Cohort Size and the Marriage Market: Explaining Nearly a Century of Changes in U.S. Marriage Rates

Mary Ann Bronson and Maurizio Mazzocco

Cohort Size and the Marriage Market:
Explaining Nearly a Century of Changes in U.S. Marriage Rates[*][†]

Mary Ann Bronson[‡]   Maurizio Mazzocco[§]

Current Draft, December 2017

Abstract

Marriage rates have important implications for many economic outcomes. Yet only a limited share of the variation in marriage formation over time and across geographies has been explained thus far. In this paper, we provide an explanation for changes in U.S. marriage rates that holds empirically over nearly a century. We document that a single variable, changes in cohort size, explains around 50% of the variation in marriage rates since the 1930s. Whenever the size of cohorts increases for consecutive years, the share of individuals marrying from those cohorts decreases. Whenever cohort size decreases, marriage rates of those cohorts in turn increase. This relationship holds systematically both over time and across states. Using plausibly exogenous variation across states in access to oral contraceptives, and consequently the number of births, we provide evidence that the relationship between changes in cohort size and changes in marriage rates is causal.

1 Introduction

Why do marriage rates vary over time? For both economists and policy-makers, this is a question of significant interest, as a large body of evidence suggests that marriage rates have important implications for many economic variables. Such variables include fertility rates, children's welfare, children's education, labor force participation, hours of work, income

[‡]Georgetown University, Dept. of Economics, Washington, D.C. Email: mary.ann.bronson@gmail.com.

[§]UCLA, Department of Economics, Bunche Hall, Los Angeles, CA. Email: mmazzocc@econ.ucla.edu.

inequality, the fraction of individuals on government aid, population growth, and workers' productivity.[1]

Given the importance of marriage for many economic outcomes, it is noteworthy that only a limited share of the variation in marriage formation over time and across geographies has been explained thus far. For example, expanded access to the "the Pill", a leading explanation for the dramatic drop in marriage rates observed in the 1970s in the U.S. (Goldin and Katz (2002)), accounts for only about 15% of the changes in marriage rates in the period in which it had the greatest impact. An alternative explanation that is often considered a central factor in rationalizing falling marriage rates is women's rising employment. Such an explanation appears relevant in the seventies and eighties but, from the 1940s to the 1960s, substantial increases in women's labor force participation coincided with sharp increases in marriage rates in the U.S., contradicting the theory. More generally, as we document in the next section, existing theories can only explain changes in marriage rates for specific periods or specific groups of individuals. Thus, while those explanations contribute to our understanding of the evolution of marriage rates, a large share of their variation remains unexplained.

The main contribution of this paper is to provide evidence that one variable, changes in cohort size, explains about half of the variation in U.S. marriage rates since the early twentieth century, both over time and across states. Our findings indicate that the relationship between changes in cohort size – the number of individuals from a given birth cohort – and changes in the marriage rates of the corresponding birth cohorts is strong and systematic, for both women and men. On average, a 10% increase in cohort size is associated with a 0.5 to 1 percent decrease in the share of men and women ever married by 40, with even larger effects on the share ever married at younger ages.

Two observations motivate our analysis. First, standard matching and search models of marriage decisions imply that population changes that affect the ratio of men to women in the marriage market should have substantial effects on marriage rates.[2] Second, changes in cohort size, generated for instance by baby booms or busts, constitute one of the most important sources of variation in the supply of men relative to women in the marriage market.[3] Because most marriages form either between women and men of similar ages or

---

[1]See for example Killingsworth and Heckman (1986), Moffitt (2000), Angrist and Lavy (1996), Gruber (2004), McLanahan and Percheski (2008).

[2]See, for example, Becker's seminal papers (Becker (1973), Becker (1974)), and the thorough review of this literature in Browning, Chiappori, and Weiss (2014).

[3]This has been noted by a number of demographers in the 1960s and 1970s (e.g., Akers (1967), Schoen (1983)), as well as by Bergstrom and Lam (1989).

between women and older men, increases over several years in cohort size make older men a scarce resource, leading to a rise in the ratio of women relative to men in the marriage market. Persistent declines in cohort size have the opposite effect. Changes in cohort size should therefore have significant effects on marriage rates. Despite this clear relationship, our paper is one of the first attempts to formally test and quantify the importance of changes in cohort size for explaining the overall variation in marriage rates observed in the U.S. in the last century.

We provide empirical evidence on the existence of a negative relationship between our two main variables for both women and men in three steps. First, using time-series variation, we document that short-run, year-on-year variation in cohort-size does not significantly affect marriage rates in the U.S.. However, cumulative changes even over three to four years generate significant changes in marriage rates, with increasing effects when the changes accumulate over longer periods. In the second step, we confirm the existence of the negative relationship between our two main variables using cross-sectional, state-level variation in changes in cohort size at marriageable age. Instrumenting for changes in cohort size at marriageable age using cohort size at birth, we find that states with greater growth in cohort size at birth experience larger drops in the share ever married in those cohorts, twenty to thirty years later, with effects that are of similar magnitude as those documented using time-series variation.

In the last step, we provide evidence that the relationship between changes in cohort size and changes in marriage rates is causal. Using an idea based on Bailey (2010), we employ the interaction between the 1957 introduction of *Enovid*, later known as the birth-control pill, and cross-state variation in anti-obscenity laws, which limited the use of contraception in some states until the mid-1960s, to generate exogenous variation in the number of births and therefore cohort size in those years. Our results indicate that states with limits on contraceptives experienced increases in cohort size relative to states without such limits, and that these changes generated a decline in marriage rates for women as well as for men, twenty to thirty years later. The exogenous variation, therefore, gives results that are consistent with the estimates obtained using time-series and cross-state variation.

Our findings have three main implications. First, our results indicate that changes in cohort size explain a large share of the variation in marriages, which in turn has first-order effects on important economic variables, such as aggregate labor supply, savings, and fertility (Doepke and Tertilt (2016)). Studies of marriage market dynamics should

therefore take such changes into account, especially those focused on explaining aggregate trends in marriage rates. Of course, the variable cohort size is typically not exogenous. It is affected by variables such as technological progress, child care availability, supply of housing, and changes in labor supply decisions of women. This fact, however, does not diminish the importance of our results for policy makers and social scientists who study marriage decisions.

Second, our finding that an increase in cohort size reduces the marriage rate of both women and men contradicts the predictions of the standard matching model of the marriage market introduced by Becker (1973) in his seminal paper and the predictions of standard search models. As we argue above, an increase in cohort size increases the number of women relative to men and, hence, the sex ratio. Becker's matching model predicts that an increase in the relative number of women should reduce the marriage rate of women and increase the marriage rate of men. Bronson and Mazzocco (2017a) show that a standard search model has the same predictions. The results presented here suggest that the marriage market operates in a way that is more complex than is allowed by the two most popular models of the marriage market.

The last implication is about policy. Since the 1990s, politicians and policy makers have frequently considered and at times implemented policies to increase the fraction of married individuals with the intent of reducing the poverty rate. Several examples of such policies exist, such as "First Things First" in Tennessee, "Healthy Marriages Grand Rapids," and the federal Temporary Assistance to Needy Families (TANF) program, which allows states to use a fraction of its funds to implement policies aimed at increasing the share of married individuals. Our findings indicate that these types of policies are likely to have only short-term effects, since in the medium and long run marriage rates are in large part driven by demographic changes that are difficult to control.

The paper proceeds as follows. In Section 2, we discuss related papers. In section 3, we describe the data sets used to derive the empirical results. Section 4 documents our findings. Section 5 concludes.

2   Existing Explanations

In this section, we discuss the existing literature on marriage formation. We begin with a general discussion of the explanations for which some empirical evidence has been provided. We then focus specifically on the literature that has looked at the effects of population

growth and sex composition on marriage behaviors.

2.1   General Explanations

Many factors contribute to explaining changes in U.S. marriage rates over time. The importance of income, technologies, welfare policy, men's marriageability, and women's labor force participation have all been emphasized in the literature. However, as we discuss in this section, the evidence about many theories is mixed or indicates that they have explanatory power limited to specific time periods or groups.

One prominent set of explanations for historical changes in marriage rates focuses on changes in income. Cherlin (1981) among others notes the correspondence between rising incomes after World War II and the associated marriage boom during this period. A related common insight is that low income during the Great Depression was the main factor behind the reduction in marriage rates during this period. A positive relationship between income and marriage rates, however, has not been successfully tested over different periods. Hill (2011) rejects the hypothesis that there is a positive correlation between income and marriage rates after 1960. Wolfers (2010) looks at the relationship between marriage rates and recessions for the past 150 years and rejects any pattern between marriage and periods of economic decline. These results suggest that income may generate part of the variation in marriage rates. But they also suggest that variation in income cannot be the general explanation that underlies the changes in marriage rates observed in the past century. In addition, there is a potential reverse causality problem with this theory that is not addressed: men have higher labor hours and earn more following marriage.[4] It is therefore difficult to determine whether an increase in income causes marriage rates to rise or whether an increase in marriage rates generates higher income levels.

Another important set of explanations focuses on technological changes that affect the benefits of being married relative to being single. Such technological changes include the invention and widescale legalization of fertility technologies, including oral contraceptives, or "the Pill", starting in the 1960s (Goldin and Katz (2002)), and abortion in the early 1970s (Akerlof, Yellen, and Katz (1996)). Additionally, labor-saving technological progress in the household sector has made it easier for singles to maintain their own home, increasing the value of being single and reducing the value of marriage (Greenwood and Guner (2009)). All three of these theories about technological change address in particular the period of

---

[4] A number of studies have documented this relationship. See Mazzocco, Ruiz, and Yamaguchi (2014).

very rapidly falling marriage rates starting around 1970, the time when the aforementioned fertility technologies became legal and widely available. Of course, labor-saving household technologies had already been improving for decades before 1970. However, these improvements coincided with significant increases rather than decreases in marriage rates from the 1940s to the early 1960s, counter to the predictions of the theory. Of the papers mentioned above, only Goldin and Katz (2002) use micro-data to estimate the effect of early access to the pill on marriage rates. Their estimates indicate that this explanation accounts for around 15% of the changes in marriage rates in the seventies.

A different set of well-known theories considers factors that may affect women's desire to marry, including men's marriageability and policies such as welfare aid. These theories have been widely empirically tested, with mixed results. Ellwood and Crane (1990) review papers that have tested the hypothesis proposed by Wilson (1987) that limited labor market opportunities reduce the number of marriageable men, and thus the marriage rate. While some papers provide evidence in support of this theory, others reject this hypothesis (e.g., Plotnick (1990) and Lerman (1989)). The degree to which men's employment opportunities affect marriage rates is largely still an open question. A related issue affecting especially black men's marriageability in the last three decades is the rise in incarceration. Charles and Luoh (2010) study the relationship between incarceration and marriage rates and find that higher incarceration rates decrease the fraction of married individuals both in black and white populations. Quantitatively, however, such an explanation is primarily relevant for black populations and only for the period after 1980, when drug-related policies significantly increased incarceration rates. Ellwood and Crane (1990) have also evaluated the link between welfare aid and marriage. They conclude that there is very little empirical support for the proposition that welfare benefits played a major role in marriage trends for black or white women.

