The "Wedding-Ring": An Agent-Based Marriage Model Based on Social Interaction

Francesco Billari

Institute of Quantitative Methods, Bocconi University, Milan

Alexia Prskawetz and Belinda Aparacio Diaz

Vienna Institute of Demography, Austrian Academy of Sciences

Thomas Fent

ETH Zurich, Department Management, Technology, and Economics

March 24, 2005

Abstract

In this paper we study the emergence of the age-at-marriage curve from the bottom-up, building an agent-based model of marriage based on social interaction. More specifically, we build a population of agents whose willingness to marry depends on the share of relevant others who are already married and on the availability of partners. Agents live on a circular space. Our simulations show that age-at-marriage curves of realistic shape can emerge from micro-level hypotheses and social interaction.

1   Introduction

The timing of marriage has been studied from two different perspectives in the empirical literature in the social and behavioral sciences. On the one side, demographers and sociologists have focused on explaining and modelling important stylized facts such as the typical shape of age-at-marriage curves;

their analytical strategies have usually relied on mathematical and statistical macro–level models. On the other side, psychologists and economists have focused on studying and modelling the process of partner search; micro–level assumptions have usually been at the heart of this approach. More recently, agent–based modelling has been proposed as a convenient approach to build models that account for macro–level marriage patterns starting from plausible micro–level assumptions and allowing for the interactions between potential partners that typically take place in the marriage market (see e.g. Todd and Billari, 2003; Todd and Simão, 2003; Todd et al., 2005). The study if macro–level outcomes of micro–level models containing social interactions allow to bridge the two perspectives we started from.

Partnership formation is by definition social interaction itself. Potential partners, in contemporary societies, interact socially before cohabiting or getting married. Nevertheless, there is important evidence that social interactions taking place in the marriage market are not limited to those between potential partners. The influential macro–level model of marriage patterns proposed by Hernes (1972) was built on the assumption that a diffusion process takes place within a cohort of individuals, with the share of married "peers" influencing the propensity to marry. In that case, the assumption was that members of the same cohort constitute the influential peer group. Hernes' assumption was mostly based on the casual evidence of everyday life. More recently, the study of the impact of social interaction on key decisions concerning our lives, such as getting married and having children has emerged as an important field of research (see e.g. Bongaarts and Watkins, 1999). More specifically, two distinct processes of "social learning", reflecting information exchange in a network, and "social influence", reflecting normative pressure in a network, have been identified (Montgomery and Casterline, 1996). Recent evidence from qualitative surveys shows that social influence and social learning are among the important factors in the decision to get married in a contemporary society (Bernardi, 2003). In what follows we shall refer to social interaction in general, although the role of social influence is likely to be more important than the role of social learning.

More specifically, we build a marriage model starting from the micro–level, including social influence as the key force driving the marriage process. Our aim is to let the typical macro–level shape of age-at-marriage patterns emerge "from the bottom up", as an outcome of our assumptions on individual behavior and social interactions. This approach is very widely used in agent-based models, for which the application to demographic choices has

2

been recently advocated (Billari and Prskawetz, 2003).

More specifically, we build an agent-based model in which we assume that agents belonging to the social network of an agent — the "relevant others" — influence the desire of an individual to get married. Such desire is mediated by features of the marriage market (i.e. the availability and location of potential partners). As two partners get married they may start having children. Similarly to Todd and Billari (2003) and Todd et al. (2005), the main benchmark against which we test our model is the shape of the age-at-marriage hazard function, i.e. the age-specific probability of marrying conditional on not having married by a certain birthday.

