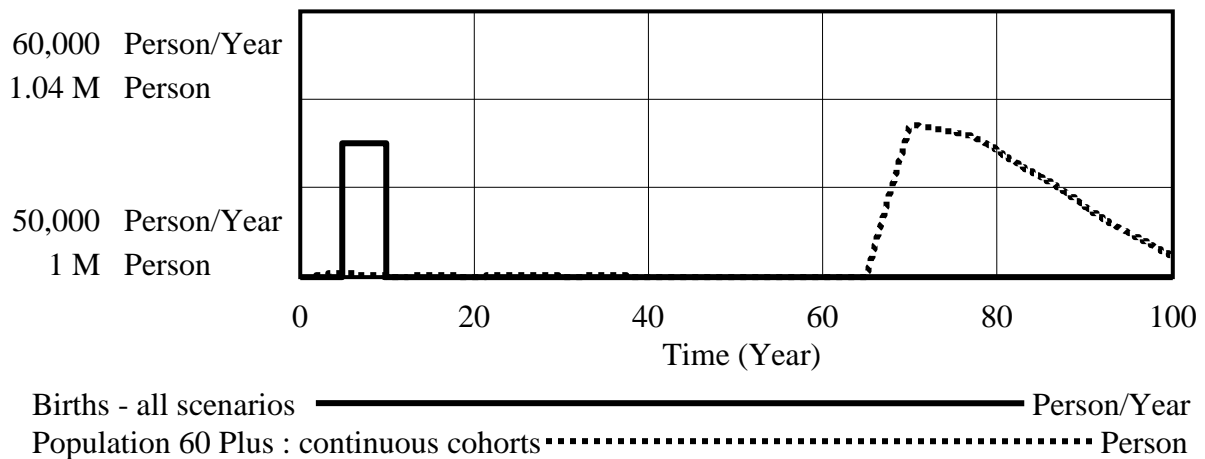


Figure 8: Simulation results for continuous cohorts (sampled every 1/16th of a year)

The behavior is almost identical to that in Figure 6. There is tiny amount of bumpiness early on that results from the initialization method described below. In practice, we found it necessary to use double precision computation to prevent numerical errors (e.g., the equilibrium initialization equations are numerically sensitive) from making that bumpiness even more pronounced.

In the real world, birth is not the only way to enter a population, and death is not the only way to leave. The COHORT CONTROL function, as we implemented it, allows secondary inflows and outflows such as immigration and emigration. All model equations are written against the model cohort groupings, and applied by the function to the finer cohorts used internally. If there is a net outflow from any of the model cohorts it takes away from the subdivided cohorts proportionally to their value. If there is a net inflow, it is decomposed evenly across the subdivided cohorts. In the case where there are significant inflows, such as immigration, a variant of the function can be used that specifies the inflow, and the inflow is spread evenly across the subdivided cohorts with the outflows removing proportionally.

When initializing the model to match historical conditions it is only necessary to specify the initial values for the model-based cohorts. The COHORT CONTROL function spreads the

values across the subdivided cohorts to smooth transition from one model-based cohort to the next.⁵

Initializing to equilibrium is somewhat more involved, and difficulties compound if the linear interpolation described above for assigning the subdivided cohort values is inappropriate. Typically this is not an issue for one-year cohorts, but five-year cohorts, whether based on equilibrium computations or historical starting conditions, it can be problematic. To overcome this, we have also implemented a version of the function that allows a finer specification of initial values but uses the larger cohort sizes for other operations (named COHORT CONTROL INIT in the external function file).⁶

When does it Matter?

We have presented a compact, comprehensive and relatively easy-to-use solution to the cohort blending problem. The solution overcomes the within-year volatility associated with discrete cohort shifting and allows analysis of dynamics in a conceptually continuous frame. Still, the question of when it is important to prevent such blending warrants more discussion. To elucidate this we present three examples. The first is taken from a study of dementia trends in Singapore and reflects on a specific and real issue. The second two are more theoretical in nature but are useful for thinking about the issue.

Thompson et al. (2012) present a demographically-based model for looking at future trends in age-related dementia of all types in Singapore, which prompted much of the analysis done in this paper. We configured an alternative version of the model which that study used to run in three ways: continuous cohorting with one year groupings, aging chains with one year groupings and

⁵ The algorithm for doing this is straightforward. We set the edge points to be the average of the two successive cohorts then use linear interpolation through the cohort's midpoint to ensure that the total quantity in the subdivided cohorts matches the quantities in the model-based cohorts.

