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

Hamish Low
Costas Meghir
Luigi Pistaferri
Alessandra Voena

Marriage, Labor Supply and the Dynamics of the Social Safety Net
Hamish Low, Costas Meghir, Luigi Pistaferri, and Alessandra Voena
NBER Working Paper No. 24356
February 2018, Revised November 2020
JEL No. H2,H31,H53,J08,J12,J18,J22

ABSTRACT

The 1996 US welfare reform introduced limits on years of welfare receipt. We show that this reduced program participation, raised employment for single mothers, and reduced divorce. A limited commitment, lifecycle model of labor supply, marriage and divorce, estimated on pre-reform data, replicates these effects. A large part of the responses occur in anticipation of benefit exhaustion, impacting primarily women with low potential earnings. The reform reduces lifetime utility of women, even allowing for the government savings, but has negligible effects on men. The expectation of marriage attenuates the losses for women and an increased probability of single-motherhood raises them.

Hamish Low
Department of Economics
University of Oxford
Oxford England
hamish.low@economics.ox.ac.uk

Costas Meghir
Department of Economics
Yale University
37 Hillhouse Avenue
New Haven, CT 06511
and IZA
and also NBER
c.meghir@yale.edu

Luigi Pistaferri
Department of Economics
579 Serra Mall
Stanford University
Stanford, CA 94305-6072
and NBER
pista@stanford.edu

Alessandra Voena
Stanford University
Landau Economics Building
579 Jane Stanford Way
Stanford, CA 94305
and NBER
avoena@stanford.edu

Introduction

Welfare programs are an important source of insurance for low-income households. If carefully targeted and designed, they can can increase overall welfare. However, they can also distort household decisions in several dimensions that interact with each other: work incentives, marriage and divorce may all be affected. These issues have been the source of continuous debate and motivated the major US welfare reform of 1996: the key innovation of this reform was to make welfare only temporary, whereas previously there were no time limits on support.[1] The aims of this paper are first, to understand the dynamic trade-offs between incentives and insurance when support becomes time-limited; and second, to evaluate the utility consequences of reform in a rich life-cycle environment that accounts for the role of family structure.

We start by empirically documenting the short-run effects of the reform on labor supply, welfare participation and marital status separately. In particular, we use data from the SIPP (1989-2007) and the March CPS (1990-2007) to estimate the effects of introducing a limit to the number of years one can receive welfare based on a quasi-experimental approach. Women whose youngest child is close enough to 18 years old would have initially remained unaffected by the introduction of time limits because eligibility has always depended on having a dependent child; by contrast, women with younger kids will be affected by the time limit in benefit eligibility. Considering each choice in isolation, we show that welfare utilization declined dramatically and persistently especially for single mothers without an actual reduction in immediate eligibility but mainly because of the anticipation of future needs: benefits were being "banked"; and that the employment of all mothers increased, partially offsetting the decline in benefit use. Alongside these effects, the prevalence of divorces declined. The results on the anticipatory decline in benefit take-up is consistent with the findings of Grogger and Michalopoulos (2003) who used a earlier experiment designed to test the impacts of time limits.

The reduced form results support the idea that the time-limits reform had a substantial effect on behavior, much of it anticipatory. But to better understand the lessons from this key event, its impact on utility and the mechanisms that underlie the longer-term impacts we

---

[1]Prior to the reform of 1996, the main welfare program for mothers of dependent children was AFDC (Aid to Families with Dependent Children), which covered families with dependent children based on income and asset requirements, but without a time limit. After reform and the imposition of time limits, this was renamed TANF (Temporary Assistance to Needy Families). Under TANF, federal block-grants cover adult recipients with dependent children for up to 5 years.

specify a life-cycle model of family formation and dissolution, welfare program participation, labor supply and savings. In our model, married couples share risk and resources, and enjoy economies of scale in a limited commitment framework, which is suitable for understanding whether policy changes can affect intra-household allocations.[2] The resulting framework has also broader implications, because it allows us to understand how behavior is shaped by the interaction of the various components of the welfare system, that we model in some detail.

