Simple model analysis of the Artificial Anasazi

By Marco Janssen, July 2008

In their PNAS paper a number of parameters were varied to calibrate the Artificial Anasazi model to the historical records (better estimates of population numbers based on room counts). The default parameter values are Harvest adjustment 1.0, Harvest variance 0.1 (there are two variances, spatial and temporal. We assume Axtell et al. vary them both the same when they talk about changing Harvest variance), Household death age (30 years), Maximum Fertility age (30 years), and fertility rate (0.125).
With these values they produce a population curve like Figure 1 in Dean et al (2000).

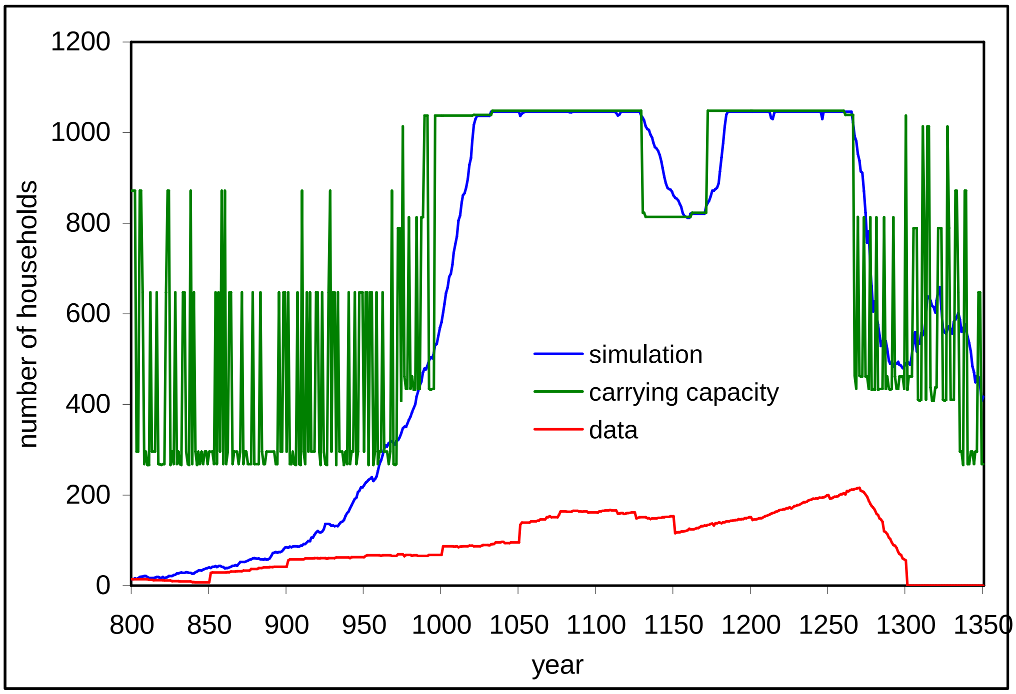

Figure 1: Results with the default parameter values in line with Dean et al. (2000).

Carrying capacity in Figure 1 is defined as the number of cells which has a base yield equal or higher than the nutrition needs. These are potential farming plots. You see fluctuations in these values due to fluctuations in PDSI values in the data, and periods in which streams and alluviums exists. The simulated number of households grows towards the carrying capacity. When the carrying capacity drops, this does not immediately lead to a drop in the population since stocks of corn lead households to hang on longer.

In their PNAS paper Axtell et al. calibrated the model by minimizing the difference of the simulated and historical data using 15 simulations. The Harvest variance became 0.4 and the Harvest adjustment rate was reduced to 0.6. The fertility rate remained the same: 0.125. They assume variation in the agent population of the Deathage (30-36) and the maximum fertility age (30-32).

We use these parameter values and run our implementation 100 times (Figure 2). In their PNAS paper they published the best fit (as we also did in the ODD description), but we see that there is some variation in the results.