To illustrate the shape of these hazard rates for marriage, we show in Figure 1 the empirically observed functions for men and women in three populations of the late twentieth century: Romania, 1998, and Norway, 1978 and 1998. In all the cases shown in the figure, notice that the rise of age-specific probabilities is faster than its decrease. Although the shape of the curve looks rather different for Norway 1998, where non-marital cohabitation is widespread, it can still be described qualitatively in a similar way. In addition, hazard rates tend to converge to a level close to zero at later ages. This typical hazard rate function can be observed for several other populations, and it is this overall pattern that we want to account for in our models.
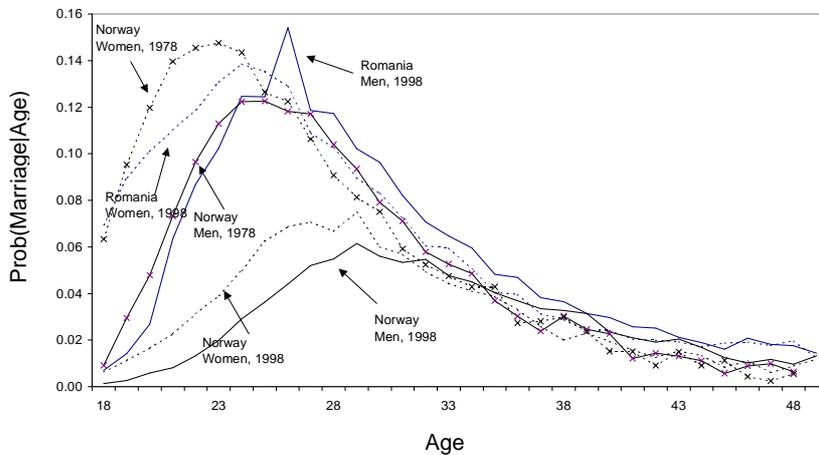

Figure 1: Hazard function for marriage in European populations. Source: own elaborations on Eurostat, New Cronos database.

This paper is structured as follows. In Section 2, we discuss the hy-

3

potheses that guide our modelling strategy. In Section 3, we present our agent-based model. In Section 4, we describe the implementation of the model. Simulation results are discussed in Section 5. Section 6 includes an outline of possible and useful extensions to the model.

2   Social interaction and marriage patterns: Theory and hypotheses

For an individual, the set of "relevant others" consists of people who are close to her/him, i.e. the member of her/his "social network". Closeness is a general feature we shall exploit in what follows. In our context, the term "close" refers to a distance that may represent a spatial distance (that is, neighbors constitute relevant others), but might as well represent a distance in terms of kinship, age, education, professional occupation, and so on. The size and characteristics of an individuals' social network may themselves depend on the individuals' characteristics. For instance, the number of relevant others increases with age during youth and adulthood, at least up to ages that are important for processes such as getting married or have children (Micheli, 2000). Moreover, education may have an impact on the spatial mobility of an individual (e.g., more frequent and long-distance trips for higher educated individuals). This could explain the spatial enlargement of the social network, while the number of persons within the network may extend also independent of mobility. E.g. an individual who is not travelling much will perhaps compensate this by having more social contacts within a local neighborhood. Education may also enhance the ability to learn new information and cause a delay in the age of marriage.

If one needs to operationalize the importance of social interaction in the decision to get married, besides the mere need to find an available partner around oneself, we can look at how relevant others behave. Relevant others provide information (i.e. social learning), and seeing relevant others behaving may trigger normative pressure (i.e. social influence). Of course, only social influence has a clear impact towards the diffusion of a certain behavior, but if we assume that marriage makes people happy at least for a while (see e.g. Clark and Oswald, 2002; Kohler et al., 2004), also social learning triggers the same effect. It is thus likely that the share of married people among the relevant others in one's social network has presumably a positive impact on

the individuals' desire to get married (Bernardi, 2003). This was the key assumption of the macro–level diffusion model developed by Hernes (1972).

More specifically, we assume that a high share of married within a certain neighborhood in a theoretical multidimensional space (spatial location, kinship, age, education, professional occupation, etc.) increases the social pressure to get married among those not yet married within that neighborhood. This implies that a diffusion process takes place, with the already married affecting the not yet married population. One important difference with respect to diffusion processes explaining the transmission of infectious diseases is that in case of marriage it is not sufficient to get "infected" — in our case, to experience social interaction effects — by married people. Obviously an infected person, that is a person experiencing a high level of pressure through social interaction and, therefore, wants to get married, also needs to find a partner who is not yet married. This partner will usually also be somebody within the individuals' social network. Hence, the highest incidence of marriage occurs within a social network exhibiting a relatively high share of both married and unmarried persons. On the other hand, in case of an infectious disease an individual who is in contact with an almost entirely infected population experiences the highest exposure to the risk of getting infected.

Moreover, not only the share of married people within the social network, but also the time span since they got married may have an influence on the social pressure. This may be due to the fact that weddings that have just taken place may make a stronger impression within the social network than weddings that have taken place decades ago. For instance, the increase in happiness related to marriage might slowly vanish over time (Clark and Oswald, 2002), or the information on married couples might become less and less relevant for singles. An unmarried person who has already been confronted with married people for a long time without getting actually married her-/himself may feel less pressure than an unmarried person witnessing the transition of the network partners from unmarried to married status. Moreover, weddings are occasions during which potential partners may meet. Another aspect to be considered is the possibility of married people to get divorced, which may cause a negative attitude toward marriage within their social network. However, to model this effect we would need to model divorce in addition to marriage. In what follows we shall consider only first marriage.

