

MARRIAGE, LABOR SUPPLY AND
THE DYNAMICS OF THE SOCIAL SAFETY NET

By

Hamish Low, Costas Meghir, Luigi Pistaferri, and Alessandra Voena

February 2018

COWLES FOUNDATION DISCUSSION PAPER NO. 2121



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

MARRIAGE, LABOR SUPPLY AND THE DYNAMICS OF THE SOCIAL SAFETY NET *

Hamish Low[†] Costas Meghir[‡] Luigi Pistaferri[§] Alessandra Voena[¶]

February 17, 2018

Abstract

The 1996 PRWORA reform introduced time limits on the receipt of welfare in the United States. We use variation by state and across demographic groups to provide reduced form evidence showing that such limits led to a fall in welfare claims (partly due to “banking” benefits for future use), a rise in employment, and a decline in divorce rates. We then specify and estimate a life-cycle model of marriage, labor supply and divorce under limited commitment to better understand the mechanisms behind these behavioral responses, carry out counterfactual analysis with longer run impacts and evaluate the welfare effects of the program. Based on the model, which reproduces the reduced form estimates, we show that among low educated women, instead of relying on TANF, single mothers work more, more mothers remain married, some move to relying only on food stamps and, in *ex-ante* welfare terms, women are worse off.

Keywords: time limits, welfare reform, life-cycle, marriage and divorce

JEL codes: D91, H53, J12, J21

*We thank Joe Altonji, Orazio Attanasio, Richard Blundell, Marianne Bruins, Mariacristina De Nardi, Felix Gray, John Kennan, Pat Kline, Robert Moffitt, Magne Mogstad, Andreas Mueller, Ben Polak, Andrew Shephard, Melissa Tartari, Matt Wiswall, James Ziliak and participants in various seminars and conferences for helpful comments. Jorge Rodriguez Osorio, Samuel Seo and Davide Malacrino provided excellent research assistance. Costas Meghir thanks the UK ESRC (grant number ES/K00624X/1), the ISPS at Yale and the Cowles foundation for financial assistance. All errors are our own.

[†]University of Cambridge, IFS and Cambridge-INET (email: hamish.low@econ.cam.ac.uk).

[‡]Yale University, NBER and IFS (email: c.meghir@yale.edu).

[§]Stanford University, NBER, CEPR and SIEPR (email: pista@stanford.edu).

[¶]The University of Chicago, CEPR, NBER and BREAD (email: avoena@uchicago.edu).

Introduction

Welfare programs constitute an important source of insurance for low-income households, particularly in an incomplete markets world in which people have little protection against income and employment shocks. If they are carefully targeted and designed to minimize work disincentives, social insurance programs can increase overall welfare. However, the potential disincentives distort household decisions in several dimensions and in both static and dynamic ways: family formation, saving and work incentives are all impacted, and with far reaching consequences. These issues have been the source of continuous debate and underlay the major US welfare reform of 1996. The key innovation of the Personal Responsibility and Work Opportunity Reconciliation Act of 1996 (PRWORA) was to introduce lifecycle time limits on receipt of welfare benefits as well as to reduce marital disincentives implicitly built into the preceding program, the Aid to Families with Dependent Children (AFDC). Indeed, the new program replacing AFDC was aptly named Temporary Assistance for Needy Families (TANF). Understanding the tradeoff between incentives and insurance for such programs and their broader effects both in the short run and the long run is a central motivation of this paper.¹

The reform has impacts potentially on four key decisions: benefit use, employment, marriage and divorce. An important issue is how much these decisions interact with each other. For example, if welfare reforms increase employment, that may lead to a decline in the benefits from marriage; however, if welfare reforms reduce single mothers' access to benefits, this will push towards less divorce and possibly increase marriage.

We start by documenting the immediate short-run effects of the reform on labor supply, welfare participation and marital status of individuals who could not have adjusted their behavior in anticipation of the changes. To this end, we use a quasi-experimental approach, following Grogger and Michalopoulos (2003), based on the different ways in which the reform was implemented across states and on how it affected various demographic groups. Women whose youngest child is close enough to 18 years old (when benefit eligibility terminates anyway) would have remained unaffected by the introduction of time limits, while women with younger kids may be affected, depending on the rules implemented by their state of residence. Based on this approach, we show that welfare utilization declined dramatically and persistently, especially for single women; and that the employment of all women increased,

¹During the same period the Earned Income Tax Credit (EITC) was expanded with the goal of increasing labor force participation among low-income individuals.

albeit only offsetting about half the decline in benefit use. Further, the number of divorces declined. On the other hand, we find no detectable effects on the flow of marriages or on fertility.

The core of the paper is to present a life-cycle model to capture the behavioral insights we document in the reduced form analysis. In this model, family formation and dissolution, welfare program participation, labor supply and savings are all endogenous choices. Married couples share risk and resources, and enjoy public goods (modeled as economies of scale) in a limited commitment framework.² We estimate our model using the method of simulated moments (McFadden, 1989; Pakes and Pollard, 1989) on pre-reform data from the Survey of Income and Program Participation (SIPP), focusing on low-educated women of working age. We analyze intrahousehold allocations and quantify the impact of the reform on within-household inequality. More generally, we quantify the long-term effects of time limits and assess their welfare implications.

A key element of our approach is the budget constraint and how this is shaped by the welfare system, and how the welfare system interacts with being married or single. We account for the structure of the welfare system that low-income households are likely to face, including AFDC/TANF, Food Stamps, and the Earned Income Tax credit (EITC). The complete picture of the welfare system allows us to understand the dynamics implied by the time limits, as well as interactions between the different programs, and more generally to evaluate how the structure of welfare affects marriage, labor supply and the allocation of resources within the household.

Our main findings are that time limits lead to a decline in benefit use as women defer claiming benefits until their children are older. About half of this decline for single women is offset by a rise in employment, and some women chose not to claim at all after the reform. On the other hand, a large fraction of those not claiming welfare do not work either. For this group of women in particular, who have to rely on food stamps or their partners, there has been a worsening of the outside option to marriage. This is reflected in a decline in Pareto weights, particularly at the time of marriage, which translates to a lower welfare share for women within marriage but also a decrease in divorce.

Our paper builds on existing work relating both to welfare reform and to life-cycle behav-

²By limited commitment in marriage, we mean that couples will renegotiate sharing rules if one person is at a binding participation constraint, but will not renegotiate otherwise. Full commitment would mean no renegotiation. No commitment would mean renegotiation in every period. For an overview of the literature on dynamic household decision making, see Chiappori and Mazzocco (2017).

ior. The literature on the effects of welfare reform is large and too long to list here. Excellent overviews are featured in Blank (2002) and Grogger and Karoly (2005). Experimental studies have highlighted that time limits encourage households to limit benefit utilization so as to “bank” their future eligibility (Grogger and Michalopoulos, 2003) and more generally are associated with reduced welfare participation (Swann, 2005; Mazzolari and Ragusa, 2012).

The literature on employment effects of welfare reform has primarily focused on single women (see, for instance, Keane and Wolpin (2010)). This is not surprising, given that both institutionally and in practice single women with children are the main recipients and targets of welfare programs such as AFDC or TANF. Chan (2013) shows that time limits associated with welfare reform are an important driver of increased labor supply in this group. Kline and Tartari (2016) examine both intensive and extensive margin labor supply responses in the context of the Connecticut Jobs First program, which imposed rather stringent time limits. An important issue, however, is that the decision to be a single woman is likely to be affected by welfare reform: this was indeed part of the motivation of the reform. We consider the joint decision of being single/married alongside benefit use and employment. In the existing literature there is some evidence on the overall effect of welfare reform on household formation and dissolution, suggesting in particular that the reform was associated with a small decline in both marriages and divorces. However, the estimated effects tend to be rather noisy (Bitler et al., 2004).

Finally, our paper draws from the literature on savings and labor supply in a life-cycle family context such as Blundell et al. (2016). We build on this literature by endogenizing both marriage and divorce and allowing intrahousehold allocations to evolve depending on changes in the economic environment and preferences. The theoretical underpinnings draw from Chiappori (1988, 1992) and Blundell, Chiappori and Meghir (2005) and its dynamic extension by Mazzocco (2007). We apply the risk sharing framework with limited commitment of Ligon, Thomas and Worrall (2000) and Ligon, Thomas and Worrall (2002*b*) as extended to the lifecycle marriage model by Mazzocco, Yamaguchi and Ruiz (2013) and Voena (2015).³⁴

In what follows, we provide detail of the 1996 welfare reform in Section 1. We present the data and the reduced form analysis of the effects of the time limits component of the

³Our paper also relates to the life cycle analyses of female labor supply and marital status (Attanasio, Low and Sanchez-Marcos, 2008; Fernández and Wong, 2014; Blundell et al., 2016; Fernández and Wong, 2017) and contributes to existing work on taxes and welfare in a static context including Heckman (1974), Burtless and Hausman (1978), Keane and Moffitt (1998), Eissa and Liebman (1996) for the US as well as Blundell, Duncan and Meghir (1998) for the UK and many others.

⁴See Persson (2014) for an example of how social policy can directly influence household formation.

PRWORA in Sections 2 and 3, respectively. In section 4 we present our life-cycle model with the estimation in section 5. The core results of the paper are the analysis of the implications of time limits and counterfactual policy simulations presented in section 6. Section 7 concludes. Further details are given in the Online Appendix, where Appendices A to F correspond respectively to sections 1 to 6.

1 Welfare Reform

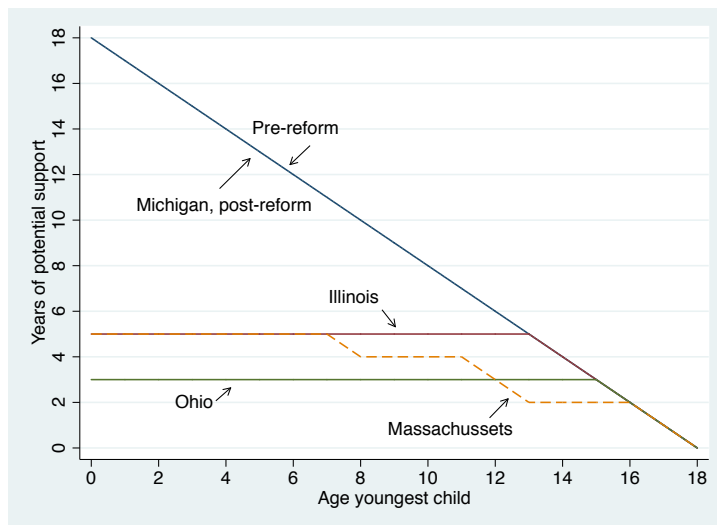
In this section, we present background on the features of the PRWORA reform of 1996, focusing on the changes in eligibility criteria. PRWORA replaced the cash welfare program known as AFDC (effectively an entitlement program) with TANF, which gave states considerable latitude in setting parameters for welfare within broad federal guidelines (Ziliak (2016), Moffitt, Phelan and Winkler (2017)). While AFDC was funded through a state-federal matching system on an unconditional basis, the funding of TANF came from federal block grants assigned to states. There were two further important elements to the reform: the introduction of time limits and the introduction of work requirements.

As regards time limits, federal funds could be used to provide assistance to family units only up to a maximum of sixty months.⁵ However, states could choose lower time limits and about one-third of them did so. States could also set longer limits but would have to cover assistance beyond the statutory limit with state-specific funds. This flexibility meant that TANF varied significantly across states.⁶ Figure 1 illustrates the variation in time limits across selected states, and highlights how the reform changed potential benefits for families with children of different ages. The vertical axis represents the number of years of potential benefits available to the household. This is plotted against the age of the youngest child on the horizontal axis. The blue solid line (Pre-reform) indicates that before the reform a household could claim benefits for as many years as the difference between 18 and the age of the youngest child, provided that it did not have additional children. Post-reform, Michigan maintained a similar regime (at least until 2008). When the post-reform line aligns with the pre-reform line, households with a youngest child of that age are unaffected. The vertical distance between the pre-reform and post-reform line gives a measure of the severity of the decline in welfare. Clearly, mothers of younger children faced deeper cuts in welfare support

⁵The size of the grant was based on a three-year average spending computed over the years preceding the reform, and was independent of the business cycle.

⁶Table A.1 in Appendix A shows how time limits differed across the 50 US states, using 2000 as a baseline.

Figure 1: Time Limits across States



than those with older children.

The second element of reform was the introduction of stronger incentives to work than existed under AFDC, although some work incentives had already been introduced by the Family Support Act of 1988. Work incentives came in two forms: the imposition of work requirements for maintaining welfare eligibility and the availability of child care assistance.⁷ While states were free to set their own policies on who must participate in work activities and what an individual must do, they needed to have at least half of the families on TANF engaged in some work activity in order to avoid fiscal penalties. Such “work” activities (20-30 hours a week) varied across states, but typically consisted of formal work, job training, job search, or educational training. Part of the difficulty with assessing the importance of the work requirement is that individuals in most cases need only to be actively looking for work, rather than being formally employed. This makes measurement and enforcement problematic. Appendix A provides more detail on work requirements. Bruins (2017) argues that these work requirements for single mothers led to an increase in poverty by cutting eligibility for TANF among those most in need. While work requirements and child care assistance may be important and also potential empirical confounders of the effect of time limits, the set of women affected by the three elements of the reform (as a function of the age of the youngest child) do not perfectly overlap. We use this less-than-perfect overlap to isolate the effect of time limits from other elements of the reform.

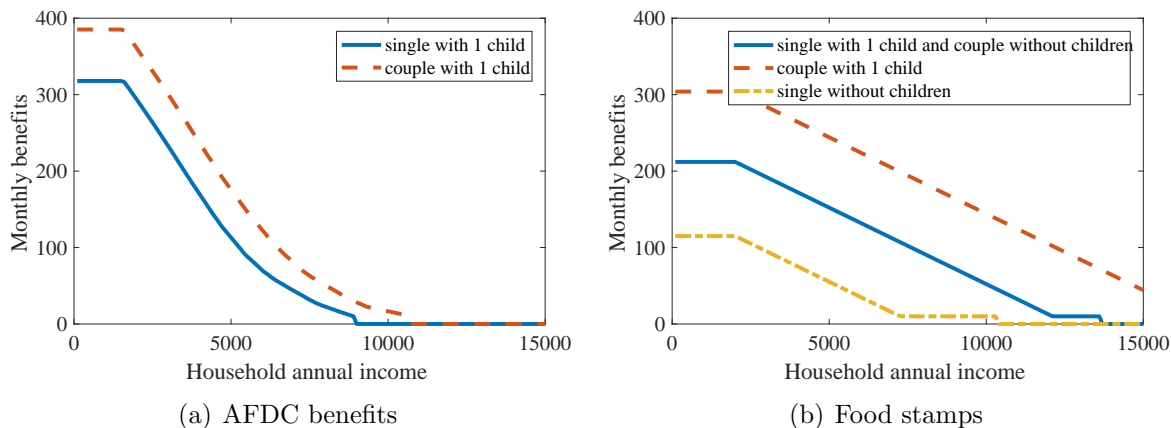
⁷Mothers of children below the age of 1 are, in most states, exempt from work requirements.

A statutory goal of PRWORA was to improve family stability. When AFDC was first introduced, benefits were limited to those who were single but over time eligibility was extended to include couples where the second adult was unemployed. TANF went further and removed any eligibility restriction to single parents or unemployed partners. Further, for most aspects, the implementation of pro-family policies was again left to the states, with little to no accountability at the federal level. These policies include the adoption of family caps, family planning provisions, step-parents' income disregards, etc. Moffitt, Phelan and Winkler (2017) discuss how TANF rules changed the incentives for marriage, and in particular, how the rules reduced the incentives for women to live with the biological father of their children.

These changes in eligibility requirements were not accompanied by appreciable changes in the amount of benefits available conditioning on eligibility. In Figure 2 we show how AFDC benefits vary by income level and by marital status in 1995 (since benefit amounts vary by states, we plot a weighted average using population shares as weights). The plot in the right panel of Figure 2 shows, for comparison, the amount available from the Food Stamps program in the same year. Most states condition AFDC eligibility on having gross income below the poverty line.⁸ A single mother with no sources of income receives approximately \$315 per month in AFDC benefits. If she were to marry a man who has no income, benefits would increase, but by less than needed to keep the adult equivalent amount unchanged. Food stamps are available to all households, irrespective of the presence and of the age of the children, although the amount varies. Eligibility and amounts are determined by household income, including all earnings and AFDC or TANF benefits (eligible households have gross income below 130% of the poverty line). The difference between Food Stamps and AFDC is that the former is a federal program and the in-kind benefits received do not vary across states, while AFDC exhibits substantial variation across states (for example, in 1995, monthly benefits for a couple with 1 child ranged from \$120 in Mississippi to \$923 in Alaska).

⁸In 1995, the poverty line was \$10030 a year for a single mother and \$12590 for a couple with 1 child

Figure 2: AFDC and Food Stamp Amounts



Notes: Monthly AFDC and food stamps benefits by household annual income.

2 Data

We use eight panels of the Survey of Income and Program Participation (SIPP) spanning the 1989-2011 period.⁹ The SIPP is a representative survey of the US population collecting information on participation in welfare and social insurance programs. In each panel, people are interviewed every four months for a certain number of waves.¹⁰ We restrict the sample to individuals between 18 and 60 who are not college graduates, and with at least one child under age 19. We focus on low-skilled individuals because they are the typical recipients of welfare programs. To avoid the well-known “seam effect” in the SIPP (Young, 1989), for each household we keep only the 4th monthly observations in a given wave. Table B.1 in Appendix B describes the sample selection in detail: our main regression includes 64,739 women who are heads or spouses of the head of their household, leading to a total of 406,370 quarterly observations.¹¹

Table 1 summarizes the data. In the first column we report statistics for our regression sample, which is the one we use in our reduced form analysis below. In the other two columns we break this sample into pre- and post-reform period. This is useful both because it gives a first glance at how summary statistics changed after the reform, and because our structural

⁹These are the 1990, 1991, 1992, 1993, 1996, 2001, 2004 and 2008 panels. We do not use the panels conducted between 1984 and 1989 because during this period most states had categorical exclusion of two-parent households from AFDC. This was changed with the Family Support Act of 1988.

¹⁰The number of waves differ by panels. For example, the 1990 panel covers eight waves, while the 1993 panel was conducted for nine waves.