⁶ To fill out the functionality needed in our research, we implemented a version of the function that allows the fine subdivisions to be implicitly addressed for inflows (COHORT CONTROL TRANSFER) as well as a way to measure populations within the fine subdivisions (COHORT COUNT). These variations are not germane to the main point of this paper but may be useful in other applications. In addition, to prevent issues in equation ordering the actual use of the COHORT CONTROL function is broken up into two separate function calls, one with a function called COHORT SHIFT. Details of this are available in documentation available from the authors.

aging chains with five year groupings. In the two cases with one-year cohorts, population is initialized by age and mortality from published one-year life tables (Singapore Department of Statistics (2010)) then held constant after 2010. In the case of the five-year cohorts the mortalities are averaged over their included ages and the populations are summed up to get the total for the cohort. A similar adjustment is made for the initial value of the dementia populations.

The model is calibrated to cohort time-series of population and deaths over the period of 2000 to 2010 to get an indirect measure of immigration and emigration by age group – phenomena that are extremely important in Singapore but for which reliable data are unavailable. The model is then run through 2050 to create a projection for four stages of dementia using established onset rates for different age groups.

To make the comparison in this experiment similar to what would result in a real project, the process of calibration was repeated with each of the different model configurations. The results of the calibrations were different, although we will not go through the details here. Because the research topic is age-related dementia, the outcome of interest is the difference in the dementia prevalence projections that result from the various formulations.

Figure 9 shows projections of the dementia population from 2010 to 2050. For the first fifteen years the results are all quite similar, but after that they begin to diverge. On reflection, the similarity in the early part of the simulation is not surprising. Cohort blending is slow process and over a relatively short period of time will be of only limited significance for slowly changing populations. The difficulty is that blending effects persist and do not correct themselves over time; instead the distortions grow as simulation time increases. This is why the issue is of such significance to simulating mortality and morbidity.

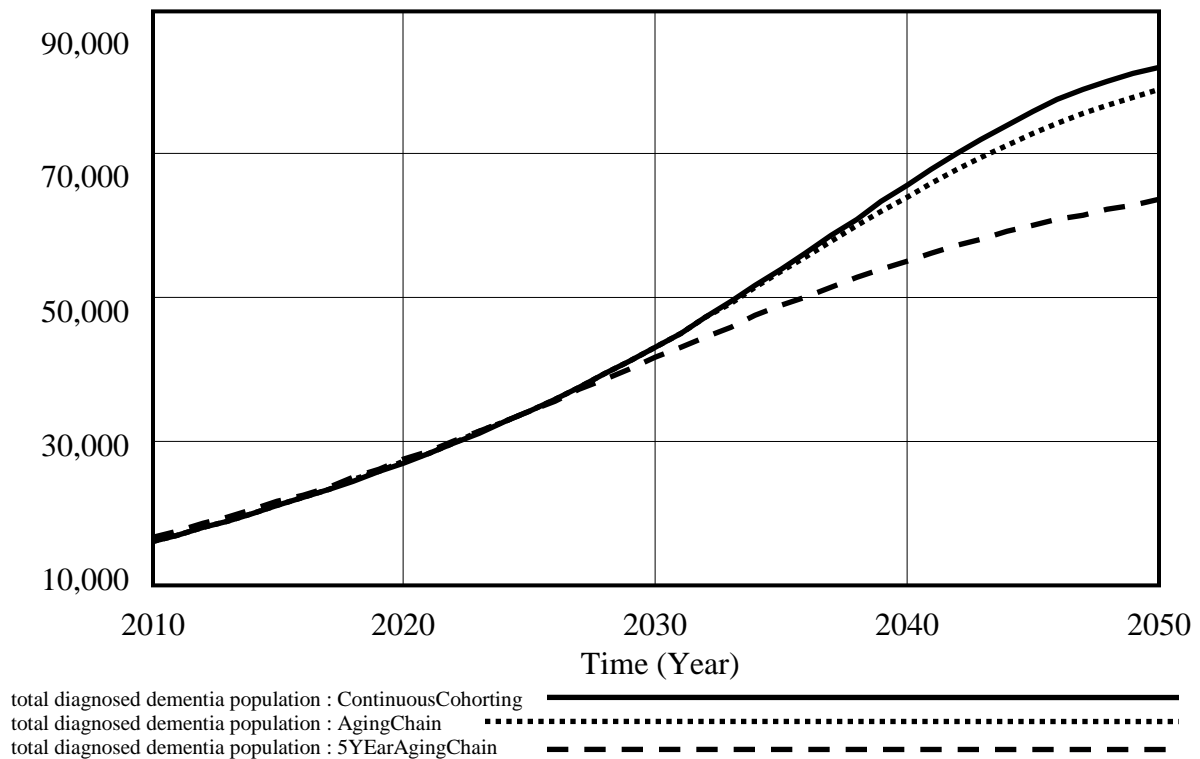


Figure 9: Dementia projections with different cohort assumptions

The continuous cohorting formulation generates the largest projection and the aging chain formulation with five-year cohorts the smallest. The aging chain formulations result in smoother curves with lower, broader peaks as a result of blending. The bigger the cohort population, the more pronounced this effect.