3   An Agent Based Model

To study the interactions among individuals and their impact on prevalence and incidence of marriage we apply an agent-based computational model. The agents live in a world which is arranged along a circular line — due to the analogy with a wedding ring, this circle may also be called a ring. Hence, the spatial location of each agent is entirely determined by its angle, $\varphi \in [0, 2\pi]$. The advantage of a circular line is the fact that each point possesses a neighborhood which goes entirely into the circular line. Thus, a ring is a simple analog to the real world, where we all live on a surface approximating a sphere.

For simplification, we use age x as the only additional coordinate. Thus, the agents are distributed within a space that can be seen as the surface of a cylinder (cf. Figure 2). The social network contains all agents located within a two-dimensional neighborhood on this surface. This two-dimensional neighborhood is always symmetric with respect to the agents spatial location. However, there is the possibility of asymmetric intervals with respect to age. This feature reflects the fact that some individuals are more accustomed to deal with younger people others are more accustomed to deal with older people. The maximum amplitude of this heterogeneity with respect to the individuals habits is determined by a numerical parameter γ (see section 4).

The share of married people among this network, rop, determines the social pressure which is given by the function (cf. Figure 3)

$$sp = \frac{\exp(\beta(rop - \alpha))}{1 + \exp(\beta(rop - \alpha))},\tag{1}$$

where α and β determine the inflection point and slope of the function. This level of social pressure, in turn, determines the length of another two-dimensional interval (acceptable range) to look for a potential partner. Thus, an increase of social pressure increases the range to search for an acceptable partner and, consequently, increases the chances to find an acceptable partner and get married. While the set of relevant others includes individuals of both sexes, only agents of opposite sex may be married.[1] If there is any unmarried agent B of opposite sex within the acceptable range of agent A, a) the two

---

[1]We neglect homosexual relationships since the mechansims of social influence and union formation for those agents might be more complex and it is questionable whether the inclusion of this group has any impact on marriage between men and women.
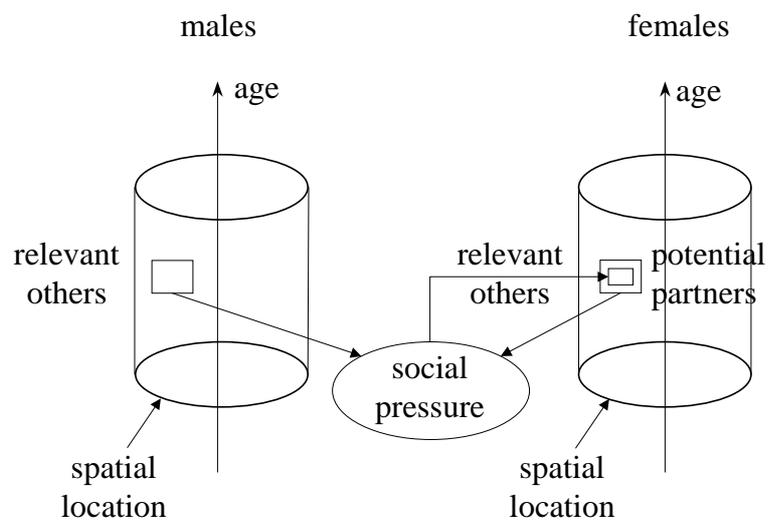
Figure 2: Implementation of the agent based model.

agents get married — in case of one-sided search or alternatively b) it is checked whether agent A is also within the acceptable range of agent B — in which case it would be a two-sided or mutual search.
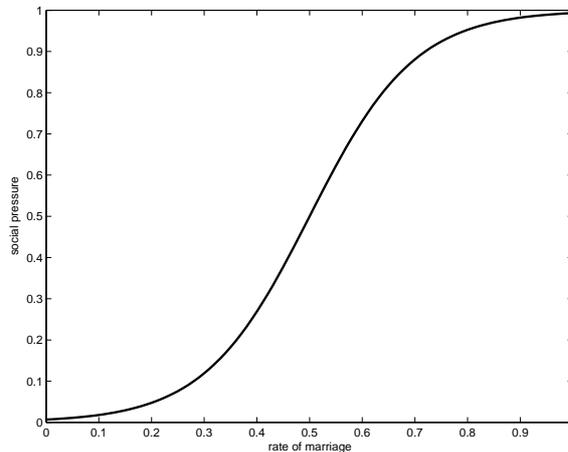

Figure 3: Functional form of social pressure.