¹¹The reason for focusing on female heads or spouses is that we can more accurately identify whether a minor in the household is the woman’s child (as opposed to, say, a sibling).

Table 1: Summary statistics

Variable	Regression sample	Pre-reform	Post-reform
	(1)	(2)	(3)
On Welfare	0.067	0.106	0.032
On Welfare (married)	0.024	0.036	0.011
On Welfare (unmarr.)	0.170	0.305	0.075
Employed	0.641	0.630	0.651
Employed (married)	0.637	0.636	0.637
Employed (unmarr.)	0.652	0.613	0.679
Divorced or separated	0.162	0.157	0.166
Div/sep if $m_{t-1} = 1$	0.009	0.009	0.010
Married	0.705	0.742	0.672
Married if $m_{t-1} = 0$	0.025	0.025	0.025
Less than high school	0.173	0.174	0.171
High school	0.459	0.484	0.437
Some college	0.368	0.342	0.391
White	0.793	0.816	0.772
Age	36.7	36.0	37.3
Number of children	2.0	2.0	2.0
Age of youngest	7.5	7.3	7.7
<i>Exposed * Post</i>	0.403	0.000	0.761
N. of obs.	406,370	191,127	215,243

Notes: Data from the 1990-2008 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. *Exposed* denotes those affected by the reform and *Post* indicates the post reform period.

estimation uses only pre-reform data (with the added restriction of focusing only on the 1960's birth cohort).¹²

Looking first at column 1, we note that the program participation rate (AFDC/TANF) is 6.7%. However, it is 2.4% for married heads of household and 17% for unmarried heads. There is a 0.9% annual divorce rate and a 2.5% annual marriage rate. The employment rate for married and unmarried women is similar: 64% and 65% respectively. At least from an unconditional point of view, after the reform we observe a reduction in welfare utilization and an increase participation among unmarried women (columns 2 and 3). There is also an increase in the stock of divorcees and a decrease in the stock of married women. In the next section we use state-level heterogeneity in time limits to analyze the impact of time limits on a number of outcomes.

3 Reduced-Form Evidence on Time Limits

3.1 Empirical strategy

We examine the relationship between the introduction of time limits and key outcome variables: welfare benefit utilization, female employment, marital status, and fertility. We compare households that, based on their demographic characteristics and state of residence, could have been affected by time limits with households that were not affected, before and after time limits were introduced. This strategy extends prior work about time limits and benefits utilization (Grogger and Michalopoulos, 2003; Mazzolari and Ragusa, 2012).¹³

We define a variable *Exposed* which takes value 1 if the household's expected benefits have changed as a result of the reform, assuming the household has never used benefits before.¹⁴ The relationship between this exposure variable and the effect of time limits becomes

¹²The structural estimation also uses earnings and wealth data on married and single men and wage data for married and single women, from all cohorts.

¹³In particular, Grogger and Michalopoulos (2003) use experimental data from the Florida Family Transition Program to test whether the introduction of time limits induced "banking" and largely find evidence consistent with time limits affecting welfare use before they become binding. Mazzolari and Ragusa (2012) use SIPP data to regress welfare utilization, employment, and other "income generating" activities against a variable measuring the stock of remaining benefits, which they impute using retrospective information on welfare use and state-specific time limits policies. Their main finding is that for families who are predicted to have hit the time limit, there is evidence that the policy was enforced. Indeed, such households experience a drop in monthly income from welfare of about \$250 on average. Since they find no evidence that such loss is offset by increases in other income sources, they conclude that time limits enforcement resulted in an increase in the rates of deep poverty for households hitting the time limits.

¹⁴For example, if a household's youngest child is aged 13 or above in year t and the state's lifetime limit is 60 months, the variable *Exposed* takes value 0, while if a household's youngest child is aged 12 or below

increasingly attenuated over time as we do not observe the actual history of welfare utilization. Moreover, in most states the reform also imposed stricter work requirements, so that a level effect on employment may be expected across both treated and control groups. However, unless work requirements interact with age of children in a complex way, our strategy still identifies the differential effect of time limits. *Exposed* takes value 0 if a household's benefits (in terms of eligibility or amounts) has not been affected in any way by the reform. Hence, *Exposed* is a function of the demographic characteristics of a household and the rules of the state in which the household resides. Figure 1 reports the potential variation for the year 2000. The value of *Exposed* may change over time because some states change their statutory time limits during the sample period and because states differ with regards to the date where the time limit clock starts to tick.¹⁵

We estimate two versions of the reduced form relationship. First, we construct a variable $Post_{st}$ to indicate the entire post-reform period, based on the timing of the introduction of time limits reported in Mazzolari and Ragusa (2012). We interact this $Post_{st}$ variable with $Exposed_{ds}$ and estimate:

$$y_{idst} = \alpha Exposed_{dst} \times Post_{st} + \mathbf{X}_{idst}\boldsymbol{\beta} + f_{st} + f_{ds} + f_s + f_t + f_d + \epsilon_{idst} \quad (1)$$

where α is the average effect of the reform.

Second, to study the dynamics of the outcome variables, we interact $Exposed_{ds}$ with dummies for each calendar year between 1990 and 2011 (excluding 1995 for scaling). We estimate pre-reform interactions with year dummies to rule out pre-reform trends across demographic groups:

$$y_{idst} = \sum_{\tau=1990}^{2011} \alpha_{\tau} Exposed_{dst} \times \mathbf{1}\{t = \tau\}_t + \mathbf{X}_{idst}\boldsymbol{\beta} + f_{st} + f_{ds} + f_s + f_t + f_d + \epsilon_{idst}. \quad (2)$$

A key question is whether this approach separates the effect of time limits from other features of welfare reform, such as work requirements, increased stigma or childcare provisions. As discussed above, as long as these components of the reform did not affect women differentially depending on the age of their youngest child, their effect would be captured by

in year t and the state's lifetime limit is 60 months, the variable *Exposed* takes value 1.

¹⁵For example, Arizona had the clock starting retroactively in November 1995, while California started the clock in July 1997.

the year fixed effects and by the state-by-year fixed effects f_{st} . However, work requirements were weaker for mothers of very young children. Identification of the effect of time limits stems from comparing employment increases of mothers of younger children with those of mothers of older children and so we will underestimate the increase in employment caused by time limits because mothers of older children were under more pressure to work. On the contrary, childcare provisions, by favoring the employment of mothers of small children, would lead us to overestimate the impact of time limits. We rule out these contamination effects by re-running regression (1) only on a sample of mothers of children between age 6 and 18, who should be facing similar static incentives but be differentially affected by the likelihood of incurring binding time limits.

3.2 Reduced Form Empirical Results

Benefits Utilization and Employment

We start by examining changes in welfare utilization. As shown in Table 1, before the reform, 11% of all households claimed benefits, and among unmarried women, the rate was 31%. These numbers fell to 3% and 8% respectively. Our regressions on utilization, reported in the top half of Table 2, show how much of this fall is due to time limits. Exposed households have a 5 percentage point (pp) lower probability of claiming benefits after the introduction of time limits. Unmarried women have a 15pp lower probability of claiming welfare. This decline among unmarried women is particularly large partly because they were the primary participants of the program before the reform.

The second panel of Table 2 reports the corresponding impact on employment. The introduction of time limits is associated with a 2.3-2.7pp increase in the employment rate of women, while the sample average employment rate is 64%. The result is mostly driven by an 8pp increase in the employment of unmarried women. The increase in employment is likely to be a direct consequence of the decline in welfare utilization, but the size of the increase is only about half of the decline in benefit utilization. We consider two main explanations for the gap: first, some previous recipients move onto alternative social insurance programs (such as Food Stamps, SSI, etc.); second, some have switched from a status of “working and on welfare” to “working and not on welfare”. We examine these two explanations in Appendix C.

First, we consider receipt of food stamps: Table C.1 in Appendix C shows that food

Table 2: Use of Benefits and Employment after Reform

Dependent Var:	AFDC/TANF Utilization					
Sample:	Whole		Married		Unmarried	
$Exposed_{dst}Post_{st}$	-0.0504*** (0.00284)	-0.0502*** (0.00275)	-0.0180*** (0.00243)	-0.0177*** (0.00234)	-0.157*** (0.0102)	-0.154*** (0.00935)
Obs	406,370	406,370	286,425	286,425	119,945	119,945
R^2	0.070	0.108	0.043	0.058	0.172	0.196

Dependent Var:	Employed					
Sample:	Whole		Married		Unmarried	
$Exposed_{dst}Post_{st}$	0.0266*** (0.00558)	0.0232*** (0.00486)	0.00277 (0.00713)	0.00122 (0.00646)	0.0820*** (0.0101)	0.0727*** (0.00976)
Obs	406,370	406,370	286,425	286,425	119,945	119,945
R^2	0.060	0.108	0.055	0.085	0.096	0.202

Basic controls	Yes	Yes	Yes	Yes	Yes	Yes
Race	No	Yes	No	Yes	No	Yes
Disability status	No	Yes	No	Yes	No	Yes
Unemp*Demog.	No	Yes	No	Yes	No	Yes

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. Data from the 1990-2008 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state fixed effects, demographics fixed effects, state-by-demographics fixed effects, state-by-year fixed effects. Standard errors in parentheses, clustered at the state level.

stamp use changes only for unmarried women, where there is a 5pp decline in use. Such a decline is to be expected given the substantial increase in the employment probability, but it does not resolve the gap between fall in numbers on welfare and the increase in numbers employed.¹⁶ Second, we report in Table C.2, Appendix C, the change of the composition of unmarried women across four groups: women who are both employed and on welfare (-3pp), women who are employed and not on welfare (+11pp), women who are not employed but are on welfare (-13pp) and women who are not employed nor on welfare (+5pp).

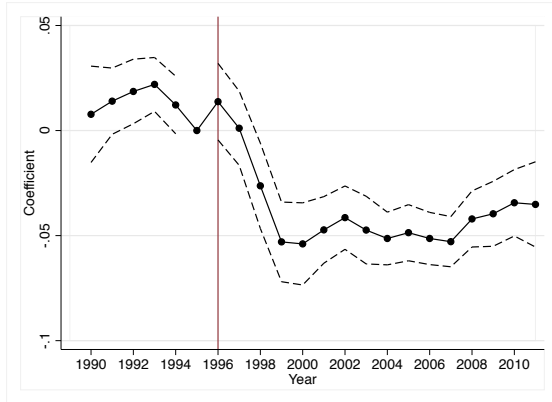
The decline in employed women on welfare reinforces the view that work requirements are unlikely to be driving the decline in welfare use. However, the fraction of single women on AFDC who report being employed pre-reform is 19.8%, and the fraction of single women on TANF who report being employed post-reform is 25.8%. What this masks is that the increase in employment is greatest for mothers of children below age 6 (from 19.0% to 26.7%) compared to mothers of older children (from 21.2% to 24.3%), and the work requirement is more likely to be enforced on the group of mothers of older children.

These average effects on utilization and on employment mask the dynamics of behavior and also whether or not time-limits induced banking of benefits. Our second regression, equation (2), addresses this. Figure 3 plots the coefficients of the variable *Exposed* interacted with year dummies. The utilization rate already begins to decline significantly in 1998, to a persistent drop of 7 percentage points by 1999. Households reduce their benefit utilization *before* anyone is likely to have run out of benefits eligibility. Similar time patterns are observed by marital status.

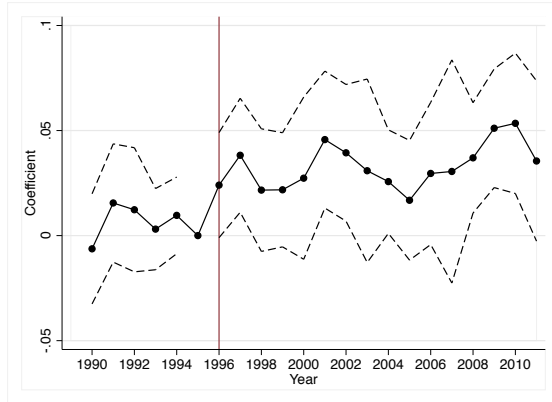
We provide additional evidence of forward looking behavior in two ways: first, we count the number of years since the introduction of time limits in each state and re-define the annual exposure dummy as $Exposed_{dst} \times \mathbf{1}\{\tau \text{ years since } TL\}_{st}$. In Figure C.1 in Appendix C, we show the decline in welfare use on the sample that excludes states with shorter time limits (less or equal to 24 months). This immediate decline is not consistent with the decline being purely mechanical and only beginning when limits actually bind. The pre-emptive decline suggests forward looking behavior (in the form of “banking” of benefits). Second, we interact the exposure coefficient with dummy variables for the age of the youngest child. Figure C.2 shows that the parents of younger children respond more strongly to the dynamic incentives created by the introduction of time limits compared to parents whose child is closer to age 13. This is consistent with the idea that the cuts in insurance induced by the introduction of

¹⁶On the other hand, there is greater use of EITC due to the work requirement of this program.

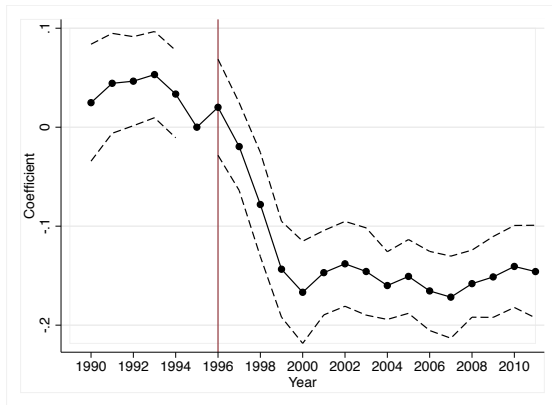
Figure 3: Program Participation and Employment Dynamics



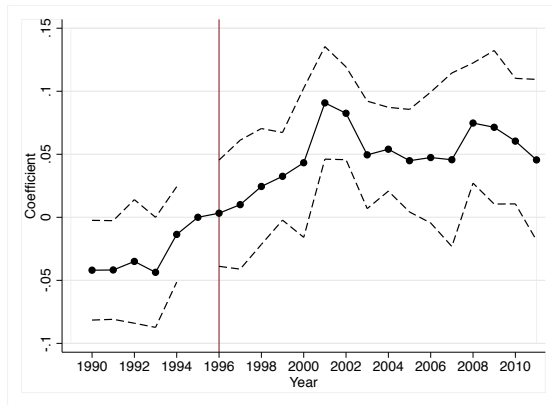
(a) Welfare: Everyone



(b) Employment: Everyone



(c) Welfare: Unmarried



(d) Employment: Unmarried

Notes: Data from the 1990-2008 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state fixed effects, demographics fixed effects, state-by-demographics fixed effects, state-by-year fixed effects, race and disability status.

time limits were deeper for women with younger children, who hence reduced their welfare use proportionally.

In Appendix C we perform a number of additional robustness checks. In Table C.3, we include age-by-year fixed effects. In Table C.4, we consider only the first 2 waves of the SIPP for every household to address potential selective attrition. In Table C.5, as a falsification exercise, we use a sample of college graduates and show that their behavior is little affected by time limits. In Table C.6, we only consider mothers of children aged 6 and above. The similarity of our findings shows that our results are not driven by mothers of children below age 6, who may respond to other features of welfare reform, like childcare provisions. Finally, in Table C.8 and C.9 and in Figure C.3, we repeat our analysis using the March Current Population Survey. Our findings are consistent throughout the robustness analysis.

Marriage formation and dissolution

A central motivation (and indeed a stated goal) of the 1996 welfare reform was to encourage “the formation and maintenance of two-parent families”. To the extent that the reform changed household formation, the results on changes in benefit use and employment for unmarried women will contain both a treatment effect and a composition effect. This issue is rarely if ever considered in typical reduced form analyses of the welfare reform. We use our model in section 4 to decompose these effects. In this section, we consider whether there is direct evidence of effects of time limits on divorce and marriage.

We first consider the impact of the welfare reform on the probability of being divorced or separated for women. In the top panel of Table 3, we look both at flows and stocks. We find a significant 0.19pp decline in the probability of transitioning into divorce conditional on being married during the previous interview. Since on average 0.9 percent of marriages end in divorce each year, this is a non-negligible effect. In terms of stocks, women exposed to time limits are 3pp less likely to be divorced or separated. Figure 4 reports the dynamics of these outcomes. The second panel of Table 3 shows that this decline in divorce was not associated with an increase in marriage.

Thus, more people appear to remain married, but at the same time no change is observed in the stock of married people, indicating a potential decline in new marriages outside of our sample of mothers. In theory, as discussed by Bitler et al. (2004), the effects of the welfare reform on household formation and dissolution are not obvious. The welfare reform, by curtailing the extent of public insurance available to low-income women, may have induced

Table 3: Marital Status

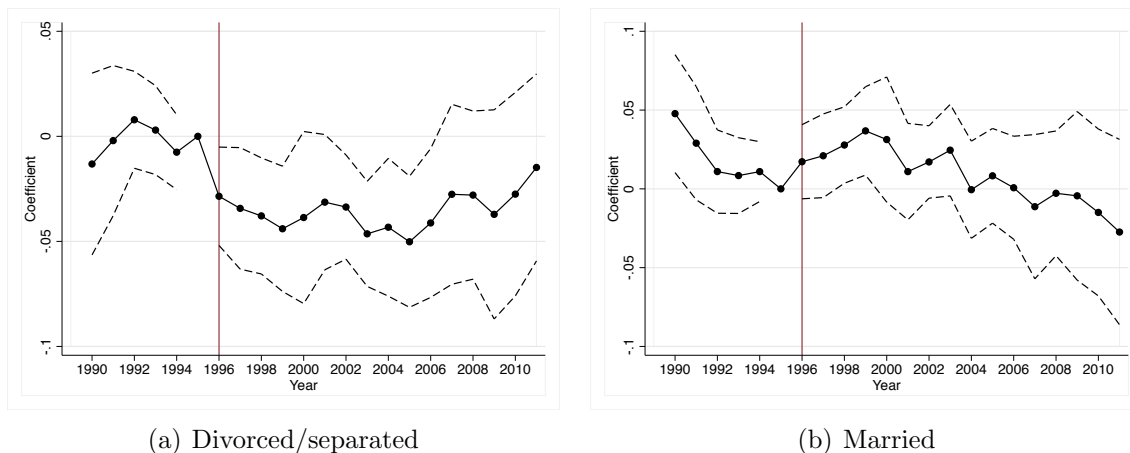
Dependent Var:	Gets Divorced/ separated		Divorce/ separation	
Sample:	$m_{t-1} = 1$		Whole	
$Exposed_{dst}Post_{st}$	-0.00186** (0.000821)	-0.00185** (0.000832)	-0.0301*** (0.00730)	-0.0288*** (0.00691)
Obs	242,963	242,963	406,370	406,370
R^2	0.007	0.007	0.011	0.020
Dependent Var:	Gets Married		Married	
Sample:	$m_{t-1} = 0$		Whole	
$Exposed_{dst}Post_{st}$	-0.00345 (0.00239)	-0.00454* (0.00238)	-0.00621 (0.00920)	-0.00862 (0.00846)
Obs	99,650	99,650	406,370	406,370
R^2	0.018	0.022	0.044	0.109
Basic controls	Yes	Yes	Yes	Yes
Race	No	Yes	No	Yes
Disability status	No	Yes	No	Yes
Unemp.*Demog.	No	Yes	No	Yes

Standard errors in parentheses clustered at the state level.

