ABSTRACT

Title of dissertation: THE EFFECT OF SAME-SEX MARRIAGE LAWS ON DIFFERENT-SEX MARRIAGE: EVIDENCE FROM THE NETHERLANDS

Mircea Ştefan Trandafir,
Doctor of Philosophy, 2009

Dissertation directed by: Professor Judith K. Hellerstein
Department of Economics

It has long been argued that the legalization of same-sex marriage would have a negative impact on marriage. My dissertation examines what happened to different-sex marriage in the Netherlands after the enactment of two laws: in 1998, a law that provided *all* couples with an institution almost identical to marriage—registered partnership, and in 2001, a law that legalized same-sex marriage for the first time in the world. The first chapter provides a brief description of the same-sex marriage debate and of the legal background in the Netherlands. In the second chapter, I analyze the marriage decision at the individual level. I construct a unique data set covering the period 1995–2005 by matching individuals from the Dutch Labor Force Survey with their marriage and residence history from official records. I estimate the first-marriage decision using a discrete-time hazard model with unobserved heterogeneity and I find that the marriage rate rose after the registered partnership law but fell after the same-sex marriage law.

In the third chapter, I study the evolution of the marriage rate in the aggregate.

I construct a synthetic control for the Netherlands as a weighted average of OECD member countries over the period 1988–2005. A comparison of the marriage rates in the Netherlands and the synthetic control confirms the findings from the individual-level analysis: the different-sex marriage rate rose after the registered partnership law and then fell after the same-sex marriage law. I also conduct a placebo test that supports the validity of the results. Finally, I examine the evolution of the different-sex union (marriages and registered partnerships) rate and I find that the different-sex union rate returns to its long-term trend after the same-sex marriage law.

My results could be explained by the combination of two effects. First, couples may learn over time about registered partnership and gradually switch from marriage to the new institution. Second, the same-sex marriage law may have caused some couples to turn away from marriage.

THE EFFECT OF SAME-SEX MARRIAGE LAWS ON DIFFERENT-SEX MARRIAGE: EVIDENCE FROM THE NETHERLANDS

by

Mircea Ştefan Trandafir

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2009

Advisory Committee:
Professor Judith Hellerstein, Chair
Professor Seth Sanders
Professor John Wallis
Professor John Ham
Professor Joan Kahn

Dedication

This dissertation is dedicated to my wife Meltem, without whom this project would have never begun nor ended, and to my parents, Ileana and Romică, who supported me every step of the way.

Acknowledgments

I am extremely fortunate to have worked with extraordinary economists during my graduate studies. Judy Hellerstein provided invaluable guidance and support, while Seth Sanders was a constant source of advice and inspiration. John Wallis taught me how to better formulate my research and John Ham guided me through the world of duration models. I also would like to thank Joan Kahn for her feedback while serving on my dissertation committee.

Finally, I owe special thanks to Jaap Abbring, Eric Bartelsman and the Tinbergen institute for their hospitality and support in obtaining the data, to Kees Waaldijk for his clear description of the Dutch legal background, and to Gerhard Meinen, Ineke Brekelsman and the staff of the Microdata lab at Statistics Netherlands for all their patience and assistance with the data.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 The same-sex marriage debate and the end-of-marriage argument

Economists have long been interested in the effects of various policies on individuals' decision to marry. Some of the policies studied are aimed directly at the institution of marriage, such as no-fault divorce laws (Allen et al., 2006; Rasul, 2006) or minimum age requirements (Blank et al., 2007). Other policies alter monetary incentives associated with marriage, such as welfare reform (Bitler et al., 2004), or reduce the cost of premarital sex arising from a reduction in the legal age for use of oral contraceptives (Goldin and Katz, 2002). The common theme in all these studies is that the policies each alter the value of marriage relative to alternative arrangements.

It has been argued recently that another policy that could affect the value of marriage and implicitly the incentives to marry, particularly for heterosexual couples, is the legalization of same-sex marriage. The "end-of-marriage argument" holds that opening marriage to same-sex couples would lead to a fall in the number of different-sex marriages because the value of marriage would be reduced: "It demeans the institution. [...] The institution of marriage is trivialized by same-sex marriage." (Rep. Henry Hyde, U.S. House of Representatives subcommittee meeting, as cited in Mohr, 1997) This argument has been mentioned frequently in the same-sex marriage debate, both in the media and in the political discourse, and was used to justify amendments to state constitutions such as Proposition 8 in California, or the Defense of Marriage Acts, laws meant to protect the federal or state governments from having to recognize a (same-sex) marriage performed elsewhere.[1]

To understand the logic behind the end-of-marriage argument, consider the following analogy. Country clubs establish a certain privileged status for their members and usually reserve the right to deny membership to persons they deem "unacceptable."[2] Now suppose the state enacts a law that forbids country clubs to discriminate against admissions in any way and effectively allow anybody who desires so to be a member. In this case, some individuals who were considering becoming members before the enactment of the law in order to gain the privileged status might choose not to apply anymore, since there is no privileged status—the value of membership was reduced for these individuals. A similar effect could be observed if the state were to create a "competing" country club, which would offer the same benefits as the private country club, but which would be open to everybody.

However, it is theoretically ambiguous whether same-sex marriage should have a negative effect on heterosexual marriage. For example, the legalization of same-sex marriage could be perceived as a move toward a more secular, less traditional

[1] By the end of 2008, Congress and forty states had enacted such acts (Stateline.org, 2009). Thirty states also had constitutional amendments that specifically defined marriage as the union between a man and a woman.

[2] A case that made the headlines was the refusal of membership by the Kansas City (Mo.) Country Club to Henry Bloch, chairman and co-founder of H&R Block Inc., because he was a Jew. This led to the withdrawal of professional golfer Tom Watson from the club (New York Times, 1990).

institution, increasing its value for some different-sex couples.[3] The final effect on the number of marriages for different-sex couples would then depend on which effect dominates: the decline in the value of marriage for some couples or the increase in the value of marriage for some other couples.

This dissertation constitutes the first analysis of the effects that same-sex marriage has on heterosexual marriage by studying its effect in the Netherlands. The Netherlands is a good candidate for such an analysis for three reasons. First, it was the earliest (2001) country to allow same-sex couples to marry, thus offering the longest period with which to examine the impact of this law.

Second, prior to legalizing same-sex marriage, the Dutch legislature formalized in 1998 the legal concept of registered partnership, an institution that is identical to marriage in almost every respect but name and tradition. Unlike in the Nordic countries, the Dutch registered partnership is also open to different-sex couples. Since this contract is virtually identical to marriage, this offers a unique opportunity to distinguish between a change in the marriage rate itself and a change in the number of couples who wish to legally formalize their relationship. It also allows for the distinction between the effects of granting same-sex couples the rights and benefits of marriage through an alternative institution and the effects of same-sex marriage itself. The short period of time between the two laws, however, also presents the challenge of distinguishing between the long-term effects of the registered partner-

[3]In a survey of Dutch couples married or in alternative institutions ran after the legalization of same-sex marriage, Boele-Woelki et al. (2006) found that almost the same fraction of married and unmarried couples pointed to the traditional value of marriage as a reason for their choice of union form. This suggests that some couples preefer marriage because of its traditional value, while other couples dislike it for the same reason.

ship law and the effects of the legalization of same-sex marriage.

Finally, the same-sex marriage debate in the Netherlands was also subject to a vigorous end-of-marriage argument. Boris Dittrich, former member and floor leader of the Dutch Parliament and a supporter of the same-sex marriage bill, reports the following episode: "I distinctly remember my former colleague, Kees van der Staaij from the Orthodox Christian Party SGP, using those arguments. He even said that God would punish those who are destroying the institution of marriage between a man and a woman. [...] That night of the debate (we were debating same-sex marriage for two full days) I drove home very carefully. I thought: if I will get into a car accident tonight, people will think God has punished me."[4]

1.2 The Dutch legal environment

The road to same-sex marriage in the Netherlands was long and bumpy.[5] At the beginning of the 1990s, gay rights organizations in the Netherlands tried to build on the success of their Danish counterparts, who had obtained the enactment of a registered partnership law in 1989, and push for the legal recognition of same-sex couples. The first move was to suggest in 1991 the creation of a symbolic registry for same-sex couples, which could have evolved into an alternative to the marriage registry and to which municipalities would participate voluntarily. Under this arrangement, same-sex couples were allowed to register their relationship with the municipality, a registration that did not involve any benefits or obligations from

[4]Translation of a previous statement (de Verdieping Trouw, 2005) provided during personal conversation.

[5]The presentation in this section draws extensively on van Velde (2005) and Merin (2002).

either the couple or the municipality or the Dutch government. This suggestion was not in conflict with the general feelings of the population, as opinion polls showed that almost 53% of the population supported same-sex marriage in 1990, a share that increased over time to around 63% in 1991 to about 73% in 1995 (van Velde, 2005). More than 100 of the 650 Dutch municipalities voluntarily decided to participate within the first year, thus paving the way for the introduction of the registered partnership system.

The government was to set up a committee of legal advisers (the Kortmann Committee) to inquire into the legal effects and the desirability of the legal recognition of same-sex couples. The committee recommended the introduction of a Danish-style partnership, a proposal warmly received by the government. A bill for the new institution was introduced in the Parliament in 1993, but held up because of the 1994 elections. After the elections, the governing coalition in the Netherlands did not include the Christian Democrats, the largest party opposing same-sex marriage, for the first time in almost eighty years. In 1995, the new cabinet presented a white paper that suggested the introduction of registered partnerships for same-sex couples and, in a departure from the Danish model, for different-sex couples as well. The argument was that the new institution was not supposed to discriminate based on sexual orientation and that heterosexuals not willing to marry should have access to this alternative contract (Merin, 2002).

Therefore, the registered partnership was designed to be almost identical to marriage and, in the case of different-sex couples, an almost perfect substitute. Waaldijk (2004) compares the two institutions based on their "levels of legal consequences" (the rights and obligations derived from a contract) and finds only three differences for heterosexual couples. First, a married man is automatically acknowledged as the father of a child born in that marriage, whereas a man in a registered partnership has to explicitly claim the child before or after birth (although this is rather a formality). Second, both contracts can be terminated in court, but registered partnerships can also be dissolved at the civil registry by mutual agreement.[6] Finally, couples in registered partnerships are prohibited from engaging in international adoption. This restriction can be circumvented by one partner adopting the child as a single individual and the other partner subsequently adopting the child as the partner of the adoptive parent.

Since the planned legislation granted same-sex couples access only to registered partnerships, there was an argument that same-sex couples would still face discrimination. A motion calling for the opening of civil marriage to same-sex couples was introduced in 1996, when the white paper was presented as a bill in the Dutch Parliament. While the registered partnership bill was making its way through the Parliament's two chambers, the government acknowledged the request to open up marriage and appointed a new panel of experts (the second Kortmann Committee) to analyze the desirability and consequences of same-sex marriage. In the meantime, the registered partnership bill was approved and signed into law, becoming effective on January 1, 1998.

[6]Starting from 2001, married couples can change their marriage to a registered partnership. Statistics Netherlands reports that, in more than 90% of the cases, this is part of a two-step procedure commonly called "flash divorce", where the partnership is dissolved by mutual agreement. This is a cheaper alternative to divorce when there is an understanding with respect to the division of common property. Therefore, the rest of the analysis will only include new partnerships rather than the total number of partnerships.

The report of the Kortmann Committee, released in October 1997, recommended the legalization of same-sex marriage while dismissing the arguments against it. In particular, it addressed the issue of a possibly negative effect on heterosexual marriage: "The argument that a large part of the population would no longer be able to identify with marriage if it were opened up, applies to an ever diminishing part of society. They can continue to identify with a marriage in church." (Kortmann Committee report, as quoted in van Velde, 2005) Still, the government opposed opening up marriage to same-sex couples in its February 1998 answer, which prompted another request in the second chamber of the Parliament for new legislation to allow same-sex couples to marry. The 1998 elections allowed the same coalition to remain in power, and the ensuing negotiations for the formation of a new cabinet led to an agreement for the introduction of same-sex marriage during that term. Finally, in 2000, the bill legalizing same-sex marriage was introduced in the Parliament and was approved in September by the House of Representatives and in December by the Senate (Merin, 2002). On April 1, 2001, the Netherlands became the first country in the world to allow same-sex couples to marry, granting them access to an institution that was traditionally available only to different-sex couples.

In conclusion, the road to the opening of marriage was long and uncertain. There was no sudden overhaul of the marriage law, but rather a series of small changes: starting from 1998, both same-sex and different-sex couples could form registered partnerships. From 2001, same-sex couples are allowed to enter marriage. Some scholars argue that the success of the gay rights movement in the Netherlands was actually due to this small-steps approach (Waaldijk, 2001). For the purpose of this paper, however, the fact that there was uncertainty with respect to the timing of the laws allows for an interpretation of them as "exogenous". Individuals could not anticipate perfectly the enaction date of each law and marriages would have already been planned by the time the laws were announced, as they are usually scheduled in advance.

1.3 Previous literature

The evidence in support of or against the end-of-marriage hypothesis is remarkably sparse and is mostly conducted by members of various advocacy groups. The arguments usually rest on the interpretation of aggregate trends or on anecdotal evidence from the Netherlands and from the Scandinavian countries. In a long series of articles focusing on the Netherlands and on the Scandinavian countries, Kurtz (2004a,b,c,d) argues that the legalization of same-sex marriage or registered partnership led to an increase in out-of-wedlock childbirth, a decline in the marriage rate and an increase in alternative family arrangements. His argument is that the legalization of same-sex marriage (and, implicitly, the introduction of registered partnership) created the idea that any family form is acceptable. As such, the countries that legalized any form of legal recognition of same-sex couples, be it marriage or registered partnership, experienced a crumbling of the family structure, with more children being born to unmarried parents and less people choosing to marry at all. His conclusions rest entirely on the graphical analysis of aggregate data and on the anecdotal evidence found in local media.

Similarly, Fagan and Smith (2004) use aggregate data from the Netherlands to calculate a number of indicators that point to a deterioration in the traditional family structure. In particular, they note an increase in the number of children born out of wedlock or who live in single-parent families, in induced abortions and in the number of people who choose not to have children. They also report a decrease in the number of married people and an increase in the number of divorced people.

The response from proponents of same-sex marriage uses the same graphical tools to rebuke the claims above. Badgett (2004a,b), for instance, uses the same Dutch and Scandinavian aggregate data to argue that there is a long-term downward trend in the marriage rate that the articles above ignore and that, once this trend is taken into account, there is no decline in the marriage rate after the legalization of same-sex marriage. Similarly, non-marital births rates did not experience a faster increase after the legalization of same-sex marriage or registered partnership. In addition, most of the parents are still married—for example, the likelihood of cohabiting parents in Sweden to break up is lower than the likelihood of married parents in the U.S. to break up (Cherlin, 2009). Eskridge and Spedale (2006) reach the same conclusions by analyzing data from all Scandinavian countries that introduced registered partnership during the 1990s. Finally, Mello (2004), Cahill (2004) and Eskridge and Spedale (2006) bring extensive anecdotal evidence to support the claim that same-sex marriage or registered partnership did not or would not affect negatively heterosexual marriage.

All these studies suffer from two major problems. First, they implicitly assume that the only reason why marriage rates (or alternative indicators of family structure) change is the legalization of same-sex marriage or the introduction of same-sex registered partnership. In reality, that there are many factors that can bring about changes in the marriage rate as the examples at the beginning of this chapter show. Ignoring these factors can lead to the contradicting conclusions offered by the media while using the same aggregate data.

Second, all these studies assume that the marriage rate prior to the legalization of registered partnership or of same-sex marriage is an appropriate counterfactual for what the marriage rate would have been in the absence of the two laws. This, again, ignores the fact that there are many changes in the environment that can lead to changes in the marriage rate. One cannot disentangle the effect of a same-sex marriage law from changes caused by population shifts only by looking at the marriage rates after the introduction of same-sex marriage.

1.4 Summary of findings

In this dissertation, I address both of the issues mentioned above. In chapter 2, I conduct my analysis at the individual level. I first construct a unique and highly confidential individual-level data set that includes demographic characteristics as well as information on the marriage decisions over the period 1995–2005 for a significant fraction of the Dutch population. I then estimate a duration model for age at first marriage. Using multiple specifications, I calculate the effect of the registered partnership law and of the same-sex marriage law on the first-marriage rate. I find that not controlling for other factors or for unobserved heterogeneity in the propensity to marry introduces significant biases in the results. My estimates from specifications with unobserved heterogeneity suggest that the marriage rate increases after the enactment of the registered partnership law and falls after the same-sex marriage law came into effect. However, this pattern is not uniform. Individuals living in the more conservative municipalities commonly called the Dutch Bible belt tend to marry less after the registered partnership law, but their marriage rate returns to the long-term trend after the same-sex marriage law. In contrast, individuals living in the four largest cities (the more liberal areas) marry less after both laws. Finally, people residing outside these two regions follow the same pattern as the overall marriage rate, marrying more after the registered partnership law and less after the same-sex marriage law.