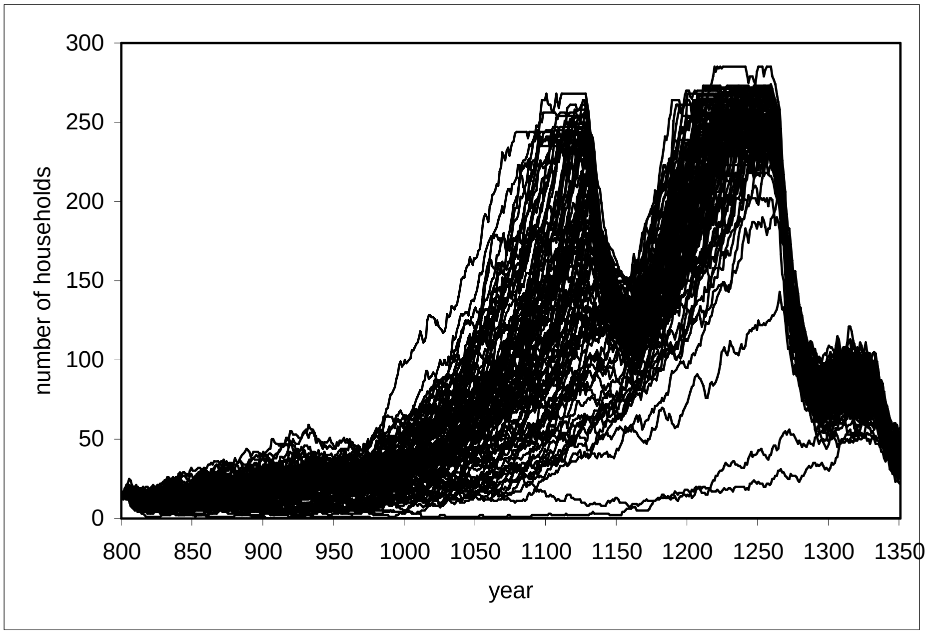

Figure 2: 100 runs with the “calibrated” Artificial Anasazi.

In Figure 3 I projected the best simulation of the 100 simulations, with the historical data, and the carrying capacity. The carrying capacity is changed due to changes in the Harvest adjustment rate (if you assume people get only 60% of their harvest this reduces the carrying capacity) as well as the harvest variation. We see again that the simulated population grows towards the carrying capacity, and is somewhat delayed when the carrying capacity drops due to storage of corn.

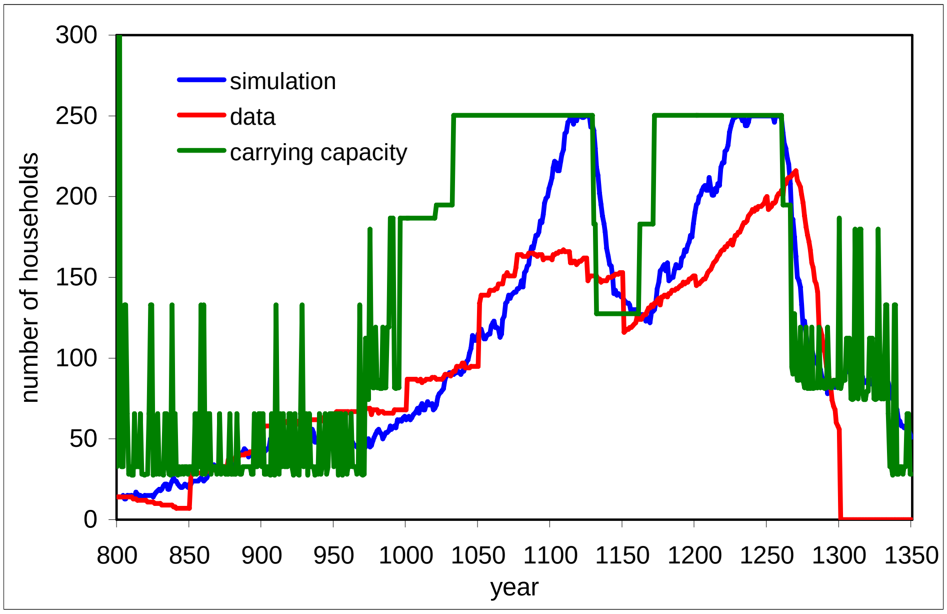

Figure 3: Simulated (“best” fit) and historical data and the carrying capacity for the parameter values from Axtell et al. (2002).

So, what leads to this good fit of the simulation with the data? We performed some sensitivity analysis of the calibrated model, but assumed no heterogeneity of the agents in fertility, Deathage or maximum fertility age. We assumed a Deathage of 30 and a maximum fertility age of 30. We ran the model 100 times for each variation of the parameter values. We changed one parameter at the time. In the figures below you see how this affect the value of the best fit (L1 stands for L1 norm, which is the sum of the absolute differences between historical and simulated data. The L2 norm, quadratic difference between data and simulation, results in the same qualitative insights). We depicted both the L1 norm for the agent based model, as well as the carrying capacity estimate alone.