Quantitative and Qualitative Differences

We have demonstrated a significant difference between using a series of cascaded levels and tracking cohorts discretely. That difference, however, is largely quantitative. An interesting question is to what extent qualitative differences in behavior can result from the different formulations.

To shed some light on the broader question we turn to one of the first models in our field, the production distribution model from Chapter 15 of *Industrial Dynamics* (Forrester (1961)). This model uses third-order delays for all material flows, justified by the thought that this seems intuitively appropriate as explained in Forrester (1961), Chapter 9. Clearly, what is intuitive is very dependent on the knowledge and inclination of the person whose intuition is being called

upon. For Jay Forrester in 1960 it was some mixing, but not as much as would result from a first-order exponential smooth. For others the first-order exponential might suffice, and still others might not expect any mixing at all. Does it matter?

To look at that question we took the third order exponential delays in the original model and replaced them with first and 9th order exponential delays and a discrete fixed length delay (continuous cohorting). The results of this experiment are shown in Figure 10 with unfilled orders at the factory graphed against final orders at retail. The first order delay is the most damped, which is exactly in accord with what we would expect because this has the lowest phase shift among the three constructs. With the exception of the very sharp corner in the case of continuous cohorting there are no significant qualitative differences. The period is longest for the first order delays and this is again consistent with what we have seen before: the higher the mixing, the more things spread out.




Figure 10: Simulation results for the production distribution model with different delay orders

There is a clear continuum of outcomes. The third-order delay moves about half way to the continuous cohorting case and the ninth-order about two thirds of the remaining way. Even the kink observed at Time 7 is closely approximated with a ninth-order delay.

We see quantitative differences but not strong qualitative differences. The question is which formulation more accurately represents observed performance. In the case of real production and

distribution, things can pass one another in transit and some orders take longer than others to complete, so there is no logical argument that a third or first-order delay is clearly wrong. In this case using the third order delay Forrester originally chose seems to be an appropriate choice.

Finally, let us return to *World Dynamics* and *The Limits to Growth* – the sources of the two structures described at the beginning of this article. Population in the *World Dynamics* model is formulated in a single stock that is increased by births and decreased by deaths. If that stock is broken down into age cohorts, the sum across cohorts will still be increased by births and decreased by deaths. Unless the multipliers on death are made to be different for different age groups there will be no change in behavior if continuous cohorting is used for population. In short, unless a model makes use of the age structure of the population, changing the manner in which cohorts are formulated will not have any impact on results.

In the *Limits to Growth* model, on the other hand, there is an age structure with distinct mortalities for the different age groups in place that is used to determine fertility and the labor force. We can retain these characteristics but formulate the model with continuous cohorting to see if preventing cohort blending makes a difference. In Figure 11 we present the results for population, which shows the greatest disparity between the two formulations. Other variables are surprisingly similar, and the overall patterns of behavior are changed very little.

As can be seen in Figure 11, population does show significant differences. Population grows more quickly with the continuous cohort formulation and reaches a higher value, peaking at a later time.




Figure 11: Simulation results for population using alternative cohort formulations

The causes of the differences are interesting. Until Time 1960 the two formulations yield nearly the same results. But in the original formulation the number of people in the age group of 15-44 spreads to the older age groups earlier, leaving fewer 15-44 year-olds. This can be seen more clearly in Figure 12, which shows just the Population aged 15 to 44. With fewer people in the fertile population group, there are fewer births with the same fertility rate. The distortion created by accelerated movement of people into older age groups is significant for rising populations, while the longer residence time in younger population groups would be significant for falling populations. This means the direction of difference is an artifact of the dominant trend in the earlier part of the simulation.




Figure 12: Simulation results for 14 to 44 year-olds with alternative cohort formulations

It is worth asking whether the *Limits to Growth* model could be improved by formulating it with continuous cohorts. Doing only that, as we have done here, the answer is likely no. When models are developed, there are ongoing adjustments made so that the whole structure gives us insight into the system being represented. By replacing part of the structure we have negated those adjustments, and one would need to do more work to find the appropriate adjustments for the new formulation. It is an open question whether a more accurate rendering of chronological aging would improve the quality of the *Limits to Growth* model.

Conclusions

The principles of continuous cohorts can be applied to any population (even populations of inanimate objects) for which true chronological aging is the best approximation to behavior.