Finally, rising income and employment opportunities for women may also affect their desire to marry. However, we do not know of papers that formally test this hypothesis over time or across geographies. One reason may be that reverse causality potentially complicates such empirical analysis: women who face poorer marriage prospects may invest more in human capital and work more.

2.2 Explanations Related to Cohort Size and Sex Composition

Next, we discuss explanations related most closely to ours, which consider the effect of changes in cohort size and changes in sex composition on marriage rates.

The idea that changes in population growth introduce variation in the sex composition of the marriage market dates back to the late 1960s. As baby boom cohorts entered marriageable age, demographers argued that women in these cohorts, and in any rapidly growing population, would experience a "marriage squeeze." The classic studies on this topic are by Akers (1967) and Schoen (1983). They analyze the effect of cohort size on marriage rates using demographic models of the marriage market and data simulated from their models. In these models, changes in cohort size modify the supply of women relative to men in the marriage market, which in turn changes the marriage rate. The results of these papers suggest that an increase in cohort size should reduce the marriage rate of women and increase the marriage rate of men. Since those findings are the outcome of the developed models and of the corresponding assumptions, they should be interpreted as theoretical insights on the possible effect of changes in cohort size on marriage decisions. Our paper tests this idea formally using nearly a century of longitudinal data as well as cross-state variation. Our findings indicate that those models are rejected by the U.S. data since increases in cohort size reduce the marriage rates for both men and women. Our results also indicate that the effects of changes in cohort size on the eventual share ever married are substantially larger than those predicted by those models.

A second well-known paper studying the effects of cohort size on sex ratios and marriage behaviors is by Bergstrom and Lam (1989). The authors observe that predictions about the effects of changes in cohort size on sex ratios may overstate the effects on marriage rates, since age differences between spouses can adjust. Even if individuals prefer older or younger spouses, they can always marry within their own cohort, which generally is composed of equal numbers of men and women. To illustrate this point, they calibrate a marriage matching model following Becker (1973) in which marriage rates across cohorts are held constant, but age differences between spouses are allowed to vary. They provide evidence using Swedish data that average age differences between spouses in Sweden do in fact adjust in response to changes in cohort size, in line with the predictions of their model. However, their model assumes that marriage rates never change in response to increases or decreases in cohort size. As we document, this assumption is empirically rejected. Our findings using U.S. data in this paper and in Bronson and Mazzocco (2017a), where we also

examine age differences between spouses, indicate that age differences adjust in response to changes in cohort size, as one might expect, but marriage rates are also strongly affected.

An extended literature related to the previous two papers makes use of changes in the sex ratio, computed as the number of men divided by the number of women in the marriage market, to understand changes in marriage rates. Almost all papers in this literature rely either explicitly or implicitly on a two-sided matching model of the marriage market, as formalized by Becker (1973). A testable implication of this model is that a reduction in sex ratio should reduce women's and increase men's marriage rates. Several papers have attempted to test the relationship between sex ratios and marriage rates. Angrist (2002), Seitz (2009), Abramitzky, Delavande, and Vasconcelos (2011), and Knowles and Vandenbroucke (2016) are examples of papers in this literature.[5]

In our paper, we do not explicitly analyze the link between sex ratios and marriage rates. However, under the standard and empirically based assumption that women primarily marry men of similar or older ages, an increase in cohort size will reduce the sex ratio, or the number of men relative to women in the marriage market. Therefore, an increase in cohort size should reduce the marriage rate of women, but increase the marriage rate of men. An important finding of our paper is that this prediction of the standard matching model is rejected if one considers U.S. data on cohort size and marriage rates over the past century.

This finding does not imply that a change in sex ratio can never generate a marriage market outcome that is consistent with the matching model. For instance, Abramitzky, Delavande, and Vasconcelos (2011) consider variation in the sex ratio due to World War I casualties in France and find that a smaller sex ratio is associated with a lower marriage rate for women and a larger marriage rate for men, in line with Becker's prediction. Our findings simply indicate that, if the change in sex ratio is generated by a change in cohort size, the standard matching model is not consistent with U.S. marriage rate data from the past 100 years.

There is one important paper in the sex ratio literature whose results are consistent with ours. Angrist (2002) uses variation in immigration rates from different European countries to the U.S. at the beginning of the twentieth century to study the relationship between sex ratios and marriage rates. He exploits the fact that the majority of migrants were men and that marriages were often formed between individuals belonging to the same ethnicity.

---

[5]Several papers have studied the relationship between sex ratios and other economic variables, including rates of single motherhood (Neal (2004)), labor force participation of women (Grossbard-Shechtman (1984)), and birth rates (Bitler and Schmidt (2011)).

Consistent with the patterns documented here, he finds that ethnicities with lower sex ratios experienced lower marriage rates for women as well as men. In our paper, we document that the positive relationship between sex ratio and marriage rates of women and men applies to a century of data and not only to the specific period considered in Angrist (2002), if the changes in sex ratios are produced by variation in cohort size. Moreover, in the companion paper Bronson and Mazzocco (2017a), we provide an explanation for the apparent inconsistency between the findings in Angrist (2002) and the findings in Abramitzky, Delavande, and Vasconcelos (2011).

One alternative explanation for the link between cohort size and marriage rates that has commonality with the one we propose is the Easterlin hypothesis. Easterlin (1987) argues that the relative size of a cohort can explain many variables that determine the economic and social outcomes of that cohort: earnings and unemployment rates, college enrollment rates, divorce, fertility, crime, suicide rates, and marriage. Easterlin's theory is composed of two parts. First, the distribution of income of a cohort is affected by its size, with larger cohorts having worse economic outcomes. Second, when income of a cohort is above its aspiration level, the individuals in that cohort will be optimistic and therefore will have better economic and social outcomes. Researchers who have attempted to test the general idea behind Easterlin's hypothesis have found mixed results (Pampel and Peters (1995)).

3   Data

In this section, we provide a brief description of the main variables used in the empirical analysis and the data employed in their construction. A more detailed discussion is given in the appendix.

Throughout the paper we rely on the following data sets: National Vital Statistics (1909-1980), Census total population counts (1910-1990), Survey of Epidemiology and End Results (SEER) population estimates (1969-2010), IPUMS CPS (1962-2015), and IPUMS USA (1940-2010). Using these data sets, we analyze changes over time and across states in two variables: marriage rates and cohort size.

Our preferred measure of marriage rates is the share ever married by a given age for a given cohort. While this measure has been used in some studies, e.g. Goldin and Katz (2002), most existing studies of changes in marriage rates typically employ one of the following two alternative variables: the number of new marriages per population; and the share of individuals currently or ever married within some age range, e.g. between the ages of 18

and 30. In Bronson and Mazzocco (2017b), we provide evidence that these two measures are not well suited to study changes in marriage rates for the following reasons. The first measure confounds changes in new marriages with changes in overall population structure. If the population undergoes any substantive growth or decline in the number of individuals of marriageable age – due, for instance, to changes in fertility or migration patterns – one will generally draw the wrong inference about changes in marriage rates. The second measure is strongly affected by changes in the age at first marriage: when examining the share of people ever married within an age range, one cannot determine whether individuals are simply delaying marriage, for example to complete their education, or whether they choose not to marry at all. In the same paper, we also provide evidence that a measure that is not affected by those limitations is the share of individuals ever married by a given age in a given cohort, as long as the age-cutoff is sufficiently high. For this reason, we use only this cohort-based measure throughout our analysis. The variable share ever married at age 30, 35, and 40 is constructed using either the decennial Censuses or the CPS. The appendix describes the exact procedure used to construct this variable.

In recent decades, cohabitation has become a popular form of household formation and a close substitute for marriage. For this reason, we extend our analysis and also consider cohabiting households by studying the variable share ever married or cohabiting at a given age. We use the variable "Relationship to household head" in the Census and CPS to record households in which a cohabiting partner is present. The Census began recording unmarried partners only in 1990, and the CPS only in 1995. Since in the longitudinal analysis we employ CPS data, we may miss cohabitations for cohorts born before 1965 when we use 30 as the age cutoff, or 1955 when we use 40 as the cutoff. In the cross-sectional analysis, since we employ the Censuses, we may similarly miss cohabitations for cohorts born before 1960 or 1950, depending on the age cutoff. In the data, however, we observe that cohabitation for early cohorts is limited. For instance, the share of individuals cohabiting at age 30 in the 1965 cohort was 2.8% and the share cohabiting at age 40 in the 1955 cohort was 0.76%.

To confirm that we do not miss a substantial number of cohabitations for the cohorts for which we do not have the cohabitation variable, we examined data from the National Survey of Families and Households (NSFH). The first wave of the NSFH (1987-1988) is nationally representative and provides retrospective data on marriage and cohabitation. We use the dataset to examine cohabitation patterns for cohorts born in 1957 or earlier. We found that cohabitation at the ages relevant for our analysis – 30, 35, and 40 – is almost non-existent for

pre-baby boom cohorts. For the 1945 to 1957 birth cohorts, the average share of individuals cohabiting at age 30 is 0.5%, and this share further decreases with age. We conclude that we only marginally underestimate the share ever married or cohabiting for the early baby boom cohorts. A limitation of the cohabitation variable in the Censuses and CPS is that we are only able to examine currently cohabiting rather than ever cohabiting individuals. However, in the NSFH we find that the share ever married or currently cohabiting at age 30 is a very close approximation for the share ever married or ever cohabiting at age 30, and by age 40 the two figures are virtually identical.

For our other main variable of interest, cohort size, we employ two measures: cohort size at birth and cohort size at marriageable age. Since we are interested in the evolution of cohort size at the ages in which individuals choose whether and whom to marry, cohort size at marriageable age (or its change over time) is the relevant independent variable for most of our analysis. Cohort size at birth is used as an instrument for cohort size at marriageable age in our cross-state regressions, and as one of the variables used to determine the effect of the introduction of the pill in states with different anti-obscenity laws. Additionally, in the analysis that uses longitudinal variation, we employ cohort size at birth as a proxy for cohort size at marriageable age. The are two reasons for this choice. First, it is possible to construct the variable cohort size at birth annually for all cohorts born in 1909 and after, whereas the variable cohort size at marriageable age is available only for cohorts born in decennial years prior to 1940. Second, as shown in Figure 1, when cohort size is computed for the U.S. population there is little difference between cohort size at birth and cohort size in adulthood, since net migration from and to the U.S. for any given cohort is limited. By using cohort size at birth in the longitudinal exercise, we can therefore consider a larger number of cohorts without significant effects on the analysis. As described in the appendix, the cohort size variables are constructed using the U.S. Vital Statistics, population counts, from the U.S. Census, and the Survey of Epidemiology and End Results (SEER).