Thus, we account for various welfare programs when modeling the budget constraint because the way individuals change behavior following welfare reform crucially depends on how the various programs interact. We also account for uncertainty in both income and family formation: the presence and source of uncertainty is of first order importance in determining individual behavior and how welfare programs affect them. This allows us to understand the dynamics implied by time limits, as well as interactions and feedback between the different programs and choices.

We estimate our model using the method of simulated moments (McFadden, 1989; Pakes and Pollard, 1989) on pre-reform data, focusing on low-educated women of working age. We analyze intra-household 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.

We find that a forward-looking model can fully account for the decline in welfare participation, the increase in employment and the decline in divorce that we observe in the data as a response to time limits. We focus our analysis on low educated women, but even within this group, the reform had the greatest effects on women in the bottom two quintiles of the individual productivity distribution, which in our model is unobserved heterogeneity for women who are not employed.

The loss in utility from time limits is only partially compensated by the government revenue saved: even after imposing revenue neutrality and redistributing the resulting lower taxes exclusively within the group of low educated women and men, introducing time limits leads to a decline in utility enjoyed by women equivalent to 0.5% to 0.7% of their lifetime consumption. By contrast, the effect on the utility of men is negligible: their reduced benefit eligibility when married is offset by gains in terms of intra-household allocations and lower

---

[2]By 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).

taxes.

Marriage prospects are very important in moderating these effects: in the absence of a marriage option, the utility loss from time limits for women is 50% to 90% larger, depending on how the tax savings are redistributed. This implies that welfare reforms that reduce its generosity may become substantially more costly when marriage becomes less common, as observed in past decades (Stevenson and Wolfers, 2007).

Introducing time limits was only one reform that could have reduced government spending on welfare. We show that an alternative reform that reduces benefit generosity but imposes no time limit, while having equal budgetary implications, is better for individual utility than the time limits reform. Finally, a comprehensive assessment of the welfare reform should account for the role of other programs that form the safety net. This assessment is complex because the various programs may act as complements or substitutes to AFDC/TANF: budget savings from imposing restrictions on one program (such as time limits) may be limited if participation in other programs increases as a consequence of the reform; or savings may be greater if the reform moves individuals into employment and lifts households above means test thresholds.

Our paper builds on existing work relating both to welfare reform and to life-cycle behavior. The literature on the effects of welfare reform is large, with excellent overviews by Blank (2002) and Grogger and Karoly (2005). Empirical studies have highlighted that time limits encourage households to limit benefit utilization so as to "bank" their future eligibility (Grogger and Michalopoulos, 2003; Grogger, 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, Fang and Keane (2004); 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 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 marital arrangements are likely to be affected by welfare reform: this was indeed part of the motivation of the reform. Further, ignoring this margin may lead to misleading conclusions about the effects of the reform on

behavior, utility and within-family insurance.[3] We consider the joint decision of being single vs. married alongside benefit use and employment. Some evidence suggests that the reform was associated with a small decline in both marriages and divorces, but estimated effects tended to be 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 decisions and allowing intra-household 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 (2002b) as extended to the lifecycle marriage model by Mazzocco, Yamaguchi and Ruiz (2013) and Voena (2015).[4]

In what follows, we provide detail of the 1996 welfare reform in Section 1. We present the data and reduced form analysis of the effects of time limits in Sections 2 and 3, respectively. In section 4 we present our life-cycle model, and provide details of the estimation in section 5. The core results are presented in counter-factual analysis in section 6 and welfare analysis in Section 7. Section 8 shows the interaction of reform to AFDC with other aspects of government support and Section 9 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 provide background on the features of the welfare reform of 1996, focusing on the changes in eligibility criteria.[5] The reform 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

---

[3]Autor et al. (2019) similarly stress the importance of within-household insurance in the context of disability insurance.