*** p<0.01, ** p<0.05, * p<0.1

Notes: Data from the 1990-2008 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state fixed effects, demographics fixed effects, state-by-demographics fixed effects, state-by-year fixed effects. Standard errors in parentheses, clustered at the state level.

Figure 4: Other dynamics



*Notes:*Data from the 1990-2008 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state fixed effects, demographics fixed effects, state-by-demographics fixed effects, state-by-year fixed effects, race and disability status.

those who were already married to attach a higher value to marriage as a valuable risk sharing tool (through male labor supply, for example), and reduced the option value of being single (and potentially claiming benefits). In other words, the reform increased the gains from marriage for those who would be eligible for welfare benefits if single and this immediately translates into lower divorce rates. However, moving from being single to married takes time because of search costs and not every match will translate into a marriage that is beneficial. As a result, any tendency for an increased marriage rate is attenuated and becomes too small to detect in the data.

The changes in insurance will not only affect the likelihood of marriage, but also women's bargaining power within marriage and allocations of resources within marriage, as we capture in our model below.

Fertility

Our empirical strategy relies on the age of the youngest child as a source of predetermined variation, and as a result is not suitable for estimating fertility outcomes, which directly affect the age of the youngest child. To examine whether time limits influenced fertility outcomes, we focus instead on the probability that a household will have a newborn (a child below age

1) in the following year, with the specification:

$$Newborn_{idst+1} = \gamma Exposed_{dst} * Post_{st} + \mathbf{X}_{idst}\boldsymbol{\beta} + f_{st} + f_{ds} + f_s + f_t + f_d + \epsilon_{idst}$$

Table C.7 in Appendix C reports the results of estimating this regression on the whole sample and on subsamples that depend on marital status. In no specification do we find that exposure to time limits influences the probability of future births, irrespective of marital status. This partly justifies our choice to treat fertility choices in the model as exogenous (albeit stochastic).

4 Life-Cycle Model

The model we now present, motivated by the earlier facts, captures the evolution over the life-cycle of family formation in a limited commitment framework. At the beginning of each period men (M) and women (F) observe their productivity realizations and women learn whether they have a child, as a function of their marital status and their age at the beginning of the period. If single, people may meet a partner of the same age group, drawn from the distribution of singles, and decide whether to get married. If they are married, they observe the realization of a shock to match quality, as well as individual specific shocks to wages, and decide whether or not to stay together. In all cases they also decide whether to work or not, and how much to save/consume. We allow for savings, both because these households do hold assets and because any analysis of the welfare effects of the reform would be incomplete without taking explicitly into account self-insurance. The aim of the model is to capture how individuals mitigate increased exposure to risk due to the reduction in the option of claiming welfare because of time limits. We have seen in the reduced form that the extra exposure increases labor supply and reduces divorce. The role of limited commitment is to allow for reallocation of household resources in response to changes in the outside option: whenever an individual has a credible threat to leave the relationship she/he can renegotiate the allocation of resources to her/his benefit. Since limited commitment (as opposed to full commitment) reduces the insurance value of marriage and hence the gains from marriage, it may lead to fewer marriages taking place. On the other hand *conditional* on marriage, limited commitment will not increase divorce, relative to the full commitment world: if there are gains to marriage the partners can negotiate a sharing rule and remain

together. Thus our interest in allowing for limited commitment is to estimate empirically the extent to which reforms affect intrahousehold allocations.

4.1 Problem of the single woman

We start by describing the problem of a single woman (F_s) with low education and who has completed her schooling choices,¹⁷ where s indicates that individual i is single, and F indicates that i is a woman. At each age t , she decides whether to work ($P_{i,t}^{F_s} \in \{0, 1\}$), how much to consume and save ($c_{i,t}^{F_s}$) and whether to claim AFDC/TANF. The decision to claim is given by $B_{i,t}^{F_s} \in \{0, 1\}$; and this leads to benefit payment $AFDC_{i,t}$. The within-period preferences for a woman, conditional on being single, are denoted by $u^{F_s}(c_{i,t}^{F_s}, P_{i,t}^{F_s}, B_{i,t}^{F_s})$, where the dependence on $B_{i,t}^{F_s}$ reflects stigma. In addition, she makes a choice to marry, which will also depend on meeting a man and whether he will agree. The decision to marry takes place at the start of the period, after all shocks are realized, but before any consumption, welfare participation, or work plan are implemented. Employment, savings and program participation decisions will be conditional on the marriage decision that occurs at the beginning of the period.

If she remains single, her budget constraint is given by

$$\frac{A_{i,t+1}^{F_s}}{1+r} = A_{i,t}^{F_s} - \frac{c_{i,t}^{F_s}}{e(k_{i,t}^a)} + (w_{i,t}^F - CC_{i,t}^a)P_{i,t}^{F_s} + B_{i,t}^{F_s}AFDC_{i,t} + FOOD_{i,t} + EITC_{i,t} \quad (3)$$

$$A_{i,t+1}^{F_s} \geq 0$$

where the subscript t denotes age, $A_{i,t}^{F_s}$ are assets and $CC_{i,t}^a$ is the financial cost of childcare paid if the woman works. The woman's wage rate $w_{i,t}^F$ follows a persistent stochastic process, detailed below. The term $e(k_{i,t}^a)$ is an equivalence scale to account for the presence of children ($k_{i,t}^a$) and their consumption cost: it is the amount of consumption bought by spending \$1. Hence children affect consumption, benefit eligibility and the opportunity cost of women's time on the labor market. We account for three important social safety net programs: food stamps (denoted $FOOD$), EITC and AFDC. The latter can be subject to time limits. The value of all programs depends on demographics and income.

¹⁷Since we are interested in the impacts of means-tested welfare benefits, such as TANF, we focus on low-education women. An important question is how education choices are themselves affected by the presence of welfare benefits (Blundell et al., 2016). Bronson (2014) studies women's education decisions in a dynamic collective model of the household with limited commitment.

The state space for a single woman is $\Omega_{i,t}^{F_s} = \{A_{i,t}^{F_s}, w_{i,t}^F, k_{i,t}^a, TB_{i,t}\}$, where $TB_{i,t}$ is the number of time periods the woman has claimed the time-limited benefits; this is only relevant when time limits apply. Food stamps and EITC are functions of the vector $\{k_{i,t}^a, w_{i,t}^F, P_{i,t}^{F_s}\}$, while AFDC is a function of the vector $\{k_{i,t}^a, w_{i,t}^F, P_{i,t}^{F_s}, TB_{i,t}\}$.

With probability λ_t , at the beginning of the period a single woman meets a man with characteristics $\{A_{j,t}^M, y_{j,t}^M\}$ (assets and exogenous earnings) and together they draw an initial match quality $L_{ij,t}^0$. If the match is formed $m_{ij,t}() = 1$, otherwise it is zero. The process governing this decision, which involves both partners, is described below in Section 4.3. We restrict encounters to be between a man and a woman with a 2-year age gap.¹⁸

We denote by $V_t^{F_s}(\Omega_{i,t}^{F_s})$ the value function for a single woman at age t and $V_t^{F_m}(\Omega_{i,t}^m)$ the value function for a married woman at age t , which we will define below. A single woman has the following value function:

$$V_t^{F_s}(\Omega_{i,t}^{F_s}) = \max_{q_{i,t}^{F_s}} \left\{ \begin{array}{l} u^{F_s}(c_{i,t}^{F_s}, P_{i,t}^{F_s}, B_{i,t}^{F_s}) \\ + \beta E_t \left[\begin{array}{l} \lambda_{t+1} \left[\begin{array}{l} (1 - m_{ij,t+1}) V_{t+1}^{F_s}(\Omega_{i,t+1}^{F_s}) \\ + m_{ij,t+1} V_{t+1}^{F_m}(\Omega_{i,t+1}^m) \end{array} \right] \\ + (1 - \lambda_{t+1}) V_{t+1}^{F_s}(\Omega_{i,t+1}^{F_s}) \end{array} \right] \end{array} \right\}$$

subject to the intertemporal budget constraint (3).

We parametrize within-period utility for the woman as follows:

$$u(c, P, B) = \frac{(c \cdot e^{\psi(m, k^a) \cdot P})^{1-\gamma}}{1-\gamma} - \eta B. \quad (4)$$

When a woman works ($P = 1$), her marginal utility of consumption changes if she has a child. The parameter η represents the utility cost from claiming AFDC/TANF benefits.¹⁹

4.2 Problem of the single man

Single men are all assumed to work, and not to have an option to receive welfare benefits (other than food stamps). Children affect the man's problem only when he is married to

¹⁸In principle, this distribution is endogenous and as economic conditions change, the associated marriage market will change, as supply and demand changes. In this paper we take this distribution as given and do not solve for it endogenously. This mainly affects counterfactual simulations. Note that solving for the equilibrium distribution in two dimensions is likely to be very complicated computationally.

¹⁹We assume there is no utility cost from claiming Food Stamps or EITC benefits because for these two programs we do not endogenize the participation decision.

the child's mother. The state space for a single man is defined by his asset holdings and his wage: $\Omega_{j,t}^{M_s} = \{A_{j,t}^M, y_{j,t}^M\}$. The state space when married, $\Omega_{ij,t}^m$ contains both the husband and wife's economic state variables. These assumptions determine $V^{M_s}(\Omega_{j,t}^{M_s})$, the man's value function when he is single and $V_t^{M_m}(\Omega_{ij,t}^m)$, the value accruing to a married man. The subscript j indexes the man, the subscript i indexes the woman, and ij indexes the corresponding couple.

The budget constraint of the single man is given by:²⁰

$$\frac{A_{j,t+1}^{M_s}}{1+r} = A_{j,t}^{M_s} - c_{j,t}^{M_s} + y_{j,t}^M + FOOD_{j,t} \quad (5)$$

$$A_{j,t+1}^{M_s} \geq 0$$

The problem for the single man is thus defined by

$$V_t^{M_s}(\Omega_{j,t}^{M_s}) = \max_{c_{j,t}^M} \left\{ \begin{array}{l} u^{M_s}(c_{j,t}^{M_s}) \\ + \beta E_t \left[\begin{array}{l} \lambda_{t+1} \left[\begin{array}{l} (1 - m_{ij,t+1}) V_{t+1}^{M_s}(\Omega_{j,t+1}^{M_s}) \\ + m_{ij,t+1} V_{t+1}^{M_m}(\Omega_{ij,t+1}^m) \end{array} \right] \\ + (1 - \lambda_{t+1}) V_{t+1}^{M_s}(\Omega_{j,t+1}^{M_s}) \end{array} \right] \end{array} \right\}$$

This problem is similar but more complex than the simple consumption smoothing and precautionary savings problem because assets affect the probability of marriage as well as the share of consumption when married.

Since single men always work, within period utility is a special case of (4) and takes the CRRA form:

$$u(c) = \frac{(c \cdot e^{\psi})^{1-\gamma}}{(1-\gamma)}$$

to guarantee symmetry with the women's utility function.

4.3 Problem of the couple

When a couple marries their assets are merged and they solve a dynamic collective problem with limited commitment (Mazzocco, 2007; Voena, 2015) which leads to a given allocation of consumption within the households, to an agreed amount of common savings, as well as

²⁰We do not consider EITC for men because the value of the program for an individual without a qualifying child is modest (for example, in 2017 the maximum annual credit for an individual without a qualifying child was \$510, as opposed to \$3,400 for those with a qualifying child).

to female labor supply. These allocations depend on the bargaining weight of each member, which in turn depend on the outside option of being single and possibly remarrying.

In the optimization problem that follows the bargaining weight of each household member is reflected in the Pareto weights $(\theta_{ij,t}^M, \theta_{ij,t}^F)$, which evolve endogenously. We first define the optimization problem for the married couple, given the overall state space $\Omega_{ij,t}^m$.²¹ We then discuss transitions between marital states and the evolution of the Pareto weights.

At the start of the next period, the couple may divorce $(d_{ij,t+1} \in \{1, 0\})$, and so the joint problem that the couple solves is

$$V_t^m(\Omega_{ij,t}^m) = \max_{\mathbf{q}_{ij,t}^m} \left\{ \begin{array}{l} \theta_{ij,t}^F u^{F_m}(c_{ij,t}^F, P_{ij,t}^F, B_{ij,t}^F) + \theta_{ij,t}^M u^{M_m}(c_{ij,t}^M) + \varsigma_{ij,t} \\ + \beta E_t \left[\begin{array}{l} (1 - d_{ij,t+1}) V_{t+1}^m(\Omega_{ij,t+1}^m) \\ + d_{ij,t+1} (\theta_{ij,t}^F V_{t+1}^{F_s}(\Omega_{i,t+1}^{F_s}) + \theta_{ij,t}^M V_{t+1}^{M_s}(\Omega_{j,t+1}^{M_s})) \end{array} \right] \end{array} \right\} \quad (6)$$

where $\mathbf{q}_{ij,t}^m = \{c_{ij,t}^M, c_{ij,t}^F, P_{ij,t}^F, B_{ij,t}^F\}$ are the choices of the ij^{th} couple, and $\varsigma_{ij,t}$ is the match quality of the couple, which evolves according to a random walk process:

$$\varsigma_{ij,t} = \varsigma_{ij,t-1} + \xi_{ij,t} \quad (7)$$

This process is initialized at the start of the marriage and $\xi_{ij,t}$ is interpreted as a ‘‘love shock’’ to the marriage.

The optimization above is subject to the budget constraint:

$$\frac{A_{ij,t+1}}{1+r} = A_{ij,t} - x(c_{ij,t}^F, c_{ij,t}^M, k_{i,t}^a) + (w_{i,t}^F - CC_t^a) P_{ij,t}^F + y_{j,t}^M + B_{ij,t} AFDC_{ij,t} + FOOD_{ij,t} + EITC_{ij,t}$$

To capture economies of scale in marriage (including public goods), we assume that the individual consumptions $c_{ij,t}^F$ and $c_{ij,t}^M$ and the equivalence scale $e(k_{i,t}^a)$ imply an aggregate household expenditure of

$$x_{ij,t} = \frac{((c_{ij,t}^F)^\rho + (c_{ij,t}^M)^\rho)^{\frac{1}{\rho}}}{e(k_{i,t}^a)}$$

The extent of economies of scale is controlled by ρ and $e(k_{i,t}^a)$. If $\rho > 1$, consumption is partially public, and the sum of spouses’ consumption exceed what they would consume if

²¹The state variables for the couple (represented by $\Omega_{ij,t}^m$), are: the combined assets, each spouses’ productivity, the number of periods of welfare benefits utilization, age of any child present ($k_{i,t}^a$), and the Pareto weights $\theta_{ij,t}^M, \theta_{ij,t}^F$.

single and spending the same amount. If the couple divorce at the start of $t + 1$, then $A_{ij,t+1}$ is divided equally.²²

There are two transitions to consider: first, for couples, whether an existing marriage continues, or ends in divorce; second, for singles, whether a meeting between a single man and single woman results in marriage.

For the marriage to continue, individual participation constraints need to be satisfied. These state that the value of marriage must be larger than the value of being single for both spouses and are given by:

$$\begin{aligned} V_{t+1}^{Fm} \left(\Omega_{ij,t+1}^m \right) &\geq V_{t+1}^{Fs} \left(\Omega_{i,t+1}^{Fs} \right) \\ V_{t+1}^{Mm} \left(\Omega_{ij,t+1}^m \right) &\geq V_{t+1}^{Ms} \left(\Omega_{j,t+1}^{Ms} \right) \end{aligned} \tag{8}$$

When married, the Pareto weights remain unchanged as long as these participation constraints are satisfied. However, the various shocks, including those to match quality and to the wages of each partner, can change the value of becoming single and remaining married. If one partner's participation constraint is not satisfied the Pareto weight moves the minimal amount needed to satisfy it. This is consistent with the dynamic contracting literature with limited commitment, such as Kocherlakota (1996) and Ligon, Thomas and Worrall (2002*a*). If it is not feasible to satisfy both spouses' participation constraints and the intertemporal budget constraint for any allocation of resources, then divorce follows.²³ Divorce can take place unilaterally, and if divorce takes place, it is efficient because there is no allocation such that each person can have a positive surplus from remaining married. This is equivalent to saying that there exists no feasible allocation and corresponding Pareto weights $\theta_{ij,t}$ which satisfy the participation constraints in equation 8.

In our context, marriage is not a pure risk sharing contract. Marriage also takes place because of complementarities (i.e., economies of scale in consumption), love (ς), and possibly also because features of the welfare system promote it. And indeed, marriage can break down efficiently if no Pareto weights imply positive gains from marriage for each partner.

²²This assumption is a good approximation of the legal position (see (Voena, 2015)). After marriage, spouses' assets merge into one value: $A_{ij,t} = A_{i,t}^{Fs} + A_{j,t}^{Ms}$ and so there is no need, in the computation, to keep track of individual assets going into the marriage.

²³We can rewrite equation 6 as the weighted sum of the value of being married for men and for women at time t . This means that in making time t decisions, the weight on time $t + 1$ outcomes is determined by the time t weights even if divorce occurs. Of course, if divorce occurs in time $t + 1$, then the man and the women only optimize over their own utility, as shown by single men and single women having their own budget constraint.