We conclude this section by discussing why we study the effect of changes in cohort size on marriage rates instead of focusing directly on the variable "sex ratio", a seemingly more appropriate measure of changes in sex composition in the marriage market. The key reason for this decision is that, despite the large number of studies on sex ratios, the literature does not agree on a specific definition for this variable. The construction of the sex ratio requires three assumptions. First, the researcher has to assume which age ranges to consider for men and women. This assumption varies considerably across papers. For example, Angrist (2002)

considers the ratio of men ages 20 to 35 to women ages 18 to 33. Abramitzky, Delavande, and Vasconcelos (2011) define the sex ratio using men ages 18 to 59 and women ages 15 to 49. Seitz (2009) considers men ages 17 to 21 and women ages 15 to 19. Second, one has to choose whether to include all individuals in the population or only single individuals, a choice that varies in the literature. Lastly, to compute the sex ratio one has to assume an age difference between spouses to determine the relevant set of men and women. The standard assumption is that women marry men who are two years older. Table 1 indicates, however, that at most 32% of couples satisfy this assumption, whereas about 34% are approximately of the same age and at least 25% of women marry a man who is at least 5 years older. The alternative assumption that there is no age difference would also be misleading, and particularly problematic in our context since, under this assumption, changes in cohort sizes would have no effects on sex ratios, as the size of a cohort is approximately equal for men and women. Choosing a fixed age difference is also questionable because, as discussed previously, age differences between spouses adjust to changes in cohort size.

Our own experimentation with previously used definitions indicates that the outcome of various empirical exercises is highly sensitive to the specific definition used in the construction of the variable sex ratio. We have therefore decided to focus on changes in cohort size. This allows us to study the empirical relationship between demographic changes and marriage rates without imposing restrictions on how changes in cohort size can affect the sex ratio. The relationship between cohort size and sex composition is made explicit in the companion paper Bronson and Mazzocco (2017a), using an overlapping generations model.

4 Empirical Results

In this section, we analyze the relationship between changes in cohort size and changes in marriage rates. The section is divided into five parts. We first provide empirical evidence using longitudinal variation. We then describe findings obtained using cross-state variation. In the third subsection, we discuss endogeneity issues that may affect the longitudinal and cross-state variation. We then provide evidence that changes in cohort size generate changes in marriage rates by using variation in early access to oral contraceptives across states as a plausible source of exogenous variation in cohort size. Lastly, we discuss possible mechanisms that can explain the observed relationship between cohort size and marriage rates.

4.1 Change in Marriage Rates Over Time

In this subsection we first use longitudinal variation to provide evidence on the general nature of the relationship between changes in cohort size and changes in marriage rates. We then try to understand whether cohort size can explain the short-run, medium-run, or long-run changes in marriage rates. As discussed in the data section, we employ variation in cohort size at birth, which provides us with annual variation in cohort size and, as documented in Figure 1, constitutes a close approximation of cohort size at marriageable age at the national level over the time period we consider.

In Figure 2, we plot cohort size and the share never married by age 30, separately for women and men, for all cohorts born between 1914 and 1981. Panels A and B describe these variables for the white and black populations, respectively. We plot the share *never* married because visually it is easier to detect a positive correlation between the two variables. Figure 2 contains one main finding. To describe it, we initially focus on cohorts in Panel A born before 1960. For those cohorts, there is a strong and positive correlation between cohort size and the share never married for both women and men. The decline in size for cohorts born in the 1920s and 1930s is associated with a similar drop in the share never married. This decline corresponds to the well-documented "marriage boom" that starts in the mid-1940s and lasts through the early 1960s, the period in which the cohorts born in the twenties and thirties were active in the marriage market. The sharp increase in the size of cohorts born between 1946 and 1959, which correspond to the post-war baby boom generations, is associated with a share never married that nearly tripled. Births and the share never married for the black population follow similar patterns.

It is left to explain why we lose the positive correlation between our two main variables for cohorts born in the 1960s and 1970s. Note that those cohorts were active in the marriage market starting from the 1980s, when cohabitation started to become a popular form of household formation and potentially a close substitute for marriage. To understand whether cohabitation can resolve the inconsistency between the early and later cohorts, in Figure 3 we plot the variables reported in the previous figure, with the exception that now cohabiting individuals are treated as married individuals instead of being treated as never-married individuals. Interestingly, once cohabiting households are accounted for, the relationship between cohort size and household formation for cohorts born in the 1960s and 1970s resembles that of the earlier cohorts. Falling cohort sizes in the 1960s and early 1970s correspond to a decline in the share never married and not currently cohabiting by 30 for

13

the male and female population. Increasing cohort sizes in the second part of the 1970s are associated with a rise in the share never married and not cohabiting. Finally, Figure 3 documents that the strong positive correlation between cohort size and share never married characterizes both the white and the black populations.

In Figures 2 and 3, we use an age cutoff of 30. The results may therefore be affected by changes in age at first marriage. To address this concern, in Figure 4 we plot cohort size and shares never married and not cohabiting by age 40. Using this new cutoff age, we find patterns that are similar to the ones observed in the first two figures: there is a positive and strong correlation between cohort size and share never married and not cohabiting.

In Table 2, we report the average relationship between marriage rates and cohort size at different age cutoffs in our longitudinal data. For ease of exposition, in the rest of the paper we will consider the effect of cohort size on the share ever married instead of the share never married. Each coefficient in the table is the outcome of a separate regression of the log share ever married or currently cohabiting on log cohort size.[6] The elasticities recorded in Table 2 are highest at 30 and gradually decrease with age, for both sexes and both races. This finding suggests that an increase in cohort size is associated with two effects: (i) a decrease in the eventual share ever married or cohabiting by 40; and (ii) an increase in the age at first marriage. It also indicates that using 40 as an age cut-off better isolates the relationship between variation in cohort size and variation in marriage rates from changes in age at first marriage. Table 2 also documents that the effects are quantitatively large. On average, an increase of 10% in cohort size is associated with a decrease in share ever married or cohabiting by 40 ranging from 0.66% for white women to 4.53% for black women. In percentage points, this amounts to a decline that is between 0.6 and 3.8 points in the share of individuals ever married or cohabiting by 40, a large effect. Lastly, the cohort size variable explains a large fraction of the variation observed in marriage rates. For instance, when we use 40 as the cutoff age, the R-squared is between 0.58 and 0.85.

Table 2 reports an average relationship between cohort size and marriage rates in levels. In the remaining part of the subsection, we consider the effect of differences in cohort sizes on differences in marriage rates. This analysis enables us to evaluate whether changes in cohort size over one or two years are sufficient to generate changes in marriage rates or

---

[6]Note that we are working with non-stationary time series, and must therefore verify that the series are cointegrated to eliminate worries about spurious regressions. A Johansen cointegration test rejects the null hypothesis that the series are not cointegrated at the one-percent level. Therefore, our OLS results are consistent and estimate a non-spurious relationship. We use Newey-West standard errors for all significance tests.

persistent changes over several years are required to produce a change in this variable. In all the analysis that follows, we focus on the white population only, for one main reason. White women and men have similar cohort size at the time of marriage, whereas black men have a significantly lower cohort size than black women because of higher mortality and incarceration rates. As a consequence, the investigation of the marriage market for blacks requires a different type of analysis which we leave for future research.

To analyze whether changes in cohort size have to cumulate for several years to generate a change in marriage rates, we regress n-year differences in marriage rates on n-year differences in cohort size, where n is set equal to 1, 2, 3, 4, 5, 7, and 10. To account for the effect of adjacent cohorts, for $n > 1$ we use differences in cumulative cohort size as our independent variable, where cumulative size for the cohort born in period t for the n-year difference is constructed by adding up cohort size from $t - n + 1$ to t.[7] We interpret the coefficient estimates on the 1-year and 2-year differences as the short-run effect, the coefficients on the 3-year, 4-year, and 5-year differences as the medium-term effect, and the coefficients on the 7-year and 10-year differences as the long-term effect. This choice is somewhat arbitrary, but it helps us focus the discussion.

The results are presented in Table 3. The first two columns report the short-run effects. The estimates for the 1-year differences indicate that a 1-year change in cohort size is not sufficient to trigger a change in marriage rates. All the coefficients are statistically equal to zero, with the exception of the coefficient for women by age 30 and for men by age 40.[8] The effects are only marginally larger when we employ 2-year differences, suggesting that in the short run cohort size has at best a weak effect on marriage rates.

The negative effect of cohort size on marriage rates is much stronger when we study the medium-term effects. Already at four-year differences, all the coefficients are large in size, negative, and statistically significant. The long-term effect of cohort size is even stronger. The coefficient estimates for the 7-year and 10-year differences suggest that, for an age cutoff of 40, an increase in cohort size of 10% generates a drop in marriage rates of 0.6-0.8%. This is a significant decline since it implies that a standard deviation increase in cohort size decreases the share ever married by 40 by a third to a half of a standard

---

[7]We have also estimated the effect of cohort size using simple n-year differences. The estimates display similar patterns, but the coefficients are generally smaller and are less precisely estimated.

[8]The estimated coefficient for men by age 40 is the only one in all of the estimations we have performed that is positive and statistically significant. This result is generated by the two spikes in births that occurred in 1942 just at the start of World War II and in 1946-1947 after World War II ended, which happen to coincide with increases in share married for those two cohorts. If one drops the observations that characterize the period around World War II, the coefficient becomes zero.

deviation. These findings indicate that cohort size can explain the medium and long term variation in marriage rates, but not the short term variation. Changes have to cumulate for longer than one or two years to generate significant fluctuations in marriage rates. It also indicates that the relationship between cohort size and marriages described in Table 2 is driven by persistent growth in cohort size, not by the size of the cohort directly. If the size of the cohort alone mattered, one should observe that short-run changes in cohort size affect marriage rates just as strongly. Instead, some persistence in the growth in cohort size is required. This result is also in line with the idea that persistent changes in cohort size affect marriage rates by modifying the sex composition of the marriage market, as we discuss in Section 2.2 and previous papers have argued (e.g., Akers (1967), Schoen (1983)).

To summarize, our results in this section indicate that there is a strong and negative relationship between marriage rates and cohort size for both men and women in the time-series data. The results also indicate that changes in cohort size account for a large fraction of the medium-run and long-run time series variation in marriage rates. In the following subsections, we further explore the empirical link between changes in cohort size and marriage rates using cross-state variation. Because cohabitation has become a close substitute for marriage in recent decades, in the rest of the paper we will continue to use the same adjusted measure of household formation: the share ever married or cohabiting by a given age. Unless we specifically note otherwise, when we use the shorthand "ever married" we refer to those ever married or currently cohabiting.

4.2 Change in Marriage Rates Across States

In this subsection we provide evidence on the relationship between changes in cohort size and changes in marriage rates using variation across states. If changes in cohort size influence marriage rates, we should observe such an effect not just across time but also across geography. Specifically, states with larger increases in cohort size should experience larger drops in marriage rates.

Changes in cohort size at marriageable age differ across states partly due to cross-state migration flows. Migration decisions are generally related to differences across states in economic and social conditions which also affect marriage rates. Moreover, since migration tends to be sex-biased, it can skew sex-ratios and affect marriage rates directly. Cohort size at marriageable age has therefore the potential to be endogenous. To address this issue, we use total births in a given cohort and state, which are arguably unaffected by the

endogeneity concerns discussed above, as an instrument for the size of the marriage market.