Figure 4 shows that reducing the amount of net harvest (by lowering the parameter harvest adjustment) reduces the misfit compared to the default case (Figure 1). Interestingly, by lowering the harvest adjustment to 0.5 the carrying capacity estimate is doing better than the agent-based model. Figure 5 shows that also the change of harvesting variance can improve the fit with the data considerable, although this might be a better fit for the carrying capacity instead of the agent based model (for harvest variance = 0.3). Figures 6-8 show that changing parameters on death age and fertility will obviously has no impact on the carrying capacity, but has modest effects on the fit between the simulated and historical data. Lowering the deathage will lead to lower populations, and a worse fit. For the maximum fertility age, and increase has no effect since the parameter deathage was held constant at 30 years. A small increase in fertility rate improves the fit slightly.

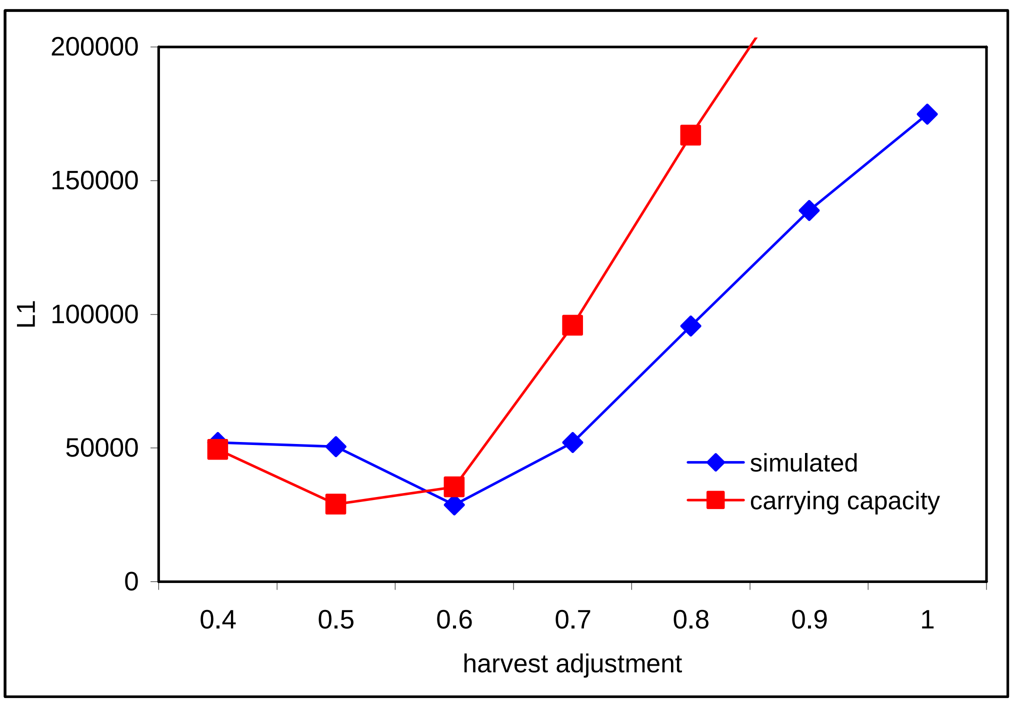

Figure 4: Average value of L1 norm for different values of the parameter harvest adjustment.