However, when marriage has the potential to be better than being single for *both* parties, overall transfers will take place and this will *de facto* lead to at least partial risk sharing. Suppose, for instance, the female wage drops relative to the male one; the husband may end up transferring resources because single life may have become relatively more attractive to the wife, say because of government transfers to single mothers.

The second transition to consider is for single individuals getting married. Whether a meeting between a single man and single woman results in marriage depends on the existence of a feasible allocation that satisfies both participation equations (equation 8). First, because of search frictions, the relevant outside option for marriage is waiting longer for an alternative partner. Second, for a number of matches there will be gains to be made over and above the outside option. The Pareto weights at the time of marriage, θ_{ij,t_0} , distribute these gains and we assume they are chosen as the solutions to a symmetric Nash bargaining game between spouses (Mazzocco, Yamaguchi and Ruiz, 2013):

$$\max_{\theta_{ij,t_0}} (V_{t_0}^{F_m}(\theta_{ij,t_0}) - V_{t_0}^{F_s}) \cdot (V_{t_0}^{M_m}(\theta_{ij,t_0}) - V_{t_0}^{M_s}).$$

The solution to this maximisation is to set the weights so that each individual receives the same mark-up over their reservation value, V_{t_0} , if (and only if) $\frac{\partial V_{t_0}^{F_m}(\theta)}{\partial \theta} = -\frac{\partial V_{t_0}^{M_m}(\theta)}{\partial \theta}$. Hence, in that case, the solution is to split the surplus of the marriage equally. More generally, the implication of this bargaining game is that if the outside option for a woman worsens, for example due to a reduction in benefits for single women, the weight on that woman will decline and so her consumption share will decline.

Ours is a context of imperfectly transferable utility, which implies that the Pareto weight affects the size of the gains to be shared. Since the outside option depends on the possibility of future marriages, the anticipated future shares of the gains will also affect the probability of marriage in this indirect way.^{24,25}

²⁴Characterizations of equilibrium in *transferable utility contexts* is given in Chiappori, Costa-Dias and Meghir (2016) for frictionless environments and by Goussé, Jacquemet and Robin (2017) for a stationary environment with frictions.

²⁵In the estimation below, we experiment with varying the share of the surplus that goes to each person. This has some effect on parameter estimates for the cost of working for married women, which need to be larger to deter participation when men have all of the surplus. However, our conclusions on the effects of time limits for behavior are unchanged.

4.4 Uncertainty

Underlying the choices described above for single women, single men and couples, there are three sources of uncertainty which we discuss in turn: female wages and male earnings; fertility; the marriage market.

Female wages and male earnings

Male earnings and female wages are respectively specified as

$$\begin{aligned}\log(y_{j,t}^M) &= a_0^M + a_1^M t + a_2^M t^2 + z_{jt}^M + \varepsilon_{jt}^M \\ \log(w_{it}^F) &= a_0^F + a_1^F t + a_2^F t^2 + z_{it}^F + \varepsilon_{it}^F \\ z_{jt}^M &= z_{j,t-1}^M + \zeta_{jt}^M \\ z_{it}^F &= z_{i,t-1}^F + \zeta_{it}^F.\end{aligned}$$

z_{it}^K ($K = F, M$) is permanent income, which evolves as a random walk following innovation ζ_{it}^K . We treat the i.i.d. shock ε_{it}^K as measurement error.²⁶ While male earnings are taken to be always positive, many women do not work. Thus, in estimation of the female wage process we will take into account endogenous selection into work to ensure unbiased estimates of the stochastic process.

Fertility

The arrival of children is stochastic and exogenous, albeit varying with the woman's marital status and age. The conditional probability of having a child is taken to be

$$Pr(k_{t+1}^a | k_t^a = 0, m_t, t). \quad (9)$$

Each woman can only have one child. This restriction is imposed for computational reasons because it limits the size of the state space, which is already large. Since the probability depends on marriage, fertility is partially endogenized through the marital decision.

²⁶One issue is the extent to which welfare reform affected the labor market and in particular human capital prices (Rothstein, 2010). Whether such general equilibrium effects are important or not depends on the extent to which the skills of those affected by the welfare reforms are substitutable or otherwise with respect to the rest of the population. With reasonable amounts of substitutability we do not expect important general equilibrium effects.

The marriage market

We parameterize the rate of arrival of meetings λ_t to vary with female age t according to the following rule to allow marriage market opportunities to vary as people become older:

$$\lambda_t = \min\{\max\{\lambda_0 + \lambda_1 \cdot (t - 1) + \lambda_2 \cdot (t - 1)^2, 0\}, 1\}.$$

When two individuals meet, at time t_0 they draw an initial match quality L_{ij,t_0} from a distribution $N(0, \sigma_0^2)$. Thereafter match quality evolves as a random walk given by equation (7). The innovations to match quality $\xi_{ij,t}$ are from a distribution $N(0, \sigma_\xi^2)$ and we allow the distribution of the initial match quality to differ from the distribution of subsequent innovations.

5 Estimation of Model Parameters

We select parameters of the model in three steps. First, some parameters are set using standard values in the literature. Second, we estimate some parameters directly from the data without imposing the model's structure. Finally, remaining parameters are estimated using the method of simulated moments, matching data and model-based simulated moments. We use moments based on pre-reform data and use post-reform data to validate the model. We describe the parameters of the model set in each of the three steps in turn.

5.1 Externally Set Parameters

Panel A of Table 4 reports parameters taken from external sources. We set the coefficient of relative risk aversion to 1.5 based on Blundell, Browning and Meghir (1994) and Attanasio and Weber (1995), the discount factor to 0.98 and the interest rate to 1.5% following Attanasio, Low and Sanchez-Marcos (2008). We set the parameter defining economies of scale in marriage from Voena (2015).

We compute parameters of the AFDC, food stamps and EITC benefit programs directly from the program rules. Eligibility for these benefits is based on a combination of economic and demographic criteria. All adult earnings within the household determine income eligibility for AFDC. We calculate AFDC benefit for different household composition and income levels by taking a population-weighted average value of the benefit across states, as reported in Figure 2.

5.2 Directly Estimated Parameters

Childcare Costs

We estimate childcare costs using information from the Consumer Expenditure Survey for the 1990-1996 period. In our model childcare costs are only incurred by working women. We use the average of total spending on day-care and babysitting for working women in the data (by child age) to be the relevant child care cost.

The Fertility Process

We allow each household to have up to one child, and compute the transition probability from no children to one child using SIPP data. We compute the Markov process for fertility by examining transition probabilities in the SIPP data as a function of a woman’s age and marital status, as in equation (9). We report the estimated transition probabilities in Figure E.1, Appendix E.

The Distribution of Characteristics of Single Men and Women

Individuals in the model use the age-dependent distribution of characteristics for partners that we observe in the data to form expectations about the matches they may be involved in. At each age men are characterized by the distribution of $\{A_t, y_t\}$. Women are characterized by the age dependent distribution for $\{\log(A_t), \log(y_t), k_t\}$, where k_t indicates whether or not she has had a child.²⁷ We also allow for additional mass for the cases in which $A_t^j = 0$. We use the same selection correction procedure described below for wages to estimate the distribution of single women’s offer wages for those single women who do not work.

Earnings processes

For women we model the hourly wage rate.²⁸ Since we do not model the intensive margin, we assume that participating women work 1530 hours per year (median hours worked in the

²⁷The bivariate normality assumption may be inappropriate as a characterization for the whole population, due to the long right tail in both assets and income, but it is less problematic for our low education sample who have low income and assets.

²⁸Female wages are the sum of the reported earnings within a year divided by annual hours. The latter are computed as the reported weekly “usual hours of work” \times the number of weeks at the job within the month \times number of months the individual reported positive earnings. We drop individuals whose hourly wage is less than half the minimum wage in the years she reported working and we drop observations whose percentage growth of average hourly earnings is lower than -70% or higher than 400% .

data). This avoids labelling fluctuations in hours worked as uncertainty: productivity shocks are the only source of uncertainty in earnings. We also need to address selection into employment. We implement a two-step Heckman selection correction procedure, described in Appendix E. The exclusion restrictions in the employment equation are “simulated” welfare benefits, following Low and Pistaferri (2015). In particular, we use state, year and demographic variation in simulated AFDC, EITC and food stamps benefits for a single mother with varying number of children who works part-time at the federal minimum wage. Since we also control for time and state effects the instruments capture differential changes in policy over time and states. The first stage showing the strength of the instruments is reported in Table E.1 in Appendix E. We use the selection correction to estimate the age profile of a woman’s wage and to account for non-participation when estimating the variance of the productivity innovations.

For men, we estimate the variance of the permanent component of log annual earnings ($\sigma_{\zeta_M}^2$) and the variance of the measurement error (σ_ε^2), using GMM as described in Appendix E. We do not correct for selection for men.

In Panel B of Table 4 we report wage process parameters. Both male and female earnings are subject to relatively high variance of permanent shocks (0.027 and 0.038 respectively). Initial heterogeneity is large, with a variance of initial wages for men and women of approximately 0.18 and 0.15 respectively, implying large initial dispersion in productivities. This also reflects differences in schooling among our (non-college graduate) group. Male and female wages have a concave lifecycle profile as expected.

Initial conditions

Our lifecycle model starts at age 21. By that age some women have already experienced marriage, divorce or childbirth. Therefore, we estimate demographic and fertility initial conditions directly from the SIPP: proportions of women married and divorced at age 20, and the proportion of women age 20 who have a child, separately for married and single women. These are reported in Panel C of Table 4.

Table 4: Externally Set and Directly Estimated Parameters of the Model

Parameter	Value/source
<i>Panel A: Externally Set Parameters</i>	
Relative risk aversion (γ)	1.5
Discount factor (β)	0.98
Economies of scale in marriage (ρ)	1.23
Welfare program parameters	Statutory rules
<i>Panel B: Directly Estimated Parameters</i>	
Childcare costs (CC^a)	CEX (see text)
Fertility process	SIPP (Figure E.1)
Distribution of single characteristics	SIPP (see text)
Variance of men's unexplained earnings in period 1	0.18
Variance of women's unexplained wages in period 1	0.15
Variance of men's earnings shocks	0.027
Variance of women's wage shocks	0.038
Life cycle profile of log male earnings (a_0^M, a_1^M, a_2^M)	9.76, 0.043, -0.001
Life cycle profile of log female wages (a_0^F, a_1^F, a_2^F)	1.96, 0.022, -0.0003
<i>Panel C: Initial conditions</i>	
% married at age 20	24.35%
% divorced at age 20	3.90%
% with one child at age 20 (married)	55%
% with one child at age 20 (single and living alone)	48%

5.3 Estimation: Method of Simulated Moments

We estimate the remaining parameters of the model by the Method of Simulated Moments (McFadden, 1989; Pakes and Pollard, 1989), where we minimize:

$$\min_{\mathbf{\Pi}} (\hat{\phi}_{data} - \phi_{sim}(\mathbf{\Pi}))' \mathcal{F} (\hat{\phi}_{data} - \phi_{sim}(\mathbf{\Pi})). \quad (10)$$

The vector $\mathbf{\Pi}$ contains the remaining unknown parameters: the disutility from working for unmarried women without children (ψ^{00}), married women without children (ψ^{01}), married women with a child (ψ^{11}), and unmarried women with a child (ψ^{10}); the variance of match quality at marriage (σ_0^2); the variance of innovations to match quality (σ_ξ^2); the parameters characterizing the probability of meeting a partner over the life cycle ($\lambda_0, \lambda_1, \lambda_2$); and the utility cost of being on welfare (η).

Empirical moments ϕ_{data} are calculated from the 1960-69 birth cohort in the pre-reform period (1990-95). These women are between age 21 and 35, ages for which we have a sufficiently large number of observations. We annualize data by considering the marital status, fertility, employment status and welfare participation status that women had for more than half of the calendar year. Simulated moments ϕ_{sim} are computed using the full numerical solution of the model. We use the inverse of the variance-covariance matrix of the empirical moments as weighting matrix \mathcal{F} , computed using the bootstrap method.²⁹

We calculate the standard errors of our parameter estimates using the asymptotic formula:

$$\text{Var}(\mathbf{\Pi}) = \left(1 + \frac{1}{S}\right) (J' \mathcal{F} J)^{-1}$$

where S is the number of sample simulations (McFadden, 1989). We calculate the derivatives of each moment with respect to each parameter, arranged in matrix J , by taking a 2.5% step of the parameter in each direction, taking the difference and then dividing by 5% of the parameter.

Selection of Moments

We consider three sets of moments. The first set of moments includes conditional moments for labor supply, i.e., the fraction of women employed by marital and fertility status (which

²⁹It is important to use the full variance-covariance matrix because of the correlation between different moments as discussed below.

pin down the disutility from work for the different types) and the proportion of single mothers on welfare (which helps identifying the utility cost from welfare participation). The second set of moments includes the profile of the probability of being ever married between ages 21 and 35. These moments jointly contribute to pinning down the variance of initial match quality, as well as the parameters characterizing the probability of meeting a potential partner. The third set of moments includes the probability of divorcing between ages 26 and 35.³⁰ These moments mostly pin down the variance of the initial marriage draw and of the innovations to match quality.

Table 5: Empirical and Simulated Moments: Employment and Welfare Use

Moment Description	Data Mean %	(s.e. in %)	Model Mean %
Employment Rate			
Married women without children	85.45	(0.01)	84.65
Unmarried women without children	87.96	(0.01)	87.20
Married women with children	60.02	(0.00)	60.22
Unmarried women with children	55.78	(0.01)	56.18
Welfare Use			
Unmarried women with children	37.53	(0.01)	36.69

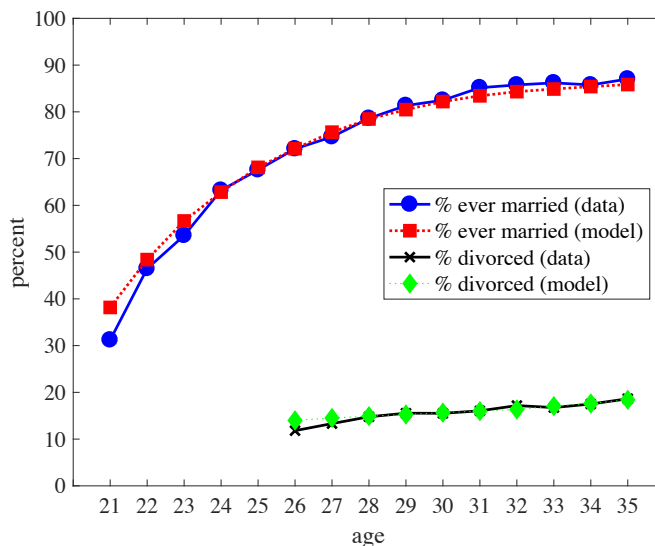
We show how the model fits the moments for employment and welfare use in Table 5 and the fit of the evolution of marriage and divorce over the life-cycle in Figure 7. Table E.2 in Appendix E shows the complete set of moments and the resulting fit.

Estimated Model Parameters

Table 6 reports the estimates of the structural parameters. In the data, unmarried women have low employment rates, despite the absence of income from a spouse. To match these low employment rates shown in Table 5, the estimation requires a large disutility of work for single women (similarly to Blundell et al. (2016)). For married women, the employment rate

³⁰The reason why we do not consider these moments for earlier ages is related to initial conditions: divorces in these early years are concentrated among people who married *before* age 21, for which we do not know the actual distribution of the match quality realizations, as the marriages occur before the model begins. This in turn was done so that we can condition on completed education.

Figure 5: Empirical and Simulated Moments: Marriage and Divorce



can be matched with a lower disutility of work because of the explicit presence of spousal income which discourages female employment.

The utility cost of welfare benefits is high, and is identified by the women who are not claiming benefits while eligible given their income. In the pre-reform period, there was no intertemporal tradeoff to claiming benefits, and hence we can attribute not claiming to utility or other costs of claiming. Our estimate of η implies that a single woman with annual consumption of \$6000 requires minimum monthly benefits of \$400 to overcome the utility cost. Figure E.2 in Appendix E gives further details. In the counterfactual simulations, for the post reform period, the intertemporal tradeoff will add to this cost, which makes it important to identify the utility cost from the pre-reform period.

The variance of innovations to match quality is less than half that of initial match quality. Finally, the table reports the parameters of the function reflecting arrival rates of partners. The implied arrival rate declines with age, but at a decreasing rate. This will prompt people to marry earlier rather than wait.