A limitation of the individual-level analysis is that I cannot construct a counterfactual—what the marriage rate would have been in the absence of same-sex marriage laws. To address this concern, in chapter 3 I turn to aggregate data and I use the method developed by Abadie and Gardeazabal (2003) to construct a synthetic control for the Netherlands. This synthetic control is a weighted average of the marriage rates of the OECD member countries such that its evolution before 1998 (the year when the registered partnership law was enacted) matches the evolution of the marriage rate in the Netherlands. The weights for each country are data-driven, calculated by matching the values of the marriage rate and its determinants in the synthetic control to the corresponding values in the Netherlands for the period 1988–1997. A comparison of the Dutch marriage rate to the synthetic marriage rate confirms the average findings from the individual-level specifications: different-sex couples marry more after the registered partnership law, but less after the same-sex marriage law.

Another advantage of the aggregate data is that it allows me to analyze both different-sex marriages and different-sex unions, i.e. marriages *and* registered partnerships. The results suggest that the rate of different-sex unions increases after the registered partnership law (not surprisingly, since the marriage rate increases) and then falls after the enactment of the same-sex marriage law, so that the rate of different-sex unions after 2001 remains close to the rate predicted by the synthetic control.

In conclusion, the introduction of registered partnership, both for same-sex and for different-sex couples, does not seem to have negative effects on different-sex marriage in the short term—indeed, there is an increase in the different-sex marriage rate between 1998 and 2000. However, the different-sex marriage rate falls after the legalization of same-sex marriage.

There can be at least two explanations for the decline in the diferent-sex marriage rate after 2001 and for the fact that the different-sex union rate returns to the long-term trend during the same period. The first explanation is that couples learn over time about registered partnership and shift gradually from marriage to registered partnership. The long-term effect of the introduction of (different-sex) registered partnership would therefore be to sort couples across the two institutions without changing the total number of couples willing to formalize their relationship.

The second explanation is the end-of-marriage argument: the same-sex marriage law changes the value of marriage for some couples, who choose not to marry anymore. The fact that some of these couples choose cohabitation over registered partnership would be offset by the fact that some other couples who were not considering marriage would enter registered partnership, such that the total number of different-sex unions stays more or less the same (around the long-term trend).

Finally, note that it is practically impossible to disentangle the long term effects of the registered partnership law from the (short-term) effects of the same-sex marriage law. In fact, additional evidence from the survey conducted by Boele-Woelki et al. (2006) suggests that the post-2001 evolution of the marriage rate might be due to both effects. It is also infeasible to gauge the relative magnitude of each effect, i.e., it is impossible to attribute the decline in the marriage rate after 2001 to either the long-term effect of the registered partnership law or to the short-term effect of the same-sex marriage law.

Chapter 2

Individual-level analysis

2.1 Introduction

In this chapter, I use individual-level data to analyze marriage behavior before and after the enactment of each law. This type of analysis goes two important steps further than the graphical analysis conducted until now on this topic. First, although I still refer to the period before the enactment of the laws as the counterfactual, I control for a variety of factors that could have affected the marriage rate during this period and that the previous literature ignores. Second, by using individual data I am able to uncover different responses in different groups in the population, which allows me to better assess the validity of the end-of-marriage argument.

I construct a unique and highly confidential individual-level data set by combining administrative data from the Dutch Municipal Records with the Dutch Labor Force Survey. The resulting data set includes demographic characteristics as well as information on the marriage decisions over the period 1995–2005 for almost 10 percent of the Dutch population. I focus on young individuals in order to eliminate selection bias and I apply a duration model for age at first marriage. An important aspect in duration models is unobserved heterogeneity. A model that does not consider the possibility that people can differ on dimensions hidden to the researchers can be severely biased. I start by estimating models without unobserved heterogeneity and I contrast them with models including unobserved heterogeneity. I model the heterogeneity as a discrete distribution, similar to Heckman and Singer (1984). The discrete distribution is particularly intuitive in this case (as opposed to a continuous distribution such as Gamma) because it separates the population into individuals who are more or less likely to marry.

As expected, I find that not controlling for unobserved heterogeneity in the propensity to marry introduces significant biases in the results. My estimates from specifications with unobserved heterogeneity suggest that the marriage rate is not negatively affected by the enactment of the registered partnership law, but it falls after the same-sex marriage law came into effect. However, this pattern is not uniform. I use the fact that different regions are characterized by different patterns in marriage, fertility, divorce, cohabitation and voting to separate individuals into three groups based on their residence. The first group includes individuals living in the more conservative municipalities commonly called the Dutch Bible, where the marriage and fertility rates are higher than the average and the divorce and cohabitation rates are lower than the average (Sobotka and Adigüzel, 2002; Statistics Netherlands, 2003). At the opposite extreme are the four largest cities: Amsterdam, Utrecht, Rotterdam and the Hague, where the marriage and fertility rates are lower than the average and the divorce and cohabitation rates are higher than the average. Finally, the rest of the country is considered all together as the third group. Using specifications that separate the effect of the laws across these three regions, I find that individuals in the Bible belt tend to marry less after the registered partnership law, but their marriage rate returns to the long-term trend after the same-sex marriage law. In contrast, individuals living in the four largest cities (the more liberal areas) marry less after each law. Finally, people residing outside these two regions follow the same pattern as the overall marriage rate, marrying more after the registered partnership law and less after the same-sex marriage law.

My results indicate that individuals living in the more conservative regions are not negatively affected by the same-sex marriage law. However, individuals residing in the more liberal areas marry less after each law, which suggests that they learn over time about the alternative to marriage, registered partnership, and they increasingly switch to it from marriage. Finally, the rest of the Netherlands and the country as a whole experience no negative impact of the registered partnership law, but a decline in the marriage rate after the same-sex marriage law. This pattern is also consistent with the learning scenario, but also with a scenario where the same-sex marriage law did indeed alter the value of marriage. I will return to these possible alternatives in more detail later in this chapter and in chapter 3.

In the next section, I detail the empirical strategy used and the particular issue it raises. I next describe the data and the construction of the sample. The estimation results are presented in section 2.4, and section 2.5 concludes.

2.2 Empirical strategy

2.2.1 Specification of the hazard function

Duration models are regularly used in economics when the outcome measures the length of stay in a particular state. Common examples include the study of unemployment spells (Ham and Rea, 1987; Meyer, 1990; Addison and Portugal, 2003) or of strike duration (Kennan, 1985; Gunderson and Melino, 1990; Campolieti et al., 2005). Although the research question in this paper is not one of the usual suspects, a duration model is particularly useful for at least two reasons. First, this type of models can easily handle individuals who do not marry while under observation (censored spells), a useful feature given that the fraction of people who are not married is rather high in the Netherlands: on average, 33 percent of the marriage-age (18 year old and above) men and 25 percent of the marriage-age women had never married between 1995 and 2005. Second and most importantly, duration models allow for time-varying variables such as the enactment of the two laws.

Thus, I will use the framework of a duration model to analyze age at first marriage. This choice of outcome variable is justified by several arguments. First, the evolution of the marriage rate is largely driven by first marriages. As figure 2.1 shows, the variation in the number of marriages which are the first for either one of the spouses closely tracks the variation in the total number of marriages. Second, people who marry for the first time are on average 12–13 years younger that people who re-marry. This makes them both potentially more impressionable to changes in the perception of the institution of marriage, and more likely to be in the sample (I will return to this last aspect at the end of this section and again in section 2.3). Finally, previously-married people are likely to already have a formed opinion on marriage and may not be likely to change it just because of the enactment of the registered partnership law or of the same-sex marriage law. On the other hand, all of these also imply that the conclusions of this analysis cannot be directly extended to the rest of the population, especially those in older age groups.

Let T_i be the random variable representing the age at first marriage age of individual i, measured in full years, and define $h_i(t)$ as the probability that an individual marries between ages t and $t+1$ conditional on never having married by age t. The function $h_i(\cdot)$ is the discrete time hazard of marriage. As in Ham and Rea (1987), I will assume that it has the logit form[1]

$$h_i(t) = \frac{1}{1 + \exp\{-y_i(t)\}},$$

where

$$y_i(t) = \theta + X_i(t)'\beta_0 + p_1(t)\beta_1 + p_2(t)\beta_2 + \gamma(t).$$

In this equation, θ is a constant and $X_i(t)$ is a vector of potentially time-varying individual characteristics. The function $\gamma(t)$ represents the form of duration dependence, i.e. the common way age influences the probability of marriage for any given person. The actual form of duration dependence is driven by the data in the sense that $\gamma(\cdot)$ is the highest degree polynomial in $\ln(t)$ supported by the data. To determine the degree of $\gamma(\cdot)$, I keep adding higher order terms until they become insignificant, as suggested by Eberwein et al. (2002).

The coefficients of interest are β_1 and β_2, corresponding to $p_1(\cdot)$ and $p_2(\cdot)$,

[1] Alternatively, the hazard function can be assumed to have the extreme value (complementary log-log) form

$$h_i(t) = 1 - \exp\big\{-\exp\{y_i(t)\}\big\},$$

which can be derived from an underlying continuous-time proportional hazards model (Prentice and Gloeckler, 1978; Meyer, 1990). The results from both specifications are qualitatively and quantitatively similar, but the logit specification performed better by yielding higher values of the log-likelihood (hence lower values of the Bayesian information criterion). Moreover, the likelihood function with unobserved heterogeneity is much less well-behaved in the extreme value case.

two indicator variables for the period 1998–2000 and after 2001, respectively. They capture the change in the hazard (probability) of marriage during these two periods as compared to the period before the enactment of either law. There are a few reasons why these coefficients might not cleanly identify the effect of the two laws. First, the laws could have delayed effects as people's attitudes toward marriage might not change instantaneously and marriages are planned ahead of time. In this case, the coefficients would capture the short-term effect of the two laws. Second, the long-term effect of the registered partnership law cannot be distinguished from the short-term effect of the same-sex marriage law, simply because of the overlap in the period when both laws are in effect. Finally, the effect of the same-sex marriage law, both in the short term and in the long term, is unidentifiable if there is learning about registered partnership: it is virtually impossible to estimate the effect of this law in the absence of the registered partnership law.

2.2.2 Likelihood function and unobserved heterogeneity

The contribution to the likelihood function of person i, who is observed to marry at age a_i before the end of the study period, is the probability of marriage at age a_i:

$$P(T_i = a_i) = h_i(a_i) \cdot \prod_{t=1}^{a_i-1} [1 - h_i(t)] \tag{2.1a}$$

This can be interpreted as being the product of the conditional probability of marriage at age a_i and the probabilities of not having married at each age prior to a_i. If the person does not marry by the end of the observation period, then the observation is censored. Let a_i be the last observed age of the individual. The contribution to the likelihood function in this case is the survivor function, i.e. the probability of marriage at an age higher than a_i:

$$P(T_i > a_i) = [1 - h_i(a_i)] \cdot \prod_{t=1}^{a_i - 1} [1 - h_i(t)] \tag{2.1b}$$

The likelihood function for a sample of N individuals is obtained from the combination of equations (2.1a) and (2.1b):

$$L = \prod_{i=1}^{N} [P(T_i = a_i)]^{\delta_i} [P(T_i > a_i)]^{1-\delta_i}, \tag{2.2}$$

where δ_i equals one if person i is observed to marry and zero otherwise. As long as $h(\cdot)$ does not include an unobserved heterogeneity term, this likelihood can be estimated using standard programs for logit specifications.[2]

A duration model without unobserved heterogeneity, however, will suffer from a serious flaw that can be seen from the following example. Suppose people fall in two categories: some who are more likely to marry young and some who are more likely to delay marriage. Over time, individuals of the first type will marry and exit the sample at a faster rate than individuals of the second type and the sample will increasingly become a selected sample of people who are more likely to delay marriage. Failure to account for this selection could severely bias the

[2]In the case of a extreme value hazard, the likelihood function can be estimated using standard complementary log-log programs. In both cases, the only requirement is to create for each individual one observation for each period they are observed (for example, if an individual is observed for 5 years, there should be 5 observations in the data, one for each year).

estimated results. A natural extension is to introduce unobserved heterogeneity through the term θ (Ham and Rea, 1987). Following Heckman and Singer (1984), I will assume that θ follows a discrete distribution with K points of support $\theta_1, \ldots, \theta_K$ and corresponding probabilities $\pi_1, \ldots, \pi_K$ (where, obviously, $\sum_{k=1}^{K} \pi_k = 1$). In this case, the contribution to the likelihood function of an individual observed to marry at age a_i is

$$P(T_i = a_i) = \sum_{k=1}^{K} \left\{ \pi_k h_i(a_i; \theta_k) \prod_{t=1}^{a_i - 1} [1 - h_i(t; \theta_k)] \right\} \tag{2.3a}$$

and that of a censored observation is

$$P(T_i > a_i) = \sum_{k=1}^{K} \left\{ \pi_k \prod_{t=1}^{a_i} [1 - h_i(t; \theta_k)] \right\}, \tag{2.3b}$$

where

$$h_i(t; \theta_k) = \frac{1}{1 + \exp\{-y_i(t; \theta_k)\}}$$

and

$$y_i(t; \theta_k) = \theta_k + X_i(t)'\beta_0 + p_1(t)\beta_1 + p_2(t)\beta_2 + \gamma(t).$$

In this setup, initial conditions are a real concern. People become "at risk of marriage" when they turn 18, the legal age of marriage for both men and women. Following the same argument as above, a sample that includes persons who were observed for the first time (entered the sample) when they were older than 18 is a selected sample because people who are less likely to marry are over-represented.

The contribution of individual i to the likelihood function should be conditional on them not having married by age a_{0i}, their age at entry into the sample:

$$P(T_i = a_i | T_i \geq a_{0i}) = \frac{P(T_i = a_i)}{P(T_i \geq a_{0i})}, \tag{2.4a}$$

$$P(T_i > a_i | T_i \geq a_{0i}) = \frac{P(T_i > a_i)}{P(T_i \geq a_{0i})}. \tag{2.4b}$$

If the model does not include unobserved heterogeneity, the conditional probabilities in equation (2.4) depend only on the *observed* data from entry into the sample:

$$P(T_i = a_i | T_i \geq a_{0i}) = \frac{h_i(a_i) \cdot \prod_{t=1}^{a_i - 1} [1 - h_i(t)]}{\prod_{t=1}^{a_{0i} - 1} [1 - h_i(t)]} = h_i(a_i) \cdot \prod_{t=a_{0i}}^{a_i - 1} [1 - h_i(t)],$$

$$P(T_i > a_i | T_i \geq a_{0i}) = \frac{\prod_{t=1}^{a_i} [1 - h_i(t)]}{\prod_{t=1}^{a_{0i} - 1} [1 - h_i(t)]} = \prod_{t=a_{0i}}^{a_i} [1 - h_i(t)].$$

In the case of unobserved heterogeneity, however, an estimation based on the unconditional contributions (2.3a) and (2.3b) will lead to incorrect estimates unless $P(T_i \geq a_{0i})$ is somehow known. One possible approach is to make some additional assumptions about the distribution of unobserved heterogeneity in the period before entry into the sample (Ridder, 1984). However, the data on the evolution of the time-varying elements of X over this period is not always available. To avoid these issues, I restrict the analysis to a "flow sample" of individuals for whom the denominator in the equations (2.4) above is close to one. I return to this issue in the next section.

2.2.3 Marginal effects

One last aspect in the estimation is the calculation of marginal effects. This is relatively trivial in the specifications without unobserved heterogeneity, as the marginal effect is simply the marginal from the logit regressions. The issue is slightly more complicated in the specifications including unobserved heterogeneity. As we are interested in the effect of the two laws on the marriage rate, which is nothing but the hazard function $h_i(\cdot)$, I focus on the *average* marginal effect. For the parameter β_j, this effect is

$$AME_j = \mathbf{E}\left[\frac{\partial h(t)}{\partial \beta_j}\right].$$

The sample counterpart is

$$\widehat{AME}_j = \frac{1}{M}\sum_{i=1}^{N}\sum_{t=1}^{a_i}\frac{\partial \hat{h}_i(t_i)}{\partial \beta_j},$$

where N is the number of individuals in the sample and $M = \sum_{i=1}^{N} a_i$ is the total number of individual-years.