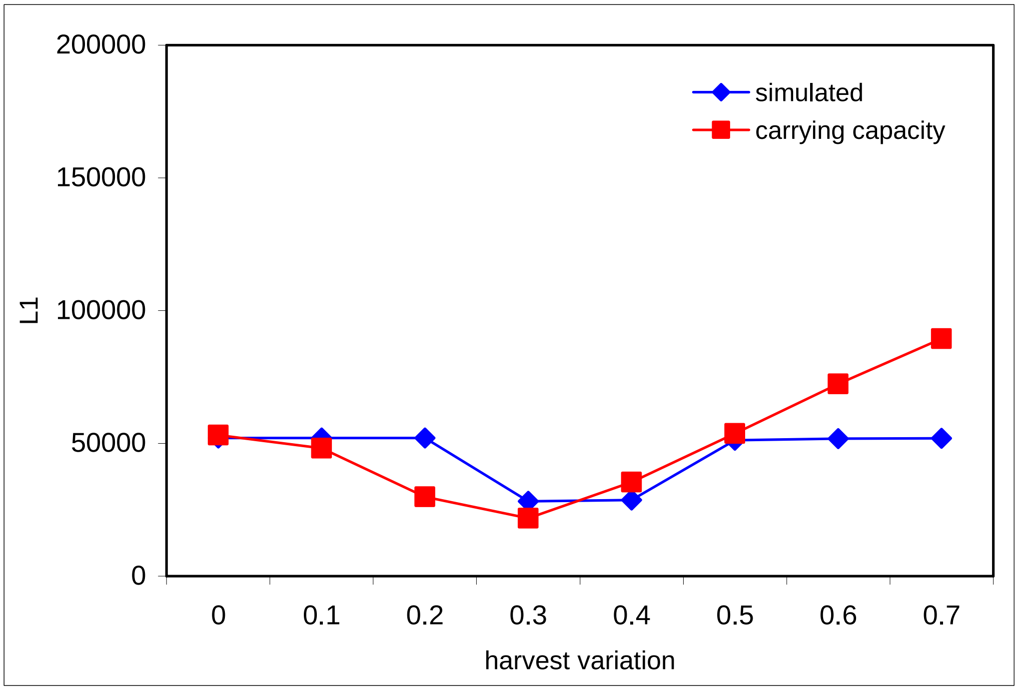

Figure 5: Average value of L1 norm for different values of the parameter harvest variation.

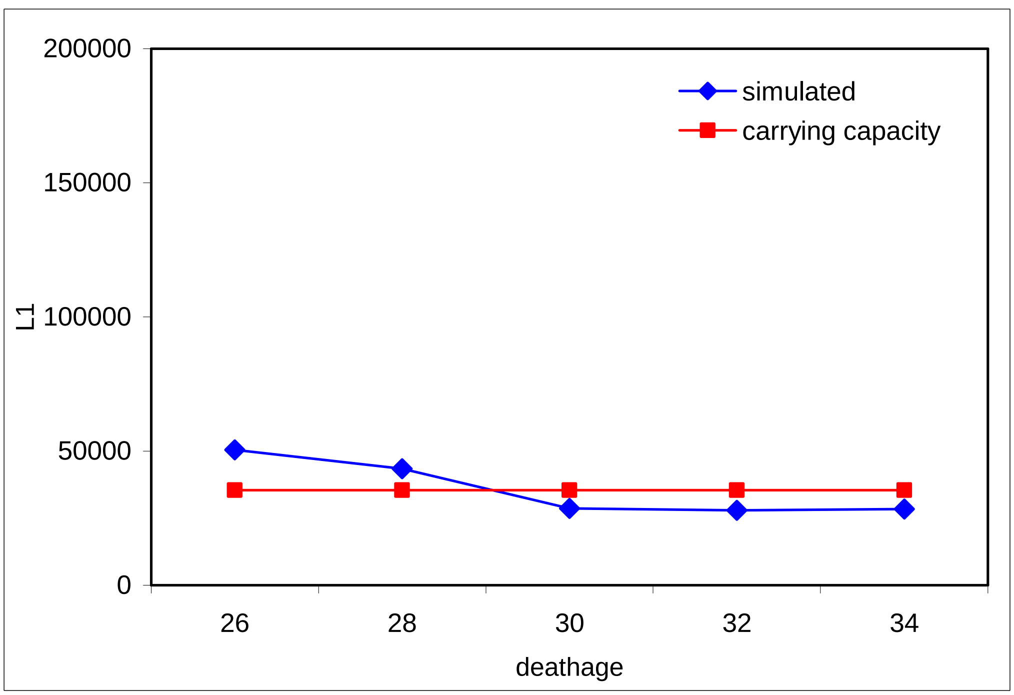

Figure 6: Average value of L1 norm for different values of the parameter deathage.