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

[5]The reform, signed by Bill Clinton in August 1996, was named the Personal Responsibility Work Opportunity Reconciliation Act (PRWORA).

(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. During the same period the Earned Income Tax Credit (EITC) was expanded to increase labor force participation among low-income individuals.

The introduction of time limits was a central element of the reform. Federal funds could be used to provide assistance to family units only up to a maximum of sixty months.[6] 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.[7]

Another element of reform was the introduction of stronger incentives to work than existed under AFDC. Work incentives came in two forms: the imposition of work requirements for maintaining welfare eligibility and the availability of child care assistance.[8] Work requirements 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.[9] While work requirements and child care assistance may be important and also potential empirical confounders of the effect of time limits (our main focus), 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 variation to isolate the effect of time limits from other elements of the reform.

A statutory goal of the welfare reform 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 extends eligibility to all households that satisfy the income and asset tests. For most aspects, the implementation of pro-family policies was again left to the states. 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

---

[6]The 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.

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

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

[9]Appendix A provides more detail on work requirements. Bruins (2017) argues that work requirements for single mothers increased poverty by cutting TANF eligibility among those most in need.

biological father of their children.

Figure 1: AFDC and Food Stamp Amounts

(a) AFDC benefits
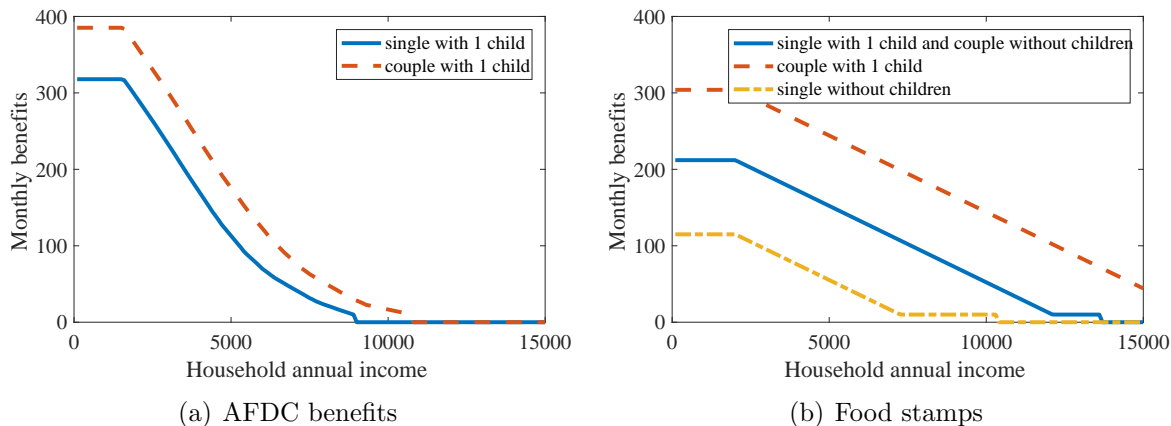
(b) Food stamps

Notes: Monthly AFDC and food stamps benefits by household annual income. AFDC benefits are a population-weighted state-level average. Amounts are in 1995 real U.S. dollars.

These changes in eligibility requirements were not accompanied by appreciable changes in the amount of benefits available conditioning on eligibility. In Figure 1 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 1 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.[10] 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. By contrast, for example, 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). As the figure clearly shows, different programs account differently for household composition and marital status, and hence have different implications for how program participation can influence or be influenced by marital status. [11]

---

[10]In 1995, the poverty line was $10030 a year for a single mother and $12590 for a couple with 1 child.

[11]Another 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).

2　Data

To study the dynamic effects of the welfare reform, we use seven panels of the Survey of Income and Program Participation (SIPP) spanning the 1989-2007 period and the March Current Population Survey (CPS) for years 1990 to 2007.[12] The SIPP is a representative survey of the US population collecting extensive information on participation in welfare and social insurance programs. Starting with the 1996 panel, the SIPP has been re-designed to include an oversampling of households from high-poverty areas, and hence is no longer nationally representative. In each panel, people are interviewed every four months.[13] The CPS is primarily a labor force survey which is conducted monthly, the March supplement includes additional information on family characteristics and program participation and is nationally representative. The so-called CPS Outgoing Rotation Group provides a limited longitudinal sample. We restrict the sample to women 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.