I approximate the derivatives in the formula above using numerical derivatives. If the parameter β_j refers to the continuous variable Z^j, the average marginal effect is approximated by the change in the estimated hazard from an infinitesimal change in the variable:

$$\widehat{AME}_j = \frac{1}{M}\sum_{i=1}^{N}\sum_{t=1}^{a_i}\frac{\hat{h}_i(t; Z_i^j + \Delta, Z_i^{-j}) - \hat{h}_i(t; Z_i^j, Z_i^{-j})}{\Delta},$$

where $Z = (X(t), p_1(t), p_2(t))$, Z_i^{-j} is the set of variables for individual i other than the variable Z^j, Δ is an arbitrarily small number, and $\hat{h}_i(\cdot)$ is the estimated hazard integrated over the distribution of unobserved heterogeneity:

$$\hat{h}_i(t; Z) = \sum_{k=1}^{K} \frac{\pi_k}{1 + \exp\left\{-\hat{y}_i(t; \hat{\theta}_k, Z)\right\}},$$

$$\hat{y}_i(t, \hat{\theta}_k, Z) = \hat{\theta}_k + X_i(t)'\hat{\beta}_0 + p_1(t)\hat{\beta}_1 + p_2(t)\hat{\beta}_2 + \hat{\gamma}(t)$$

In the case of a dummy variable Z^j, the average marginal effect is approximated by the change in the estimated hazard due to a move from a value of zero for Z^j to a value of one:

$$\widehat{AME}_j = \frac{1}{M} \sum_{i=1}^{N} \sum_{t=1}^{a_i} \left[\hat{h}_i(t; 1, Z_i^{-j}) - \hat{h}_i(t; 0, Z_i^{-j})\right].$$

This last formula needs to be applied cautiously when the dummy variable Z^j is part of a set of categorical dummy variables (e.g., ethnicity dummies or the two period dummies). In this case, the comparison should be made with respect to the omitted category, for which all the other dummy variables from the set are equal to zero. One way to ensure the correct calculation of marginal effects is to restrict the sample to the individual-years for which all the other corresponding categorical dummies are equal to zero (obviously, in this case M is the corresponding number of individual-years). For example, to calculate the marginal effect of $p_1(\cdot)$, the 1998–2000 period dummy, I restrict the sample to individual-years for the period 1995–2000, for which $p_2(\cdot)$, the 2001–2005 period dummy, is equal to zero.

2.3 Data

I create the data using ten waves of the restricted version of the Dutch Labor Force Survey (1996–2005) and the January 2006 snapshot of the highly confidential Dutch Municipal Records. The Labor Force Survey (*Enquête Beroepsbevolking*, or EBB) is an annual cross-sectional random survey of the population 15 years of age and older. It includes information on educational attainment, ethnicity, employment and other demographic and labor market characteristics at the time of the interview. In addition, the restricted version provides an identification number that can be used to match the individuals to other data sets maintained by Statistics Netherlands. The ten waves of the survey combined yield information on almost 950,000 individuals, or approximately six percent of the average population over the period 1995–2005.

The Dutch Municipal Records (*Gemeentelijke Basis Administratie*) include detailed information on changes in the marital status and residence of the entire resident population for the period between January 1, 1995 and January 1, 2006. Statistics Netherlands made available to me information on the individuals included in the ten waves of the Labor Force Survey. Using the identification number, these individuals are matched to their full marriage and residence history, both before and after their survey interview. The result is a longitudinal data set for the period 1995–2005 including information on ethnicity, marital status and residence over the whole period, and educational attainment and school enrollment at the time of the Labor Force Survey interview. Finally, the data is augmented with the yearly unemployment rate at the regional level.[3]

The variables included in the analysis measure the attractiveness of an individual on the marriage market (age, education and ethnicity), the thickness of the market (location and ethnicity), and business cycle fluctuations (the regional unemployment rate). As long as the variables from the Labor Force Survey are time-invariant, their inclusion in the final longitudinal data set will not cause problems. However, there are two cases when this does not hold. First, about 15 percent of the sample was still enrolled in some form of education at the time of the survey (approximately 9 percent full-time and 6 percent part-time). To increase the probability that the highest educational level reported, either completed or in progress, is the actual highest level attained, I restrict the initial sample to individuals who were at least 20 years of age in the year of the interview.[4]

Second, some variables such as age, residence or the unemployment rate can change continuously. These variables need to be aggregated because I am using a discrete-time approach. I do this on a calendar year basis rather than according to the birth date of individuals for two reasons. First, it would be practically impossible to measure the regional unemployment rate on a different scale than the calendar year. Second, there is a strong seasonal pattern in marriages that suggests that people make marriage decisions based on the calendar year rather than their birth

[3]The Netherlands is divided into twelve provinces: Drenthe, Flevoland, Friesland, Gelderland, Groningen, Limburg, Noord-Brabant, Noord-Holland, Overijssel, Utrecht, Zeeland and Zuid-Holland.

[4]A small fraction of the individuals still in school were enrolled in a lower educational level than their highest level completed (for example, persons with a college degree in science enrolled in professional economics or business courses). The highest of the two educational levels was used for these cases. The Dutch education system, the flows among different educational attainments and the grouping of educational attainments used are shown in figure 2.A.

dates. For example, about 61 percent of marriages between 1995 and 2005 were performed during the "summer months" (from May to September).

Therefore, I will measure age as the age in full years at the *end* of the calendar year, so that 17 year-old persons at the beginning of the year who get married by the end of the year (after turning 18) are included in the sample. In contrast, I will consider the residence at the *beginning* of the year, under the assumption that most marriage decisions are made in advance and thus the location at the beginning of the year is likely to influence the marital decision. Finally, the regional unemployment rate is the *average* over the calendar year provided by Statistics Netherlands.

Recall that, as explained in section 2.2, using data from all the individuals can be problematic due to initial conditions. Ideally, I would discard all the individuals who were older than 18 when first observed, but this can cause more problems. The average age at first marriage increased during this period from 29.6 to 32.4 years for men and from 27.4 to 29.7 years for women. If the sample includes only individuals who turned 18 in or after 1995, the oldest person in the sample would be 29 years old in 2005. This is below the average age at first marriage for both men and women, so the number of observed marriages will be low and so will be the power of the estimation. The compromise is to include individuals in an age group that accounts for a small fraction of the total number of marriages, i.e., keep individuals such that $P(T \geq a_{0i})$ in equations (2.4) is close to one. Men 18–24 years old account for about 10 percent of first marriages, so I restrict the sample of men to individuals 18–24 years old in 1995 or when first observed. In other words, men enter the sample in the calendar year they turn 18 or in 1995 if they were never-married at the beginning of the year and were 18–24 year old. I restrict the sample in similar fashion to women 18–22 years old, who also account for about 10 percent of first marriages.

Finally, I conduct the analysis separately by gender because men and women seem to have different attitudes toward marriage. For example, men tend to marry later than women: between 1995 and 2005, the average age at first marriage in the Netherlands is consistently higher by about 2 years for men than for women (29.5 to 32.5 years of age for men, compared to 27.5 to 30 years of age for women). Also, the fraction of women who marry previously-married men is higher than the fraction of men who marry previously-married women. During the sample period, 10.3 to 11.6 percent of all women marrying for the first time and 8.6 to 9.7 percent of all men marrying for the first time had a partner who had been married.

The Municipal Records provide very detailed data and, like many administrative data sets, reduce or eliminate the measurement error in the variables. However, they also present several disadvantages. First, no distinction is made between same-sex and different-sex marriages—they are both coded as "marriage". Second, due to the sensitive nature of the data, I have no data on individuals who did not participate in the Labor Force Survey. Thus, there is no information on the spouse of an individual unless he or she also participated in the Labor Force Survey. Finally, the coding of addresses changed over time and was aggregated at the street address level since 2003, making it practically impossible to identify the spouse of an individual even if the information for that person were made available. As a result, I am unable to separate individuals entering a different-sex marriage from those entering a same-sex marriage. This induces an upward bias in the estimate of the different-sex marriage rate after 2001 and is problematic only if I find that the effect of the same-sex marriage law is positive (or insignificant). I further address this issue in chapter 3 by using aggregate data, where I can distinguish between different-sex and same-sex marriages.

A second disadvantage of the data is that the information on labor market outcomes applies to only one point in time. The highest level of education attained is plausibly constant over the eleven years of the study period, but this is less likely to be the case with employment status, industry or occupation. Therefore, I cannot use any of these variables from the Labor Force Survey, except for what is implicitly included in the regional unemployment rate.

The final sample includes 70,718 men and 53,803 women. The higher number of men is simply due to the selection process. In addition, since women tend to marry younger, there will be more never-married men then women for every single age group. Descriptive statistics for the sample, separately by sex, are listed in table 2.1. All the statistics and the subsequent analysis use the sample weights provided in the Labor Force Survey.[5]

The youth of the sample is evident from the first statistic. The average age at first marriage is around 27 years for men and 25 years for women, significantly lower than the average age at first marriage in the entire population during this period. Censored individuals (persons who do not marry until the end of 2005) have a similar age distribution, with only slightly higher average age. This is not surprising given

[5]The sample weights refer to the year of the interview. Under the assumption that the structure of the population did not change significantly during the period under study, I rescale the weights to represent the probability of interview relative to the entire sample of ten waves of the Labor Force Survey.

that figure 2.2, which plots the Kaplan-Meier estimates of the probability of being single by age (the survival function), shows that about half of the oldest individuals had not married by 2005. Overall, only 26.29 percent of men and 33.30 percent of women married by the end of 2005.

The main difference in the two subsamples is due to the initial age restriction. While the bulk of the individuals in the sample were born between 1970 and 1984, the age distribution is different for men because the sample includes older people. The 1970–1974 birth cohort consists of men who were between 21–24 years old in 1994, but only 21 and 22 year-old women, which explains why it accounts for 41.04 percent of men and only 23.79 percent of women.

Note also that the distribution of education is skewed toward higher levels of education. Only about 28 percent of men and 23 percent of women have at most a high school (general secondary) degree. Almost 40 percent of both men and women have some post-secondary vocational training, while higher vocational training and college degrees account for 23.16 and 9.05 percent of men and 28.39 and 8.67 percent of women, respectively. Approximately 83 percent of the sample are natives and almost 8 percent are Western immigrants, i.e. people from Europe (except Turkey), North America, Oceania (including Australia and New Zealand), Japan, and Indonesia. Immigrants from potentially more conservative areas such as the predominantly Muslim countries Turkey and Morocco, or Dutch current and former territories Aruba and Suriname, account for about 6 percent of both men and women.

The Netherlands is one of the most urban countries in Europe, as approximately 63 percent of the sample resides in urban areas. However, there is geographic heterogeneity with respect to people's attitudes toward marriage and cohabitation. One area that I will focus on comprises the four largest cities (Amsterdam, the Hague, Rotterdam and Utrecht) and exhibits low fertility and marriage rates and high non-marital birth and divorce rates, as well as low frequency of church-going. Another area of interest is the so-called Dutch Bible belt (*De Bijbelgordel*), a set of municipalities characterized by relatively high church participation, high fertility rates, low cohabitation and divorce rates, low non-marital birth rates, and high frequency of church-going (Sobotka and Adigüzel, 2002; Statistics Netherlands, 2003). These two areas are identified on the map in figure 2.3.

Summary statistics for each region are shown in table 2.2. At exit from the sample, i.e. when marrying or at the end of 2005 for individuals still single, the fraction of people in the four largest cities is much larger than the fraction of people in the Bible belt—15.16 percent compared to 2.71 percent for men and 16.26 percent compared to 2.73 percent for women. Yet the fraction of marriages is disproportionately high in the Bible belt: 4.39 percent of first-marrying men and 4.42 percent of first-marrying women reside in the Bible belt, while only 12.11 percent of first-marrying men and 12.03 perecent of first-marrying women reside in the four largest cities.[6] The difference between the two regions is also evident from the percentage of their residents who marry, which is more than twice higher in the Bible belt (42.60 percent against 21 percent for men and 53.93 percent against 12.03 percent

[6]One explanation for this is that single people tend to move to cities, where the marriage markets are thicker, but (married) couples tend to move outside of the cities, where housing is cheaper (Gautier et al., 2005). Thus, the number of single individuals in the cities tends to be higher than in the more rural areas that comprise the Bible belt.

for women).

Based on this evidence and the patterns of church-going, fertility and divorce mentioned above, I will consider the four largest cities to represent mostly liberal individuals and the Bible belt municipalities to include mostly conservative people. This distinction can be used to determine the impact of the two laws on individuals based on the individuals' degree of conservatism.

2.4 Results

2.4.1 Baseline regressions

I first estimate the discrete-time duration model in section 2.2 without unobserved heterogeneity. Recall that the coefficients represent the effect of the corresponding variables on the probability that a given individual marries during the calendar year. In other words, they represent the effect of the variables on the marriage rate in a given year. Recall also that the sample of men and of women are not the opposite sides of an accounting equation: the spouses of the men in the sample who marry are not necessarily in the sample of women and vice versa. As a result, there are differences in the coefficient estimates between the two genders and in the estimated marriage rates.

The results from the baseline regressions for men are provided in table 2.3. To show the importance of taking individual-level characteristics and aggregate factors into account, I start with a simple set of covariates and then discuss the change in the coefficients as new variables are added. Some of the variables are common to all the specifications. First, I always include the two dummy variables that represent each of the two new law regimes. I also add a linear trend to capture changes in marriage behavior not accounted for by the other explanatory variables, such as the increased secularization of the Dutch society. Thus, the coefficients on the two period dummies should be interpreted as deviations from the long-term trend in the marriage rate in comparison to the period before 1998. Also present in all specifications is a sixth-degree polynomial in the natural log of age minus 17 (since the minimum legal age of marriage is 18) to account for duration dependence. Finally, a set of dummy variables representing five-year birth cohorts is included in all but the first specification.

The specification in columns 1 and 5 controls only for duration dependence and is similar to the type of aggregate analysis present in the media that takes into account the long-term trend in the marriage rate. For both sexes there is a secular decline in their first-marriage rate, as shown by the estimated negative trend terms. In the sample of men, both laws were followed by a higher propensity to marry, which is basically Badgett's (2004b) argument that the marriage rate did not fall after the enactment of same-sex marriage laws once the downward trend is taken into account. Women, however, are estimated to marry less after the enactment of the same-sex marriage law.

The difference in the pattern of the marriage rate still holds after adding controls for demographic characteristics (education, ethnicity and birth cohort) as shown in columns 2 and 6: men are again estimated to marry more after both laws, while women marry less after the same-sex marriage law. This is not surprising because none of the variables added changes over time and therefore they should have no impact on the coefficients of time-varying variables such as the trend or the two law regime dummies. What does emerge from this exercise is that the relationship between education and marriage is also different between men and women. The estimates show the propensity to marry relative to the omitted category, which is an intermediate level of education between the general secondary and higher professional levels.[7] They suggest an almost inverted U-shaped pattern for men, with the least and the most educated having the lowest propensities to marry. For women, on the other hand, the relationship is more skewed, lower educated women having higher propensities to marry but higher educated women being much less likely to marry. These patterns are consistent with a scenario of female hypergamy (women "marry up") and male hypogamy (men "marry down"), such that lower educated women have a better chance of finding a match, and higher educated people have lower propensities to marry, in general.

Ethnicity influences the decision to marry in an expected way. Compared to natives, both men and women with a Western background or from Suriname and Aruba are less likely to marry in the Netherlands, while people from the more conservative Turkey and Morocco are much more likely to marry. Men from other countries seem to have higher tendency to marry, while women they seem to be similar to natives in this regard.

Finally, the addition of demographic characteristics improves the fit of the

[7]Keep in mind that the numbers are not comparable across equations because logit coefficients are measured with respect to the standard deviation of the dependent variable.

model. The values of the log-likelihood can be directly compared since the specifications are nested. For men, the log-likelihood in column 2 is higher by 0.88% than the one in column 1 (from −9973.84 to −9885.94), while for women the log-likelihood in column 6 is is 2.11% higher than the one in column 5 (from −8966.37 to −8777.34).[8]

Columns 3 and 7 add the regional unemployment rate, an indicator of the business cycle that varies both over time and cross-sectionally. Higher unemployment is associated with lower marriage rates for both men and women, which could be due to couples delaying marriage during economic downturns. Its inclusion does not affect the estimated effects of education and ethnicity, which are the time-invariant variables. It does, however, change the estimated trend in the marriage rate and the period dummies. Unlike in the first two columns, column 3 shows that both men and women are estimated to marry less after each law, while the long-term trend becomes less negative.

Finally, columns 4 and 8 add information on the residence of individual which, as mentioned before, can and does change over time. Living in an urban area is associated with a lower propensity to marry, even lower if the individual resides in one of the four largest cities. In contrast, individuals who live in the Bible belt are estimated to marry significantly more. The effects of education and ethnicity are virtually unaffected, as are the effect of the unemployment rate and the long-term trend. The estimates of the period dummies change yet again, becoming even

[8]The values of the log-likelihood are divided by 1,000 throughout the paper in order to improve readability.

more negative. This suggests that the marriage patterns were influenced by the migration patterns described in section 2.3. Over time, people are more likely to move to urban areas or to one of the largest cities in order to find a better match. People (already) living in these urban areas are less likely to marry because of lower religiosity (Sobotka and Adigüzel, 2002). These two effects offset each other and lead to overestimated marriage rates in the later periods represented by the two period dummies. Once controls for the migration patterns are included, the marriage rate is estimated to fall by even more after the enactment of two laws. Note also that including the location variables increases the likelihood function by much more than the inclusion of the regional unemployment rate (from -9884.76 to -9846.89 for men and from -8776.94 to -8732.63 for women), one more indication of the importance of these variables.

