Using Microsimulation to Estimate the Impact of Crisis Mortality on the Marriage Market. An Application to the Historical Demography of Roman Italy

Saskia Hin[1,*] **and Emilio Zagheni**[2,*]

[1]Max Planck Institute for Demographic Research
[2]Queens College & CUNY Institute for Demographic Research

Please do no not cite this extended abstract without permission of the authors

Abstract

Demographic microsimulation has been used very successfully to model the dynamics of kinship structure in historical and contemporary populations. One of the main limitations of current approaches is the calibration process, which typically requires ad hoc solutions. We propose a method to calibrate demographic microsimulation to historical records and we apply it to the case of ancient Roman Italy. We evaluate the impact that the Punic wars had on population growth via both mortality of soldiers and the indirect effect on fertility rates resulting from a form of marriage squeeze. Preliminary results indicate that when male mortality perturbed marriage markets, the share of Roman women unable to find a partner increased substantially. As a result, fertility for the women bound to get married during the war decreased dramatically. Although the impact on population growth was transitory, it took several decades before the population size returned to the pre-war levels.

*Authors are in alphabetical order reflecting equal contribution to the paper.

Introduction

Demographic microsimulation has been used very successfully to model the dynamics of kinship structure in historical and contemporary populations (e.g., Wachter, Hammel and Laslett 1978; Wachter 1997; Wachter, Knodel and VanLandingham 2002). One of the most widely used demographic microsimulators is SOCSIM, a computer program that was originally developed by Eugene Hammel and Kenneth Wachter, at UC Berkeley in the 1970s (e.g., Hammel, Mason and Wachter 1990). The simulator has been designed to model very detailed sub-groups of a population, and to address a wide range of research questions.

One of the most understudied areas of demographic microsimulation (and, more broadly, agent-based models) is the calibration process. In principle, perfect knowledge of demographic rates should lead to an unbiased reconstruction of the kinship network through demographic microsimulation. The only uncertainty associated to the simulated kinship structure would be related to the stochasticity of the microsimulation. In practice, knowledge of vital rates is far from being perfect. Kinship reconstruction and forecasting demand a level of detail for demographic rates that is often missing in available data sets. For instance, transition rates from one marital status to another one are usually not readily available and estimates may not be very accurate. Fertility rates are usually not broken down by marital status or parity. In most cases, demographic rates that are used as input to the microsimulation need to be estimated from various data sources with different sampling errors. Even when reliable data sources exist to compute demographic rates broken down by the categories of interest, the heterogeneity of the population's rates within the tabulated categories, and the complexity of the matching process in the marriage market, constrain the accuracy of the microsimulation.

Traditionally, the problem of calibrating SOCSIM has been addressed using *ad hoc* tuning. Input rates are adjusted on a trial and error basis in order for the output of the simulation to match key summary demographic measures obtained from a population census or sample surveys. The development of methods to calibrate SOCSIM has been mainly constrained by the limitation of computer power in the past. Traditional methods have heavily relied on minimizing the number of simulation runs by using expert judgement in adjusting the input demographic rates in a consistent and appropriate way.

With increasing computer power, there has been a growing interest for the development of methods to calibrate simulation models. Bayesian methods, in particular, have proved useful to formalize the process of calibration and statistical inference. These approaches have not been developed or used for SOCSIM. However, some of the key features of these approaches can be applied to the context of demographic microsimulation. In this paper, we propose a method to calibrate demographic microsimulation inspired by recent developments in this area pioneered by Adrian Raftery and his colleagues (Poole and Raftery 2000; Sevcikova, Raftery and Waddell 2007). We offer an application to the historical demography of ancient Rome. More specifically, we analyze the consequences of the Punic wars, and the related mortality of soldiers, on marriage rates, number of births, and short- and long-term population dynamics. The main idea is that, in a society where practically only married women give birth, the impact of a long conflict does not have only a direct effect on population growth (through the mortality of men), but also an indirect effect (through a 'marriage squeeze' that generates posponement of childbearing).

The Historical Debate

Questions regarding population size, population trend and demographic structures are fundamental for our understanding of Republican history. How we ought to judge what voting rights meant in the Republic, how the performance of the Roman economy compared to that of later pre-industrial counterparts and what impact it had on the living standards of individuals, how the course of Roman history might have affected by the availability of military manpower and the resilience of free labourers - these and similar questions cannot be answered without at least an approximate idea of population dynamics (cf. Scheidel 2008). It does not come as too much of a surprise then that recent years have witnessed a revival of intense debate on the demographic background against which political and social developments during the period between Hannibal and Augustus took place. Did Italy by the end of the Republic count about 6 million inhabitants, or rather about 16 million? And was its free citizen population dwindling as a result of continued natural birth deficits, or growing forcefully instead?