Whenever two agents A and B get married this has an impact on the social pressure of those agents who consider A or B as parts of their relevant others. Afterwards, two married agents may get children. These children are randomly located somewhere within the neighborhood of their parents, of course starting at age zero. Therefore, the social network of the parents may be partially inherited.

In this model, social pressure is the only cause why agents start looking for a partner. However, in reality it may happen that people who have already found a girlfriend/boyfriend in the past decide to get married because of social pressure. In order to investigate this effect as well we would need to incorporate unmarried relationships into our model. However, since data about unmarried relationships are not that comprehensive it may be a feasible simplification to neglect the transition from a casual relationship to marriage.

4 Implementation

We implement the wedding ring model using the software package NetLogo, which is a programmable modelling environment for building and exploring

multilevel systems. NetLogo was developed at the "Center for Connected Learning and Computer Based Modeling" at NorthWestern University. The NetLogo environment enables researchers to give rules to individual agents in a simulation and observe the collective result of all the agents' behavior (Wilensky [14]).

Each agent possesses the following characteristics: a numeric identifier i, the year of birth b, age x, sex s, spatial location φ, length of the symmetric interval in which an agent searches for a potential partner d, social pressure (as it depends on the relevant others) sp (cf. equation (1)), marital status m, identifier of the partner if married (missing value for not yet married individuals and for individuals of initial population) j, duration of the marriage md (0 for not yet married individuals), relevant others rop and potential partners pop that include all agents of opposite sex within the search interval. Note that except the numeric identifier, the year of birth and sex, all other characteristics are time varying.

We initialize the simulation with a starting population of N individuals with an age distribution approximating the population of the United States in 1995. We choose sex and marital status randomly assuming a sex ratio at birth of 1.048 and the age and sex specific marital status of the U.S. population in 1995. The duration of marriage is 0 for unmarried agents and equal to 1 at the year of marriage. Thereafter it increases every calendar year by one year. For couples of the initial population the duration of marriage is randomly chosen within the interval $[0, x - 16]$.[2]

To find the relevant others of an agent we consider five different kinds of agents: (a) those who are influenced by younger and older agents similarly, those who are only (b) or mostly (d) influenced by younger agents and those of the same age and those that are only (c) or mostly (e) affected by older agents and those of the same age (cf. Figure 4). To implement the choice of relevant others we first choose the type of the agent, drawing a random number among the discrete distribution $(1, 2, 3, 4, 5)$ which denotes the five possible shapes of age intervals illustrated in Figure 4. Next we randomly choose a parameter $\gamma \in [0, \bar{\gamma}]$ which determines the midpoint of each age interval. In case of an agent of type (a) we do not need to choose γ since the interval will be located symmetrically around the age of the agent. The width of the interval is determined by choosing another random variable

---

[2]Since in the initial population we miss the partners we need to assume an exogenous setting of the marriage duration.

$a \in [0, \bar{a}]$. The interval for the spatial dimension is symmetric around the spatial location φ of the agent and we assume that it depends on the number of initial people in order to avoid a dependence of the number of relevant others on the size of the total population. Among this set of agents (within the chosen age interval and the space interval) the agents choose a random number of agents to be his relevant others. Once we have defined the interval for relevant others, the share of married agents in this interval will determine the social pressure sp as given in equation (1).

In a final step we need to determine the space that includes potential partners. Essentially it is given by transforming the value of the social pressure into a distance $d = sp * m$. The factor m depends on the number of initial people and avoids that the probability of finding a partner is influenced by the population size. The range for potential partners is then equal to $[\varphi - d, \varphi + d]$ along the spatial dimension. The parameter m, which is given as $180 * 500/N$, assures that an agent in a small population (e.g. a starting population of 500 agents) searches the whole ring for a partner as soon as his social pressure becomes 1 (or even earlier in an even smaller population) whereas an agent in a big world never uses the whole ring for his partner search. The range for potential partners along the age dimension is equal to $[x - sp * 50, x + sp * 50]$.