Although the form of duration dependence is rather flexible, it might be argued that a non-parametric approach could be more accurate. In order to make sure that the polynomial form chosen is a good approximation of the actual form of duration dependence, I also ran the specification in columns 4 and 8 by replacing the polynomial in log-age with a set of age dummies. The results, listed in table 2.A (for comparison, columns 1 and 3 are the same as columns 4 and 8 in table 2.3), show that most of the coefficients are virtually unchanged. A comparison of the fit of the model with the previous case can be performed using the Bayesian (or Schwartz) information criterion (for which lower values indicate better fit). The BIC is almost the same for the models in columns 1 and 2, and 3 and 4, respectively, suggesting that the polynomial form is a good approximation for the actual form of duration dependence.

The magnitude of the estimated effects on the marriage rate is relatively large, considering the marginal effects calculated for the specification in column 4. The marriage rate of men over the 1995–2005 period is, on average, 2.99 percent and is estimated to fall by 0.06 percentage points after the registered partnership law and by 0.16 percentage points after the same-sex marriage law, compared to a long-term downward trend of 0.05 percentage points per year. In the case of women, the average marriage rate is 4.07 percent and the decline is 0.14 percentage points and 0.65 percentage points, respectively, while the downward trend is 0.05 percentage points per year.

2.4.2 Unobserved heterogeneity

As mentioned in section 2.2, a duration model without unobserved heterogeneity can yield severely biased estimates. I assume that the unobserved heterogeneity term follows a discrete distribution with two mass points.[9] I use the parametric form of duration dependence for because models with both nonparametric unobserved heterogeneity and nonparametric duration dependence have been shown to suffer from severe bias (Baker and Melino, 2000). Moreover, as seen in the previous section, using non-parametric duration dependence does not significantly change the results.

Columns 2 and 4 in table 2.4 report the results from the models including

[9] Previous research found that two mass points are in general sufficient (Narendranathan and Stewart, 1993; Ham and Rea, 1987). A distribution with three mass points produced basically unchanged estimates and increased the log-likelihood function only by about 0.5%, but the likelihood function converged with much more difficulty.

unobserved heterogeneity. For comparison, columns 1 and 3 report the specifications without unobserved heterogeneity (columns 4 and 8 in table 2.3). For both men and women, the registered partnership law (period 1) is now associated with an increase in the marriage rate, while the same-sex law (period 2) is followed by a decline. The reduction in the marriage rate after 2001 is, however, much smaller than estimated by the model without unobserved heterogeneity (columns 4 and 5).

The inclusion of unobserved heterogeneity leaves some of the estimates virtually unchanged, such as the regional unemployment rate or the long-term trend. It does, however, alter some of the other coefficients. For instance, the relationship between education and marriage changes slightly. The pattern remains more or less the same for men, but it becomes steeper and more linear for women. This lends even more support to the idea that women tend to marry up, making it more difficult for higher educated women to find a match. Ethnicity seems to influence the propensity to marry in similar ways for both men and women, unlike the results with no unobserved heterogeneity. In particular, women of non-Western background seem to marry more than natives, whereas before they were estimated to marry less. Residence seems to influence the marriage rate even more than before: living in an urban area or in one of the four largest cities is associated with even lower propensities to marry. In contrast, residence in the Bible belt increases the probability of marriage by even more.

The magnitude of the marginal effects is much smaller than in the case when unobserved heterogeneity was ignored. The marriage rate of men increases by 0.02 percentage points after the registered partnership law and falls by 0.13 percentage points after the same-sex marriage law, while the downward trend is 0.12 percentage points per year. Women experience an increase in the marriage rate after 1998 of 0.13 percentage points, but a larger decline after the same-sex marriage law at -0.59 percentage points. The downward trend in their marriage rate is estimated at 0.09 percentage points per year.

Finally, controlling for unobserved heterogeneity improves the fit of the regression. The BIC in the models with unobserved heterogeneity is smaller than the BIC in the corresponding model with unobserved heterogeneity.

2.4.3 Heterogeneous effects by location

The results presented in the previous section suggest that the residence of individuals is related to the timing of marriage. In this section, I analyze in more detail this relationship by modifying the specification of the hazard function, still using the logit form

$$h_i(t) = \frac{1}{1 + \exp\{-y_i(t)\}},$$

but where

$$\begin{aligned} y_i(t) = \theta + X_i(t)'\beta_0 + p_1(t)BB_i(t)\beta_1^{BB} + p_2(t)BB_i(t)\beta_2^{BB} + p_1(t)LC_i(t)\beta_1^{LC} \\ + p_2(t)LC_i(t)\beta_2^{LC} + p_1(t)[1 - LC_i(t) - BB_i(t)]\beta_1^{OTH} \\ + p_2(t)[1 - LC_i(t) - BB_i(t)]\beta_2^{OTH} + \gamma(t), \end{aligned}$$

$X(\cdot)$ includes the same set of variables as before (including location in a large city or in the Bible belt), and $BB_i(t)$ and $LC_i(t)$ are dummy variables for individual i residing in the Bible belt or in one of the four largest cities in year at the beginning of year t. In this specification, β_1^{BB} and β_2^{BB} represent the change in the propensity to marry among individuals living in the Bible belt after the enactment of the registered partnership law and of the same-sex marriage law, respectively. Similarly, β_1^{LC} and β_2^{LC} represent the change in the marriage rate among individuals residing in one of the four largest cities following the legalization of registered partnership and of same-sex marriage, respectively. Finally, β_1^{OTH} and β_2^{OTH} capture similar changes for individuals living outside the Bible belt or the four largest cities.

I start again by estimating a model with no unobserved heterogeneity. The results are shown in columns 1 and 3 in table 2.5. As before, residence in the Bible belt is associated with a higher probability of marriage and residence in the four largest cities is associated with a lower marriage rate as compared to the rest of the Netherlands, for both men and women. The registered partnership law was followed by reductions in the marriage rate in both areas. However, a one-sided F-test indicates that the decline in the marriage rate between 1998 and 2001 in the four largest cities is significantly larger than the decline in the Bible belt (F-statistic = 19.67 for men and 19.02 for women). In contrast, the reduction in the marriage rate after the same-sex marriage law was larger in the Bible belt than in the four largest cities (F-statistic = 359.74 for men and 32.84 for women). Within each region, the marriage rate fell even more after the introduction of same-sex marriage, with the exception of men in the four largest cities.[10] Finally, individuals living outside these two areas seem not to be affected by the introduction of registered partnership, but they do marry less after the legalization of same-sex marriage.

Columns 2 and 4 show the results from the specification including unobserved heterogeneity. While location still has the same overall effect on the marriage rate—higher in the Bible belt and lower in the four largest cities—its evolution after the enactment of the two laws is now significantly different. The marriage rate in the Bible belt fell after the registered partnership law came into effect, by 0.15 percentage points for men and 0.19 percentage points for women. However, it *increased* (relative to the long-term downward trend) after the enactment of same-sex marriage law (the F-statistic for the test of equality of the coefficients is 66.82 for men and 3.47 for women, rejecting the hypothesis in both cases). People residing in the four largest cities also married less after the registered partnership law by approximately -0.60 percentage points for both men and women, but they married even less after the enactment of the same-sex marriage law (F-statistic = 183.82 for men and 5604.64 for women, rejecting the hypothesis that the estimate for period 2 is the same as the estimate for period 1). Moreover, the initial decline in the four largest cities is larger than in the Bible belt, with an F-statistic of 130.92 for men and 47.38 for women. Finally, both men and women residing outside of the Bible belt or the four largest cities married more after the registered partnership law (by 0.12 percentage points for men and 0.25 for women) and less after the same-sex

[10]The F-statistics for the one-sided tests comparing the estimated coefficients for period 1 and period 2 are: 1044.35 (men, Bible belt), 1885.10 (women, Bible belt), 0.005 (men, four largest cities) and 2251.92 (women, four largest cities). In all cases but the last they fail to reject the hypothesis that the decline in the marriage rate after 2001 is larger than during 1998–2000.

marriage law (by 0.03 and 0.33 percentage points, respectively), consistent with the results from the previous section.

These results indicate that there is variation in the response to the enactment of the two laws. Individuals residing in the more conservative municipalities included in the Bible belt seem to have been affected by the registered partnership law, but their marriage rate recovers after 2001. A possible explanation for this pattern is that these individuals had strong beliefs about the institution of marriage that were not easily changed by the legalization of same-sex marriage.

On the other hand, individuals living in the more liberal four largest cities marry relatively less after the enactment of the registered partnership law and even less after the enactment of the same-sex marriage law. The first effect may be driven by the existence of more individuals on the margin (who are swayed away from marriage toward registered partnership) in the four largest cities than in the Bible belt. The second effect is more intriguing, as one would expect people living in more liberal areas to be unaffected by the same-sex marriage law. However, recall that it is practically impossible to separately identify the influence of the two laws separately on marriage decisions because of their overlap. Figure 2.4 suggests that there might have been a learning process involved, as the number of new registered partnerships contracted by different-sex couples increased in each year after 2000. This could explain the fall in the propensity to marry after 2001, as more couples consider this alternative to marriage.

Finally, individuals residing outside of the Bible belt or the four largest cities tend to marry more after the enactment of the registered partnership law, but less after 2001, when the same-sex marriage law was enacted. This is a similar pattern to the one found in the previous section for the overall marriage rate and could be explained under two scenarios. First, the registered partnership law might have a positive effect on the marriage rate in the short run, but as people learn about the new institution over time, there is a gradual shift from marriage to registered partnership and the effect of the registered partnership law becomes negative in the long run. Second, the same-sex marriage law could have a negative effect on the marriage rate that offsets any short- and long-term positive effects of the registered partnership law. However, as described in section 2.2.1, it is practically impossible to disentangle these two effects.

2.5 Conclusion

In this chapter, I analyze the validity of the claim that the institution of marriage is negatively affected when opened to same-sex couples. I focus on the Netherlands, the first country to legalize same-sex marriage in 2001, which also introduced in 1998 an alternative institution, identical in almost all ways to marriage—registered partnership. I study the marriage behavior at the individual level using a discrete-time duration model, with and without unobserved heterogeneity. The estimates indicate that the overall marriage rate increases after the enactment of the registered partnership law, only to decline after the same-sex marriage law comes into effect. In other words, the introduction of same-sex registered partnership does not affect marriage negatively. It is impossible to establish, however, whether the decline in the marriage rate after 2001 is due to the legalization of same-sex marriage or to learning about registered partnereship. I will return to this issue in the next chapter, where I will also present some more (indirect) evidence.

I also find that the results differ by location, where the residence of individuals is considered as an indicator of their liberalism. Individuals in the more liberal area (the four largest cities) seem to marry less after each of the two laws, which is consistent with them learning about registered partnership over time and switching to the new institution. In contrast, individuals in the more conservative areas (the Bible belt) marry less after the introduction of registered partnership but their marriage rate increases after the legalization of same-sex marriage, returning to levels similar to the period before 1998.

One concern with the method used in this chapter is that it considers the period before the enactment of both laws as the counterfactual. If the set of variables $X_i(t)$ is large enough to control for any potential factor that could have affected marriage patterns after 1998, then the results in section 2.4 are unbiased. In chapter 3, I turn to aggregate data and I construct a more accurate counterfactual which will allow me to assess better the accuracy of the results in this chapter.

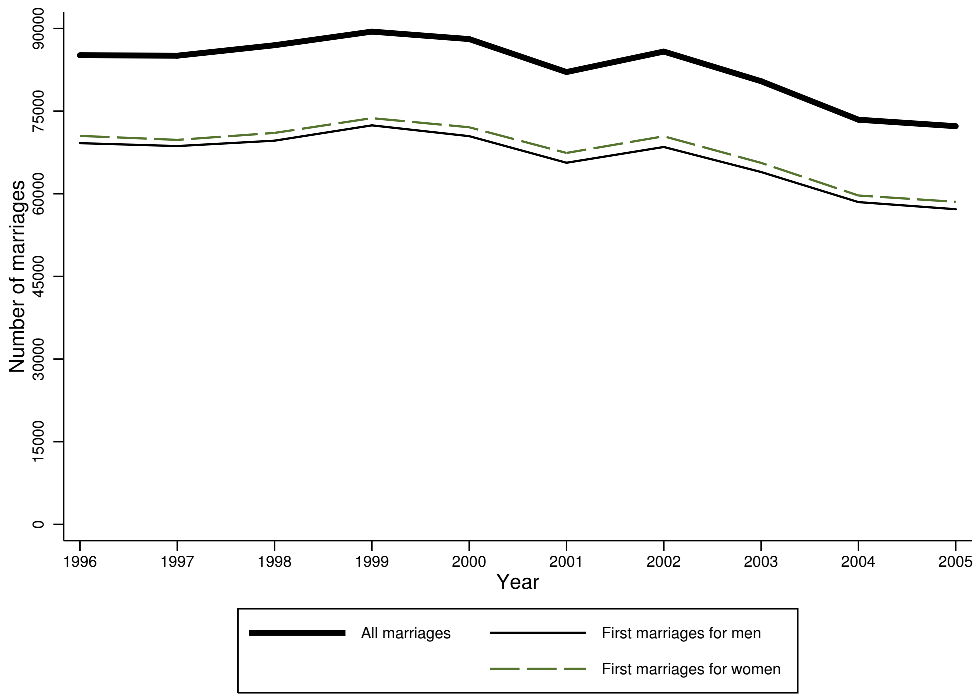

Figure 2.1: The evolution of all marriages and first marriages for one of the spouses

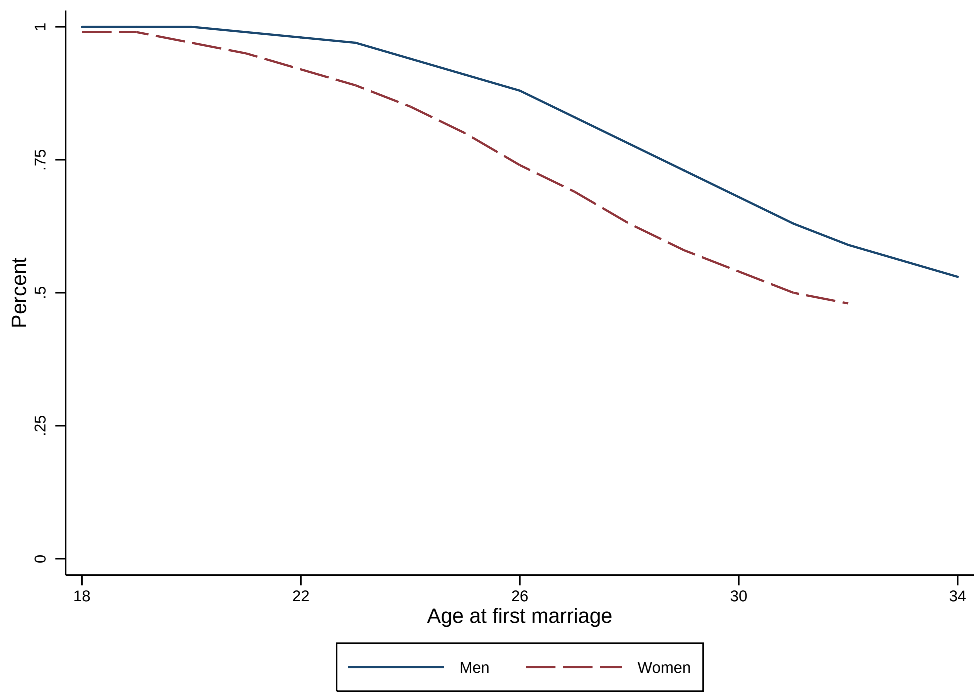

Figure 2.2: Kaplan-Meier estimates of the survival function (the probability of still being single by age)