To perform the analysis at the state level, we rely on the decennial Census over the entire period of interest since it is the only dataset with sufficiently large sample sizes for all states. In Section 3, we have provided details about how we construct the three main variables needed for the analysis: cohort size at birth, cohort size at marriageable age, and share ever married for each state and cohort. Using the Census data we can only use the decennial cohorts born between 1910 and 1970, since the share ever married at 30 and 40 can only be computed for them. We therefore perform the empirical analysis by first constructing ten-year differences for the log of share ever married, the log of cohort size at birth, and the log of cohort size at the time of marriage for each state and cohort. We then pool all cross-sections, because of the small number of observations, and regress the differences in log share ever married on the differences in log cohort size where log cohort size is instrumented using log cohort size at birth. We add year fixed effects to the regression to control for general time trends.

Table 4 presents the results of the cross-state regressions separately by gender. The first column presents the estimates obtained using a standard OLS regression. Similarly to the findings obtained using longitudinal variation, the estimated coefficient is negative and statistically significant for the two age cutoffs we consider, for both men and women. As argued above, the OLS estimates may be biased because of migration decisions. In particular, if migration is partially driven by the desire to find a spouse, the estimates will be positively biased. A similar result applies if migration is partially driven by job-related reasons. In this case, on average one would expect individuals to move to states with higher earning opportunities. If individuals with higher income are also more likely to marry, this would generate a positive correlation between cohort size growth and marriage rates, biasing the results toward zero away from the strongly negative relationship we predict.

The first stage results obtained using the IV estimator are reported at the bottom of the Table 4. They suggest that cohort size at birth explains a large fraction of cohort size at marriageable age. But the coefficient on log cohort size at birth, estimated to be 0.44, is well below 1, suggesting that cross-state migration significantly affects changes in cohort size at marriageable age. Our second stage estimates, which are reported in column two, are all negative, statistically significant and, as expected, larger in size than the OLS estimates. For the share ever married by 30, the coefficient is estimated to be -0.104 for women and -0.123 for men. When we use the age cutoff of 40, the coefficient drops to -0.058 for women

and -0.047 for men.

The sizes of the IV coefficients imply large effects that can explain a significant share of the changes in marriage rates. For example, in the two decades corresponding to the 1940 and 1950 cohorts, and the 1950 and 1960 cohorts, cohort size increased by substantial amount, 42% and 17%, respectively. Our estimates imply that this corresponded to a 3.4 and 3.8 percentage point decrease in the share of women and men ever married by 30 from 1940 to 1950, and a 1.5 and 1.4 percentage point decrease from 1950 to 1960. This is equivalent, on average, to about 55% and 39% of the actual changes in share ever married by 30 observed for women and men over this time period.

In the third column, we add time-region fixed effects to the IV specification to make sure that our findings are not driven by systematic differences in trends across regions. The coefficient estimates are similar to the ones presented in column 2. The only difference is that we lose significance for men when we use an age cutoff of 40 because of an increase in the standard errors.

4.3 Potential Endogeneity Concerns

The findings obtained using cross-state regressions strongly corroborate the negative relationship between changes in cohort size and changes in marriage rates observed for both women and men when longitudinal variation is used. However, there are reasons that prevent a causal interpretation of the relationship. While using number of births as an instrumental variable allows us to reasonably avoid reverse causality problems as well as important endogeneity concerns due to migration, one may nevertheless worry about omitted variables: state-level characteristics that drive changes in birth rates for a particular cohort as well as changes in marriage decisions of individuals that belong to that cohort 20 to 30 years later. Such omitted variables would have to be highly persistent shocks that affect growth in births in a given year as well as growth in subsequent marriage rates about 20 to 30 years later.

It is not easy to think of variables that fit this description. Potential examples include highly persistent productivity shocks which cause wages to grow more rapidly in some states over time, affecting both birth rates at the time of the initial shock and marriage rates two or three decades later. A positive trend in men's earnings in some states fits this description. If children are a normal good, states with such positive trends may see increased births in 1950 relative to 1940 as well as a greater number of marriageable men in 1980 relative to

1970. Alternatively, states may differ in their religiosity or preferences for forming a family. In states with weaker preferences for family, one might expect depressed birth rates as well as lower marriage rates in the future.

Note that in these and most credible cases we would typically expect an increase in both births and subsequent marriage rates or a decline in both variables. This bias would work against our favor and would result in a positive coefficient on cohort size, which is not what we find. Nevertheless, without exogenous variation in cohort size we cannot entirely eliminate the possibility that some biases could work in our favor.

4.4 Instrumental Variables Strategy

To address the potential endogeneity issues outlined in the previous subsection, we construct an instrument which is based on an idea first proposed by Bailey (2010). The idea is to use the interaction between the introduction of *Enovid* in 1957, later known as the birth control pill, and cross-state variation in anti-obscenity laws, which limited the use of contraception, to generate exogenous variation in the number of births and therefore cohort size.

In 1873, the U.S. Congress enacted the Comstock Act which had two main objectives. The direct goal was to ban the interstate mailing, shipping, and importation of products and printed materials that were considered to be "obscenities". Since the Act considered anything employed for the prevention of conception an obscenity, it outlawed any interstate transaction involving contraceptives. The indirect goal was to "incite every State Legislature to enact similar laws" as stated by U.S. Representative John Merriman during an interview with the New York Times on March 15, 1873. The Comstock Act was highly successful in achieving this goal. By 1900, 42 states had approved anti-obscenity laws and by 1943 the number of states had increased to 48.

Because the specific language of these statues varied significantly across states, the laws had different effects on the introduction of the pill in different states. As suggested by Bailey (2010), the states can be grouped into four categories depending on the type of law they enacted. The first group consists of the seventeen states that explicitly banned the sale, advertisement, and distribution of information of any product for the prevention of conception. The second group consists of seven states that had the same ban as the first group of states, but added an exception for physicians and pharmacists who were allowed to sell, advertise, and distribute information on products related to birth-control methods. The third category is comprised of six states that explicitly only banned the advertisement and

distribution of information about contraceptives, but did not outlaw their sale. The final group is composed of eighteen states that approved a law banning the sale, advertisement, and dissemination of information of obscene products, without explicitly classifying the prevention of conception as obscene. In our analysis, we refer to states in the first two groups as having a sales ban on contraceptives. We control explicitly for whether or not a state had a physician exception.

An important question is whether some states enacted stricter anti-obscenity laws because they had more conservative constituencies or because of other observable cross-state differences. Bailey (2010) provides evidence that this is not the case. For instance, among the states that adopted sales bans of contraceptives one can find both typically conservative and typically liberal states. California and Washington, two of the states that repealed anti-abortion laws before the Roe v. Wade decision, enacted the strictest version of the bans whereas Alabama, a generally conservative state, adopted a statute that did not explicitly categorize the prevention of conception as obscene.

These anti-obscenity laws lasted until the sixties when they were repealed or struck down by the individual states or by the 1965 U.S. Supreme Court's decision in *Griswold v. Connecticut*. Specifically, two states repealed their anti-obscenity statutes in 1961, one state in 1962, four states in 1963, and Connecticut in 1965 after the U.S. Supreme Court's decision. *Griswold v. Connecticut* expedited the repeal of anti-obscenity statutes in all the remaining states between 1965 and 1971. In the empirical part, we follow Bailey (2010) and use the period between 1957 and 1965, the year of the U.S. Supreme Court decision, as the period in which the introduction of *Enovid* interacted with the anti-obscenity laws to generate what is arguably exogenous variation in total births, and therefore cohort size.

Before presenting our results, it is important to emphasize a difference between this paper and Bailey (2010). Bailey is interested in the relationship between the anti-obscenity laws and birth rates of married women after the pill was introduced. She finds that states with the sales ban experienced a marital birth rate that was 8% higher than the remaining states during the period 1957-1965. In this paper we are interested in the link between the anti-obscenity laws and the following two variables: cohort size at birth and cohort size at marriageable age.

We start by giving some graphical evidence on the effect of the source of variation described above on our variables of interest. We follow Bailey (2010) and present the results separately for the four Census regions. In Figure 5, Panel A, we report the difference

in growth of cohort size at birth between states with the sales ban and the remaining states from 1950 to 1970. Two features generate the hump shape that characterizes all census regions. First, in all regions, after the introduction of the pill, states with the ban experienced larger growth in cohort size at birth. In the South, states with the ban were already increasing relative to states without a ban before 1957, but the process was expedited by the introduction of the pill. Second, in all regions, when states started to outlaw the sales bans on contraceptives, the growth in cohort size at birth in states with the ban started to converge to the growth in states with no ban. The convergence continues until 1965 when the *Griswold v. Connecticut* decision took place, at which point the two groups of states have similar rates of growth in cohort size.

In Figure 5, Panel B, we replace cohort size at birth with cohort size at age 25, which represents a measure of cohort size at marriageable age. The figure displays patterns that are similar to the ones observed in Panel A, with the differences in growth in cohort size following the same familiar hump shape between 1957 and 1965. Thus, the differences in cohort size at birth driven by the Comstock laws and the introduction of the Pill persisted to generate differences in cohort size at marriageable age.

We now formally use the introduction of the pill interacted with the sales bans as an instrument for cohort size at marriageable age. To that end, we construct two dummy variables: the first one, ban_s, is equal to one for all cohorts from state s that adopted a sales ban on contraceptives and zero otherwise; the second dummy variable, $ban * pill_{c,s}$, takes a value of one if a cohort c was born between 1957 and 1965 in a state s that enacted a contraceptive ban. In the first stage, we then regress the n-year difference in log cohort size at the time of marriage on the two dummy variables and a set of controls, i.e.

$$\log \frac{y_{c,s}}{y_{c-n,s}} = \alpha + \beta_1 ban_s + \beta_2 ban * pill_{c,s} + \sum_{c,r} \pi_{c,r} + X'\gamma + \varepsilon_{c,s}, \qquad (1)$$

where $y_{c,s}$ is the size of cohort c in state s, $y_{c-n,s}$ is the same variable for cohort $c - n$, $\pi_{c,r}$ are cohort-region fixed effects, and X' is a set of control variables that includes an indicator equal to 1 if the state had a physician exception, the physician indicator interacted with $pill_{c,s}$, and an indicator equal to 1 if the state enacted an advertising ban on contraception.

The results of the first stage are presented in Table 5, where we report the effect of the anti-obscenity laws on 1-year, 3-year, 5-year, and 7-year differences in log cohort size. After the introduction of the pill and before the repeal of the Comstock laws, the sales ban on contraceptives had a positive and statistically significant effect on cohort size at

marriage age in all cases. Consistent with Figure 5, Panel B, the effect is larger as we go from a 1-year difference to a 5-year difference, while the coefficients for the 5-year and 7-year differences are similar in size. Our estimates on the 5-year and 7-year differences indicate a 4% higher growth in cohort size in states with the sales ban between 1957 and 1965. The F-tests to evaluate the strength of the instruments are between 10.11 and 19.22 in our four specifications. An additional noteworthy result is that the coefficient on ban_s is always small and statistically insignificant suggesting that the sales ban had no effect on cohort size before the introduction of the pill. This finding is consistent with Bailey's results which indicate that the Comstock laws had no effect on other forms of contraception.

In the second stage, we use a specification similar to the one employed with the cross-state variation, i.e.

$$\log \frac{mar_{c,s}}{mar_{c-n,s}} = \beta_0 + \beta_1 \log \frac{size_{c,s}}{size_{c-n,s}} + \sum_{c,r} \pi_{c,r} + X'\gamma + \varepsilon_{c,s},$$

except that now we instrument cohort size growth with ban_s and $ban * pill_{c,s}$.