Traditionally, macro-demographic debates have been strongly rooted in literary evidence on the Republican and Augustan censuses. Whereas some of the recent publications have drawn on this literary evidence to propose new demographic scenarios (notably Lo Cascio 1994, De Ligt 2004, and Hin 2008), others have turned to archaeological survey evidence to seek an answer to demographic questions (e.g. Launaro 2011, De Ligt 2012).

The widely divergent population sizes adhered to in this so-called 'low count high count' debate rest on implicit assumptions regarding a range of demographic parameters, including population growth rates, marriage ages, life expectancies and fertility levels. Theoretically, these underlying assumptions can be tested against historical demographic evidence. This would allow us to evaluate the respective plausibility of diverse population development scenarios. In practice, such an approach is only possible through microsimulation: only then can demographic assumptions underlying various population scenarios be made explicit.

Were there sufficient men of marriageable age for Roman women to find a partner during and after the war, which can serve as a case study of shock mortality? Did phases of intense conquest and military loss have long-term consequences for population size? Could the Roman citizenry maintain itself as a population, and even grow, without changing behavioural preferences and strategies in response to mortality shocks? Or did historical conditions require adaptive behaviour to survive? And could any feasible combination of adaptive strategies lead to rapid population growth? Our paper contributes to the debate in historical demography by evaluating the short- and long-term impacts of the wars and, more generally, of high mortality rates in a warring society on marriage dynamics, fertility, and population trends.

Calibration of the Microsimulation Model

Figure 1 shows a schematic representation of the approach that we use to calibrate the microsimulation model. The microsimulator requires demographic rates as input, and generates population files as output. The population files can then be analyzed, and summary demographic measures can be extracted. These quantities can be compared with the respective summary statistics obtained from historical records, like the Roman Egypt census or Roman commemorative shift inscriptions.

The goal of the calibration process is thus to find a set of input rates which are associated to simulated populations whose key characteristics closely match the ones extracted from historical records. We may rescale some age-specific demographic input rates using a set of parameters θ. We express our uncertainty about the rescaling parameters using a probability distribution for the parameters; that is the prior distribution. The choice of a particular set of parameters from the prior distribution yields a population output for which we can compute the likelihood, based on a comparison with historical records. The combination of the prior distribution and the likelihood gives the posterior distribution for the rescaling parameters. The posterior means for the parameters of interest are the final choice for the rescaling parameters.

We assumed that the shape of the estimated age-specific rates from historical records is fairly accurate and thus most of the uncertainty is associated to the scale or level of the parameters. Based on this premise, we chose a set of parameters that rescale age-specific fertility rates and marriage rates. We refer to these parameters as θ.

For the set of rescaling parameters θ, the likelihood is computed as follows:

1. Run SOCSIM n times, with the same set of chosen parameters, but different seeds, and store the outputs.

2. From the n simulation outputs, compute the mean vector μ and the variance-covariance matrix Σ for a set of five k summary quantities.

3. Based on a normal approximation, the probability of observing the historical records for the key summary quantities, x, given the chosen set of parameters is:

$$f_X(x) = \frac{1}{(2\pi)^{k/2}|\Sigma|^{1/2}} exp(-\frac{1}{2}(x-\mu)'\Sigma^{-1}(x-\mu)) \quad (1)$$

   This is the likelihood for the specific set of parameters

Given a prior distribution and a parameterization to compute the likelihood, the posterior distribution for the rescaling parameters is obtained using the SIR algorithm (Rubin 1987, 1988):

1. Sample with replacement a number m of parameter vectors from the prior distribution.

2. For each sampled vector, compute the sampling importance weight, which is proportional to the likelihood for the sampled vector.

3. Sample with replacement from the m parameter vectors with probabilities proportional to the weights to approximate the posterior distribution.

A standard choice for the point estimate of the parameters is the mean of the posterior distribution.

Preliminary Results

Figure 2 shows simulated population pyramids of Roman citizens before the conflict and at the end of the first Punic war. The population pyramids provide an idea of the large impact that the war had on the population of Roman male citizens. Figure 3 shows the simulated proportions of Roman citizens never married, by sex, over time. When male mortality perturbed marriage markets, the share of Roman women unable to find a partner increased substantially. As a result, fertility for the women bound to get married during the war decreased dramatically. Figure 4 shows the overall impact that the Punic wars had on population size. Although the impact on population growth was transitory, it took several decades before the population size returned to the pre-war levels.