During each simulation year, the agent ages by one year, and dies off if he has reached age 100. Agents who reach the marriageable age of 16 search for relevant others. The share of married couples among this network determines the social pressure function which then determines the region in which unmarried adults look for potential partners. In case an agent finds a potential partner, it is checked if the agent himself is among the set of potential partners of his partner. If the latter condition holds, the two agents get married.

Married women can give birth to new agents with a probability of $rom(a) * asfr(a) * F$ where $rom(a)$ is the proportion of married women at the age a and $asfr(a)$ is the U.S. age specific fertility rate taken from the US Bureau of Census. The parameter F is adjusted each simulation step in order to keep the population stationary. New born agents are randomly located within an interval which is twice as large as the mothers interval of relevant others. Note, that except year of birth and sex which is randomly chosen assuming a sex ratio at birth of 1,048, all other characteristics are missing during the childhood. At age 16 these characteristics are initialized similar to the initial population.
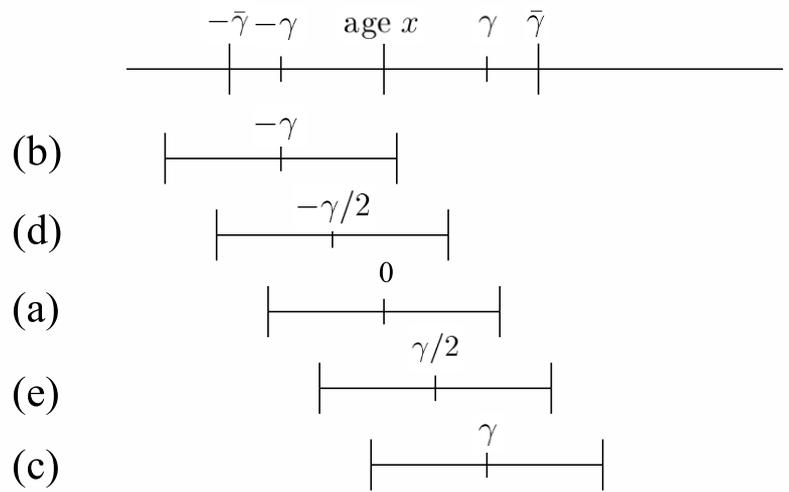
Figure 4: Determination of relevant others.

5   Simulation Results

In this section we discuss the results we obtained by running simulations with a population size of $N = 800$. We are mostly interested in the hazard function for marriage in order to compare the results obtained from the simulation with empirical data. Since the population contains only 800 agents it is obvious that the hazard functions of the artificial population exhibit rather erratic patterns. In order to smooth the curves we collect the data of 75 consecutive cohorts to compute the hazard function and take the average of ten simulation runs.

We set the length of the age-interval to look for relevant others equal to two years ($\bar{a} = 2$). For the function determining the social pressure we use $\alpha = 0.7, \beta = 5$ as the benchmark. For the heterogeneity of the agents with respect to the age interval that determines their network of relevant others we choose $\bar{\gamma} = 2$.

Applying the benchmark setting of parameters our simulations yield the hazard function as given in Figure 5.a (case A) and Figure 5.b (case A) for women and men respectively. The fact that age at first marriage peaks at young ages is in contrast to the typical right-skewed bell distribution observed in empirical data (cf. Figure 1) . As Todd et al. (2005) have shown one may obtain a closer fit to the empirically population wide observed hazard function if we introduce heterogeneity in the underlying mate search process at lower ages. We implement a similar mechanism and add a lower acceptable age at marriage for each agent that is chosen randomly from a normal distribution. Adding heterogeneity in the marriage process at lower ages results in the empirically observed right-skewed bell distribution of ages at marriage (case B, in Figure 5a and Figure 5b). To demonstrate that social influence is an important mechanism (in addition to the heterogeneity introduced at the lower acceptable ages of marriage) we run two further simulations where we kept social pressure constant for the whole age interval (case C) and alternatively where we assume that social pressure decreases with age and neglect heterogeneity in the initial distribution of marriage ages (case D). (Figure 6 plots the alternative functional forms of the social pressure that are applied in the simulations.) Ignoring social influence results in a population wide marriage hazard function that exhibits unrealistic peaks at higher ages while assuming that social pressure decreases with age results in an unrealistic fast decay of the marriage hazard with age. The fact that the hazard rates are lower for men compared to women can be explained