To construct the share ever married we must use the decennial Censuses since the CPS does not have enough state-level observations. In principle, the share ever married can be computed for each cohort born in a particular state if one observes in the Census data a recall variable measuring the age at first marriage. Unfortunately, after 1980 this variable is not available in the Censuses. Without this recall variable, the share ever married cannot be computed directly for each cohort, because in each decennial Census we only observe a particular cohort at a particular age. To address this limitation of the Censuses, we rely on the following strategy. In each Census, we first consider all individuals between the ages of 25 and 45. We then compute the share ever married for each cohort born in a particular state. Notice that we cannot use this variable directly in our regressions because it is affected by the age at which we observe a particular cohort in a particular Census. To deal with this issue, we regress the computed share ever married on age, state, cohort, and cohort-region dummies. We then remove the effect of age by subtracting the estimated coefficient on the age dummy multiplied by the dummy itself. Finally, we use the constructed variable in our regressions.

The second stage results are reported in Table 6 for men and women separately. The coefficient estimates have the expected negative sign, are statistically different from zero, and large in magnitude for both women and men. They indicate that during the period considered a 1% increase in cohort size at marriageable age generated a reduction in marriage

rates between 0.24 and 0.44 percent. The point estimates in the IV regressions are somewhat larger in size than the corresponding estimates obtained using longitudinal or cross-sectional variation. Note, however, that we would expect the IV coefficients to be negative and larger in magnitude given the discussion in Section 4.3 on potential endogeneity concerns, since the most plausible omitted variables would bias the coefficients positively toward zero. The IV findings are therefore consistent with the results obtained using longitudinal and cross-state variation. They indicate that there is a causal negative relationship between changes in cohort size and changes in marriage rates.

We conclude this section with a discussion of a potential threat to our IV strategy. It is possible that the negative relationship between cohort size and share ever married we find in our IV regressions is generated by some type of selection process governing who becomes a mother in states without the ban after the introduction of the pill. The most serious hypothesis that could confound the interpretation of our results is that, after the introduction of the pill, mothers in states without the ban give birth to fewer children who are positively selected along some dimension. If those children are more likely to marry, as the literature suggests, our IV regressions will estimate a negative relationship between our two main variables.

To evaluate this hypothesis, we follow Ananat and Hungerman (2012), who study the characteristics of children born to women under age 21 in states with different laws governing young women's access to oral contraceptives. The authors find that expanded access to the pill for young women increased the share of children born with low weight, a strong predictor of future economic outcomes. Following their empirical strategy, we test whether children born in states where the pill was initially banned are more or less likely to have low birth weight. We employ the same specification we use in the first stage of the IV estimation except that, to make our results comparable with Ananat and Hungerman (2012)'s findings, we use levels instead of differences for the following two new dependent variables: the share of children born with extremely low birth weight, which is defined as a birth weight below 1500 grams, and the share of children with low birth weight, which is a birth weight below 2500 grams.

The estimation results are reported in Table 7. Using both dependent variables, the estimated coefficient on the interaction between the ban and the introduction of the pill is small and statistically insignificant. We find therefore no evidence that early access to the pill had an effect – positive or negative – on the fraction of children born with low weight.

23

These findings are in line with those documented in Bailey (2013), who similarly finds no evidence, using a slightly different set of specifications, that the birthweight of infants born in the 1960s changed differentially in states where selling the Pill was legal from states where it was not. Our results differ from those obtained in Ananat and Hungerman (2012) for two reasons. First, we focus on contraceptive bans in the early 1960s that primarily affected women who had already reached the age of majority,[9]. Here, we consider all women who gave birth in the period of interest, the majority of whom are married, whereas Ananat and Hungerman (2012) study the behavior of single women younger than 21. The different results can therefore be rationalized by a more widespread early use of the pill by married women relative to young single women.

Ananat and Hungerman (2012) also find weak evidence that increased access to the pill to single women under 21 had the effect of increasing the share of children born in poor families, suggesting that the young women in their sample period who reduced their unwanted pregnancies using the Pill may have been positively selected. Bailey (2013) tests for such selection during our relevant sample period in the 1960s using data from the Integrated Fertility Survey Series (IFSS). The IFSS includes socioeconomic measures such as race and education, as well information about the wantedness and timing of pregnancies and births of female respondents. Bailey documents that Pill usage in the early 1960s was concentrated among women in married households, as expected. However, Bailey does not find evidence that reductions in unwanted births were higher for highly-educated women.

Finally, while we find no evidence of an increase in the share of children born in poor families in states with contraceptive bans, we note that an increase of the kind documented in Ananat and Hungerman (2012) would be a threat to our IV estimates only if children born in low income families are more likely to marry. In this case, the negative selection would generate the negative relationship between cohort size and marriage rates we observe in the data. But the literature on household formation appears to rule out this alternative.[10]

To summarize, we have provided evidence on two main results. First, there is a causal negative relationship between changes in cohort size and changes in marriage rates of women.

---

[9]Legal age of majority in the 1960s was typically 21. Even in those states that did not have contraceptive bans, women under the age of majority either had to have a guardian's consent or already be married to obtain oral contraceptives. while Ananat and Hungerman focus specifically on laws in the late 1960s and early 1970s that reduced the age of majority below 21 and increased access to the pill for minors. A second and related reason for the different findings is that we study a different population

[10]For instance, the handbook chapter by Black and Devereux (2010) indicates that there is a positive intergenerational correlation in income and education. Moreover, Stevenson and Wolfers (2007) find no difference in the share ever married by education and therefore income. These two results suggest that children of low income parents are not more likely to marry.

Second, there is a similar causal negative link between changes in cohort size and changes in marriage rates of men.

4.5 Cohort Size and Marriage Rates: Possible Mechanisms

The documented negative relationship between changes in cohort size and changes in marriage rates for women follows the predictions of a standard matching model. However, the fact that one observes the same negative relationship for men is more difficult to explain using such a model. Indeed, this is a surprising finding that previous papers studying the potential effects of population growth on marriage outcomes in the 1970s in the U.S. (e.g., Akers (1967), Schoen (1983)) and over longer periods of time in Sweden (e.g., Bergstrom and Lam (1989)) have predicted incorrectly or overlooked.

In the companion paper Bronson and Mazzocco (2017a), we explore possible mechanisms that can explain the relationship between changes in cohort size and marriage rates for both men and women. We first develop and test a standard dynamic search model of the marriage market. We find that the search model is rejected by the data for the same reason the standard matching model is rejected: in the search model an increase in cohort size always increases the marriage rate of men. The search model we develop and the standard matching model have one common feature: changes in cohort size affect marriage rates exclusively through the number of women in the marriage market relative to the number of men and the corresponding probabilities of meeting a spouse. The rejection of these two models indicates that the relationship between cohort size and marriage rates cannot be explained by changes in meeting probabilities alone.

Using this insight, in Bronson and Mazzocco (2017a) we develop two variations of the search model and show that they are both able to generate the empirical results documented in this paper. In the first version, a man can undertake a pre-marital investment that makes him more attractive to potential partners and, hence, increases his probability of meeting a potential spouse. This pre-marital investment also increases the expected marital surplus of the potential couple, and therefore their probability of marrying. In this alternative formulation, men are more likely to invest when cohort size decreases because, in this case, the probability of meeting a potential spouse without investment is too low. In the second version, the model is modified to allow the value of being single to increase with cohort size. The economic idea behind this new feature of the model is that it is more enjoyable to be single when cohort size is large because there are more individuals of the same age with

whom to perform leisure activities.

In the companion paper, we show that both versions of the search model can generate the observed pattern that a rise in cohort size generates a decline in the marriage rate of men. But, interestingly, the two alternative models generate this additional effect in different ways. The first version does this by increasing the value of marriage when cohort size declines, whereas the second version achieves this by increasing the value of being single when cohort size rises. This difference generates a testable implication based on the relationship between cohort size and divorce rates. The first version of the search model predicts that an increase in cohort size should generate an increase in divorce rates. The second version has the opposite prediction and generates a negative relationship between those two variables. In Bronson and Mazzocco (2017a), we use this implication to test one model against the other. We document that, in the data, there is a strong and positive relationship between changes in cohort size and changes in divorce rates. We can therefore reject the model in which the value of being single increases with cohort size in favor of the investment model. This result suggests that the negative relationship between cohort size and marriage rates can be explained by changes in cohort size that simultaneously affect the relative supply of women in the marriage market and the marital surplus.

5 Conclusions

Changes in cohort size over time or across geographies constitute one of the most important potential sources of variation in the supply of women relative to men in the marriage market, as a number of previous studies have noted. Though a large literature has studied the effects of sex composition on marriage behaviors in specific contexts, the importance of this factor for explaining overall variation in marriage rates for the majority of the U.S. population thus far has not been quantified nor formally tested. In this paper, we provide such a test and quantification. Using time-series variation, cross-state variation, and quasi-random variation in the adoption of the pill across states we provide evidence in support of the following two results. First, an increase in cohort size reduces marriage rates for both women and men and a decrease has the opposite effect. Second, changes in cohort size can explain about 50% of the variation in U.S. marriage rates.

Our findings help shed light on the drivers of marital changes over the last century in the U.S. Of course, they are also relevant in other settings. For example, large increases in births immediately after World War II affected growth in cohort size and therefore sex

composition in virtually every developed country. This may explain at least partly why large declines in marriage rates throughout the 1970s – when these cohorts reached marriageable age – were observed in all of those countries. This is true even in Japan, which legalized oral contraceptives – a predominant alternative explanation behind falling marriage rates in the 1970s – only in the late 1990s. Finally, our findings help clarify why complementary explanations about changes in marriage rates account in many periods for a relatively limited share of overall variation in marriage rates, and in some periods appear irrelevant. For example, the significant increase in women's labor supply during and after World War II (Acemoglu et al. (2004)) and the improvements in home technologies in the 1940s and 1950s were accompanied by large increases in marriage rates, counter to the predictions of those theories. This does not necessarily indicate that these theories are wrong, but rather that other factors – in particular the large demographic changes affecting sex composition in the marriage market in those decades – had a greater effect and pushed marriage rates upward.

Our findings also have policy implications. Both at the state and the federal level, politicians and policy makers have discussed and implemented policies that attempt to improve the well-being of low income families by increasing the fraction of married individuals. For instance, during the Bush Administration, such proposals allocated up to 1.5 billion dollars to implement and evaluate policies aimed at promoting marriage.[11] Our results suggest that these policies may prove largely ineffective since a significant part of the changes in marriage rates is generated by forces that are mostly outside the control of policy makers.

---

[11]Seefeld and Smock (2004) provide a nice discussion of the recent interest of policy makers in marriage as a policy tool.

References

Abramitzky, Ran, Adeline Delavande, and Luis Vasconcelos. 2011. "Marrying Up: The Role of Sex Ratio in Assortative Matching." *American Economic Journal: Applied Economics* 3 (3): pp. 124–157.

Akerlof, George A., Janet L. Yellen, and Michael L. Katz. 1996. "An Analysis of Out-of-Wedlock Childbearing in the United States." *The Quarterly Journal of Economics* 111 (2): pp. 277–317.