As we develop our models, we plan to evaluate different scenarios and the effect of behavioral responses to the mortality shock. We will also be able to make statements about the uncertainty of our results. Microsimulation, and, in particular, statistical calibration of simulation models, will become increasingly relevant in historical demography, as more and more historical census micro-data become available.

References

[1] De Ligt, L. 2004. Poverty and demography: the case of the Gracchan land reforms, *Mnemosyne* 57.6: 725-757.

[2] Hammel, E.A., C. Mason and K.W. Wachter. 1990. SOCSIM II, a Sociodemographic Microsimulation Program, rev. 1.0, Operating Manual: Graduate Group in Demography Working Paper No. 29. Berkeley, California, University of California, Institute of International Studies, Program in Population Research.

[3] Hin, S. 2008. *Counting Romans*, in People, land and politics. Demographic developments and the transformation of Roman Italy, 300 BC - AD 14, eds. L. de Ligt and S. Northwood. Leiden: 187-238.

[4] Lo Cascio, E. 1994. The size of the Roman population: Beloch and the meaning of the Augustan census figures, *Journal of Roman Studies* 84: 24-40.

[5] Poole D. and A.E. Raftery. 2000. Inference for Deterministic Simulation Models: the Bayesian Melding Approach. *Journal of the American Statistical Association* 95(452):1244-1255.

[6] Rubin, D. 1987. Comment on "The Calculation of Posterior Distributions by Data Augmentation", by M. Tanner and W.H. Wang. *Journal of the American Statistical Association* 82:543-546.

[7] Rubin, D. 1988. Using the SIR Algorithm to Simulate Posterior Distributions. *Bayesian Statistics* 3 (J.M. Bernardo, M.H. DeGroot, D.V. Lindley and A.F.M. Smith, eds.) 395-402. Clarendon Press, Oxford, U.K.

[8] Saller, R. P. 1994. *Patriarchy, property and death in the Roman family.* Cambridge.

[9] Scheidel, W. 2008. *Roman population size. The logic of the debate*, in People, land, and politics. Demographic developments and the transformation of Roman Italy, 300 BC - AD 14, eds. L. de Ligt, and S. Northwood. Leiden: 17-70.

[10] Sevcikova H., A.E. Raftery and P.A. Waddell. 2007. Assessing Uncertainty in Urban Simulations Using Bayesian Melding. *Transportation Research*. Part B 41:652-669.

[11] Wachter K.W. 1997. Kinship resources for the elderly. *Philosophical transactions of the royal society of London - Series B: Biological sciences* 352(29).

[12] Wachter, K.W., E.A. Hammel and P. Laslett. 1978. *Statistical Studies of Historical Social Structure*. New York, Academic Press.

[13] Wachter K.W., J.E. Knodel and M. VanLandingham. 2002. AIDS and the elderly of Thailand. *Demography* 39(1):25-41.

Figures

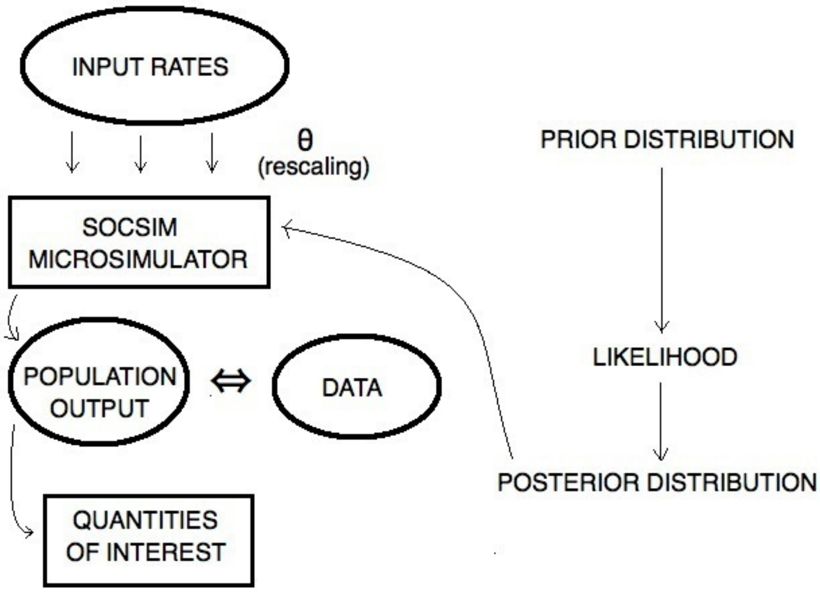

Figure 1: A schematic representation of the approach that we propose to calibrate the microsimulation model.