by the fact that the sex ratio at birth implies more men than women and therefore more men stay unmarried in our simulations.
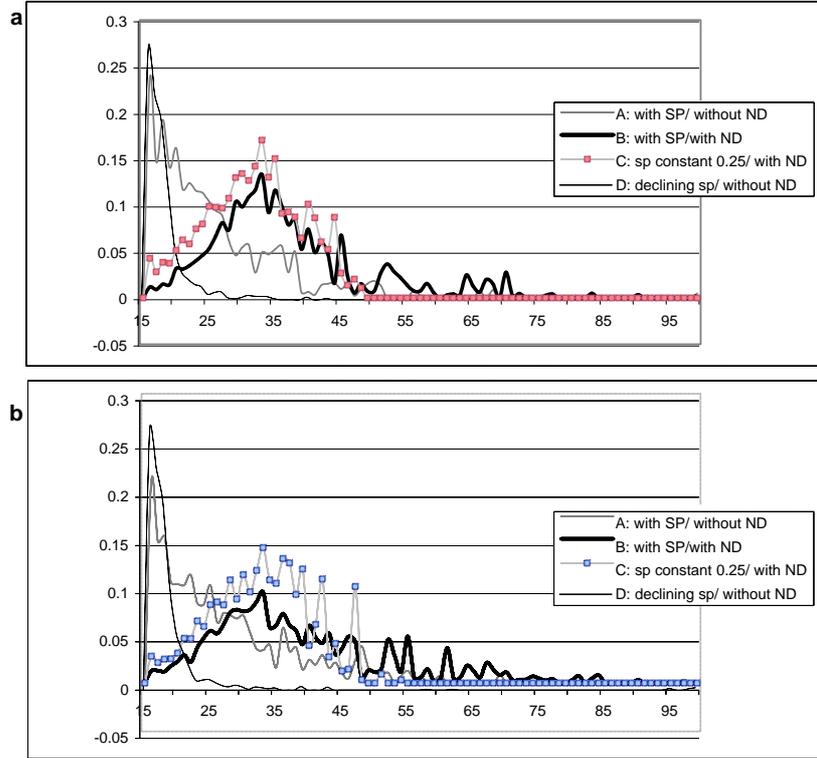

Figure 5: Hazard functions for marriage in a population of simulated agents with alternative settings for social pressure and heterogeneity in the initial distribution of age at marriage. a. Women. b. Men

We next test the sensitivity of the aggregate hazard function of age at marriage when we vary the asymmetry (as represented by the parameter $\bar{\gamma}$) in the age interval that determines an agent's network of relevant others (Figure 7a and Figure 7b). We choose the simulation setting case B in Figure 5 as our benchmark, i.e. we allow for heterogeneity in the initial distribution of ages at marriage and assume a social pressure function that increases with age. In comparison to the benchmark settings we alternatively choose $\bar{\gamma} = 0$ and $\bar{\gamma} = 5$, i.e. we exclude any asymmetry and alternatively
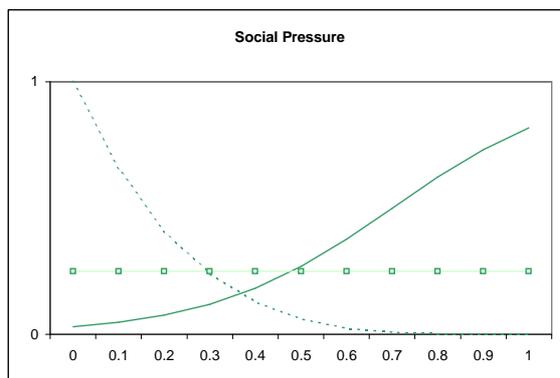
13

Figure 6: Functional form of social pressure applied in Figure 5a and Figure 5b.

increase the asymmetry. Our results indicate that the qualitative shape of the hazard function for marriage is preserved for different parameter settings of $\bar{\gamma}$.

To test the sensitivity of our results with respect to changes in the functional form of the social pressure function we have run simulations where we change the inflection point and alternatively the slope of the social pressure function. The alternative set of experiments is summarized in Figure 8a and Figure 8b for women and men. The alternative forms of the social pressure function are plotted in Figure 9. As the benchmark settings of our simulations we choose again case B in Figure 5 (i.e. $\beta = 5, \alpha = 0.7$). Our results indicate that if the inflection point of the social pressure function shifts to the left, i.e. social pressure increases at each age, the resulting marriage hazard increases mostly at younger ages. A change in the slope of the social pressure function keeping the inflection point the same implies a reduction in the marriage hazard at lower ages while the likelihood of marriage increases at higher ages.