Akers, Donalds. 1967. "On Measuring the Marriage Squeeze." *Demography* 4 (2): 907–924.

Ananat, Elizabeth Oltmans, and Daniel M. Hungerman. 2012. "The Power of the Pill for the Next Generation: Oral Contraception's Effects on Fertility, Abortion, and Maternal and Child Characteristics." *The Review of Economics and Statistics* 94 (1): pp. 37–51.

Angrist, Josh. 2002. "How Do Sex Ratios Affect Marriage and Labor Markets? Evidence from America's Second Generation." *The Quarterly Journal of Economics* 117 (3): pp. 997–1038.

Angrist, Joshua D., and Victor Lavy. 1996. "The Effect of Teen Childbearing and Single Parenthood on Childhood Disabilities and Progress in School." *NBER Working Paper*, no. 5807.

Bailey, Martha. 2013. "Fifty Years of Family Planning: New Evidence on the Long-Run Effects of Increasing Access to Contraception." *NBER Working Paper*, no. 19493.

Bailey, Martha J. 2010. ""Momma's Got the Pill": How Anthony Comstock and Griswold v. Connecticut Shaped US Childbearing." *The American Economic Review* 100 (1): pp. 98–129.

Becker, Gary S. 1973. "A Theory of Marriage: Part I." *Journal of Political Economy* 81 (4): pp. 813–846.

———. 1974. "A Theory of Marriage: Part II." *Journal of Political Economy* 82 (2): pp. S11–S26.

Bergstrom, Theodore C., and David A. Lam. 1989. "The Effects of Cohort Size on Mariage Market in Twentieth Century Sweden." *Center for Research on Economic and Social Theory, Michigan* 91, no. 6.

Bitler, Marianne, and Lucie Schmidt. 2011. "Marriage Markets and Family Formation: The Role of the Vietnam Draft." *University of California at Irvine Working Paper.*

Black, Sandra E., and Paul J. Devereux. 2010. "Recent Developments in Intergenerational Mobility." *Prepared for the Handbook in Labor Economics 2010.*

Bronson, Mary Ann, and Maurizio Mazzocco. 2017a. "Cohort Size and the Marriage Market: What Explains the Negative Relationship?" *CCPR Working Paper.*

———. 2017b. "An Evaluation of Three Ways of Measuring Marriage Rates." *Journal of Demographic Economics, Forthcoming.*

Browning, Martin, Pierre-André Chiappori, and Yoram Weiss. 2014. *Economics of the Family.* Cambridge University Press.

Charles, Kerwin Kofi, and Ming Ching Luoh. 2010. "Male Incarceration, the Marriage Market, and Female Outcomes." *The Review of Economics and Statistics* 92 (3): pp. 614–627.

Doepke, Matthias, and Michelle Tertilt. 2016. "Families in Macroeconomics." *Handbook of Macroeconomics, Forthcoming.*

Easterlin, Richard A. 1987. *Birth and Fortune: The Impact of Numbers on Personal Welfare.* Chicago: Chicago University Press.

Ellwood, David T., and Jonathan Crane. 1990. "Family Change Among Black Americans: What Do We Know?" *The Journal of Economic Perspectives* 4 (4): pp. 65–84.

Goldin, Claudia, and Lawrence F. Katz. 2002. "The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions." *Journal of Political Economy* 110 (4): pp. 730–770.

Greenwood, Jeremy, and Nezih Guner. 2009. "Marriage and Divorce since World War II: Analyzing the Role of Technological Progress on the Formation of Households." *NBER Macroeconomics Annual* 23 (1): pp. 231–276.

Grossbard-Shechtman, Amyra. 1984. "A Theory of Allocation of Time in Markets for Labour and Marriage." *The Economic Journal* 94 (376): pp. 863–882.

Gruber, Jonathan. 2004. "Is Making Divorce Easier Bad for Children? The Long-Run Implications of Unilateral Divorce." *Journal of Labor Economics* 22 (4): pp. 799–833.

Hill, Matthew. 2011. "Love in the Time of Depression: The Effect of Economic Downturns on the Probability of Marriage." *UCLA Working Paper.*

Killingsworth, Mark R., and James J. Heckman. 1986. "Female Labor Supply: a Survey." *Handbook of Labor Economics, Edited by O. Ashenfelter and R. Layard,* 1:pp. 103–204.

King, Miriam, Steven Ruggles, J. Trent Alexander, Sarah Flood, Katie Genadek, Matthew B. Schroeder, Brandon Trampe, and Rebecca Vick. 2010. *Integrated Public Use Microdata Series, Current Population Survey: Version 3.0. [Machine-readable database].* Minneapolis: University of Minnesota.

Knowles, John A., and Guillaume Vandenbroucke. 2016. "Fertility Shocks and Equilibrium Marriage-Rate Dynamics: Lessons from World War 1 in France." *Working Paper.*

Lerman, Robert I. 1989. "Employment Opportunities of Young Men and Family Formation." *The American Economic Review, Papers and Proceedings* 79 (2): pp. 62–66.

Mazzocco, Maurizio, Claudia Ruiz, and Shintaro Yamaguchi. 2014. "Labor Supply and Household Dynamics." *American Economic Review: Papers & Proceedings* 104 (5): 354–59.

McLanahan, Sara, and Christine Percheski. 2008. "Family Structure and the Reproduction of Inequalities." *The Annual Review of Sociology* 34:pp. 257–76.

Moffitt, Robert A. 2000. "Welfare Benefits and Female Headship in U.S. Time Series." *The American Economic Review* 90 (2): pp. 373–377.

Neal, Derek. 2004. "The Relationship between Marriage Market Prospects and Never-Married Motherhood." *The Journal of Human Resources* 39 (4): pp. 938–957.

Pampel, Fred C., and H. Elizabeth Peters. 1995. "The Easterlin Effect." *Annual Review of Sociology* 21:pp. 163–194.

Plotnick, Robert D. 1990. "Welfare and Out-of-Wedlock Childbearing: Evidence from the 1980s." *Journal of Marriage and Family* 52 (3): pp. 735–746.

Ruggles, Steven, J. Trent Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schroeder, and Matthew Sobek. 2010. *Integrated Public Use Microdata Series: Version 5.0 [Machine-readable database].* Minneapolis: University of Minnesota.

Schoen, Robert. 1983. "Measuring the Tightness of a Marriage Squeeze." *Demography* 20 (1): pp. 61–78.

Seefeld, Kristin S., and Pamela J. Smock. 2004. "Marriage on the Public Policy Agenda: What Do Policy Makers Need to Know from Research?" *PSC Research Report, University of Michigan*, no. 04-554.

Seitz, Shannon. 2009. "Accounting for Racial Differences in Marriage and Employment." *Journal of Labor Economics* 27 (3): 385–437.

Stevenson, Betsey, and Justin Wolfers. 2007. "Marriage and Divorce: Changes and Their Driving Forces." *The Journal of Economic Perspectives* 21 (2): pp. 27–52.

Wilson, William J. 1987. *The truly disadvantaged: the inner city, the underclass, and public policy.* Chicago: Chicago University Press.

Wolfers, Justin. 2010. "What Is Going on With Marriage?" *New York Times Op-Ed and http : //www.freakonomics.com/2010/10/13/what − is − going − on − with − marriage/, October 12.*

Tables and Figures

Table 1: Average Age Differences Between Spouses

Age Difference between Husband and Wife	Percent of couples, 1962-2015
Husband is more than one year younger	8.1%
Husband and wife are approximately the same age	34.4%
Husband is 2 to 4 years older	31.8%
Husband is 5 to 6 years older	10.8%
Husband is 7 to 9 years older	7.6%
Husband is 10+ years older	7.3%

Notes: Because CPS data does not provide information about exact age or month of birth, we include in one category couples born in the same or adjacent birth years. Source: IPUMS CPS, 1962-2015.

Table 2: Time Series Regression of Log Share Ever Married on Log Cohort Size

	White Men	Black Men	White Women	Black Women
Ever married	-0.294***	-0.592***	-0.193***	-0.870***
by age 30	(0.037)	(0.048)	(0.026)	(0.064)
R^2	0.50	0.70	0.46	0.74
Ever married	-0.182***	-0.440***	-0.111***	-0.560***
by age 35	(0.015)	(0.031)	(0.011)	(0.043)
R^2	0.71	0.77	0.65	0.74
Ever married	-0.107***	-0.322***	-0.066***	-0.453***
by age 40	(0.008)	(0.019)	(0.007)	(0.026)
R^2	0.77	0.84	0.58	0.85

*** Significant at 1%. Newey-West standard errors in parentheses. Each coefficient is the outcome of a separate regression. Regressions include cohorts born after 1914 until the most recent cohort observed at a given age in 2015. The number of observations in each regression is equal to 71 for the share ever married by 30, 66 for the share ever married by 35, and 61 for the share ever married by 40. Sources: IPUMS CPS 1962-2015, IPUMS USA 1960-1970.

Table 3: Regression of Change in Log Share Ever Married on Change in Log Cumulative Cohort Size

	1-Yr.	2-Yr.	3-Yr	4-Yr	5-Yr.	7-Yr.	10-Yr.
Men	0.049	0.017	-0.157**	-0.220***	-0.233***	-0.238***	-0.277***
By Age 30	(0.037)	(0.062)	(0.061)	(0.049)	(0.045)	(0.040)	(0.047)
Men	-0.054	-0.057	-0.090**	-0.101***	-0.098**	-0.115***	-0.163***
By Age 35	(0.054)	(0.044)	(0.040)	(0.034)	(0.026)	(0.023)	(0.017)
Men	0.073***	0.037	-0.018	-0.048**	-0.054**	-0.052**	-0.085***
By Age 40	(0.016)	(0.038)	(0.029)	(0.022)	(0.023)	(0.023)	(0.021)
Women	-0.064***	-0.077***	-0.101***	-0.095***	-0.103***	-0.125***	-0.156***
By Age 30	(0.024)	(0.037)	(0.014)	(0.019)	(0.021)	(0.024)	(0.028)
Women	-0.005	-0.053**	-0.070***	-0.084***	-0.081***	-0.090***	-0.113***
By Age 35	(0.056)	(0.021)	(0.011)	(0.012)	(0.015)	(0.016)	(0.015)
Women	-0.011	-0.005	-0.008	-0.022*	-0.036**	-0.056***	-0.073***
By Age 40	(0.021)	(0.021)	(0.012)	(0.013)	(0.014)	(0.010)	(0.006)

* Significant at 10%. ** 5%. *** 1%. See notes in Table 1. Newey-West standard errors in parentheses.