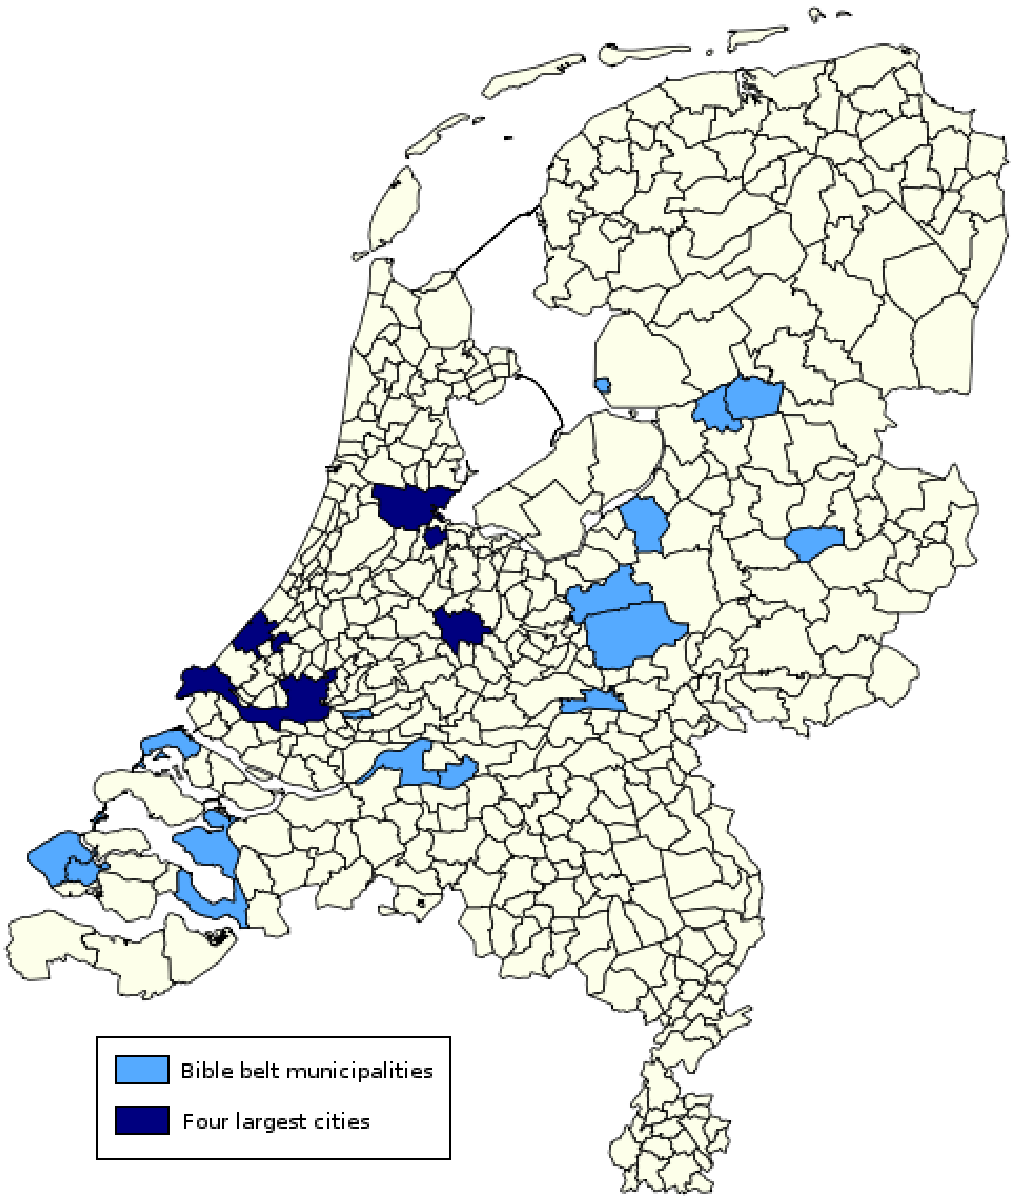

Figure 2.3: The four largest cities and the Bible-belt municipalities.

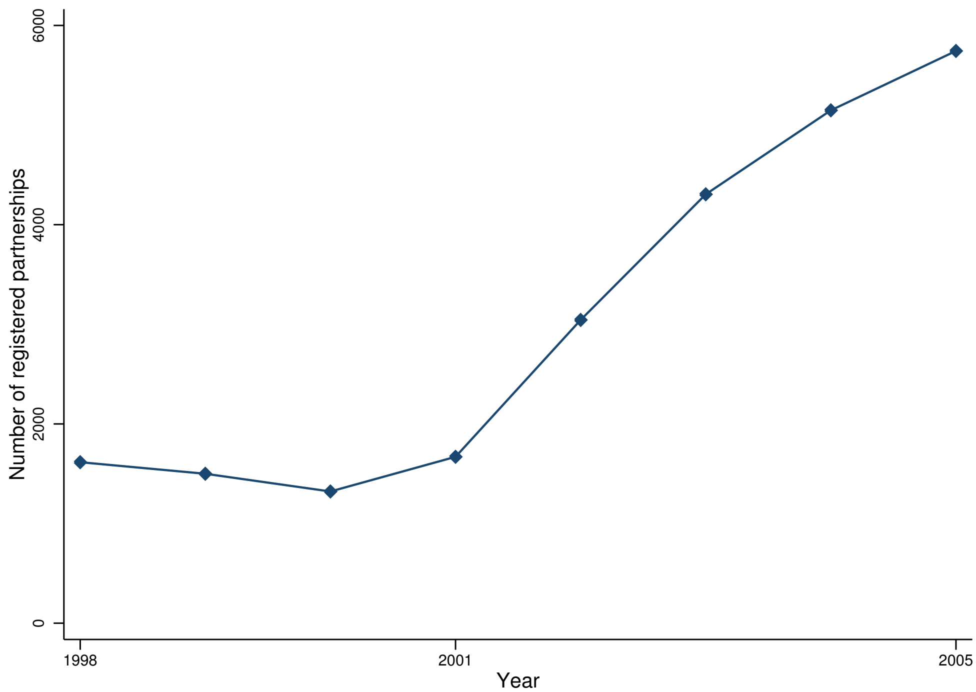

Figure 2.4: The number of new different-sex registered partnerships in the Netherlands, 1998–2005

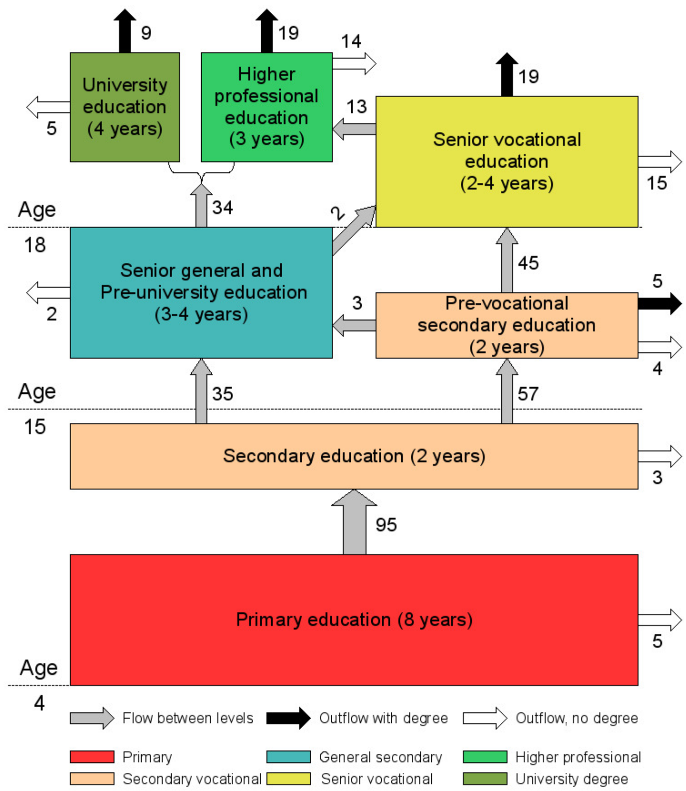

Note: Numbers next to arrows represent percentages of a cohort.

Source: Dutch Ministry of Education and Science (2003)

Figure 2.A: The education system in the Netherlands and the definition of educational attainment used, 2002

Table 2.1: Summary statistics

	Men (%)	Women (%)
First marriages		
Percent	26.29	33.30
Average age (years)	27.37	25.25
	(2.97)	(2.96)
Censored observations		
Percent	73.71	66.70
Average age (years)	28.19	26.96
	(3.70)	(3.23)
Birth cohort		
1970–74	41.04	23.79
1975–79	39.54	51.19
1980–84	18.43	23.79
1985–89	0.99	1.23
Education		
Primary education	4.24	3.03
Secondary vocational	16.83	12.95
General secondary	6.83	7.10
Senior vocational	39.88	39.85
Higher professional	23.16	28.39
College	9.05	8.67
Ethnicity		
Natives	83.11	82.84
Western immigrants	7.71	7.80
Turks/Moroccans	3.21	3.43
Surinamese/Arubans	3.01	3.26
Other non-Western immigrants	2.97	2.66
Urban area	62.86	63.80
Number of individuals	70,718	53,883

Note: All statistics weighted using sample weights.

Table 2.2: Summary statistics, four largest cities and the Dutch Bible belt

	Four largest cities (%)	Bible belt (%)
	Men ($N = 70,718$)	
Location at exit from sample	15.16	2.71
Percent of total marriages	12.11	4.39
Percent of residents marrying	21.00	42.60
	Women ($N = 53,803$)	
Location at exit from sample	16.26	2.73
Percent of total marriages	12.03	4.42
Percent of residents marrying	24.64	53.93

Note: All statistics weighted using sample weights.

Table 2.3: Discrete-time duration model for age at first marriage

	Men ($N = 70,718$)				Women ($N = 53,803$)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Period 1 (1998–2000)	0.094*	0.095*	−0.031*	−0.040*	0.031*	0.039*	−0.034*	−0.052*
	(0.003)	(0.003)	(0.004)	(0.004)	(0.003)	(0.003)	(0.004)	(0.004)
Period 2 (2001–2005)	0.070*	0.074*	−0.089*	−0.101*	−0.127*	−0.116*	−0.210*	−0.234*
	(0.004)	(0.004)	(0.006)	(0.006)	(0.005)	(0.005)	(0.006)	(0.006)
Linear trend (1995=0)	−0.057*	−0.051*	−0.035*	−0.034*	−0.029*	−0.029*	−0.020*	−0.017*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Education (omitted category: Senior vocational)								
Primary education		−0.272*	−0.271*	−0.258*		−0.027*	−0.027*	0.011*
		(0.004)	(0.004)	(0.004)		(0.004)	(0.004)	(0.004)
Secondary vocational		−0.055*	−0.055*	−0.053*		0.094*	0.094*	0.098*
		(0.002)	(0.002)	(0.002)		(0.002)	(0.002)	(0.002)
General secondary		−0.422*	−0.423*	−0.342*		−0.344*	−0.345*	−0.265*
		(0.003)	(0.003)	(0.003)		(0.003)	(0.003)	(0.003)
Higher professional		−0.220*	−0.221*	−0.169*		−0.449*	−0.450*	−0.397*
		(0.002)	(0.002)	(0.002)		(0.002)	(0.002)	(0.002)
College		−0.289*	−0.292*	−0.172*		−0.786*	−0.786*	−0.647*
		(0.002)	(0.002)	(0.002)		(0.003)	(0.003)	(0.003)
Ethnicity (omitted category: Natives)								
Western immigrants		−0.201*	−0.202*	−0.157*		−0.255*	−0.255*	−0.194*
		(0.003)	(0.003)	(0.003)		(0.003)	(0.003)	(0.003)
Turks/Moroccans		1.156*	1.153*	1.270*		1.469*	1.467*	1.608*
		(0.003)	(0.003)	(0.003)		(0.003)	(0.003)	(0.003)
Surinamese/Arubans		−0.289*	−0.288*	−0.155*		−0.519*	−0.518*	−0.330*
		(0.005)	(0.005)	(0.005)		(0.005)	(0.005)	(0.005)
Other non-Western immigrants		0.050*	0.050*	0.141*		−0.130*	−0.130*	−0.009*
		(0.004)	(0.004)	(0.004)		(0.005)	(0.005)	(0.005)
Regional unemployment rate			−0.029*	−0.032*			−0.017*	−0.022*
			(0.001)	(0.001)			(0.001)	(0.001)
Urban indicator				−0.171*				−0.238*
				(0.001)				(0.002)
Bible belt				0.665*				0.661*
				(0.003)				(0.004)
Four largest cities				−0.248*				−0.268*
				(0.002)				(0.002)
Log-likelihood / 1000	−9973.84	−9885.94	−9884.76	−9846.89	−8966.37	−8777.34	−8776.94	−8732.63
BIC / 1000	19947.86	19772.28	19769.93	19694.24	17932.91	17555.08	17554.29	17465.73

Notes: All specifications are weighted using sample weights and include 5-year birth cohort dummies (except for columns 1 and 5) and a sixth-degree polynomial in $\ln(age - 17)$. Standard errors in brackets. BIC is the Schwartz (Bayesian) Information Criterion. Starred coefficients are significant at the 1 percent level.

Table 2.4: Discrete-time duration model for age at first marriage, with and without unobserved heterogeneity

	Men ($N = 70,718$)		Women ($N = 53,803$)	
	(1)	(2)	(3)	(4)
Period 1 (1998–2000)	−0.040*	0.005	−0.052*	0.028*
	(0.004)	(0.004)	(0.004)	(0.004)
Period 2 (2001–2005)	−0.101*	−0.038*	−0.234*	−0.130*
	(0.006)	(0.006)	(0.006)	(0.006)
Linear trend (1995=0)	−0.034*	−0.035*	−0.017*	−0.020*
	(0.001)	(0.001)	(0.001)	(0.001)
Education (omitted category: Senior vocational)				
Primary education	−0.258*	−0.127*	0.011*	0.545*
	(0.004)	(0.006)	(0.004)	(0.007)
Secondary vocational	−0.053*	0.035*	0.098*	0.400*
	(0.002)	(0.003)	(0.002)	(0.003)
General secondary	−0.342*	−0.501*	−0.265*	−0.418*
	(0.003)	(0.004)	(0.003)	(0.004)
Higher professional	−0.169*	−0.331*	−0.397*	−0.691*
	(0.002)	(0.003)	(0.002)	(0.003)
College	−0.172*	−0.395*	−0.647*	−1.101*
	(0.002)	(0.004)	(0.003)	(0.004)
Ethnicity (omitted category: Natives)				
Western immigrants	−0.157*	−0.180*	−0.194*	−0.173*
	(0.003)	(0.004)	(0.003)	(0.004)
Turks/Moroccans	1.270*	2.282*	1.608*	2.822*
	(0.003)	(0.006)	(0.003)	(0.006)
Surinamese/Arubans	−0.155*	−0.079*	−0.330*	−0.164*
	(0.005)	(0.006)	(0.005)	(0.008)
Other non-Western immigrants	0.141*	0.229*	−0.009*	0.175*
	(0.004)	(0.006)	(0.005)	(0.007)
Regional unemployment rate	−0.032*	−0.033*	−0.022*	−0.022*
	(0.001)	(0.001)	(0.001)	(0.001)
Urban indicator	−0.171*	−0.201*	−0.238*	−0.324*
	(0.001)	(0.002)	(0.002)	(0.002)
Bible belt	0.665*	0.904*	0.661*	1.004*
	(0.003)	(0.005)	(0.004)	(0.005)
Four largest cities	−0.248*	−0.376*	−0.268*	−0.370*
	(0.002)	(0.003)	(0.002)	(0.003)
Unobserved heterogeneity	no	yes	no	yes
Log-likelihood / 1000	-9846.89	-9802.94	-8732.63	-8641.61
BIC / 1000	19694.24	19606.25	17465.73	17283.57

Notes: All specifications are weighted using sample weights and include a sixth-degree polynomial in $\ln(age - 17)$ and 5-year birth cohort dummies. Standard errors in brackets. Unobserved heterogeneity is modeled as a discrete distribution with two mass points. BIC is the Schwartz (Bayesian) Information Criterion. Starred coefficients are significant at the 1 percent level.

Table 2.5: Discrete-time duration model for the age at first marriage, by location

	Men ($N = 70,718$)		Women ($N = 53,803$)	
	(1)	(2)	(3)	(4)
Bible belt				
Main effect	0.968*	0.922*	0.994*	1.021*
	(0.008)	(0.008)	(0.007)	(0.008)
Period 1 (1998–2000)	−0.232*	−0.044*	−0.283*	−0.041*
	(0.010)	(0.011)	(0.010)	(0.011)
Period 2 (2001–2005)	−0.483*	0.029	−0.640*	−0.022
	(0.010)	(0.012)	(0.010)	(0.012)
Four largest cities				
Main effect	−0.029*	−0.166*	0.076*	−0.047*
	(0.006)	(0.006)	(0.006)	(0.006)
Period 1 (1998–2000)	−0.283*	−0.185*	−0.333*	−0.127*
	(0.007)	(0.008)	(0.007)	(0.008)
Period 2 (2001–2005)	−0.283*	−0.260*	−0.580*	−0.556*
	(0.008)	(0.009)	(0.008)	(0.009)
Rest of the Netherlands				
Period 1 (1998–2000)	0.006	0.035*	0.006	0.055*
	(0.004)	(0.004)	(0.004)	(0.005)
Period 2 (2001–2005)	−0.053*	−0.009	−0.158*	−0.072*
	(0.006)	(0.006)	(0.006)	(0.006)
Linear trend (1995=0)	−0.035*	−0.035*	−0.017*	−0.020*
	(0.001)	(0.001)	(0.001)	(0.001)
Constant	−6.784*	−6.128*	−4.765*	−8.233*
	(0.020)	(0.020)	(0.009)	(0.013)
Unobserved heterogeneity	no	yes	no	yes
Log-likelihood / 1000	−9845.00	−9802.21	−8729.29	−8638.51
BIC / 1000	19690.54	19604.85	17459.12	17277.43

Notes: All specifications are weighted using sample weights and include controls for education, ethnicity, regional unemployment rate, urbanization, 5-year birth cohort and a sixth-degree polynomial in $\ln(age - 17)$. Standard errors in brackets. BIC is the Schwartz (Bayesian) Information Criterion. Starred coefficients are significant at the 1 percent level.