In summary, the simulation results so far indicate that among those sets of numerical parameters the modifications with respect to social pressure have a stronger impact than those with respect to the asymmetry in the age interval that determines the agents network. Moreover, the steeper increase of the social pressure increases the fluctuations of the curves but also the probabilities of getting married at higher ages. Consequently the probability of getting married at younger ages is lower.
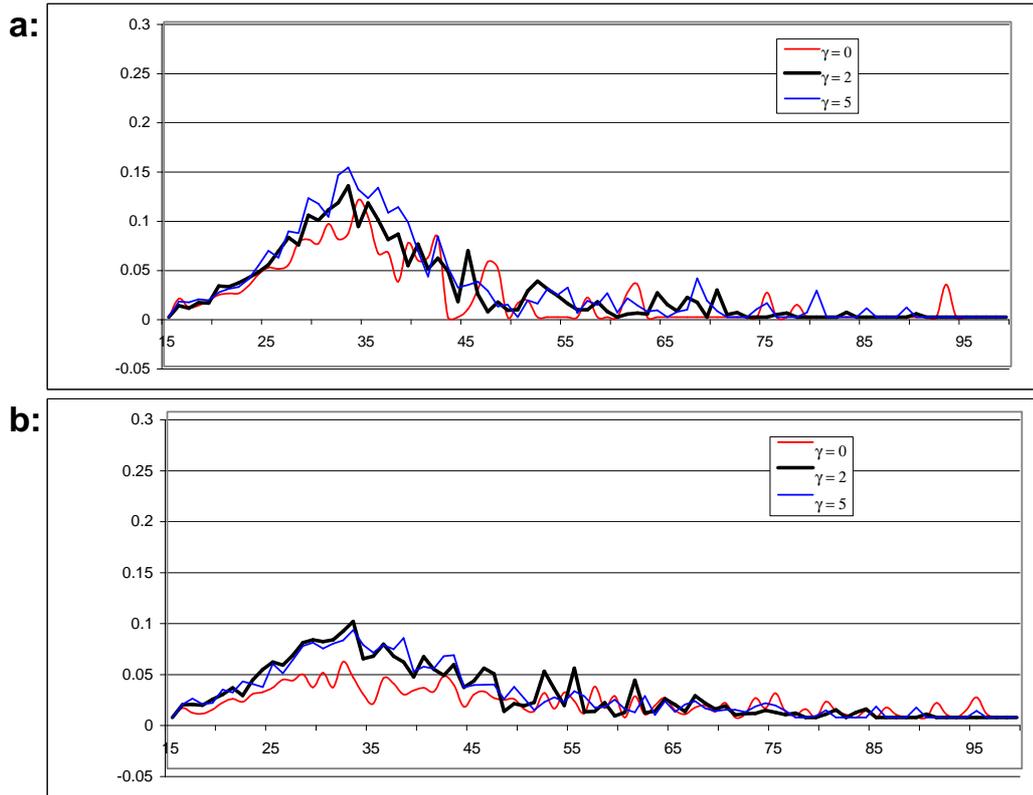
14