Table 4: Cross-Sectional Regression of Log Share Ever Married by 30 or 40

Dependent Variable: 10-Yr. Difference in Log Share Ever Married

10-Yr. Dfference in Log Cohort Size	OLS	IV (1)	IV (2)
Men	-0.041**	-0.123***	-0.168***
By Age 30	(0.019)	(0.041)	(0.056)
R^2	0.807	0.791	0.819
Men	-0.029***	-0.047**	-0.040
By Age 40	(0.010)	(0.018)	(0.028)
R^2	0.555	0.551	0.595
Women	-0.041***	-0.104***	-0.115***
By Age 30	(0.010)	(0.023)	(0.036)
R^2	0.555	0.765	0.796
Women	-0.032***	-0.058***	-0.048**
By Age 40	(0.010)	(0.017)	(0.023)
R^2	0.539	0.528	0.605
First Stage Results			
Log Cohort Size at Birth			0.440*** (0.073)
F-test			36.44
R^2			0.831

** Significant at 5%. *** 1%. Robust std. errors in parentheses. N=288. Each coefficient is the outcome of a separate, population-weighted regression. We control for cohort fixed effects, and for cohort-region fixed effects in IV-(2). Includes all decennial cohorts born between 1910 and 1970, in all states except HI and AK. Sources: US Population Counts, 1910-1990; IPUMS USA, 1940-2010.

Table 5: Comstock Laws and N-Year Differences in Log Cohort Size

Dependent Variable: N-Yr. Difference in Log Cohort Size

	1-yr	3-yr	5-yr	7-yr
$Ban * Pill_{c,s}$	0.012**	0.035***	0.041***	0.040**
	(0.005)	(0.011)	(0.015)	(0.017)
Ban_s	-0.002	-0.003	0.006	0.018
	(0.005)	(0.010)	(0.012)	(0.014)
N	1248	1152	1056	960

* Significant at 10%. ** 5%. *** 1%. Robust standard errors in parentheses. Regressions are weighted by population and include controls for physician exception, physician exception interacted with "pill," advertising bans, and cohort-region fixed effects. Sources: IPUMS USA, 1980-2000, NIH SEER Population Counts.

Table 6: N-Year Differences in Log Share Ever Married and N-Year Differences in Log Cohort Size

Dependent Variable: N-Yr. Difference in Log Share Ever Married (Men)

	1-yr	3-yr	5-yr	7-yr
N-yr Difference	-0.389*	-0.241**	-0.239***	-0.260***
in Log Cohort Size	(0.206)	(0.101)	(0.091)	(0.080)
N	1248	1152	1056	960

Dependent Variable: N-Yr. Difference in Log Share Ever Married (Women)

	1-yr	3-yr	5-yr	7-yr
N-yr Difference	-0.441***	-0.311***	-0.302***	-0.304***
in Log Cohort Size	(0.161)	(0.093)	(0.084)	(0.076)
N	1248	1152	1056	960

* Significant at 10%. ** 5%. *** 1%. See note in Table 5.

Table 7: Comstock Laws and Birth Weight

	Share with Birth Weight < 1500	Share with Birth Weight < 2500
$Ban * Pill_{c,s}$	0.0001025	-0.0004241
	(.0002037)	(.0007311)
Ban_s	-0.000462	0.0006242
	(0.0003064)	(0.0006155)
N	1006	1006

See note in Table 5.

Figure 1: Cohort Size at Birth and at Age 30
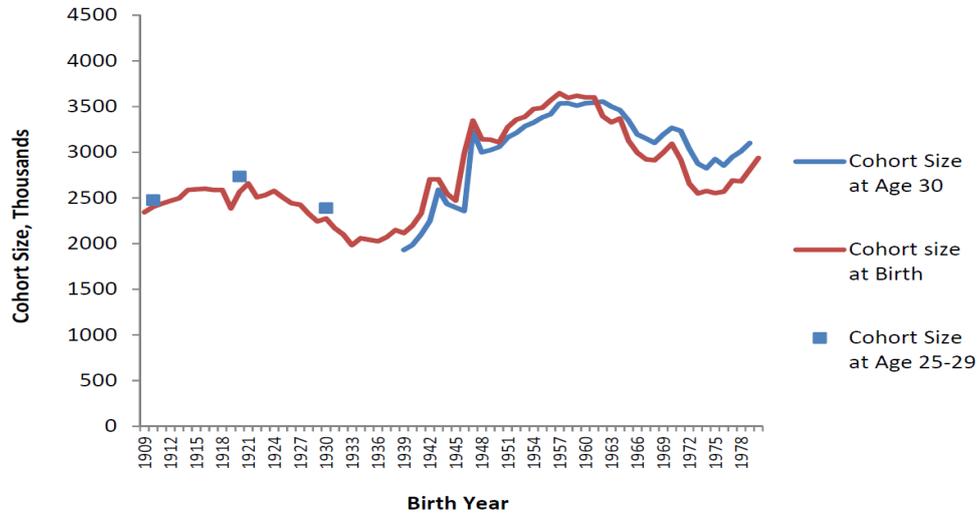
White only. Sources: Vital Statistics of the U.S.; NIH SEER population estimates (1969-2010); U.S. Census 1940-1960.

Figure 2: Share Never Married By 30

(a) White
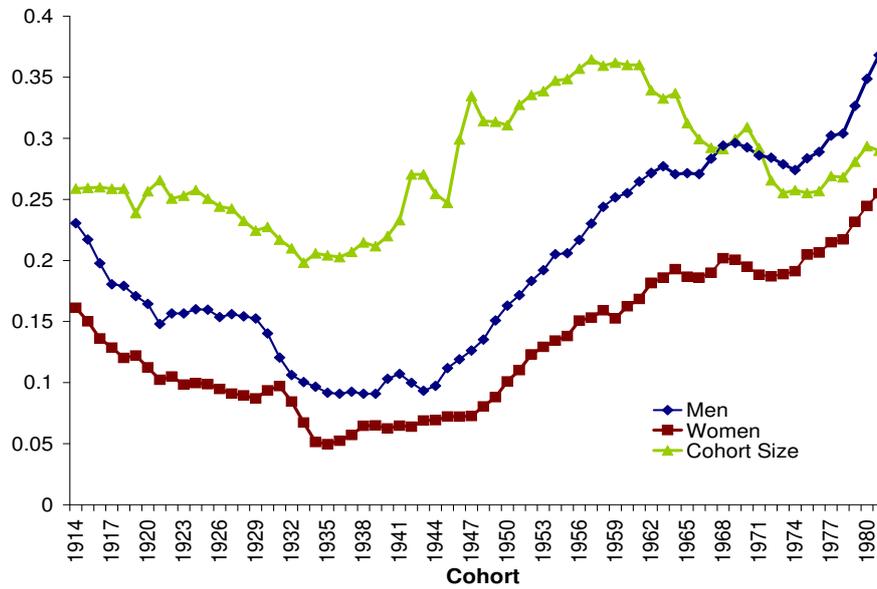
(b) Black
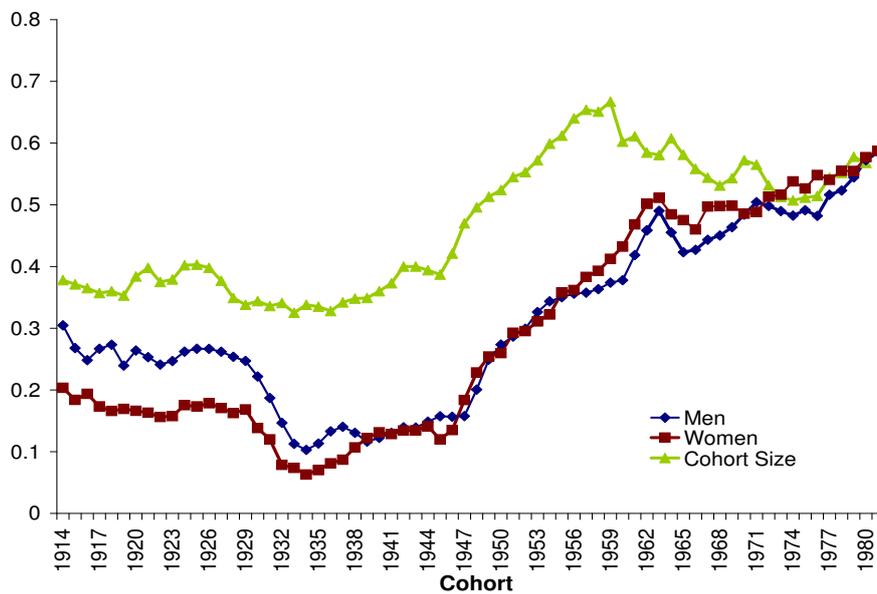
Notes: The vertical axis represents both the percentage of individuals ever married as well as normalized cohort size. In Panel A we normalize cohort size by dividing by 10,000,000; in Panel B by 1,000,000. For share ever married, we graph three-year moving averages. Sources: Vital Statistics of the United States; IPUMS CPS, 1962-2015; IPUMS USA, 1960-1970.

**Figure 3:** Share Never Married and Not Cohabiting By 30

(a) White
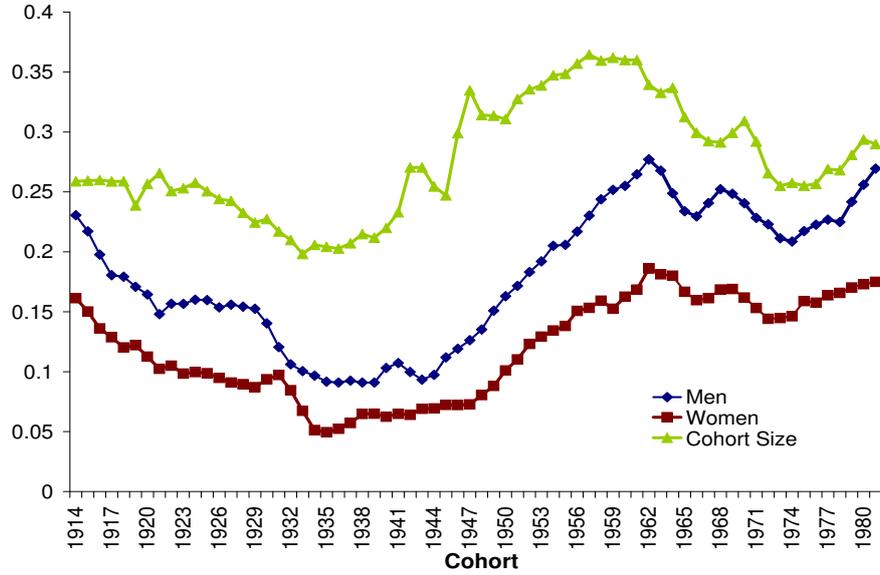
**(b)** Black
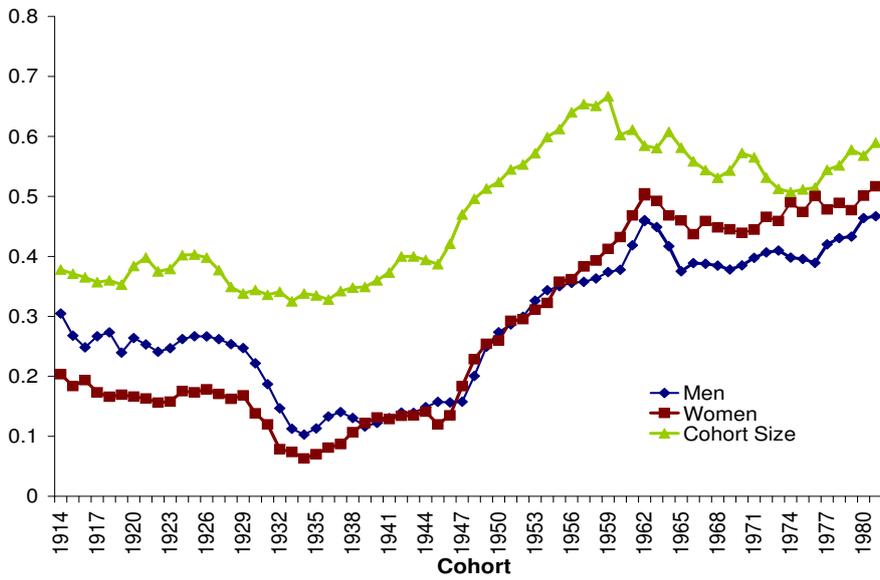
* See note in Figure 2.