5.4 Implications for intrahousehold Allocations

The model has implications for the distribution of resources in the household under AFDC. The mean Pareto weight for women is about one half of the weight for men. This is in line with estimates from the literature on collective household models for the United States, the

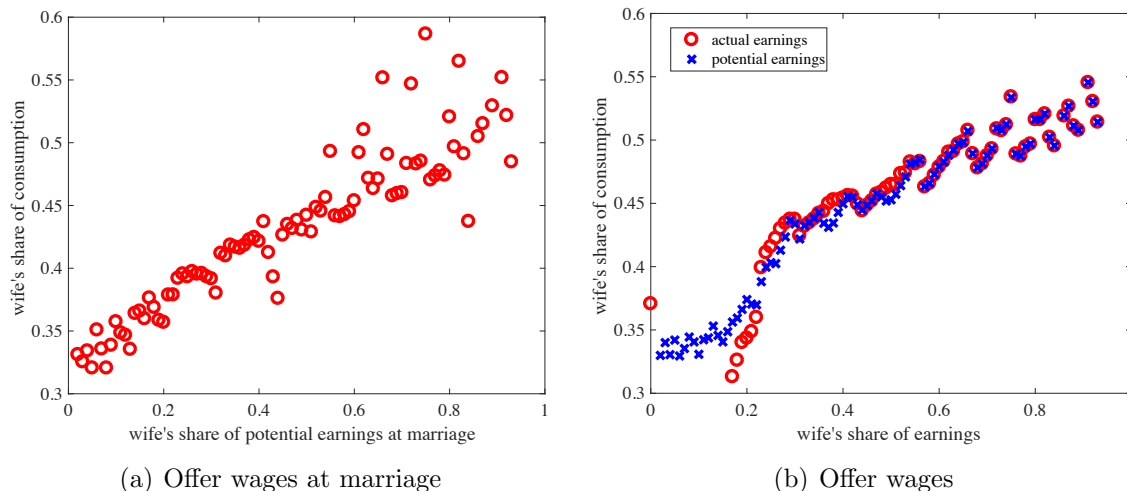
Table 6: Parameters Estimated by Method of Simulated Moments

Parameter		Estimate	(s.e.)
Cost of work			
Unmarried, no children	$\exp\{\psi^{s0}\}$	0.33	(0.036)
Married, no children	$\exp\{\psi^{m0}\}$	0.45	(0.042)
Unmarried, with child	$\exp\{\psi^{s1}\}$	0.32	(0.023)
Married, with child	$\exp\{\psi^{m1}\}$	0.36	(0.028)
Cost of being on AFDC	η	0.0071	(0.0002)
Match quality			
Variance at marriage	σ_0^2	0.0226	(0.0069)
Variance of innovations	σ_ξ^2	0.0107	(0.0029)
Probability of meeting partner by age			
	λ_0	0.324	(0.021)
	λ_1	-0.031	(0.003)
	λ_2	0.0007	(0.0002)

Notes: The cost of work is expressed as the amount of consumption if not working that is equivalent to working and consuming one unit of consumption: from the utility function 4 this is $\exp\{\psi\}$. The standard error reported is on the underlying parameter, ψ . The underlying parameter vector for ψ is given by $\{\psi^{s0}, \psi^{m0}, \psi^{s1}, \psi^{m1}\} = \{-1.12, -0.80, -1.14, -1.01\}$

United Kingdom, and Japan (Lise and Seitz, 2011; Mazzocco, Yamaguchi and Ruiz, 2013; Voena, 2015; Lise and Yamada, 2014). The left hand graph of Figure 6 plots the average of the woman’s share of consumption against the woman’s share of potential earnings, defined as her offered wage at the time of marriage. This is shown for couples who get married, and so for whom there is a gain to marriage. An increase in the woman’s share of potential earnings by 1 percentage point leads to an increase of 0.25 percentage points in the consumption share. However, women only achieve a 50% share of consumption when the potential earnings share reaches 80%. The consumption share only captures one aspect of the utility because changes in potential earnings will also change labor force participation which directly changes utility for women.

Figure 6: Consumption Allocation in the Household



Notes: Simulations from the estimated model.

The right hand graph shows what happens within the marriage. The blue crosses show how the share of consumption changes with offered wages in each period. The red circles show how the share of consumption changes with actual earnings once the labor supply decision has been made. The red circle when the woman’s actual earnings are zero is the average consumption share for non-participating women. This graph highlights the selection effect into participation: no woman whose share of earnings is less than 0.19 chooses to work.

Consumption shares remain highly persistent after marriage. This arises because even under limited commitment Pareto weights (and hence the sharing rule) are only renegotiated if there is a credible threat and one of the participation constraints given by equation (8) becomes binding. These constraints can bind following positive shocks to wages or earnings

which increase individual outside options more than the value of marriage. Shocks to the match quality can also cause the participation constraints to bind. However, in practice the constraints rarely bind following shocks: across all couples, the Pareto weights do not adjust in 87% of marriages. On the other hand, in some cases they do bind and indeed the shocks become large enough to induce divorce.

6 The Effect of Time Limits

6.1 Short-Run Response to Time Limits

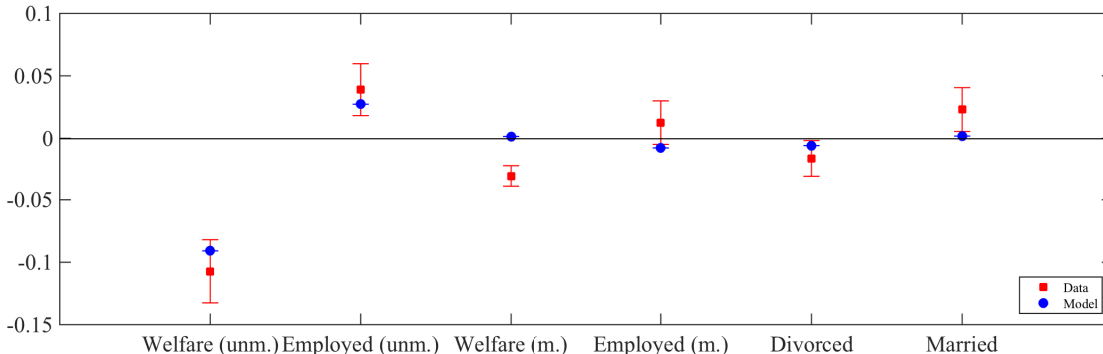
We simulate the introduction of time limits for women at different ages at the time the benefit reform took place. We match the age distribution at the introduction of time limits to the 1996 age distribution in the SIPP. We then use simulated data to estimate short run effects equivalent to those estimated by the difference-in-difference specification in section 3; this allows us first to check that our model can replicate the reduced-form evidence above.

Figure 7 reports the estimated coefficients on the simulated data and in the SIPP, focusing on the sample of women aged 21 to 53. The SIPP data specification replicates the specification in Appendix C, table C.3 (i.e., it includes year-by-age fixed effects). The simulated difference-in-differences estimates are qualitatively and quantitatively close to the empirical ones. We match the large decline in welfare utilization and the smaller (in absolute value) positive effect on employment among unmarried women. We also find that time limits lead to a reduction in the number of divorces and an increase in the stock of marriages, although both effects are small. On the other hand, the model does not replicate the small decline in welfare utilization and rise in employment among married women. In our model, married couples are very unlikely to qualify for AFDC even without time limits because all men work and so there is very limited scope for changes in behavior on reform.³¹

These responses raise the question of to what extent they reflect forward looking behavior and the banking of benefits. To address this question, we simulate the counterfactual of how would individuals have behaved if they had been myopic with respect to the time limits. By “myopic”, we mean individuals who behave as if the introduction of time limits had not occurred (until they actually run out of benefits), but are forward looking in terms of other behaviors. Figure 8 compares welfare utilization by years since the reform in the data with

³¹Married women whose husbands do not work respond in a similar way to single women.

Figure 7: Difference-in-Differences Estimates: Simulated and Actual Data



Notes: Effects estimated using data 7 years before the reform and 14 years afterwards in both the simulated model and the SIPP data. We mimic the difference-in-difference estimation of section 3.

the full model and under myopia. Myopic individuals do not cut their use of welfare when the reform takes place. By contrast, in the model and in the data, there is a decline directly after the reform.

6.2 Long-Run Effects of Time Limits

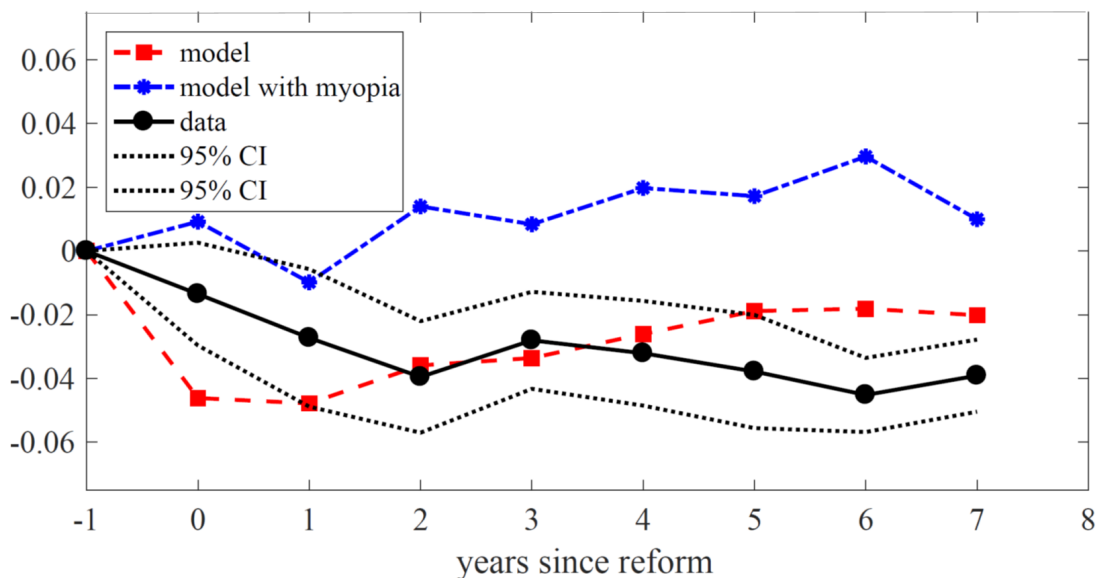
We address the long term effects of time limits on women by simulating the effects for 33 years after the reform. We consider time limits of varying length up to 5 years, and also the removal of AFDC completely.³² To summarize the effects of welfare reform we estimate the following equation on the simulated data:

$$y_{it} = \beta \cdot Reform_{it} + f_{it}^{age} + u_{it}.$$

The results are in Figure 9, where we report the effect on welfare utilization in panel (a), and employment in panel (b). For the standard five year time limit, welfare utilization by single women declines in the long run by 14.1% and employment increases by 6.1%. Even in the scenario when AFDC is removed completely and so welfare use falls by 35% to zero, employment only increases by 20%. Figure 10 shows that the effects on the marriage market are more muted. With a 5 year time limit, divorce declines by a small amount, and marriage

³²Introducing time limits, which lead to a reduction in benefit use and increase in employment, generates tax revenue. We present the results without reducing the tax rate. When we make the experiment revenue neutral by cutting the payroll tax, the results change only minimally: welfare use falls slightly further, and employment rises slightly more.

Figure 8: Dynamic response of welfare utilization to time limits

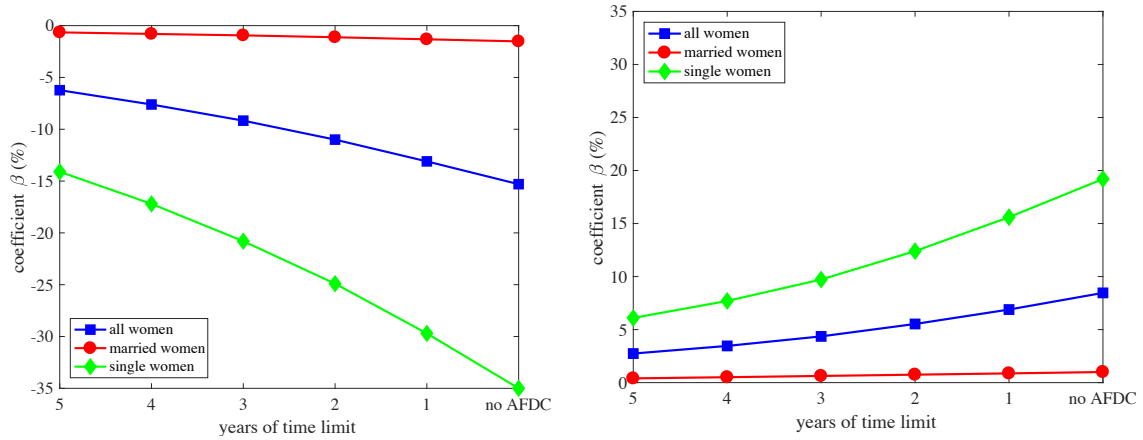


Notes: Effects estimated using data 7 years before the reform and 14 years afterwards. By *Model with myopia* we mean individuals who behave as if the introduction of time limits had not occurred (until they actually run out of benefits), but are forward looking in terms of other behavior.

increases similarly. On the other hand, women’s position within marriage worsened and there is a decline in the Pareto weight of married women.

While Figure 9 reports changes to the level of benefit use, Figure 11 show at what stage this decline occurs, focusing on the introduction of the 5 year limit. Under AFDC, the average welfare user is on welfare for 5.76 years. Under a 5-year time limit, the average utilization among welfare users drops to 3.34 years. This decline is partly mechanical in that time limits, once introduced, are binding for 10.5% of all women, and we observe significant bunching at 5 years once the limit is introduced. However, the decline also reflects an increase in individuals who never claim and reduced claiming even by those who have not hit their time limits. Figure 11 shows how changes in fraction of women receiving welfare and the fraction who have hit the 5 year limit against the age of their child, with and without time limits. Even right after birth, the use of benefits is lower under time limits, though the reduction in use is greatest when the child is aged 9-12. Reflecting this banking of benefits, under time limits, women reach 5-years of benefit use much later in their child’s life. By child-age 18, 12.5% of women have reached 5 years of benefits use, whereas without time limits, the same fraction reach 5 years of benefit use before their child is age 9.

Figure 9: Long-Run Effects of Time Limits on Welfare and Employment

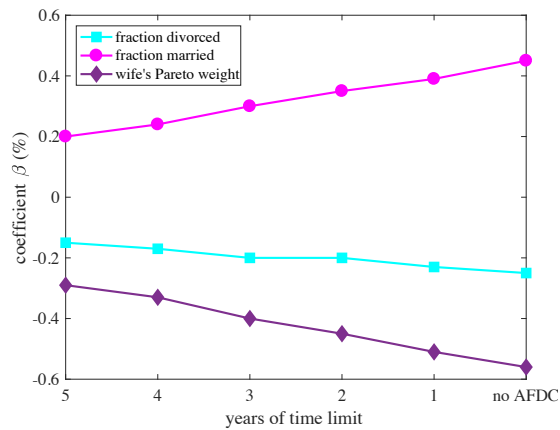


(a) Welfare utilization

(b) Employment

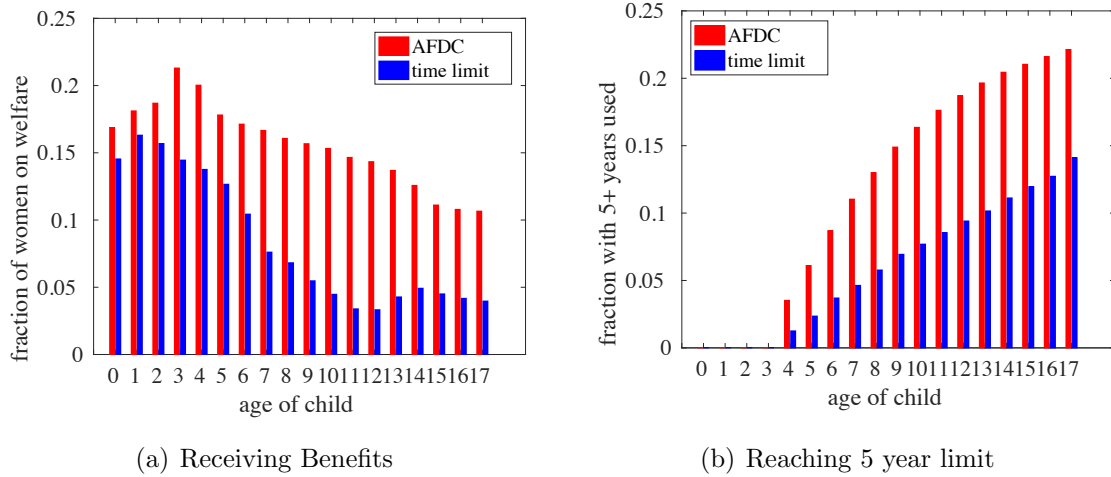
Notes: Simulated effects up to 33 years after the reform for mothers.

Figure 10: Long-Run Effects of Time Limits on Marriage and Divorce



Notes: Simulated effects up to 33 years after the reform for mothers.

Figure 11: Welfare Utilization by Age of Child



Notes: Simulation from estimated model.

Intrahousehold Allocations

We now consider how the introduction of time limits can affect intrahousehold allocations. Time limits reduce the value of being single because welfare becomes less generous. This weakens their bargaining position within marriage. In Figure 12 we consider separately the effect of this change on those already married from those who are entering marriage. For those who are already married, the limited commitment model means that weaker bargaining power will only lead to a change in Pareto weights when the man's participation constraint binds. As we discussed above, this only happens in 13% of cases in the baseline, which implies that only few families will be on the margin when the reform is implemented. Thus the Pareto weights will shift for only a small group of people and by a slight amount when they do change. This explains the empirical finding that the increase in divorce is small.

On the other hand, Figure 12 shows that there is a larger negative effect on wives' Pareto weights for newly formed couples. This partly reflects the worse position at the point where marriage happens and translates to a lower share of the gains in marriage.

There are two takeaway points: first, the main impact of time limits on household bargaining and sharing is at the time of marriage. After marriage, there is little renegotiation in practice and there is close to full commitment for most. Second, while the reform may have reduced divorce and increased marriage, it is not clear that this was welfare increasing. We explore the impact on utility explicitly in section 6.4.

Figure 12: Changes in Wives' Pareto Weights from Time-Limits



Notes: Simulation from estimated model.

6.3 Alternative Welfare Reform

In Table 7 we perform a number of counterfactual experiments to explore alternative welfare reforms and to study interactions between the various welfare programs. In panel A, we consider removing the food stamps program. In panel B, we turn to the EITC. In column 1, for both panels, we report again the effects of imposing a 5-year time limit on AFDC (as in Figure 9). In column 2 we report the effect of removing a specific feature of the welfare system (food stamps or EITC) while leaving AFDC unaltered (i.e., with no time limits). In column 3 we simultaneously remove a specific feature of the welfare system and impose a 5-year time limit on AFDC. Finally, in column 4 we impose a 5-year time limit on AFDC to an environment in which a specific feature of the welfare system was absent both at baseline and after the reform.