Table 2.A: Discrete-time duration model for age at first marriage, specifications with parametric and non-parametric duration dependence

	Men ($N = 70,718$)		Women ($N = 53,803$)	
	(1)	(2)	(3)	(4)
Period 1 (1998–2000)	−0.040*	−0.028*	−0.052*	−0.046*
	(0.004)	(0.004)	(0.004)	(0.004)
Period 2 (2001–2005)	−0.101*	−0.087*	−0.234*	−0.230*
	(0.006)	(0.006)	(0.006)	(0.006)
Linear trend (1995=0)	−0.034*	−0.035*	−0.017*	−0.017*
	(0.001)	(0.001)	(0.001)	(0.001)
Education (omitted category: Senior vocational)				
Primary education	−0.258*	−0.258*	0.011*	0.011*
	(0.004)	(0.004)	(0.004)	(0.004)
Secondary vocational	−0.053*	−0.053*	0.098*	0.098*
	(0.002)	(0.002)	(0.002)	(0.002)
General secondary	−0.342*	−0.342*	−0.265*	−0.265*
	(0.003)	(0.003)	(0.003)	(0.003)
Higher professional	−0.169*	−0.169*	−0.397*	−0.397*
	(0.002)	(0.002)	(0.002)	(0.002)
College	−0.172*	−0.172*	−0.647*	−0.647*
	(0.002)	(0.002)	(0.003)	(0.003)
Ethnicity (omitted category: Natives)				
Western immigrants	−0.157*	−0.157*	−0.194*	−0.194*
	(0.003)	(0.003)	(0.003)	(0.003)
Turks/Moroccans	1.270*	1.270*	1.608*	1.608*
	(0.003)	(0.003)	(0.003)	(0.003)
Surinamese/Arubans	−0.155*	−0.155*	−0.330*	−0.330*
	(0.005)	(0.005)	(0.005)	(0.005)
Other non-Western immigrants	0.141*	0.141*	−0.009*	−0.009*
	(0.004)	(0.004)	(0.005)	(0.005)
Regional unemployment rate	−0.032*	−0.030*	−0.022*	−0.022*
	(0.001)	(0.001)	(0.001)	(0.001)
Urban indicator	−0.171*	−0.170*	−0.238*	−0.238*
	(0.001)	(0.001)	(0.002)	(0.002)
Bible belt	0.665*	0.665*	0.661*	0.661*
	(0.003)	(0.003)	(0.004)	(0.004)
Four largest cities	−0.248*	−0.247*	−0.268*	−0.268*
	(0.002)	(0.002)	(0.002)	(0.002)
Duration dependence	poly	np	poly	np
Log-likelihood / 1000	−9846.89	−9846.28	−8732.63	−8732.27
BIC / 1000	19694.24	19693.22	17465.73	17465.14

Notes: All specifications are weighted using sample weights and include 5-year birth cohort dummies. Duration dependence is modeled either as a sixth-degree polynomial in $\ln(age-17)$ ("poly") or as a set of age dummies ("np"). Standard errors in brackets. BIC is the Schwartz (Bayesian) Information Criterion. Starred coefficients are significant at the 1 percent level.

Chapter 3

Aggregate analysis

3.1 Introduction

Th simple graphical analysis of the marriage rate present in the media assumes that the only possible influence on the marriage rate is the introduction of registered partnership or the legalization of same-sex marriage. Moreover, by comparing the marriage rate before to the marriage rate after each of these legal changes, the implicit assumption is that the "before" period is an appropriate counterfactual for the "after" period. In most empirical cases, however, this is not the case.

The most common strategy to estimate the effects of an intervention is the "differences-in-differences" method (Wooldridge, 2002). This method requires observations from two groups: one that was affected by the intervention, called the "treated" group, and one that is similar to the treated group but was not affected by the intervention, called the "control" group. The method then consists of comparing the difference in outcomes before and after the intervention in the treatment and the control groups. This is done through an estimating equation of the form

$$Y_{it} = X_{it}'\beta + T_i\delta_1 + P_i\delta_2 + T_iP_i\delta_3 + \epsilon_{it}, \tag{3.1}$$

where Y_{it} is the outcome of interest for unit i at time t, X_{it} is a set of characteristics,

T_i is a dummy variable for individual i being in the treated group, and P_i is a dummy variable for period t being after the intervention (the "post" period). The estimated effect of the intervention is then δ_3, the differential change in the outcome in the treated group compared to the control group between the periods before and after the intervention.

This estimation method relies on the identification of an appropriate control group. The control group is used as the "counterfactual," the evolution of the outcome in the treated group in the absence of the intervention. Thus, the treated and control groups should have experienced similar trends in their outcomes in the absence of the intervention. The control group is usually chosen such that the pre-intervention trends in the two groups are similar, indicating that the post-intervention trends would have been similar as well in the absence of the intervention.

It should be clear now that the choice of a control group is crucial, since using the wrong counterfactual can lead to severely biased inference. In some applications, the control group is apparent for geographic, demographic or historical reasons (see, for example, Card and Krueger, 1994; Almond et al., 2005; Tyler et al., 2000). In other applications, such as my study, there is no obvious choice for the control group. The major concern with the individual-level approach in the previous chapter is that it uses the period before the enactment of the registered partnership law as the counterfactual. In other words, since the two laws apply to all Dutch residents, there is no group in the population that was not affected by the laws (i.e., there is no control group) and the change in marriage rates after each law is interpreted as resulting directly and only from the laws themselves. Inference based on this approach is biased if there were other factors that influenced the marriage decision and that came into effect during the same period as the laws.

The alternative is to look for a country or a set of countries that can be used as a control group. For historical and demographic reasons, the best control group would have been Belgium. However, Belgium followed closely in the footsteps of the Netherlands and enacted a registered partnership law in 2000 and a same-sex marriage law in 2003. In the absence of an obvious control group, Abadie and Gardeazabal (2003) and Abadie et al. (2007) suggest an alternative: a synthetic control. The synthetic control is a weighted average of potential "donors" such that the weighted average of their outcomes and of their determinant variables (the X's) exhibit a similar trend to the treated group in the period before the intervention. This artificial data series provides a credible counterfactual because it takes into account the evolution of both the outcome and its determinants.

In this chapter, I construct a synthetic control for the marriage rate in the Netherlands as a weighted average of OECD member countries. The results support the conclusions reached in chapter 2. I also conduct several placebo tests that support the validity of the method and of the results. Finally, I provide an interpretation of these results using additional information from the survey of married and registered couples conducted by Boele-Woelki et al. (2006).

3.2 Construction of the synthetic control and estimation method

Similar to Abadie et al. (2007), let subscript 1 indicate the Netherlands and let $W = (w_2, \ldots, w_{J+1})$ be a vector of weights assigned to the J potential donor countries. Without any restrictions on the weights, a sufficiently large number of potential donor countries and of determinant variables will lead to a synthetic control that matches perfectly the evolution of the marriage rate in the Netherlands prior to the introduction of the two laws. However, negative weights or weights larger than one would be difficult to interpret. Hence, the weights are restricted to lie in the unit interval ($0 \leq w_j \leq 1$ for all j) and to sum up to one ($\sum_{j=2}^{J+1} w_j = 1$), which results in a synthetic control that will likely not match perfectly the trend in the marriage rate before the two laws.

For the synthetic control, the marriage rate m_1^* and its determinants X_1^* are calculated as weighted averages of the corresponding variables in the donor countries:

$$m_{1t}^* = \sum_{j=2}^{J+1} w_j m_{jt},$$
$$X_{jt}^* = \sum_{j=2}^{J+1} w_j X_{jt}.$$

Let T_0 be the number of available periods before 1998 and let the vector $K = (k_1, \ldots, k_{T_0})$ define a linear combination of the pre-1998 marriage rates for any country i:

$$\overline{m}_i^K = \sum_{t=1}^{T_0} k_t m_{it}.$$

Now consider M such linear combinations for the Netherlands: $\overline{m}_1^{K_1}, \ldots, \overline{m}_1^{K_M}$, and define $Z_1 = \left(X_1', \overline{m}_1^{K_1}, \ldots, \overline{m}_1^{K_M}\right)'$ as the vector obtained by combining the determinants of the marriage rate prior to 1998 (T_0) and these M linear combinations of the pre-1998 marriage rate in the Netherlands. Next, consider the matrix Z_0 constructed by combining similar vectors for the J potential donors, such that the j-th column of Z_0 is $\left(X_j', \overline{m}_j^{K_1}, \ldots, \overline{m}_j^{K_M}\right)'$, where X_j is the set of determinants of the marriage rate prior to 1998 in country j.

In principle, the linear combinations $(K_1, \ldots, K_M)$ are arbitrary. In practice, Abadie et al. (2007) suggest choosing $M = 1$ and $K = \frac{1}{T_0}$, such that the linear combinations amount to the average over the period before the intervention:

$$\overline{m}_i^K = \frac{1}{T_0} \sum_{t=1}^{T_0} m_{it}$$

Thus, the vector of data for the Netherlands becomes $Z_1 = (X_1', \overline{m}_1)'$, while the corresponding matrix Z_0 for the donor countries has columns of the form $\left(X_j', \overline{m}_j\right)'$ for the j-th donor country.

Given this structure of the Z matrices, let V be a diagonal matrix of loadings corresponding to all the variables (both the determinants X and the marriage rate m). The optimal set of weights is the one that minimizes the weighted distance between Z_1 and Z_0:

$$W^*(V) = \text{argmin}\sqrt{(Z_1 - Z_0W)'V(Z_1 - Z_0W)}.$$

The matrix V can be arbitrary, but a natural choice is the one that minimizes the mean squared error of the marriage rate in the synthetic control relative to the actual marriage rate in the Netherlands (Abadie et al., 2007):

$$V^* = \text{argmin}\sqrt{[m_1 - m_0 W^*(V)]'[m_1 - m_0 W^*(V)]},$$

where m_1 is a $(T_0 \times 1)$ vector containing the marriage rate in the Netherlands and m_0 is a $(T_0 \times J)$ matrix of marriage rates of the potential donors in the pre-1998 period. This ensures that the marriage rate in the synthetic control constructed using the resulting weights $W^*(V^*)$ is the best match to the marriage rate in the Netherlands in the period before 1998.

In conclusion, the synthetic control is constructed by assigning a set of data-driven weights to potential "donor" countries such that the weighted average of their marriage rates and determinant variables closely match the marriage rate and the determinant variables in the Netherlands during the "pre-treatment period" (before the enactment of the first law, the registered partnership law). These weights are calculated via an iterative two-step maximization problem. In the first step, each variable (both in the set of determinants X and the marriage rate) is assigned a loading and the country weights are calculated as a function of these loadings so as to minimize the (weighted) distance between the synthetic control and the Netherlands in terms of all the variables. In the second step, the variable loadings are chosen so that the marriage rate in the control group matches the marriage rate in the Netherlands as closely as possible. Finally, the two steps are repeated until convergence is achieved.

3.3 Data

The list of potential donors includes all the OECD member countries that did not enact a registered partnership or same-sex marriage law during the period 1988–2004. After excluding Mexico and the Slovak Republic, for which not enough data are available, the sample of potential donors consists of 17 countries: Australia, Austria, Czech Republic, Germany, Greece, Hungary, Ireland, Italy, Japan, Korea, New Zealand, Poland, Portugal, Switzerland, Turkey, United Kingdom, United States. The data on marriage rates and their determinants come from the OECD, European Union's Eurostat, World Bank's World Development Indicators, or national statistical offices. The exact source and definition of each variable are provided in appendix 3.A.

The ideal outcome variable would be the marriage rate measured among the "population at risk," i.e. single individuals legally allowed to marry. Unfortunately, statistics agencies commonly report a different measure, the crude marriage rate, defined as the total number of marriages per 1,000 inhabitants. Figure 3.1 plots these two different indicators using data on different-sex marriages and population from Statistics Netherlands over the period between 1988 and 2004 (the two vertical dotted lines correspond to 1998 and 2001, the years when registered partnership and same-sex marriage were introduced). Since the marriage rate among unmarried individuals of legal age to marry (the "correct" marriage rate) is much higher than the crude marriage rate, the two measures are plotted on separate but proportional axes. It is clear from the figure that although the "correct" marriage rate drops relatively more sharply than the crude marriage rate, the two measures exhibit extremely similar patterns. Therefore, the crude marriage rate can provide an accurate indication of the evolution of the "correct" marriage rate.

The figure also shows that changes in the crude marriage rate understate changes in the "correct" marriage rate. Indeed, the percentage change in the crude marriage rate between 1989 and 1997 was smaller by an average factor of 1.19 than the percentage change in the "correct" marriage rate. The difference between the evolution of the two indicators is much smaller after the enactment of the two laws, the average factor during the period 1998–2005 being 1.02. The crude marriage rate is thus a better indicator of the "correct" marriage rate after the enactment of the two laws.

The data available for Statistics Netherlands allows for a breakdown of marriages into different-sex and same-sex marriages. In addition, it includes information on new registered partnerships, again separately for same-sex and different-sex couples. This makes it possible to study in turn the evolution of three different indicators: the different-sex marriage rate, the overall marriage rate, and the different-sex union rate (marriages and registered partnerships), all defined as the corresponding number of contracts per 1,000 individuals.

As in the individual-level analysis, the variables included in the vector of determinants X can be classified in three groups. First, there are the variables that describe the number of people at risk of marriage and the probability that they will meet, or the thickness of the marriage market. This group includes the fraction of the population in the 25–44 age group, the fraction of population living in urban areas, the ratio of women to men in the population, and the age at first marriage of both men and women. The second set of variables describes the attractiveness of individuals in the marriage market in terms of their current or potential earnings and in terms of fertility. These variables are the labor force participation of both men and women aged 25–34, the total fertility rate (the average number of children that would be born by women of bearing age), and girls' enrollment share in tertiary education. Finally, the unemployment rate of individuals in the 25–34 age group describes business cycle fluctuations.

The first two columns in table 3.1 list the mean of each variable for the Netherlands and for the potential donors as a group (unweighted average) for the period between 1988 and 1997. The differences in the numbers range from very small (labor force participation of men) to relatively large (unemployment), indicating that there are some significant differences between the potential donors as a group and the Netherlands. Therefore, the unweighted average of the potential donors might not be an appropriate control group and the construction of a synthetic control is required.

3.4 Results

3.4.1 Crude marriage rate

The procedure described in section 3.2 produces three quantities of interest: a diagonal matrix V^* of optimal loadings for each determinant variable X, a vector W^* of optimal weights for each potential donor, and the synthetic control constructed as a weighted average of marriage rates in the potential donors by applying the optimal weights W^*.

Recall that the matrix V includes loadings for both the determinant variables X and the average marriage rate for the period before the intervention. In order to assess the importance of each variable in X in the construction of the synthetic control, column 4 in table 3.1 lists the loadings for the determinant variables in X rescaled so as to sum to one. The table indicates that the variables with the most predictive power are the fertility rate and inflation, followed by the age at first marriage of women and the labor force participation rate of men between 25 and 34 years of age.

The means of the determinant variables for the synthetic control are listed in column 3 of table 3.1. It is apparent that the differences between the Netherlands (column 1) and the synthetic control are much smaller than the differences between the Netherlands and the unweighted average of potential donors. Indeed, the largest percentage difference between columns 1 and 3 is about half the largest difference between column 1 and 2, supporting the idea that the synthetic control is a more appropriate control group than the group of potential donors.

The weights assigned to each country in the synthetic control are listed in column 1 of table 3.2. Note that the non-negativity restriction on the weights leads, in general, to corner solutions: 10 of the 17 countries have zero weight. Of the seven countries with non-zero weights, Switzerland has almost three times more weight than any other country and together with Italy, Austria and New Zealand accounts for more than 75 percent of the synthetic control. Finally, the mean squared error of the marriage rate in the synthetic control relative to the actual marriage rate in the Netherlands is 0.18, or 3 percent of the average crude marriage rate of 5.84 over the same period (before 1998).

Figure 3.2 plots the marriage rate in the Netherlands and the synthetic control. The two lines are relatively close for the period 1988–1997, before the registered partnership law was enacted, a reassuring finding since this is the period when the treated and the control group have to be similar. Between 1998 and 2000, when registered partnership was made available to both same-sex and different-sex couples, the marriage rate in the Netherlands is slightly higher than in the synthetic control, though still relatively close. However, the marriage rate in the Netherlands falls rapidly after 2001 but it increases slightly in the synthetic control.

One relatively straightforward way to gauge the decline in the marriage rate is to compare the largest gap between the actual marriage rate in the Netherlands and the synthetic control in each of the three periods: before 1998, between 1998 and 2000, and after 2001. Column 1 in table 3.3 lists these numbers. Note that the largest absolute difference between the marriage rate in the Netherlands and the synthetic control after 2001 occurs in 2005 and is equal to 0.61, or approximately 13.8% of the crude marriage rate in 2005. In contrast, the largest relative difference between the two measures during 1988–1997 is 0.28, or 4.8% of the average crude marriage rate during this period, and between 1998 and 2000 it is 0.34, or 6.10% of the average crude marriage rate during this period. This suggests that the decline in the marriage rate after 2001 is rather significant, being at least twice as large (relatively) than any difference between the synthetic control and the real marriage rate in the previous periods.