Figure 4: Share Never Married and Not Cohabiting By 40

(a) White
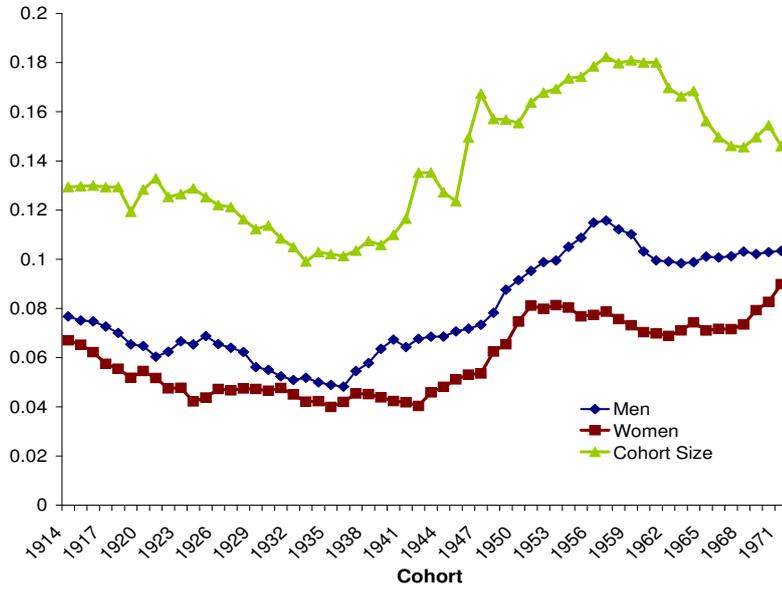

(b) Black
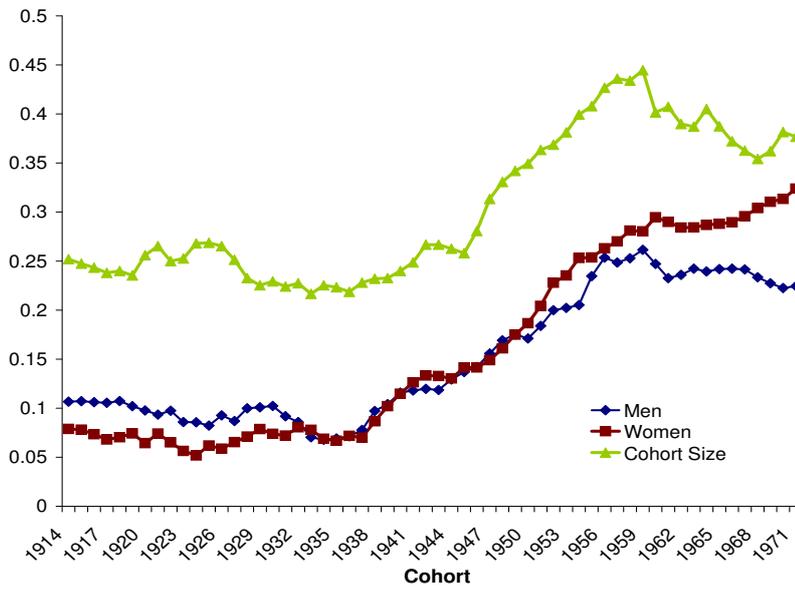
* See note in Figure 2.

Figure 5: Growth by Region: States with Sales Bans - States without Sales Bans

(a) Growth of Total Births
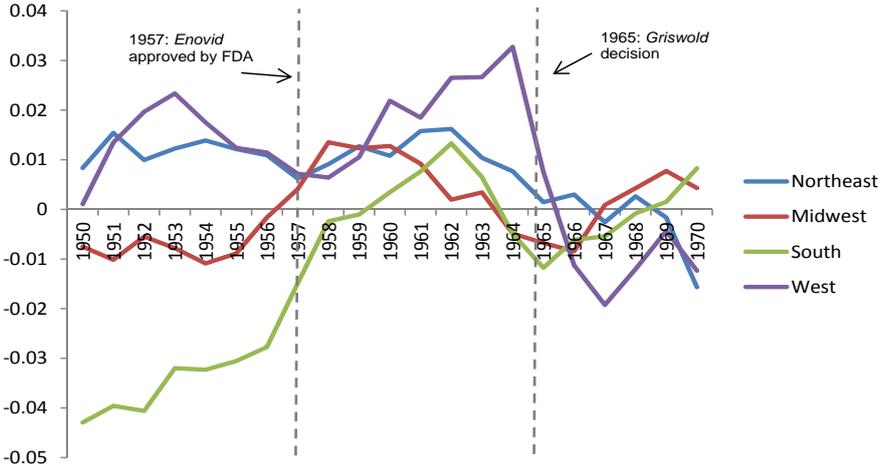
(b) Growth of Total Adult Population, at Age 25
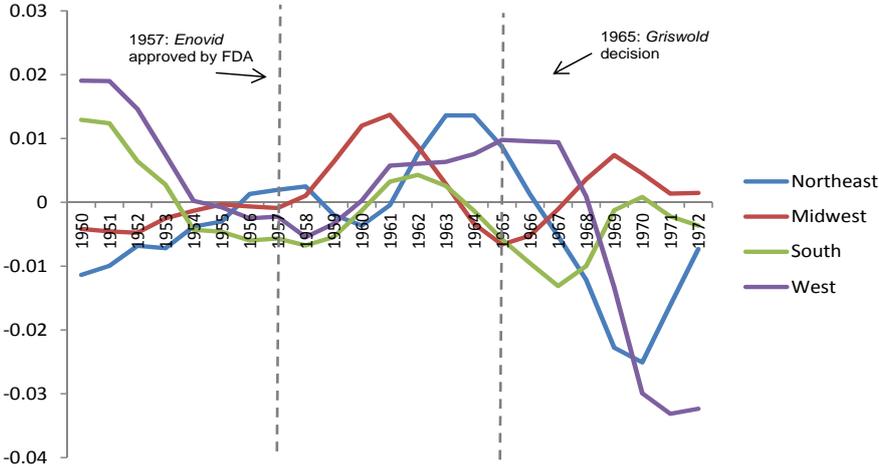
* See note in Figure 2.

Appendix: Data Description

Table A.1 provides a summary of the datasets employed in the construction of the main variables of interest. In the rest of the appendix, we give additional details about how we construct the variables share ever married, cohort size at birth, and cohort size at marriageable age.

The variable share ever married is constructed using a different procedure depending on whether we use longitudinal or cross-state variation. With longitudinal variation, we employ a combination of the CPS, which covers the period 1962-2015, and of the decennial Censuses. In the CPS, we observe the age and the marital status of each respondent. We can therefore easily compute the share ever married by age 30, 35, or 40 for each cohort born after a particular year. For instance, for the variable share ever married by age 30, we can use the CPS for all cohorts born on or after 1932; for the variable share ever married by age 40, we can use the CPS for all cohorts born on or after 1922. For cohorts born before those years, we use the 1960, 1970, and 1980 Censuses, which contain information on the marital status and the age at first marriage, a recall variable. Using these two variables, we construct the share ever married by age 30 and 35 for different cohorts by considering all individuals who in a given Census are between the ages of 30 and 45. We use a maximum cutoff age of 45 to avoid potential measurement errors due to differential mortality rates of married and non-married individuals. For the share ever married by age 40, we use the same procedure with a maximum cutoff age of 50. With cross-state variation, for all cohorts we only use information from the Censuses, as sample sizes in the CPS are too small to provide reliable estimates at the state level. In the longitudinal variation, we can construct the share ever married annually only for cohorts born after 1914 because the 1960 Census is the first one that records the age at first marriage.

As discussed in Section 3, in the paper we use two different measures of cohort size: cohort size at marriageable age and cohort size at birth. Cohort size at marriageable age is used as the main independent variable in the cross-state regressions and in the regressions that use the introduction of the pill as an instrument. There are two datasets that can be used to measure this variable: the decennial Censuses and the SEER population estimates. SEER records cohort sizes at different ages starting from 1969. Hence, using this dataset we would be able to construct cohort size at marriageable age only for some of the decennial cohorts born between 1910 and 1970, which are the ones we consider in the cross-state regressions. For consistency, we use the decennial Census for all years in the cross-state

analysis. The decennial Census population counts are published for 5-year age groups. We therefore construct cohort size at marriage age by recording the number of individuals between the ages of 20 and 24 in the decennial Censuses 1930-1990. We also experimented with the 5-year age group 30-34 in the 1940-2000 Censuses with similar results. In the regressions that use the pill and the anti-obscenity laws as instruments, we use cohorts born between 1945 and 1970 which are all observed in SEER at age 25 or older. We therefore measure cohort size at marriage using the information in SEER at age 25.

Cohort size at birth is used in three ways: as the main independent variable when we employ longitudinal variation; as an instrument for cohort size at marriage age in the cross-state regressions; and as one of the variables used to determine the effect of the introduction of the pill in states with different anti-obscenity laws. As indicated in Table A.1, in the longitudinal analysis cohort size at birth is constructed using the U.S. Vital Statistics which provide information on this variable by race starting in 1909. For the cross-state regressions, there are two data sets that can be used to measure cohort size at birth: the U.S. Vital Statistics which record births by race and by state from 1940; and the decennial Censuses which provide information on population counts by race and state from the beginning of the twentieth century to 2010. For consistency, rather than combining two different datasets, we use the Censuses over the entire period of interest. As mentioned earlier, decennial Censuses publish population counts for 5-year age groups. From each decennial Census, we therefore record the number of individuals between the ages of 0 and 4 and use it to construct the cohort size at birth. Our results do not change if we use data from the U.S. Vital Statistics for the 1940 to 1970 cohorts. In the regressions that use the introduction of the pill as an instrumental variable, we consider cohorts born between 1945 and 1970. We can therefore use the U.S. Vital Statistics to compute cohort size at birth for all of them.

Table A.1: Data Sets Used in the Construction of the Main Variables

Variable	Variation of Interest	Dataset
Cohort Size at Birth	National, Yearly	Vital Statistics of the United States (1909-1980)
	State Level, Decennial Years	U.S. Decennial Census, 1910-1970
Cohort Size at Adulthood	National and State Level, Yearly (Around the Introduction of the Pill)	NIH/SEER Population Estimates (1969-2010)
	State-Level, Decennial Years	U.S. Decennial Census, 1930-1990
Share Ever Married	National, Yearly	CPS (1962-2015), U.S. Decennial Census, 1960-1970
	State Level, Decennial Years	U.S. Decennial Census, 1940-2000