We begin by examining the role of food stamps (Panel A). This program has a strong income effect on both single and married women, so its removal increases the probability of employment by 19pp and 7.5pp respectively (column 2). Because of the increase in the probability of working, we observe a decline in AFDC utilization. Removing food stamps also increases the probability of being married and reduces the probability of being divorced, with substantially larger effects compared to introducing time limits. This is because removing food stamps causes a greater reduction in outside options for women. Removing food stamps and imposing a 5-year limit on AFDC at the same time has broadly additive effects (column 3 as opposed to the sum of columns 1 and 2). Nevertheless, column 4 shows an important source

Table 7: Long-Run Effects of Alternative Welfare Reforms

<i>Panel A: Reform of Food Stamps</i>				
	Impose 5yr limit	Remove FS	Remove FS Impose 5yr limit	Impose 5-year limit (No FS available)
AFDC utilization	-0.062	-0.055	-0.098	-0.042
Employed	0.027	0.122	0.163	0.039
AFDC utilization (m.)	-0.007	-0.002	-0.006	-0.005
Employed (m.)	0.004	0.075	0.082	0.006
AFDC utilization (unm.)	-0.141	-0.128	-0.226	-0.096
Employed (unm.)	0.061	0.191	0.281	0.088
Divorced	-0.001	-0.006	-0.006	-0.001
Married	0.002	0.006	0.009	0.002
Wife's Pareto Weight	-0.003	-0.005	-0.007	-0.002
<i>Panel B: Reform of EITC</i>				
	Impose 5yr limit	Remove EITC	Remove EITC Impose 5yr limit	Impose 5-year limit (No EITC available)
AFDC utilization	-0.062	0.076	-0.025	-0.100
Employed	0.027	-0.042	-0.006	0.035
AFDC utilization (m.)	-0.007	0.006	-0.003	-0.010
Employed (m.)	0.004	0.051	0.059	0.007
AFDC utilization (unm.)	-0.141	0.164	-0.060	-0.220
Employed (unm.)	0.061	-0.169	-0.096	0.071
Divorced	-0.001	0.001	0.000	-0.002
Married	0.002	-0.008	-0.004	0.004
Wife's Pareto weight	-0.003	-0.001	-0.005	-0.004

Notes: Simulated effects up to 33 years after the reform for mothers. Panel A considers the removal of food stamps and Panel B the removal of EITC. Column 1 imposes 5-year time limits on AFDC. Column 2 removes each of Food Stamps and EITC separately. Column 3 removes a specific feature of the welfare system and impose a 5-year time limit on AFDC. Column 4 introduces a 5-year time limit without Food Stamps and without EITC. “m.” is married and “unm.” unmarried.

Table 8: Long-Run Effects of Reducing Stigma Costs

	Impose 5yr limit	Set $\eta = 0$	Set $\eta = 0$ Impose 5yr limit	Impose 5-year limit (Always $\eta = 0$)
AFDC utilization	-0.062	0.249	0.083	-0.166
Employed	0.027	-0.178	-0.069	0.108
AFDC utilization (m.)	-0.007	0.123	0.049	-0.075
Employed (m.)	0.004	-0.053	-0.019	0.033
AFDC utilization (unm.)	-0.141	0.416	0.130	-0.283
Employed (unm.)	0.061	-0.350	-0.139	0.208
Divorced	-0.001	0.004	0.000	-0.004
Married	0.002	-0.009	-0.001	0.007
Wife's Pareto weight	-0.003	0.007	-0.001	-0.008

Notes: Simulated effects up to 33 years after the reform for mothers. Column 1 imposes 5-year time limits on AFDC. Column 2 sets $\eta = 0$. Column 3 sets $\eta = 0$ at the same time as imposing a 5-year time limit on AFDC. Column 4 introduces a 5-year time limit when η always equals 0. “m.” is married and “unm.” unmarried.

of interaction between these programs that helps us interpret our reduced-form estimates. In particular, imposing a 5-year time limit on AFDC in an environment without food stamps leads to a reduction in AFDC utilization that is close, in absolute value, to the increase in employment. By contrast, when food stamps are available, as we have seen in the data and in the model, the reduction in AFDC use is twice as large as the increase in employment. This result confirms that food stamps represent an important source of insurance for mothers who drop out of AFDC because of time limits and are not employed.

Turning to EITC (Panel B), consistently with our empirical findings in table E.1, we see that EITC has primarily a substitution effect for single mothers: by supplementing their earnings, it increases their probability of employment. Indeed, the removal of EITC reduces the probability that a single woman is employed by 17pp and these women move onto AFDC. Overall, the program has a sizable effect on marital status, larger than AFDC or food stamps: its removal increases the incentive to be single, reduces marriage probabilities by 0.8pp and increases divorce. The removal of EITC combined with a 5-year time limit has broadly additive effects (column 3). Again, though, this welfare feature is important for generating the effect of time limits that we observe in the data. Had we not modeled EITC, we would

have obtained larger gaps between the decline in AFDC use and the increase in employment associated with time limits on AFDC than those we observe in the data (column 4).

A final reform we consider concerns stigma, with the elimination of all stigma costs to participating in AFDC (reported in Table 8). There is no clear policy equivalent to eliminating stigma costs. On the other hand, the parameter η captures both stigma as well as more general transaction costs and we show that these costs can have substantial effects: removing these costs boosts participation in the program substantially by 12.3pp for married mothers and 41.6pp for unmarried ones, as well as reducing employment to a somewhat smaller extent. The cheaper access to AFDC primarily benefits single mothers, and this improves the outside option to marriage: marriage falls, divorces rise and women's weight in household decision making increases. Further, there is a large interaction effect with the imposition of time limits: in the absence of stigma, time limits are more likely to be binding and so have a large impact on welfare use and employment.

6.4 The Welfare Costs of Time Limits

Our results so far document a rich array of behavioral responses to the introduction of time limits. Time limits reduce the use of benefits and increase employment. On the other hand, we find that half of those coming off AFDC/TANF do not move into employment and so their consumption declines. Further, divorce rates fall and this is because outside options gets worse. By reducing the outside option for women of living off benefits, the reform weakens their bargaining position within marriage, as well as reducing the value of insurance directly. In this section we use our model to calculate the net effect of all these offsetting consequences of the reform.

In computing the welfare consequences, we adjust the payroll tax to hold constant the government deficit. This is because the fall in welfare use and the rise in employment induced by the introduction of time limits lead to savings on government spending on transfers and to an increase in tax revenues. Hence, revenue neutrality means that the tax rate needs to be cut when the time limit is introduced. We implement this through a payroll tax/subsidy that falls only on the women of our group; hence we transfer the revenue gains to them in terms of lower taxes and abstract from redistribution across groups. We describe the procedure in full in Appendix F.

To define the welfare cost or benefit of introducing time limits, we calculate the proportion of consumption, π^s , that an individual is willing to pay *ex-ante* to be indifferent between

a new scenario, for example, $s = 5TL$ (a 5 year time limit) and the baseline, *AFDC*. We compute the lifetime expected utility of a woman in our model as:

$$E_0 [U (s, \tau)] |_{\pi} = \sum_{i=1}^N \sum_{t=0}^T \beta^t \left(\frac{\left((1 - \pi^s) c_{i,t}^s \cdot e^{\psi(m_{i,t}, k_{i,t}^a) \cdot P_{i,t}^s} \right)^{1-\gamma}}{1 - \gamma} - \eta B_{i,t}^s + L_{i,t} m_{i,t}^s \right) \quad (11)$$

where $\{c_t^s, P_t^s, B_t^s, m_t^s, \}$ refer to the implied consumption, labor supply, benefit status and marital status in economy s , and τ is the revenue neutral tax rate. E_0 represents the expectation at the beginning of working life, before initial conditions are known. We solve for π such that³³

$$E_0 U (5TL, \tau_w) |_{\pi} = E_0 U (AFDC, 0), \quad (12)$$

We find that women are willing to pay 0.5% of lifetime consumption to avoid time limits and remain in the original AFDC regime, despite directly benefiting (in our simulation) from the redistribution of the government revenue saved. This is an important calculation because it shows that - “behind the veil of ignorance” - the 1996 reform, while achieving its main goals of breaking the culture of welfare dependence and emphasizing one of self-sufficiency through work, induced a net welfare loss by reducing insurance available to low-income women. In contrast, without the revenue-neutrality adjustment increasing income, the willingness to pay to avoid time limits for these women is equal to 1% of lifetime consumption.

7 Conclusions

This paper addressed the broad dynamic implications of reforming the safety net in the US, accounting simultaneously for the various possible responses through welfare use, employment and marital status. The overarching goal of the reform was to curb disincentive effects, such as those related to the decision to work or form stable marital relationships, while preserving insurance provided to families with young children. We focus on reforms that limit the life-time use of benefits. In addition to assessing the specific policy, our analysis highlights the extent to which dynamic incentives really matter for welfare recipients on

³³In varying π we do not reoptimize and take as given the choices of consumption, participation, benefit use, marriage and divorce.

a broad set of outcomes. Our key finding is of substantial forward-looking behavior because benefit utilization is shown to decline immediately after the reform, despite the fact that time limits would only bind after a number of years.

We find this result of forward-looking behavior where individuals “bank” benefits, in both our quasi-experimental evidence looking at each possible response in isolation, as well as in our structural life-cycle model which endogenizes the various possible responses. Welfare use falls substantially for single women: pre-reform, 31% of low educated single mothers receive AFDC. Post-reform, this falls to 8%. Our finding is that time limits alone induced a fall to 16%. On the other hand, only half of the women who no longer receive benefits are employed. For this group in particular, the implications are serious: reliance on Food Stamps is greater and the cost of being single is much higher. This latter effect drives the reduction in divorce rates. It also changes bargaining power within relationships: we show that at the time of marriage, Pareto weights for women decline significantly post-reform. On the other hand, if women are already married, we find little changes in Pareto weights. This arises because of the limited commitment framework: the Pareto weights adjust only if participation constraints bind and this happens only very infrequently.

We use our model to consider welfare implications, alternative benefit reforms and long run consequences. Welfare analyses reveal that the groups mostly targeted by the reform (i.e., women) suffered a net welfare loss, despite the increase in self-sufficiency brought about by increasing employment. The reform caused this substantial welfare loss, even when we account for the government revenue raised. Although time limits encouraged marriage, the effect of EITC on family formation seems to have been much larger. This perhaps emphasizes again the importance for both insurance and incentives of welfare-to-work type benefits.

References

- Attanasio, O., H. Low, and V. Sanchez-Marcos.** 2008. “Explaining changes in female labor supply in a life-cycle model.” *The American Economic Review*, 98(4): 1517–1552.
- Attanasio, Orazio P, and Guglielmo Weber.** 1995. “Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey.” *Journal of Political Economy*, 103(6): 1121–1157.
- Bitler, Marianne P, Jonah B Gelbach, Hilary W Hoynes, and Madeline Zavodny.** 2004. “The impact of welfare reform on marriage and divorce.” *Demography*, 41(2): 213–236.

- Blank, Rebecca M.** 2002. “Evaluating Welfare Reform in the United States.” *Journal of Economic Literature*, 40(4): 1105–1166.
- Blundell, Richard, Alan Duncan, and Costas Meghir.** 1998. “Estimating labor supply responses using tax reforms.” *Econometrica*, 827–861.
- Blundell, Richard, Martin Browning, and Costas Meghir.** 1994. “Consumer Demand and the Life-Cycle Allocation of Household Expenditures.” *The Review of Economic Studies*, 61(1): 57–80.
- Blundell, Richard, Monica Costa Dias, Costas Meghir, and Jonathan M Shaw.** 2016. “Female labour supply, human capital and welfare reform.” *Econometrica*, 84(5): 1705–1753.
- Blundell, Richard, Pierre-André Chiappori, and Costas Meghir.** 2005. “Collective labor supply with children.” *Journal of political Economy*, 113(6): 1277–1306.
- Bronson, Mary Ann.** 2014. “Degrees are forever: Marriage, educational investment, and lifecycle labor decisions of men and women.” *Unpublished manuscript*, 2.
- Bruins, Marianne.** 2017. “TANF, Child Care and Child Wellbeing in Sole Parent Families.” *University of Oxford mimeo*.
- Burtless, Gary, and Jerry A Hausman.** 1978. “The effect of taxation on labor supply: Evaluating the Gary negative income tax experiment.” *The Journal of Political Economy*, 1103–1130.
- Chan, Marc K.** 2013. “A Dynamic Model of Welfare Reform.” *Econometrica*, 81(3): 941–1001.
- Chiappori, P.A., M. Costa-Dias, and C. Meghir.** 2016. “The Marriage Market, Labor Supply, and Education Choice.” *NBER WP 21004, Journal of Political Economy, Forthcoming*.
- Chiappori, Pierre-André.** 1988. “Rational household labor supply.” *Econometrica: Journal of the Econometric Society*, 63–90.
- Chiappori, Pierre-André.** 1992. “Collective labor supply and welfare.” *Journal of Political Economy*, 437–467.
- Chiappori, Pierre-Andre, and Maurizio Mazzocco.** 2017. “Static and intertemporal household decisions.” *Journal of Economic Literature*, 55(3): 985–1045.
- Eissa, Nada, and Jeffrey B Liebman.** 1996. “Labor supply response to the earned income tax credit.” *Quarterly Journal of Economics*, 111(2): 605–637.
- Fernández, Raquel, and Joyce Cheng Wong.** 2014. “Divorce Risk, Wages and Working Wives: A Quantitative Life-Cycle Analysis of Female Labour Force Participation.” *The Economic Journal*, 124(576): 319–358.

- Fernández, Raquel, and Joyce Cheng Wong.** 2017. “Free to Leave? A Welfare Analysis of Divorce Regimes.” *American Economic Journal: Macroeconomics*, 9(3): 72–115.
- Goussé, Marion, Nicolas Jacquemet, and Jean Marc Robin.** 2017. “Marriage, Labor Supply and Home Production.” *Econometrica, Forthcoming*.
- Grogger, Jeffrey, and Charles Michalopoulos.** 2003. “Welfare Dynamics under Time Limits.” *Journal of Political Economy*, 111(3).
- Grogger, Jeffrey, and Lynn A. Karoly.** 2005. *Welfare Reform*. Harvard University Press.
- Heckman, James.** 1974. “Shadow prices, market wages, and labor supply.” *Econometrica: journal of the econometric society*, 679–694.
- Keane, Michael, and Robert Moffitt.** 1998. “A structural model of multiple welfare program participation and labor supply.” *International Economic Review*, 553–589.
- Keane, Michael P, and Kenneth I Wolpin.** 2010. “The Role of Labor and Marriage Markets, Preference Heterogeneity, and the Welfare System in the Life Cycle Decisions of Black, Hispanic, and White Women*.” *International Economic Review*, 51(3): 851–892.
- Kline, Patrick, and Melissa Tartari.** 2016. “Bounding the Labor Supply Responses to a Randomized Welfare Experiment: A Revealed Preference Approach.” *The American Economic Review*, 106.
- Kocherlakota, N.R.** 1996. “Implications of efficient risk sharing without commitment.” *The Review of Economic Studies*, 63(4): 595–609.
- Ligon, E., J.P. Thomas, and T. Worrall.** 2002a. “Informal insurance arrangements with limited commitment: Theory and evidence from village economies.” *The Review of Economic Studies*, 69(1): 209–244.
- Ligon, Ethan, Jonathan P Thomas, and Tim Worrall.** 2000. “Mutual insurance, individual savings, and limited commitment.” *Review of Economic Dynamics*, 3(2): 216–246.
- Ligon, Ethan, Jonathan P Thomas, and Tim Worrall.** 2002b. “Informal insurance arrangements with limited commitment: Theory and evidence from village economies.” *The Review of Economic Studies*, 69(1): 209–244.
- Lise, Jeremy, and Ken Yamada.** 2014. “Household sharing and commitment: Evidence from panel data on individual expenditures and time use.” IFS Working Papers.
- Lise, Jeremy, and Shannon Seitz.** 2011. “Consumption inequality and intra-household allocations.” *The Review of Economic Studies*, 78(1): 328–355.
- Low, Hamish, and Luigi Pistaferri.** 2015. “Disability risk, disability insurance and life cycle behavior.” *American Economic Review*, 105(10): 2986–3029.
- Mazzocco, Maurizio.** 2007. “Household intertemporal behaviour: A collective characterization and a test of commitment.” *The Review of Economic Studies*, 74(3): 857–895.

- Mazzocco, M., S. Yamaguchi, and C. Ruiz.** 2013. “Labor supply, wealth dynamics, and marriage decisions.”
- Mazzolari, Francesca, and Giuseppe Ragusa.** 2012. “Time Limits: The Effects on Welfare Use and Other Consumption-Smoothing Mechanisms.” *IZA discussion paper No.6993*.
- McFadden, Daniel.** 1989. “A method of simulated moments for estimation of discrete response models without numerical integration.” *Econometrica: Journal of the Econometric Society*, 995–1026.
- Moffitt, Robert A, Brian J Phelan, and Anne E Winkler.** 2017. “Welfare rules, incentives, and family structure.” National Bureau of Economic Research.
- Pakes, Ariel, and David Pollard.** 1989. “Simulation and the Asymptotics of Optimization Estimators.” *Econometrica*, 57(5): 1027–1057.
- Persson, Petra.** 2014. “Social insurance and the marriage market.” *Unpublished manuscript*.
- Rothstein, Jesse.** 2010. “Is the EITC as Good as an NIT? Conditional Cash Transfers and Tax Incidence.” *American Economic Journal: Economic Policy*, 2(1): 177–208.
- Swann, Christopher A.** 2005. “Welfare reform when recipients are forward-looking.” *Journal of Human Resources*, 40(1): 31–56.
- Voena, Alessandra.** 2015. “Yours, Mine and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?” *American Economic Review*.
- Zabel, Jeffrey E.** 1998. “An analysis of attrition in the Panel Study of Income Dynamics and the Survey of Income and Program Participation with an application to a model of labor market behavior.” *Journal of Human Resources*, 479–506.
- Ziliak, James P.** 2016. “Temporary Assistance for Needy Families.” *Economics of Means-Tested Transfer Programs in the United States*, 1: 303–394.

Online Appendix

Appendices A to F correspond respectively to sections 1 to 6 of the main paper.

Appendix A: Time Limits

Table A.1: Time Limits in the year 2000

Type of limit	Duration	State
No limit	n.a.	Michigan, Vermont, Maine
Benefit reduction	60	California, Maryland, Rhode Island
Benefit reduction	24	Indiana
Periodic	24/48	Nebraska
Periodic	24/84	Oregon
Periodic	24/60	Arizona, Massachussets
Periodic	36/60	Ohio
Lifetime	60	Alabama, Alaska, Colorado, D.C., Hawaii, Illinois, Iowa, Kansas, Kentucky, Louisiana, Minnesota, Mississippi, Missouri, Montana, Nevada, New Hampshire, New Jersey, New Mexico, New York, North Carolina, North Dakota, Oklahoma, Pennsylvania, South Carolina, South Dakota, Tennessee, Texas, Virginia, Washington, West Virginia, Wisconsin, Wyoming
Lifetime	48	Georgia, Florida
Lifetime	36	Delaware, Utah
Lifetime	24	Montana, Idaho, Arkansas
Lifetime	21	Connecticut

Notes: Source: Welfare Rules Database (<http://wrdb.urban.org>). States with benefit reduction rules continue to provide benefits after the time limit is reached, but only to the children in the household unit. States with periodic limits of x/y months provide benefits for at most x months over a period of y months (and cap the overall time limit at y months).