The aggregate analysis above suggests that the marriage rate did not decline after the introduction of registered partnership, but it did after the legalization of same-sex marriage. This pattern is exactly the same as the one found in section 2.4, supporting the validity of the individual-level analysis. In particular, there is no negative effect of the registered partnership law on the marriage rate, but there is such an effect for the same-sex marriage law. This result is robust to the inclusion of additional variables, such as the difference in life expectancy between women and men or the growth rate of real GDP. I also find the same pattern if I include in the vector X of determinants values post-1998 for variables which are not likely to be affected by the two laws (e.g., the ratio of women to men or GDP growth).

3.4.2 Placebo test

Although the results using both methods are similar, there is still the concern that the synthetic control method might not produce an appropriate control because, for example, the set of determinants X is not well chosen. Abadie and Gardeazabal (2003) and Abadie et al. (2007) suggest two ways to conduct placebo tests that would confirm or reject the choice of determinant variables. The first type of placebo test is to choose a period of analysis prior to the intervention and assign an artificial "intervention year" during this period. The synthetic control constructed in this way should not differ from the treated group either before or after the artificial intervention. Unfortunately, there is not enough historical data on all the variables included in the analysis to conduct such a placebo test.

The second type of placebo test consists of choosing some of the countries with the highest weight in the synthetic control and assume that they actually experienced the same type of intervention at the same time as the treated group. A synthetic control can be constructed for each of these experiments using the rest of the donor countries. These synthetic controls should not be different from the "treated" countries since there is no intervention.

I conduct this second type of placebo test to confirm the validity of the method. I focus on Switzerland, the country with the largest weight in the synthetic control for the Netherlands. I eliminate it from the pool of potential donors and I construct a synthetic control for Switzerland in the same way as before, using the method in section 3.2 and data for the period between 1988 and 1997. The weights of each country in the synthetic Switzerland are listed in column 2 of table 3.2. As before, only a few of the potential donors have non-zero weights and three countries (Germany, Korea and Ireland) account for more than 86 percent of the synthetic control. Figure 3.3 plots the marriage rate in Switzerland and the corresponding synthetic control. Unlike the case of the Netherlands, the two lines are remarkably similar for the whole period of analysis. This is reassuring, as there should have been no effect of the placebo laws on the marriage rate in Switzerland.

As before, we can compare the largest absolute differences between the crude marriage rate in Switzerland and its synthetic control for the three periods. The corresponding numbers are listed in column 4 of table 3.3. Note that the synthetic control for Switzerland is a relatively poorer match than the one for the Netherlands during the period used for its construction: the absolute difference between 1988 and 1997 is 0.51 in Switzerland, or 8.03% of the average marriage rate during the period. In contrast, the absolute difference in the Netherlands during the same period is only 4.81% of the average marriage rate. However, the relative differences for Switzerland get only smaller in the subsequent periods: 0.22 (4.05% of the mean) during 1998–2000 and 0.32 (6.08% of the mean) after 2001.

In conclusion, the placebo test suggests that the synthetic control method constructs an appropriate control group and that the conclusions in section 3.4.1 are valid.

3.4.3 Different-sex marriage and union rates

Recall that the end-of-marriage argument holds that the different-sex marriage rate would fall after the legalization of same-sex marriage or even the introduction of same-sex registered partnership. In the individual-level analysis, I could not separate different-sex from same-sex marriages as there was no information on the gender of the spouse. However, Statistics Netherlands provides aggregate data for each type of marriage, allowing a separate analysis of different-sex marriage. Since the crude marriage rate used above includes both different-sex and same-sex marriages after 2001, it is an overstatement of the (crude) different-sex marriage rate. Indeed, as it can be seen in figure 3.4 (compared to figure 3.2), the fall in the marriage rate after 2001 is even greater if only different-sex marriages are considered.[1] Column 2 in

[1]Note that the synthetic control is the same as before because the counterfactual is the same: what the marriage rate would be in the Netherlands if the two laws were not enacted, i.e. if only different-sex couples were allowed to marry.

table 3.3 lists the largest absolute differences between the different-sex marriage rate and the synthetic control for each period. The differences for the periods 1988–1997 and 1998–2000 are the same as in column 1, since the different-sex marriage rate is just the crude marriage rate prior to 2001. The absolute difference in 2005 is 0.68, 15.62% of the marriage rate that year, larger than in column 1. In other words, the decline in the different-sex marriage rate is even larger than suggested by the crude marriage rate.

One aspect that was not taken into account until now is that different-sex couples have access to an alternative institution after the introduction of registered partnership. If registered partnership is perceived as a reasonable alternative to marriage, then it is possible that some couples might choose it over marriage. In a world where marriage and registered partnership are equivalent, we would expect couples to select randomly into an institution and thus have approximately half of them choosing marriage and half registered partnership. In practice, the legal differences between the two institutions, however minor, and the difference in traditional values ensure the fact that the distribution of couples across institutions is not even. Thus, one can make the argument that what matters is the total number of unions, i.e. marriages and registered partnerships, rather than just marriages. Figure 3.5 plots this different-sex union rate in the Netherlands and the marriage rate from the synthetic control. Again, the synthetic control is the same as above. Note also that the different-sex union rate is the same as the different-sex marriage rate and the overall marriage rate prior to 1998 since the institution available to couples was different-sex marriage.

As before, we can compare the differences between the different-sex union rate and the synthetic control over the three periods. The numbers are shown in column 3 of table 3.3. While the difference before 1998 is the same as in the previous two columns, it is higher between 1998 and 2000 at 0.44, or 7.67% of the mean union rate. This is not surprising since the union rate is the marriage rate plus the registered partnership rate and the marriage rate was already higher than the synthetic control during this period. The difference becomes smaller after 2001, however, to 0.33 (6.98% of the union rate in 2005 or 6.55% of the average union rate over the period), which is relatively smaller than both the corresponding differences in columns 1 and 2 and the difference between 1998 and 2000. Moreover, the difference is comparable to the difference prior to the introduction of registered partnership, which suggests that there is not much change in the total number of unions in recent years relative to the baseline period (prior to the registered partnership law), but rather sorting across the two institutions.

3.4.4 Additional evidence

These results are consistent with two explanations. One the one hand, there might be a learning process: people learn over time about registered partnership and start switching away from marriage. This seems to be suggested by the evolution of the registered partnership rate described in figure 2.4. Under this scenario, the same-sex marriage law could have no effect and the observed decline in the marriage rate is simply due to couples choosing registered partnership over marriage. Alternatively, the legalization of same-sex marriage could have changed the value of marriage relative to registered partnership. Recall that the only major difference between the two institutions is in the traditional value of marriage. If same-sex marriage reduces the traditional value of marriage, then it is possible that more people switch to registered partnership. Under this scenario, the decline in the marriage rate is directly due to the legalization of same-sex marriage and the registered partnership law could have no negative effects.

As long as the different-sex union rate is not larger than the marriage rate in the synthetic control after 2001, there has to be some sorting of couples out of marriage and into registered partnership. The number of different-sex registered partnerships almost quadruples between 1998 and 2005, but most of the increase occurrs after 2001, which suggests that the legalization of same-sex marriage might have accelerated the learning process, the sorting process, or both. It is impossible to distinguish between the two alternatives with the data available. Fortunately, some additional evidence is provided by an evaluation study of the two institutions commissioned by the Dutch Ministry of Justice in 2005. Although there is no historic information included, the results of the survey presented in Boele-Woelki et al. (2006) can help shed some light on these two scenarios.

First, there is clear evidence of learning as some of the couples interviewed who were in registered partnership report finding out about the institution after its introduction, usually from a notary.[2] Second, 57 percent of the couples in registered

[2]The couples report being told about registered partnership by a notary they visited for drawing up or renewing cohabitation agreements (privately-drawn contracts between cohabiting partners) or for inheritance issues (Boele-Woelki et al., 2006).

partnership acknowledge not having considered marriage as an option. If at least some of these couples entered registered partnership after 2001 and if the number of different-sex unions is the same as in the absence of registered partnership and same-sex marriage, as suggested by figure 3.5, then some of the couples who would have gotten married in the absence of the laws choose not to formalize their relationship anymore. Since the trend in the marriage rate and in the registered partnership rate accelerate after 2001 in different directions, this supports the idea that the same-sex marriage law might have a negative effect on different-sex marriage.

In conclusion, there is suggestive evidence that confirms both scenarios. Over time, different-sex couples become increasingly aware of registered partnership and start choosing it over marriage, but some different-sex couples choose either registered partnership over marriage or no formal relationship at all after the enactment of the same-sex marriage law.[3]

3.5 Conclusion

Similar to the analysis in chapter 2, the analysis in this chapter suggests that the marriage rate increased slightly after the enactment of the registered partnership law and then fell after the enactment of the same-sex marriage law. This pattern is even stronger if only different-sex marriages are considered. However, once registered

[3] At this point, it would be useful to compare these patterns to a model of first formalization of a relationship through either registered partnership or marriage, similar to the model in section 2.2. However, Boele-Woelki et al. (2006) report that registered partnership seems to be the choice for older couples. The average age at partnership registration among the couples surveyed was over 38 years for different-sex couples and over 43 years for same-sex couples, compared to 33 and 41 years for married couples. Indeed, only 701 men and 647 women in my sample enter registered partnership, compared to 20,670 and 19,865 marriages, and the results in section 2.4 are qualitatively the same if age at first marriage is replaced by age at first registration.

partnerships are included, the rate of different-sex unions is close to the rate in the synthetic control.

This chapter contributes to the literature in two ways. First, it provides reassuring evidence on the validity of the individual-level analysis in the previous chapter. While neither method can bring in itself hard enough evidence to convince the reader, together they provide a compelling story on the effect of the two laws on the marriage rate.

Second, the aggregate data allows for a separate study of the overall (crude) marriage rate, the different-sex marriage rate and the different-sex union rate. The evolution of these different measures uncovers two interesting stories. First, the introduction of an alternative institution for different-sex couples leads to an increased different-sex union rate. Second, there is sorting across the two institutions for different-sex couples after both laws.

Finally, additional data provided by the survey conducted by Boele-Woelki et al. (2006) suggests that two forces combine to influence the evolution of the marriage rate. On the one hand, some couples become more aware of the alternative institution of registered partnership and choose it over marriage or cohabitation agreements. On the other hand, some couples who would have gotten married in the absence of the two laws choose not to get married any more. Thus, the evidence suggests that the evolution of the marriage rate after 2001 could be attributed to both the long-term effect of the registered partnership law, via learning, and to the short-term effect of the same-sex marriage law.

3.A Data sources and description

The variables used in this chapter are defined as follows:

- *Crude marriage rate* = the number of (new) marriages per 1,000 individuals;
- *Fraction of the population in an age group* = number of people in the age group divided by total population;
- *Urban population* = fraction of the midyear population of areas defined as urban in each country and reported to the United Nations in total population;
- *Female-male ratio* = ratio of the total female population to total male population;
- *Age at first marriage* = the average or median age of individuals marrying for the first time;
- *Labor force participation, by sex* = ratio of the labor force (the total number of people employed plus unemployed) for each sex to the working age population (older than 15) for each sex, expressed in percentages;
- *Total fertility rate* = number of children born to women aged 15–49;
- *Girls' enrollment share in tertiary education* = the number of girls enrolled in tertiary education, expressed as a percentage of the total number of students in tertiary education;
- *Unemployment rate of individuals in the 25–34 age group* = number of unemployed persons as a percentage of the labor force (the total number of people

employed plus unemployed) in the age group.

- *Inflation* = annual change in the consumer price index using 2000 as the base year (the value of the index is equal to 100 in 2000).

The data sources for each country are listed in table 3.A

Notes: The crude marriage rate is measured on the left axis, while the "correct" marriage rate is measured on the right axis. The two axes are proportional: a one percent change in the crude marriage rate (on the left axis) is projected onto a one percent change in the "correct" marriage rate on the right axis.

Figure 3.1: Evolution of two measures of the marriage rate in the Netherlands

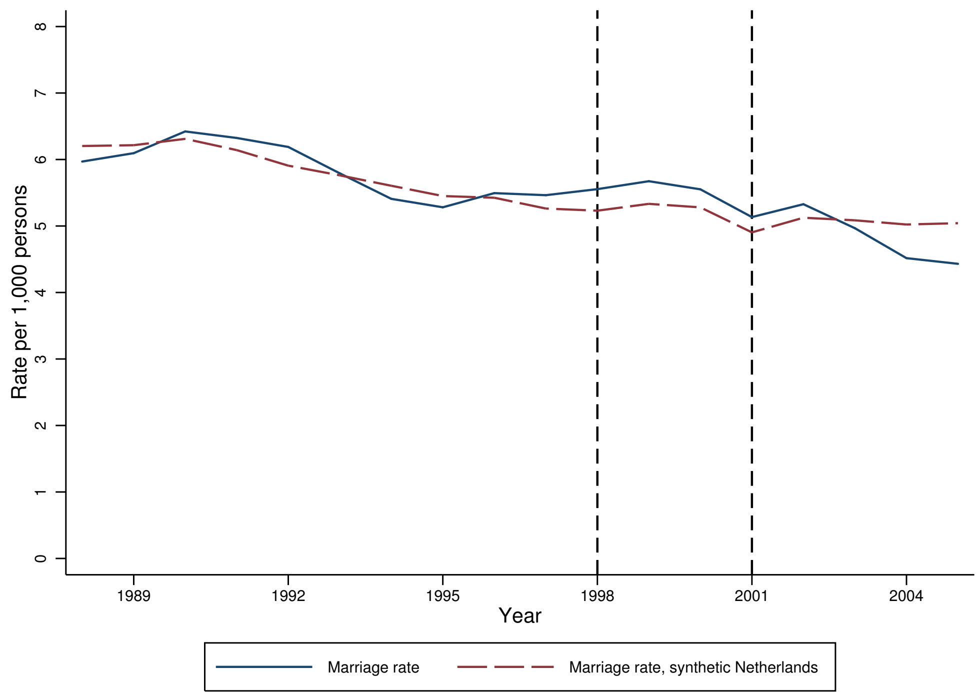

Figure 3.2: Evolution of marriage rate in the Netherlands and in the synthetic control

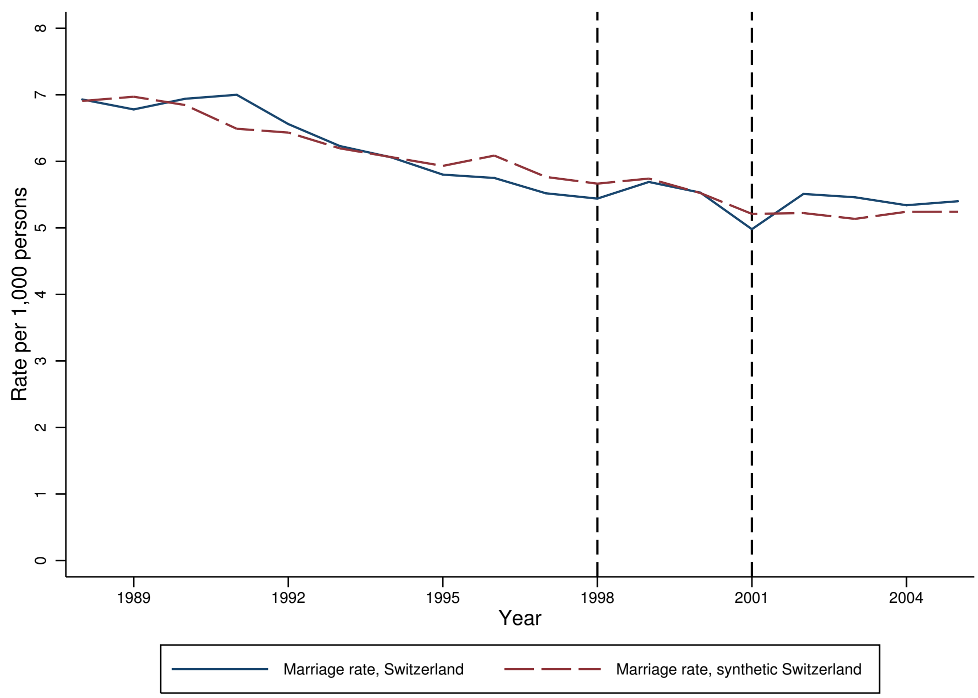

Figure 3.3: Placebo test: Switzerland

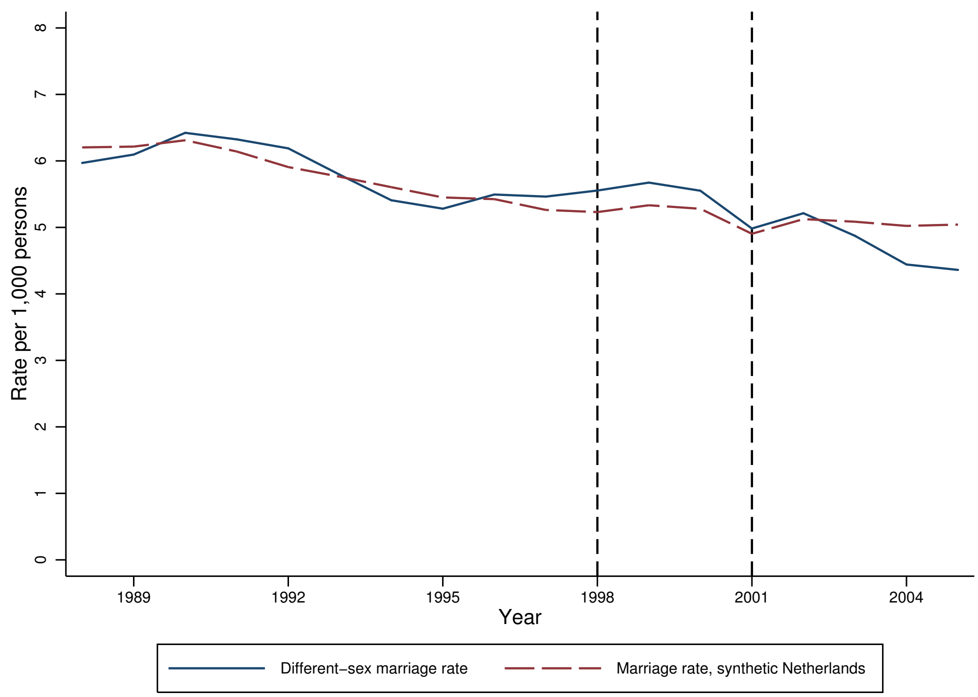

Figure 3.4: Evolution of the different-sex marriage rate in the Netherlands and in the synthetic control