The mixing of individuals that is implicit in the most commonly used outflow equations, whether first or third order, make it difficult to accurately track aging populations. While this mixing is a good approximation to many real world phenomena such as disease progression, hiring, skill development, work flow and perception, it does not make sense for chronological aging where there is no possibility of one individual passing another in age. By increasing the order of the delay process this problem can be alleviated but not eliminated. In addition, the number of levels required for a good approximation to chronological aging is surprisingly large. By tracking different cohorts and aging them discretely it is possible to overcome the mixing problem. The traditional way of doing this with annual cohorts, however, introduces initialization difficulties and anomalous within-year behavior as the cohorts adjust toward the year-end condition normally reported. We find that by refining the length of time in each cohort to be the same as the model's solution interval (dt) we can overcome this behavior and generate smooth adjustment of all model variables. We call this continuous cohorting and using this solution turns out to be particularly important for looking at incidence and prevalence of age-related conditions, such as dementia, that are more strongly pronounced in the elderly. In general the use of continuous cohorting rather than an aging chain causes quantitative, but not qualitative, changes in model behavior. Some of the quantitative changes can, however, be significant and continuous cohorting stands out as a good approach for models containing endogenous demographic elements. The use of continuous cohorting can also increase acceptance among sociologists,

demographers and geographers as the computations are closer to those customarily employed in these fields.

References

- Barney, G. O., Eberlein, R. L., Qu, W. & Sharma, P. D. Year. The Threshold 21: Sustainable Development Model. *In*: Shimada, T. & Saeed, K., eds. System Dynamics '95, 1995 Tokyo. International System Dynamics Society, 22-29.
- Forrester, J. W. 1961. *Industrial Dynamics*, Cambridge MA, Productivity Press.
- Meadows, D. H., Meadows, D. L., Randers, J. & Behrens, W. W., Iii 1971. *The Limits to Growth*, London, Pan Books, Ltd.
- Meadows, D. L., Behrens, W. W., Iii, Meadows, D. H., Naill, R. F., Randers, J. & O., E. K. 1974. *Dynamics of Growth in a Finite World*, Cambridge MA, Productivity Press.
- Preston, S. H., Heuveline, P. & Guillot, M. 2001. *Demography: Measuring and Modeling Population Processes*, Oxford, Blackwell Publishers.
- Richardson, G. P. & Pugh, A. L. I. 1981. *Introduction to System Dynamics Modeling with Dynamo*, Cambridge MA, Productivity Press.
- Statistics, S. D. O. 2010. Population Trends 2010. Singapore: Singapore Department of Statistics.
- Sterman, J. 2000. *Business Dynamics: Systems Thinking and Modeling for a Complex World with Cd-Rom*, Irwin/McGraw-Hill.
- Thompson, J. P., Riley, C. M. & Matchar, D. B. Year. System Dynamics of Diminishing Family Size on Living Arrangements of Singaporeans with Ageing-Related Dementia. *In*: The 29th International Conference of the System Dynamics Society, 2011 Washington DC. System Dynamics Society.
- Thompson, J. P., Riley, C. M. & Matchar, D. B. 2012. Future Living Arrangements of Singaporeans with Age-Related Dementia. *International Psychogeriatrics*, forthcoming.

ⁱ For a discussion of the attributes of coflows, see Sterman (2000), Ch. 12.

ⁱⁱ The aging process is decoupled from mortality and the final cohort grouping done for age 60 and over in the example shown here.

ⁱⁱⁱ It is also worth noting that the average lifetime (over all stocks) changes with the number of stocks involved. This is in contrast to a typical aging chain (without intermediate outflows) for which the average residence time is just the sum of the residence times in each stock.

^{iv} This happens because the outflows decrease with each successive cohort so that the outflow from a cohort is smaller than the average outflow from its predecessor cohorts. It is also intuitively clear because the steady state average age of the population will decrease as more stocks are added to the aging chain.

^v Because we are focusing on continuous time model conceptualization we do not explore the relationship between the solution interval (dt) and the degree of blending. Using a larger the solution interval will make the blending less pronounced. In fact, using a solution interval equal to the time in a cohort (typically both would be 1) prevents blending altogether. Mathematically, this is the same as using year-on-year difference equations and this is the standard approach used by demographers. The techniques presented in this paper allow a consistent rendering of aging without sacrificing the benefits of continuous conceptualization.

^{vi} As an example, using standard one year life tables we had constructed a model in which we were consistently getting excess survivorship to age 1. The mortality after age 1 was substantially lower than infant mortality (as it almost universally is) and this was causing persistent problems. We had formulated the population without the extra stock and the 0-1 population was continually being emptied at the end of the year and then filling back up. At the end of the year it was at the appropriate level, but over the course of the year it went from 0 to the value. This meant the average value was basically $\frac{1}{2}$ of the year-end value and that, in turn, implied that infant deaths were $\frac{1}{2}$ the measured values. This was true even through the infant death rate being instantaneously reported matched our expectations. Once we went back to sampling at the computational interval it was obvious what was happening, but to that point it was a mystery.