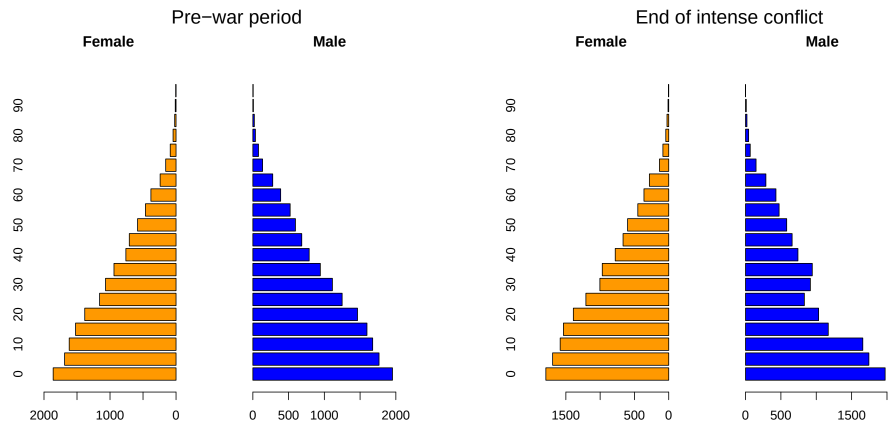

Figure 2: Simulated population pyramids of Roman citizens before the conflict and at the end of the First Punic war.

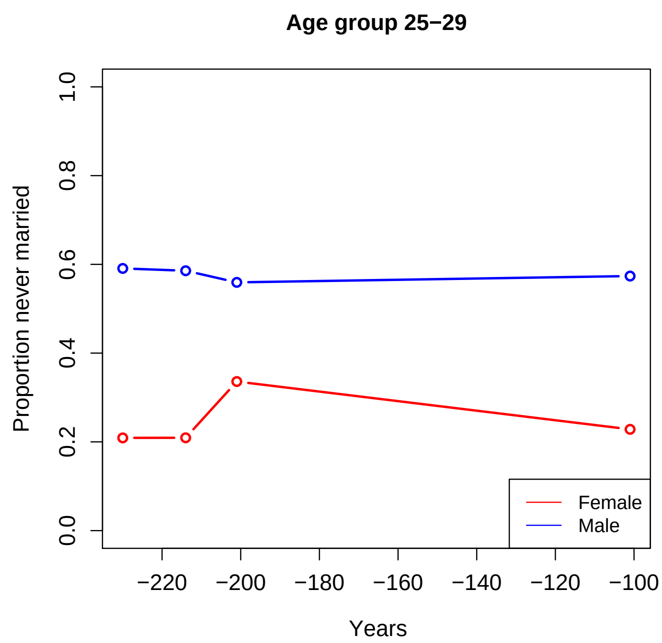

Figure 3: Simulated proportions of Roman citizens never married, by sex, over time.

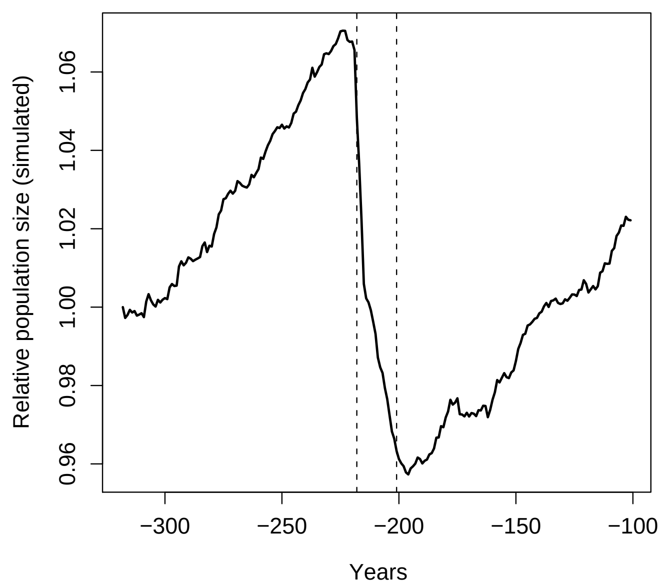

Figure 4: Simulated relative population size of Roman citizens over time. The dotted lines indicate the temporal interval of the Punic wars.