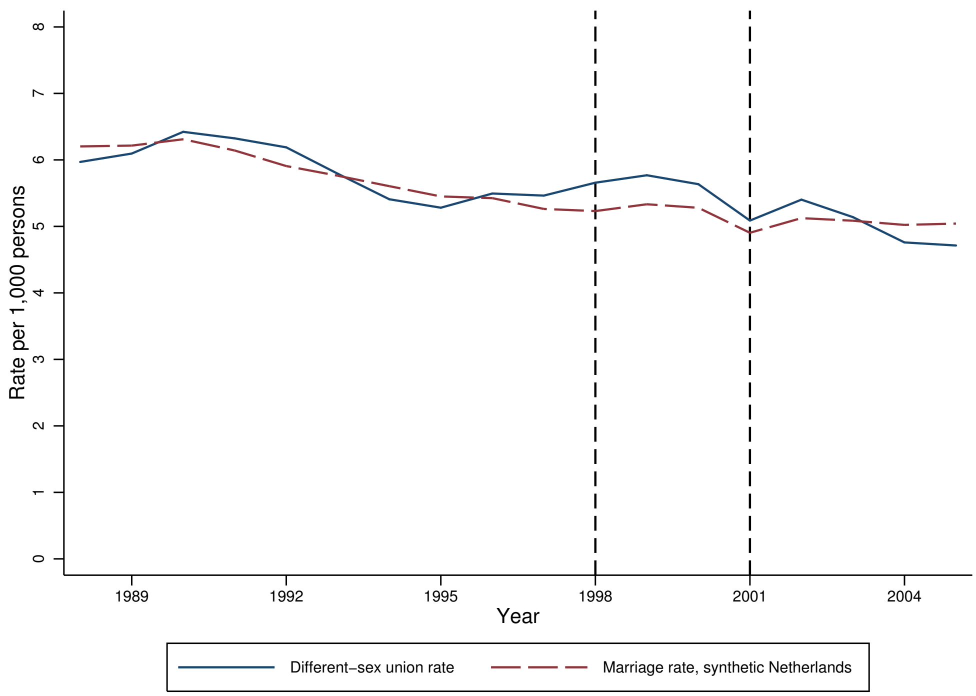

Figure 3.5: Evolution of the different-sex union rate in the Netherlands and in the synthetic control

Table 3.1: Descriptive statistics and variable weights, synthetic control

Variable	Mean, 1988–1997			Loading
	Netherlands	Potential donors	Synthetic control	
	(1)	(2)	(3)	(4)
Crude marriage rate	5.84	6.31	5.85	—
Population 25–44 y.o. (%)	32.32	29.52	30.20	0.00002
Urban population (%)	70.87	70.36	68.99	0.00001
Ratio of women to men	1.01	1.03	1.03	0.00002
Age at first marriage, men (years)	28.51	27.08	28.24	0.00001
Age at first marriage, women (years)	26.28	24.8	26.28	0.01606
Labor force participation, men 25–34 y.o. (%)	94.22	94.04	94.21	0.00869
Labor force participation, women 25–34 y.o. (%)	69.67	64.13	69.62	0.00110
Fertility rate	1.57	1.7	1.57	0.94617
Girls' enrollment share in tertiary education (%)	46.56	47.99	44.03	0.00018
Unemployment, all 25-34 y.o. (%)	6.67	8.39	6.67	0.00303
Inflation (%)	2.21	14.87	3.00	0.02470

Table 3.2: Optimal weights in the synthetic control

Country	Synthetic control for: Netherlands (1)	Synthetic control for: Switzerland (2)
Australia	0	0
Austria	0.138	0.089
Czech Republic	0	0
Germany	0.074	0.441
Greece	0	0
Hungary	0	0
Ireland	0.097	0.189
Italy	0.139	0
Japan	0.068	0.001
Korea	0	0.238
New Zealand	0.103	0
Poland	0	0
Portugal	0	0
Switzerland	0.381	—
Turkey	0	0
United Kingdom	0	0.042
United States	0	0

Table 3.3: Differences between marriage and union rates and synthetic control

	Absolute difference between synthetic control and			
	Netherlands			Switzerland
	Crude marriage rate (1)	Different-sex marriage rate (2)	Different-sex union rate (3)	Crude marriage rate (4)
1988–1997				
Max	0.28	0.28	0.28	0.51
% of mean during period	4.81	4.81	4.81	8.03
% of year when max occurred	4.54	4.54	4.54	7.29
1998–2000				
Max	0.34	0.34	0.44	0.22
% of mean during period	6.10	6.10	7.67	4.05
% of year when max occurred	6.01	6.01	7.56	4.13
2001–2005				
Max	0.61	0.68	0.33	0.32
% of mean during period	12.53	14.27	6.55	6.08
% of year when max occurred	13.78	15.62	6.98	5.95
2005 only				
Max	0.61	0.68	0.33	0.16
% of value	13.78	15.62	6.98	2.92

Notes: "Max" is the largest absolute difference between the synthetic control and the corresponding measure. "% of mean during period" refers to the ratio of the max defined above to the average of the corresponding measure during the indicated period, expressed as a percentage. "% of year when max occured" is the ratio of the max to the value of the corresponding variable during the year when the largest absolute difference is observed, expressed as a percentage. For "2005 only", only the last ratio is defined.

Table 3.A: Sources of data for the variables used in the aggregate analysis

	NL	AU	AT	CZ	DE	GR	HU	IE	IT	JP	KR	NZ	PL	PT	CH	TR	GB	US
Crude marriage rate	N	N	E	E	E	E	E	E	E	N	N	N[1]	E	E	E	E	E	N
Population 25–44	E	N	E	E	E	E	E	E	E	N	N	N[2]	E	E	E	E	E	N
Urban population	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W
Women-men ratio	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W
Age at first marriage	E	N[3]	E	E	E	E	E	E	E	N	N	N[3]	E	E	E	E	N	N[3]
Labor force participation	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O
Fertility rate	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O
Girls in tertiary education	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E
Unemployment, 25–34 age group	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O
Inflation	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O

Notes: 1. Until 1990, based on “de facto” population. Since 1991, based on resident population. 2. As of Dec 31. 3. Median rather than average. Country codes: NL = Netherlands, AU = Australia, AT = Austria, CZ = Czech Republic, DE = Germany, GR = Greece, HU = Hungary, IE = Ireland, IT = Italy, JP = Japan, KR = Korea, NZ = New Zealand, PL = Poland, PT = Portugal, CH = Switzerland, TR = Turkey, GB = United Kingdom, US = United States. Sources of data: E = Eurostat, W = WDI, O = OECD, E = edStats (World Bank), N = National statistical office. National statistical offices: Netherlands = Statistics Netherlands, Australia = Australian Bureau of Statistics, Japan = Statistics Bureau (population) or Ministry of Health, Labour and Welfare, Korea = Korea National Statistical Office, New Zealand = Statistics New Zealand, United Kingdom = Office for National Statistics, United States = Statistical Abstract of the United States (marriage rates), Intercensal estimates (population), Census Bureau (age at first marriage).

Bibliography

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller (2007), "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program," NBER Working Paper no. 12831.

Abadie, Alberto, and Javier Gardeazabal (2003), "The economic costs of conflict: A case study of the Basque Country," *American Economic Review* 93(1), 113–32.

Addison, John, and Pedro Portugal (2003), "Unemployment duration: Competing and defective risks," *Journal of Human Resources* 38(1), 156–91.

Allen, Douglas, Krishna Pendakur, and Wing Suen (2006), "No-fault divorce and the compression of marriage ages," *Economic Inquiry* 44(3), 547–58.

Almond, Douglas, Kenneth Chay, and David Lee (2005), "The costs of low birth weight," *Quarterly Journal of Economics* 120(3), 1031–83.

Badgett, M. V. Lee (2004a), "Will providing marriage rights to same-sex couples undermine heterosexual marriage?" *Sexuality Research & Social Policy* 1(3), 1–10.

Badgett, M.V. Lee (2004b), "Prenuptial jitters: Did gay marriage destroy heterosexual marriage in Scandinavia?" available online at `http://www.slate.com/id/2100884/`, last checked on Jun 16, 2009.

Baker, Michael, and Angelo Melino (2000), "Duration dependence and nonparametric heterogeneity: A Monte Carlo study," *Journal of Econometrics* 96(2), 357–93.

Bitler, Marianne, Jonah Gelbach, Hilary Hoynes, and Madeline Zavodny (2004), "The impact of welfare reform on marriage and divorce," *Demography* 41(2), 213–36.

Blank, Rebecca, Kerwin Kofi Charles, and James Sallee (2007), "Do state laws affect the age of marriage? A cautionary tale about avoidance behavior," NBER Working Paper no. 13667.

Boele-Woelki, Katharina, Ian Curry-Sumner, Miranda Jansen, and Wendy Schrama (2006), "Huwelijk of geregistreerd partnerschap? Een evaluatie van de Wet openstelling huwelijk en de Wet geregistreerd partnerschap," Universiteit Utrecht and Ministerie van Justitie.

Cahill, Sean (2004), *Same-Sex Marriage in the United States: Focus on the Facts*, Lanham, MD: Lexington Books.

Campolieti, Michele, Robert Hebdon, and Douglas Hyatt (2005), "Strike incidence and strike duration: Some new evidence from Ontario," *Industrial and Labor Relations Review* 58(4), 610–30.

Card, David, and Alan Krueger (1994), "Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania," *American Economic Review* 84(4), 772–93.

Cherlin, Andrew (2009), *The Marriage-Go-Round: The State of Marriage and the Family in America Today*, Knopf.

de Verdieping Trouw (2005), "Dossiers: Boris Dittrich," Aug 15, available online at `http://www.trouw.nl/achtergrond/Dossiers/article1431963.ece?page=1`, last checked on Jun 16, 2009.

Dutch Ministry of Education, Culture, and Science (2003), *Kerncijfers 1999–2003: Onderwijs, Cultuur en Wetenschap [Facts and Figures 1999–2003: Education, Culture and Science].*

Eberwein, Curtis, John Ham, and Robert LaLonde (2002), "Alternative methods of estimating program effects in event history models," *Labour Economics* 9(2), 249–78.

Eskridge, William, and Darren Spedale (2006), *Gay Marriage: For Better or for Worse? What We've Learned from the Evidence*, New York: Oxford University Press.

Fagan, Patrick, and Grace Smith (2004), "The transatlantic divide on marriage: Dutch data and the U.S. debate on same-sex unions," The Heritage Foundation WebMemo 577, available online at `http://www.heritage.org/research/family/wm577.cfm`, last checked on Jun 16, 2009.

Gautier, Pieter, Michael Svarer, and Coen Teulings (2005), "Marriage and the city," IZA Discussion Papers no. 1491.

Goldin, Claudia, and Lawrence Katz (2002), "The power of the pill: Oral contraceptives and women's career and marriage decisions," *Journal of Political Economy* 110(4), 730–70.

Gunderson, Morley, and Angelo Melino (1990), "The effects of public policy on strike duration," *Journal of Labor Economics* 8(3), 295–316.

Ham, John, and Samuel Rea (1987), "Unemployment Insurance and Male Unemployment Duration in Canada," *Journal of Labor Economics* 5(3), 325–53.

Heckman, James, and Burton Singer (1984), "A method for minimizing the impact of distributional assumptions in econometric models for duration data," *Econometrica* 52(2), 271–320.

Kennan, John (1985), "The duration of contract strikes in U.S. manufacturing," *Journal of Econometrics* 28(1), 5–28.

Kurtz, Stanley (2004a), "Dutch debate," National Review Online, available online at http://www.nationalreview.com/kurtz/kurtz200407210936.asp, last checked on Jun 16, 2009.

——— (2004b), "The end of marriage in Scandinavia: The 'conservative case' for same-sex marriage collapses," *The Weekly Standard* 9(20).

——— (2004c), "Going Dutch?" *The Weekly Standard* 9(36).

——— (2004d), "No explanation," National Review Online, available online at http://www.nationalreview.com/kurtz/kurtz200406030910.asp, last checked on Jun 16, 2009.

Mello, Michael (2004), *Legalizing Gay Marriage*, Philadelphia: Temple University Press.

Merin, Yuval (2002), *Equality for Same-Sex Couples*, Chicago: University of Chicago Press.

Meyer, Bruce (1990), "Unemployment insurance and unemployment spells," *Econometrica* 58(4), 757–82.

Mohr, Richard (1997), "The stakes in the gay-marriage wars," in Robert Baird and Stuart Rosenbaum, eds., *Same-Sex Marriage: The Moral and Legal Debate*, Amherst, NY: Prometheus Books.

Narendranathan, Wiji, and Mark Stewart (1993), "How does the benefit effect vary as unemployment spells lengthen?" *Journal of Applied Econometrics* 8(4), 361–81.

New York Times (1990), "An upright man named Tom Watson," Column by Ira Berkow on Dec 1., p. 46.

Prentice, Ross, and Lynn Gloeckler (1978), "Regression analysis of grouped survival data with application to breast cancer data," *Biometrics* 34(1), 57–67.

Rasul, Imran (2006), "Marriage markets and divorce laws," *Journal of Law, Economics, & Organization* 22(1), 30–69.

Ridder, Geert (1984), "The distribution of single-spell duration data," in George Neumann and Niels Westergrd-Nielsen, eds., *Studies in labor market dynamics*, Berlin, New York: Springer-Verlag.

Sobotka, Tomáš, and Feray Adigüzel (2002), "Religiosity and spatial demographic differences in the Netherlands," University of Groningen SOM Research Report no. 02F65.

Stateline.org (2009), "State of the States 2009," available online at http://archive.stateline.org/flash-data/2009StateOfTheStates.pdf, last checked on Jun 16, 2009.

Statistics Netherlands (2003), "Largest families in Urk," available online at `http://www.cbs.nl/en-GB/menu/themas/bevolking/publicaties/artikelen/archief/2003/2003-1269-wm.htm`, last checked on Jun 16, 2009.

Tyler, John, Richard Murnane, and John Willett (2000), "Estimating the labor market signaling value of the GED," *Quarterly Journal of Economics* 115(2), 431–68.

van Velde, Hans (2005), "No gay marriage in the Netherlands," available online at `http://www.gk.nl/index.php?id=9&a=bericht&bericht=279`, last checked on Jun 16, 2009.

Waaldijk, Kees (2001), "Small change: How the road to same-sex marriage got paved in the Netherlands," in Robert Wintemute and Mads Andenæs, eds., *Legal Recognition of Same-Sex Partnerships: A Study of National, European and International Law*, chap. 23, 437–64, Oxford: Hart Publishing.

——— (2004), "Major legal consequences of marriage, cohabitation and registered partnership for different-sex and same-sex partners in the Netherlands," in Kees Waaldijk, ed., *More or Less Together: Levels of Legal Consequences of Marriage, Cohabitation and Registered Partnerships for Different-Sex and Same-Sex Partners. A Comparative Study of Nine European Countries*, 137–53, Institut National d'Études Démographiques, Document de Travail No. 125.

Wooldridge, Jeffrey (2002), *Econometric Analysis of Cross Section and Panel Data*, Cambridge, MA: MIT Press.