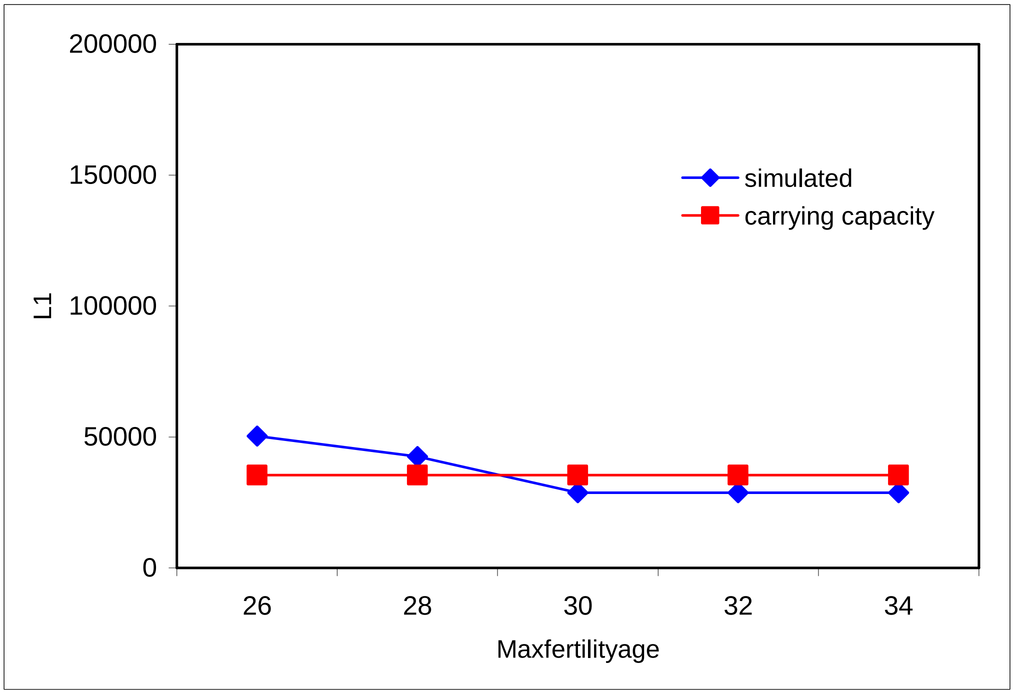

Figure 7: Average value of L1 norm for different values of the parameter maximum fertility age.

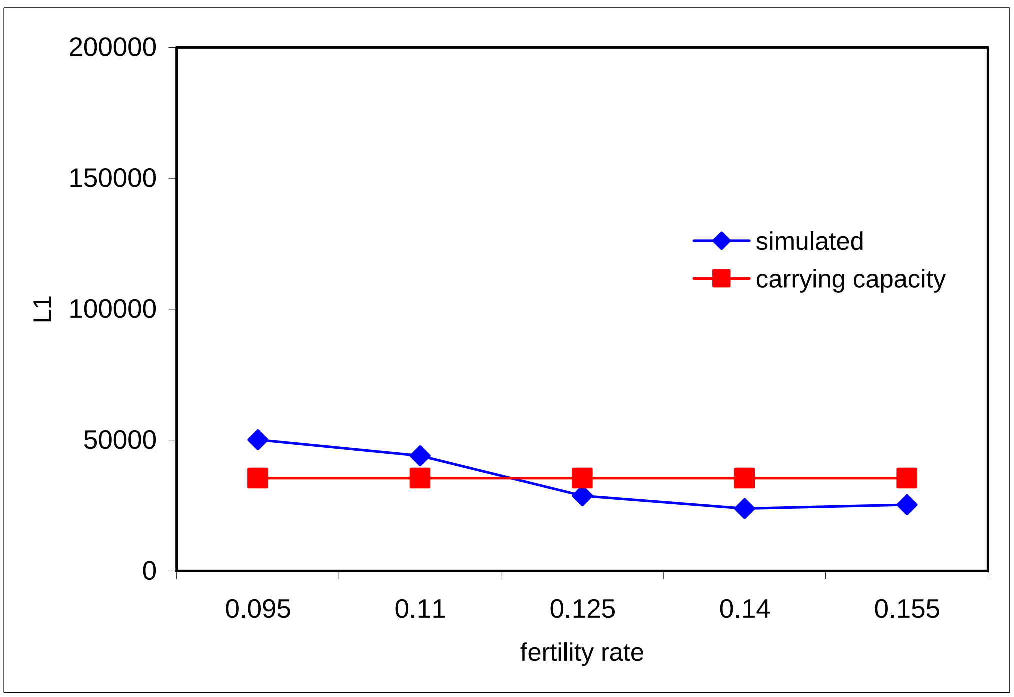

Figure 8: Average value of L1 norm for different values of the parameter fertility rate.

This analysis was done with a replication of the Artificial Anasazi model. I may have misinterpret some of the original code and therefore one has to be careful in making interpretations of the original model. I also did not look at the spatial explicit population levels. Nevertheless, this simple analysis shows that the fit between the data and model is mainly derived by changing the values affecting the carrying capacity. Variations of the agent specific behavior has limited impact on the results. Given the uncertainty of population numbers, one may wonder why not just the estimates of the carrying capacity are used, which already shows that the environment alone could not lead to the complete abandonment of the Long House Valley.

A more detailed analysis may include a systematic exploration of the parameter space that includes the interaction effects of the parameter values.