Figure 7: Hazard functions for marriage in a population of simulated agents with alternative settings for the asymmetry in the age interval that determines an agent's network. a. Women. b. Men
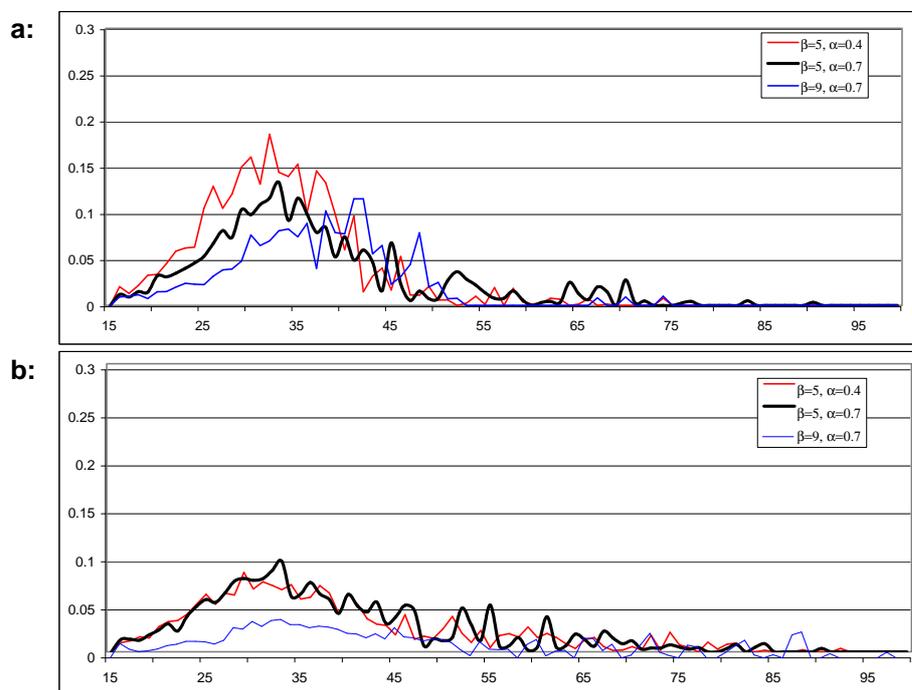
Figure 8: Hazard functions for marriage in a population of simulated agents with alternative settings for social pressure a. Women. b. Men
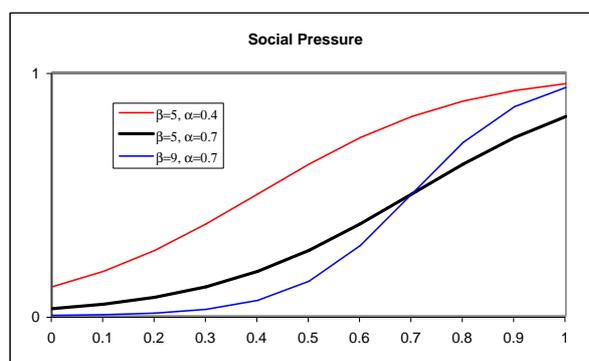

Figure 9: Functional form of social pressure applied in Figure 8a and Figure 8b.

16

References

[1] Behrman, J. R.; Kohler, H.-P.; Watkins, S. C. (2002). Social networks and changes in contraceptive use over time: evidence from a longitudinal study in rural Kenya. Demography 39 (4), 713-738

[2] Bernardi, L. (2003). Channels of social influence on reproduction, Population Research and Policy Review, forthcoming.

[3] Bongaarts, J., Watkins, S.C. (1996) Social Interactions and Contemporary Fertility Transitions, Population and Development Review 22(3), 639-682.

[4] Clark, A.E., Oswald, A.J. (2002), Well-Being in Panels. CNRS and Delta, Paris, France

[5] Hernes, G. (1972). The process of entry into first marriage. American Sociological Review 37, pp.47-82.

[6] Jager, W., Janssen, M.A. (2003). Diffusion Processes in Demographic Transitions: A Prospect on Using Multi Agent Simulation to Explore the Role of Cognitive Strategies and Social Interactions. In: Billari, F.C., Prskawetz, A., eds. Agent-Based Computational Demography. Physica-Verlag, pp. 54-72.

[7] Kohler, H.P., Behrman, J.R., Skytthe, A. (2004). Partner + Children = Happiness?. University of Pennsylvania, Philadelphia, PA.

[8] Micheli, G.A. (2000). Kinship, Family and Social Network: The anthropological embedment of fertility change in Southern Europe. Demographic Research 3, 13.

[9] Montgomery, MR., Casterline, J. (1996). Social Influence, Social Learning, and New Models of Fertility. In: Casterline, J., Lee, R., Foote, K., eds. Fertility in the United States: New Patterns, New Theories. Supplement to Population and Development Review, 22(1): 151-175.

[10] Nazio, T., Blossfeld, H.-P. (2003). The Diffusion of Cohabitation among Young Women in West Germany, East Germany and Italy. European Journal of Population 19, pp. 47-82.

[11] Simão, J., Todd, P.M. (2003) Emergent patterns of mate choice in human populations, Artifical Life 9, 403-417.

[12] Todd, P.M., Billari, F.C. (2003). Population-Wide Marriage Patterns Produced by Individual Mate-Search Heuristics. In: Billari, F., Prskawetz, A., eds. Agent-Based Computational Demography. Physica-Verlag, pp. 117-137.

[13] Todd, P.M., Billari, F.C., Simão, J. (2005). Aggregate age-at-marriage patterns from indvidual mate search heuristics. forthcoming in Demography.

[14] Wilensky, U. (1999). NetLogo. `http://ccl.northwestern.edu/netlogo`. `Center for Connected Learning and Computer- Based Modeling. Northwestern University, Evanston, IL.`