A few states have changed their limits over time. For example, Arizona moved in 2016 to a limit of just one year. Michigan started with no time limit but moved to imposing a 4 year time limit in 2008.

In addition to the introduction of time limits, the reform introduced work requirements and subsidies to child care. These subsidies were primarily to the low-income population and the administration of the program was again decentralized to states. There was a significant

increase in both spending and coverage. Federal child care funding increased in real terms from \$3 billion in 1997 to \$9 billion in 2010; and the average monthly number of low-income children under age 13 receiving subsidies increased over the same period from 1 million to 1.7 millions. Since child care assistance is relevant primarily for pre-school children, while time limits apply to all ages, we verify in our reduced form that the impact of welfare reform on welfare utilization and employment is not limited to families with young children.

Appendix B: Data

Table B.1: Sample selection for empirical analysis

	Individuals	Observations
Everyone over 18	481,327	3,306,878
Drop college graduates	303,033	1,996,570
Drop men	163,500	1,097,432
Drop if over 60	123,994	784,791
Drop if no children in household	75,938	455,514
Household heads or spouses	64,739	406,370

Notes: Data from the 1990-2008 SIPP panels.

Appendix C: Reduced Form Evidence

In this Appendix we present results for a variety of robustness checks on our reduced form evidence.

Food Stamps, Joint Welfare/Employment dynamics, and Event-study analysis

We start by presenting results commented in the main text, Section 3.

Table C.1: Food Stamps Utilization

Dependent Var:	Food Stamps Utilization					
Sample:	Whole Sample		Married		Unmarried	
$Exposed_{dst}Post_{st}$	-0.007 (0.00550)	-0.002 (0.005)	-0.004 (0.004)	0.0002 (0.004)	-0.055*** (0.011)	-0.046*** (0.010)
Basic controls	Yes	Yes	Yes	Yes	Yes	Yes
Race	No	Yes	No	Yes	No	Yes
Disability status	No	Yes	No	Yes	No	Yes
Unemp. rate*Demog.	No	Yes	No	Yes	No	Yes
Observations	406,370	406,370	286,425	286,425	119,945	119,945
R-squared	0.095	0.164	0.082	0.108	0.117	0.177

*** p<0.01, ** p<0.05, * p<0.1

Notes: Data from the 1990-2008 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state fixed effects, demographics fixed effects, state-by-demographics fixed effects, state-by-year fixed effects. Standard errors in parentheses, clustered at the state level.

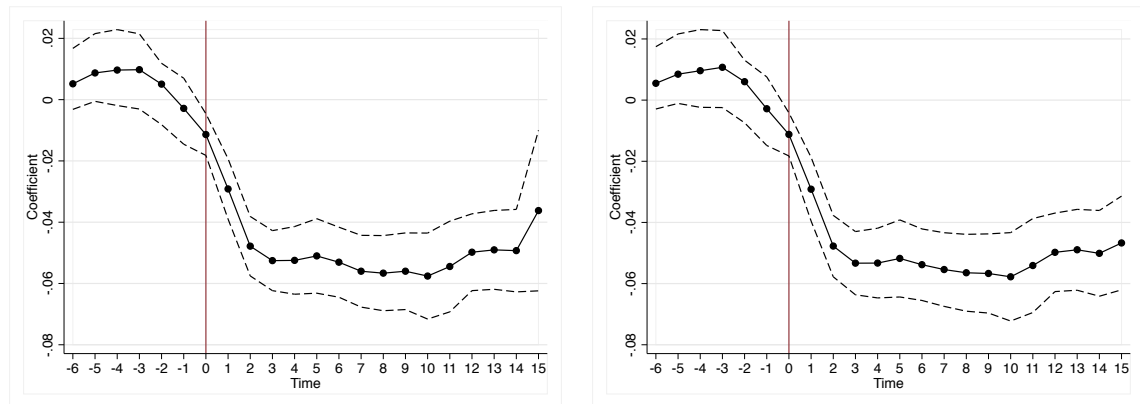
Table C.2: Joint Employment and Welfare Utilization Status of Single Mothers

Dependent Var:	Employed On Welfare	Employed Not on Welfare	Not Employed On Welfare	Not Employed Not on Welfare
$Exposed_{dst}Post_{st}$	-0.0282*** (0.00588)	0.110*** (0.00831)	-0.129*** (0.00914)	0.0468*** (0.0108)
Basic controls	Yes	Yes	Yes	Yes
Race	Yes	Yes	Yes	Yes
Disability status	Yes	Yes	Yes	Yes
Unemp. rate*Demog.	Yes	Yes	Yes	Yes
Observations	119,945	119,945	119,945	119,945
R-squared	0.033	0.105	0.155	0.075

*** p<0.01, ** p<0.05, * p<0.1

Notes: Data from the 1990-2008 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state fixed effects, demographics fixed effects, state-by-demographics fixed effects, state-by-year fixed effects. Standard errors in parentheses, clustered at the state level.

Figure C.1: Program Participation following the Introduction of Time Limits

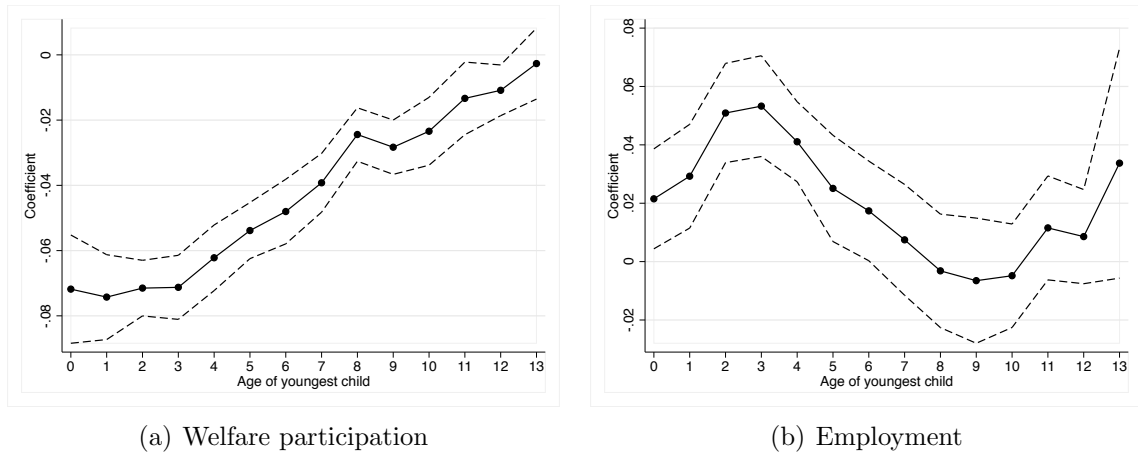


(a) All women

(b) Women in states with limit above 24 months

Notes: Data from the 1990-2008 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state fixed effects, demographics fixed effects, state-by-demographics fixed effects, state-by-year fixed effects, race and disability status. Only states where time limits were fixed at over 24 months are included in panel (b).

Figure C.2: Program Participation and Employment Dynamics by Child Age



Notes: Data from the 1990-2008 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state fixed effects, demographics fixed effects, state-by-demographics fixed effects, state-by-year fixed effects, race and disability status. Child age is defined as the age of the youngest child,

Including age-by-year fixed effects

Next, we replicate our main regression tables (Tables 2 and 3) including controls for age-by-year fixed effects. Hence, we compare women of the same age within the same year, more fully isolating the heterogenous effect due to the age of the youngest child. These demanding controls leave our main findings unchanged, but attenuated (table C.3).

Table C.3: OLS Regressions including Age-by-Year Fixed Effects

Dependent Var:	AFDC/TANF		Employed		Div/Sep	Married
Sample:	Whole	Unmarried	Whole	Unmarried	Whole	Whole
$Exposed_{dst}Post_{st}$	-0.0464*** (0.00409)	-0.111*** (0.0114)	0.0172*** (0.0059)	0.0279** (0.0112)	-0.0127* (0.0074)	0.0137 (0.0086)
Basic controls	Yes	Yes	Yes	Yes	Yes	Yes
Race	Yes	Yes	Yes	Yes	Yes	Yes
Disability status	Yes	Yes	Yes	Yes	Yes	Yes
Unemp.*Demog.	Yes	Yes	Yes	Yes	Yes	Yes
Age-by-year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	406,370	119,945	406,370	119,945	406,370	406,370
R-squared	0.110	0.201	0.108	0.205	0.020	0.110

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: Data from the 1990-2008 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state fixed effects, demographics fixed effects, state-by-demographics fixed effects, state-by-year fixed effects. Standard errors in parentheses, clustered at the state level.

Attrition in the SIPP sample

To address concerns regarding the high rate of attrition in the SIPP (Zabel, 1998), we limit our analysis to the first two waves of each SIPP panel. In Appendix table C.4 we show that this adjustment leaves the results unaffected.

Table C.4: OLS Regressions with First Two Waves of Each SIPP panel

Dependent Var:	AFDC/TANF		Employed		Div/Sep	Married
Sample:	Whole	Unmarried	Whole	Unmarried	Whole	Whole
$Exposed_{dst}Post_{st}$	-0.0622*** (0.00557)	-0.188*** (0.0156)	0.0314*** (0.00908)	0.106*** (0.0190)	-0.0299*** (0.00983)	-0.0137 (0.0100)
Basic controls	Yes	Yes	Yes	Yes	Yes	Yes
Race	Yes	Yes	Yes	Yes	Yes	Yes
Disability status	Yes	Yes	Yes	Yes	Yes	Yes
Unemp.*Demog.	Yes	Yes	Yes	Yes	Yes	Yes
Observations	52,101	15,745	52,101	15,745	52,101	52,101
R-squared	0.115	0.216	0.110	0.204	0.020	0.114

Standard errors in parentheses clustered at the state level

*** p<0.01, ** p<0.05, * p<0.1

Notes: Data from the 1990-2008 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state fixed effects, demographics fixed effects, state-by-demographics fixed effects, state-by-year fixed effects. Standard errors in parentheses, clustered at the state level.

College graduates sample

Our sample excludes college graduates because they are unlikely to be affected by the reform, given lower rates of participation in welfare. To verify this conjecture, we replicate our regressions for welfare use, employment and marital status using the same sample of college graduates. We find very small effects on welfare utilization (-0.7pp in the whole sample compared to -5pp in our main sample, again concentrated among singles) and no effects whatsoever on employment and marital status (Appendix Table C.5).

Table C.5: Effects of Time Limits on College Graduates

Dependent Var:	AFDC/TANF		Employed		Div/Sep	Married
Sample:	Whole	Unmarried	Whole	Unmarried	Whole	Whole
$Exposed_{dst}Post_{st}$	-0.00732*** (0.00185)	-0.0658*** (0.0111)	0.00380 (0.00909)	0.0217 (0.0206)	-0.00191 (0.0116)	-0.00782 (0.0134)
Basic controls	Yes	Yes	Yes	Yes	Yes	Yes
Race	Yes	Yes	Yes	Yes	Yes	Yes
Disability status	Yes	Yes	Yes	Yes	Yes	Yes
Unemp.*Demog.	Yes	Yes	Yes	Yes	Yes	Yes
Observations	180,155	26,556	180,155	26,556	180,155	180,155
R-squared	0.028	0.122	0.041	0.171	0.022	0.058

Standard errors in parentheses clustered at the state level

*** p<0.01, ** p<0.05, * p<0.1

Notes: Data from the 1990-2008 SIPP panels. Sample of female heads of household who are college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state fixed effects, demographics fixed effects, state-by-demographics fixed effects, state-by-year fixed effects. Standard errors in parentheses, clustered at the state level.

Exclude mothers of young children

A potential concern is that our results are driven by changes in the behavior of households with small children after welfare reform as a result of the more generous childcare provisions in the PRWORA.³⁴ Appendix table C.6 shows that the results are robust to excluding

³⁴The welfare reform eliminated federal child care entitlements and replaced them with a childcare block grant to the states. Under these changes, states became more flexible in designing their childcare assistance

households in which the youngest child is below the age of 6. Note that this is a sample where the decline in welfare benefits is less deep. Not surprisingly (in the light of our model), the employment effects are smaller than in the whole sample. Another important component of the 1996 welfare reform was the introduction of work requirement. The only threat to identification is that work requirement were less stringent for mothers of very young children (below age one). This should lead our estimates for employment to be downward biased. However, this is unlikely to represent a significant bias given the size of the population exempted.

Table C.6: Women with Children above age 5

Dependent Var:	AFDC/TANF		Employed		Div/Sep	Married
Sample:	Whole	Unmarried	Whole	Unmarried	Whole	Whole
$Exposed_{dst}Post_{st}$	-0.0290*** (0.00293)	-0.0879*** (0.00911)	0.00695 (0.00556)	0.0488*** (0.0130)	-0.0253*** (0.00738)	0.00238 (0.00880)
Basic controls	Yes	Yes	Yes	Yes	Yes	Yes
Race	No	Yes	No	Yes	No	Yes
Disability status	No	Yes	No	Yes	No	Yes
Unemp.*Demog.	No	Yes	No	Yes	No	Yes
Observations	228,729	68,695	228,729	68,695	228,729	228,729
R-squared	0.089	0.163	0.141	0.101	0.027	0.104

Standard errors in parentheses clustered at the state level

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Notes: Data from the 1990-2008 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state fixed effects, demographics fixed effects, state-by-demographics fixed effects, state-by-year fixed effects. Standard errors in parentheses, clustered at the state level.

programs. In practice, the total amount available for state-level childcare programs could increase or decrease depending on the state's own level of investment.

Fertility

Table C.7: Fertility

Dependent Var:	Newborn in $t + 1$					
Sample:	Whole		Married		Unmarried	
$Exposed_{dst}Post_{st}$	-0.0007 (0.00125)	-0.0009 (0.0013)	-0.0007 (0.0017)	-0.0009 (0.0018)	0.0032 (0.0028)	0.0034 (0.0028)
Basic controls	Yes	Yes	Yes	Yes	Yes	Yes
Race	No	Yes	No	Yes	No	Yes
Disability status	No	Yes	No	Yes	No	Yes
Unemp. rate*Demog.	No	Yes	No	Yes	No	Yes
Observations	233,944	233,944	167,065	167,065	66,879	66,879
R-squared	0.010	0.010	0.014	0.014	0.021	0.021

Standard errors in parentheses clustered at the state level

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Notes: Data from the 1990-2008 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state fixed effects, demographics fixed effects, state-by-demographics fixed effects, state-by-year fixed effects. Standard errors in parentheses, clustered at the state level.

Replication in the March CPS data

We use data from the March CPS between 1990 and 2011 to replicate all our main specifications in the sample that excludes college graduates. As reported in tables C.8 and C.9, all the findings in the SIPP carry through in the CPS. The dynamics suggest a substantial amount of benefits banking and an immediate response of all outcome variables (figure C.3), just as we have documented in the SIPP.

Table C.8: Effects of Time Limits in the CPS

Dependent Var:	AFDC/TANF Utilization			Employed		
Sample:	Whole	Married	Unmarried	Whole	Married	Unmarried
$Exposed_{dst}Post_{st}$	-0.0392*** (0.00287)	-0.00833*** (0.00134)	-0.188*** (0.00728)	0.0150** (0.00575)	-0.0177*** (0.00595)	0.122*** (0.0107)
Basic controls	Yes	Yes	Yes	Yes	Yes	Yes
Race	No	Yes	No	Yes	No	Yes
Disability status	No	Yes	No	Yes	No	Yes
Unemp.*Demog.	No	Yes	No	Yes	No	Yes
Observations	362,994	265,100	97,894	362,991	265,097	97,894
R-squared	0.039	0.019	0.098	0.062	0.056	0.096

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: Data from the March CPS. Sample of female heads of household who are college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year fixed effects, state fixed effects, demographics fixed effects, state-by-demographics fixed effects, state-by-year fixed effects. Standard errors in parentheses, clustered at the state level.

Table C.9: Effects of Time Limits in the CPS on Marriage

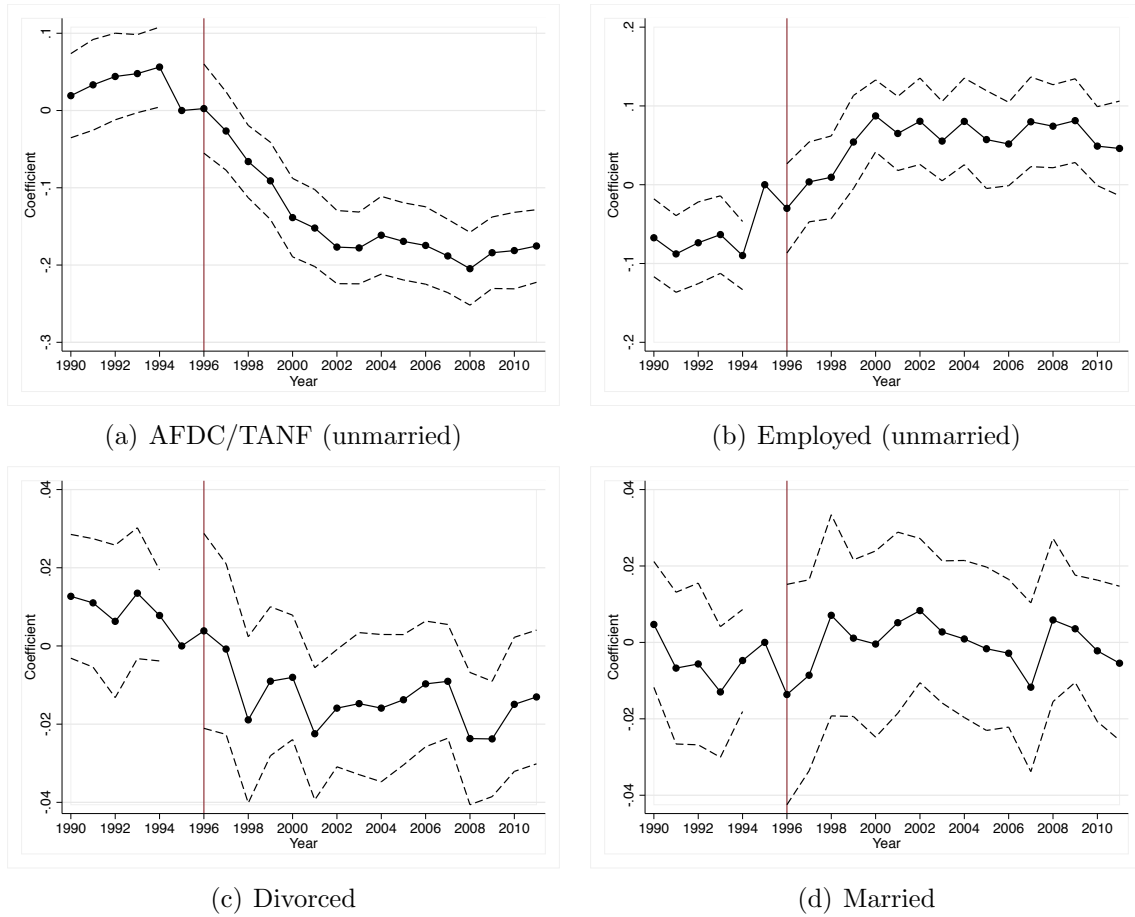
Dependent Var:	Divorced/Separated		Married	
	Whole	$m_{t-1} = 1$	Whole	$m_{t-1} = 1$
<i>Exposed_{dst}Post_{st}</i>	-0.0225*** (0.00337)	-0.00323** (0.00149)	0.00548 (0.00586)	-0.00179 (0.00594)
Basic controls	Yes	Yes	Yes	Yes
Race	Yes	Yes	Yes	Yes
Disability status	Yes	Yes	Yes	Yes
Unemp. rate*Demog.	Yes	Yes	Yes	Yes
Observations	362,994	121,081	362,994	43,152
R-squared	0.011	0.010	0.040	0.033

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: Data from the March CPS. Sample of female heads of household who are college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year fixed effects, state fixed effects, demographics fixed effects, state-by-demographics fixed effects, state-by-year fixed effects. Standard errors in parentheses, clustered at the state level.

Figure C.3: Dynamic effects of time limits in the CPS



Notes: Data from the March CPS. Sample of female heads of household who are college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year fixed effects, state fixed effects, demographics fixed effects, state-by-demographics fixed effects, state-by-year fixed effects. Standard errors in parentheses, clustered at the state level.

Online Appendix D: Life-Cycle Model

The baseline model solves for the Pareto weights on husband and wife's welfare through the solution to a symmetric Nash bargaining game between spouses (Mazzocco, Yamaguchi and Ruiz, 2013):

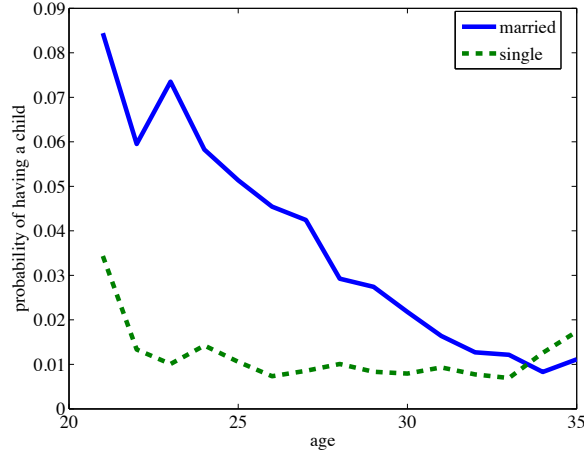
$$\max_{\theta_{ij,t_0}} (V_{t_0}^{Fm}(\theta_{ij,t_0}) - V_{t_0}^{Fs}) \cdot (V_{t_0}^{Mm}(\theta_{ij,t_0}) - V_{t_0}^{Ms}).$$

Marriages will occur if there is a positive surplus to share. This suggests that how the surplus is shared will not affect the marriage decision. However, the sharing rule that would be used in future marriages will affect outside options and so whether each person is better off continuing to search. If the surplus of future marriages was allocated exclusively to the woman, then it is worthwhile for the woman to search for the match with the highest surplus. If almost none of any future surplus was going to be allocated to the woman, there is little incentive for her to continue searching. We experiment varying the share of the surplus that goes to each person. This has some effect on parameter estimates for the cost of working for married women, which need to be larger to deter participation when men have all of the surplus. However, our conclusions on the effects of time limits for behavior are unchanged.

Appendix E: Structural Estimation Details

Fertility

Figure E.1: Probability of having a first child by woman's age and marital status



Source: Data from SIPP panels 1990-2008.

Earnings Process for Men

We use GMM to estimate the variance of the permanent component of log annual earnings ($\sigma_{\zeta^M}^2$) and the variance of the measurement error ($\sigma_{\varepsilon^M}^2$), based on the following moment conditions:

$$E[\Delta u_t^2] = \sigma_{\zeta^M}^2 + 2\sigma_{\varepsilon^M}^2 \quad (13)$$

$$E[\Delta u_t \Delta u_{t-1}] = -\sigma_{\varepsilon^M}^2 \quad (14)$$

where u_t is the residual log earnings obtained after regressing earnings on dummies for age, disability status, and year.

Wage Process for Women

We implement a two-step Heckman selection correction procedure. Wages are given by

$$\log w_{it} = \mathbf{X}_{it}\boldsymbol{\beta} + u_{it}. \quad (15)$$

Table E.1: Employment status Probit regressions - Women

VARIABLES	(1) coeff.	(2) marg. eff.
Average AFDC payment (\$100)	-0.064*** (0.007)	-0.021*** (0.003)
Average food stamps payment (\$100)	-0.002 (0.095)	-0.008 (0.031)
Average EITC payment (\$100)	0.183*** (0.054)	0.060*** (0.018)
Age dummies		Yes
State dummies		Yes
Year dummies		Yes
Controls		Yes
Observations		69,832

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Data from the 1990-2008 SIPP panels. Sample of non-college graduates. Annualized data.

In \mathbf{X}_{it} we include age dummies, disability status, race, state dummies and year dummies. Wages are observed only when the woman works ($P_{it} = 1$), which happens under the following condition:

$$\mathbf{Z}_{it}\boldsymbol{\gamma} + \nu_{it} > 0,$$

In \mathbf{Z}_{it} we include \mathbf{X}_{it} and a vector of exclusion restrictions, assumed to explain employment decision but not wages. These exclusion restrictions are “simulated” welfare benefits, as described in Low and Pistaferri (2015). In particular, we use state, year and demographic variation in simulated AFDC, EITC and food stamps benefits for a single mother with varying number of children who works part-time at the federal minimum wage. Since we also control for time and state effects the instruments capture differential changes in policy over time and states. The first stage is reported in table E.1 and clearly demonstrates the strength of the instruments.

We use the selection correction to estimate the age profile of a woman’s wage and to account for non-participation when estimating the variance of the productivity innovations. The GMM estimates of the variance of the permanent component of log income ($\sigma_{\zeta_F}^2$) is

obtained by solving the following moment conditions:

$$\begin{aligned}
E[\Delta u_t \mid P_t = 1, P_{t-1} = 1] &= \sigma_{\nu\zeta^F} \left[\frac{\phi(\alpha_t)}{1 - \Phi(\alpha_t)} \right] \\
E[\Delta u_t^2 \mid P_t = 1, P_{t-1} = 1] &= \sigma_{\zeta^F}^2 + \sigma_{\nu\zeta^F}^2 \left[\frac{\phi(\alpha_t)}{1 - \Phi(\alpha_t)} \alpha_t \right] + 2\sigma_{\varepsilon^F}^2 \\
E[\Delta u_t \Delta u_{t-1} \mid P_t = 1, P_{t-1} = 1, P_{t-2} = 1] &= -\sigma_{\varepsilon^F}^2
\end{aligned}$$

where $\alpha_t = -\mathbf{Z}_t\boldsymbol{\gamma}$ and $\phi(\cdot)$ and $\Phi(\cdot)$ are the standard normal density and distribution function. We ignore selection correction for the first order covariance in order to reduce noise.

Moments

The match between data moments and simulated moments is reported in Table E.2.

Parameter Estimates

We can put the stigma cost, η , into context by asking what is the minimum extra consumption needed to compensate for the fact that the source of the extra consumption comes from welfare payments. We calculate the amount of yearly consumption \hat{c} that would be needed to make someone on welfare indifferent between claiming and not claiming monthly benefits π for 12 months. In a simple static setup, we calculate π for a given \hat{c} using the expression: $\frac{(\hat{c}+12\pi)^{1-\gamma}}{1-\gamma} - \eta \approx \frac{\hat{c}^{1-\gamma}}{1-\gamma}$. In Figure E.2 we plot the amount π against \hat{c} .

Implications of the Model: Women's Earnings over the Life-Cycle

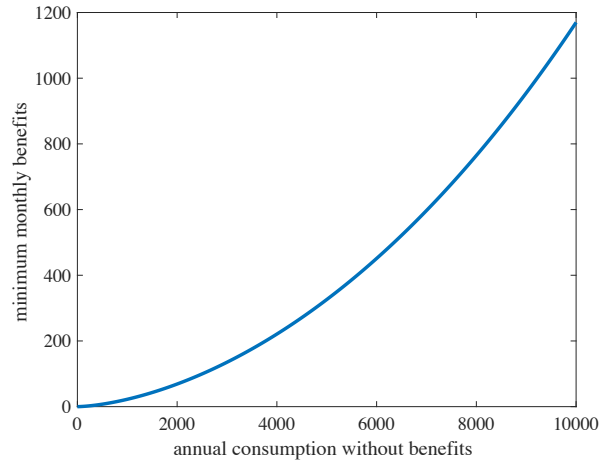
Figure E.3 shows the evolution of observed earnings over the life cycle for women. The model replicates both the level and the concavity of the profiles for all women and, crucially, for unmarried women. While we estimate a concave offered wage profile (as documented in Table 6), the fit is not mechanical because the employment decision of the woman is endogenous. In other words the model captures both the shape of the profile and the nature of selection into work.

Table E.2: Target moments

No.	Moment Description	Data		Model
		Mean %	(s.e. in %)	Mean %
1	% employed (married without children)	85.45	0.01	84.65
2	% employed (unmarried without children)	87.96	0.01	87.20
3	% employed (married with children)	60.02	0.00	60.22
4	% employed (unmarried with children)	55.78	0.01	56.18
5	% on AFDC (unmarried with children)	37.53	0.01	36.69
6	% ever married at age 21	31.20	0.03	38.72
7	% ever married at age 22	46.46	0.02	48.55
8	% ever married at age 23	53.54	0.01	56.31
9	% ever married at age 24	63.22	0.01	62.79
10	% ever married at age 25	67.58	0.01	68.06
11	% ever married at age 26	72.08	0.01	72.25
12	% ever married at age 27	74.65	0.01	75.57
13	% ever married at age 28	78.60	0.01	78.35
14	% ever married at age 29	81.35	0.00	80.55
15	% ever married at age 30	82.49	0.00	82.05
16	% ever married at age 31	85.16	0.00	83.46
17	% ever married at age 32	85.76	0.00	84.52
18	% ever married at age 33	86.20	0.00	85.19
19	% ever married at age 34	85.76	0.01	85.72
20	% ever married at age 35	87.02	0.01	85.96
21	% divorced at age 26	11.81	0.00	13.79
22	% divorced at age 27	13.32	0.00	14.36
23	% divorced at age 28	14.78	0.00	14.79
24	% divorced at age 29	15.54	0.00	15.05
25	% divorced at age 30	15.52	0.00	15.38
26	% divorced at age 31	16.05	0.00	15.73
27	% divorced at age 32	17.17	0.01	16.40
28	% divorced at age 33	16.75	0.01	16.98
29	% divorced at age 34	17.51	0.01	17.76
30	% divorced at age 35	18.63	0.01	18.54

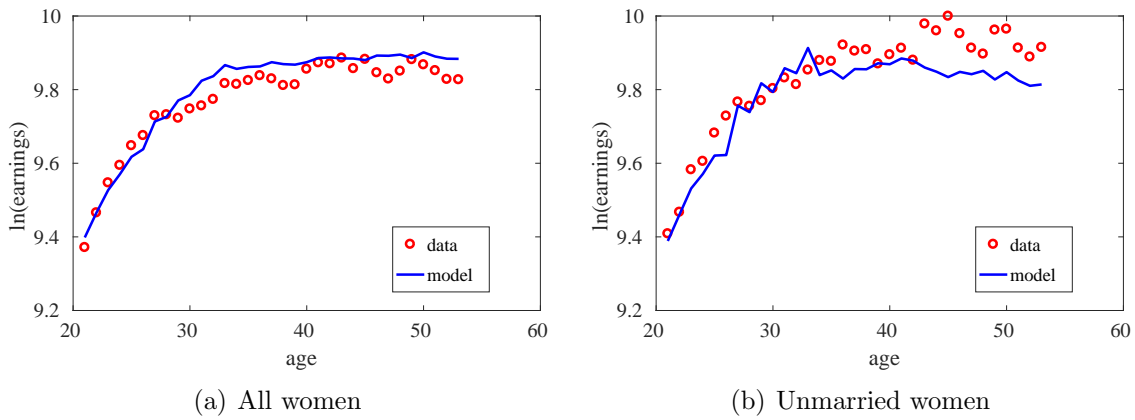
Notes: SIPP data, panels 1990-2008, but restricted to pre-reform data, which varies by state between 1995 and 1998. Sample of women born in the 1960s and aged 21-35 without college degrees. Annualized data.

Figure E.2: Static cost of claiming welfare implied by the estimated utility cost of claiming AFDC



Notes: Simulations from the model for $\eta = 0.0071$. For each c on the horizontal axis, the figure plots π that solves $\frac{(c+12\pi)^{1-\gamma}}{1-\gamma} - \eta = \frac{c^{1-\gamma}}{1-\gamma}$.

Figure E.3: Life-cycle profiles of log-wages for women in the data and in the model



Notes: Data from 1990-2008 panels of SIPP. Data from the 1960-69 birth cohorts, pre-reform.

Appendix F: Effect of Time Limits

Counterfactual distributions of singles' characteristics in the marriage market

After time limits are introduced in a counterfactual exercise, single women are characterized by the vector $\{\log(A_t), \log(y_t), k_t, TB_t\}$, where $TB_t \in [0, 5]$ represents the years of welfare that the woman has used since the reform. Because TB_t is not observed in our data after reform, we use an iterative procedure. First, we assume a uniform distribution for TB_t , solve the model and simulate the reform. Then, we compute the simulated conditional distributions of TB_t for each asset, income and fertility type. Last, we solve the model again with these updated conditional distributions and use the resulting policy function to perform the counterfactual exercises.

Revenue Neutrality

In our baseline model, the payroll tax rate on labor is set to 0, and hence we let the government run a deficit \bar{D} . When introducing time limits, we hold the government budget deficit \bar{D} constant. This is achieved by adjusting the proportional payroll tax on women's earnings, τ_w , such that the present discounted value of net revenue flows remains constant:

$$\sum_{i=1}^N \sum_{t=1}^T \frac{1}{(1+r)^{t-1}} [FS_{it} + EITC_{it} + b_{it}] = \sum_{i=1}^N \sum_{t=1}^T \frac{1}{(1+r)^{t-1}} \tau_w w_{it} P_{it} + \bar{D}$$

where b captures the payment through AFDC or time limits. This calculation can be carried out using realized payments. By doing this we are able to evaluate the welfare implications of the reform by allowing the same population to benefit from (or pay for) the resulting changes in government deficits and thus abstracting from redistribution from other groups. Moreover any distortionary taxes needed for this calculation are accounted for.³⁵

In practice, we first calculate the left hand side in the baseline. This gives the size of $\bar{D} = 147$ per woman/year when $\tau_w = 0$. Second, we change the policy rule into a 5-year limit and recalculate the LHS. This gives the new deficit if the tax rate remains at zero ($\bar{D}' = 78.5$ per woman/year). Third, we iterate on τ_w so that the deficit under the new policy is equal

³⁵Note that the summation is taken over the women only. This is because our simulations do not keep track of men unless they are in a relationship. So, benefits are being spent on the women and the extra tax to cover those benefits is being taken from the women. Hence, the amount of taxes raised from the men is held constant and budget balance comes only from women.

to \bar{D} . As τ_w adjusts, the choices individuals make will change and so the model needs to be solved again at each iteration. However, the final iteration gives the behavior of individuals in the new policy regime holding revenue constant, which occurs when $\tau_w^* = -0.7\%$ for the case of a 5 year time limit.

Welfare Calculation

To define the welfare cost or benefit of introducing time limits, we compute the lifetime expected utility of a woman in our model as

$$E_0 U(s, \tau) = E_0 \sum_{t=1}^T \beta^{t-1} \left(\frac{(c_t^s \cdot e^{\psi(M, k^a) \cdot P_t^s})^{1-\gamma}}{1-\gamma} - \eta B_t^s + L_t m_t^s \right),$$

where $\{c_t^s, P_t^s, B_t^s, m_t^s\}$ refer to the implied consumption, labor supply, benefit stream and marital status in the baseline economy ($s = AFDC$) or in an alternative economy with different welfare parameters (e.g. $s = 5TL$ (*5 year limit*)) and τ is the revenue neutral tax rate. E_0 represents the expectation at the beginning of working life, before initial conditions are known.

We calculate the proportion of consumption an individual is willing to pay ex-ante to be indifferent between environment $s' = 5TL$ and $s = AFDC$.³⁶

$$E_0 U(5TL, \tau_w) |_{\pi} = \sum_{i=1}^N \sum_{t=0}^T \beta^t \left(\frac{((1-\pi) c^s \cdot e^{\psi(M, k^a) \cdot P^s})^{1-\gamma}}{1-\gamma} - \eta B_t^s + L^t m_t^s \right) \quad (16)$$

We solve for π such that

$$E_0 U(5TL, \tau_w) |_{\pi} = E_0 U(b), \quad (17)$$

where π can be interpreted as the consumption cost of going from AFDC to a 5-year limit.³⁷

³⁶Ex-ante, no one knows the sequence of shocks that will be realised and so since there are no aggregate shocks, realized discounted lifetime utility averaged across all individuals will be equal to expected utility.

³⁷In varying π we do not reoptimize.