Our analysis includes 51,509 women in the SIPP who are heads or spouses of the head of their household, leading to a total of 336,129 quarterly observations. Our CPS sample includes 86,813 women and a total of 153,498 annual observations.[14] Table B.1 in the Appendix provides summary statistics for the SIPP and the CPS, pre- and post-reform.

3　Reduced-Form Evidence on Time Limits

We use the SIPP and CPS data to examine the relationship between the introduction of time limits and key outcome variables: welfare benefit utilization, women's employment, marriage, divorce, and fertility. In this section, we consider these outcomes separately, and then we use our model to show how the outcomes interact with one another.

---

[12]The SIPP panels used are the 1990, 1991, 1992, 1993, 1996, 2001, and 2004 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.

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

[14]The 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).

3.1 Empirical Strategy

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).[15]

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.[16] Mothers of younger children faced more consequential cuts in welfare support than those with older children because the time period over which they could claim is longer.

Figure 2: Time Limits Across States (2000)

The relationship between this exposure variable and the effect of time limits becomes

---

[15]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 and such households experience a drop in monthly income from welfare of about $250 on average. They find no evidence that such loss is offset by increases in other income sources and conclude that time limits enforcement resulted in an increase in the rates of deep poverty for households hitting the time limits.

[16]For example, if a households'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 households's youngest child is aged 12 or below in year t and the state's lifetime limit is 60 months, the variable *Exposed* takes value 1.

increasingly attenuated over time as we do not observe the actual history of welfare utilization, which, furthermore, is an endogenous variable. 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 other features of the welfare reform other than time limits interact with age of children in a complex way, our strategy can identify the 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. Note that Figure 2 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 when the time limit clock starts to tick.[17]

We estimate two versions of the reduced form relationship. First, we construct a variable $Post_{st}$ to indicate the post-reform period, based on the timing of the introduction of time limits.[18] We look at the main impacts a few years after the reform (up to 2002) as well as the longer-term effects (the entire period covered by our data, which end in 2007). Given that we do not observe how much of their benefits people have used, we focus mostly on effects that arise in the period immediately after the reform, where the risk of treatment/control contamination is minimal. We interact the $Post_{st}$ variable with $Exposed_{ds}$ and estimate:

$$y_{idst} = \alpha \, Exposed_{dst} \times Post_{st} + \mathbf{X}_{idst}\boldsymbol{\beta} + f_{st} + f_{ds} + f_{mt} + \epsilon_{idst} \qquad (1)$$

where α is the average effect of the reform, identified off of changes in the behavior of household with a youngest child of different ages, $\mathbf{X}$ is a vector of individual and state-level time-varying controls, f_{st} are state-by-year fixed effects, f_{ds} are state-by-age-of-youngest-child fixed effects, and f_{mt} are month-by-year fixed effects.

Second, to study the dynamics of the outcome variables in more detail, we interact $Exposed_{ds}$ with dummies for each calendar year between 1990 and 2007 (excluding 1995 for scaling). We estimate pre-reform interactions with year dummies to rule out pre-reform trends across demographic groups:

---

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

[18]Years of time limits are reported in Mazzolari and Ragusa (2012).

$$y_{idst} = \sum_{\tau=1990}^{2007} \alpha_\tau \, Exposed_{dst} \times \mathbf{1}\{t = \tau\}_t + \mathbf{X}_{idst}\boldsymbol{\beta} + f_{st} + f_{ds} + f_{mt} + \epsilon_{idst}. \qquad (2)$$

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 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. By contrast, 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 Empirical Results

Benefits Utilization and Employment

Panel A of Table 1 reports the main effects of time limits within the first six years after 1996 (so that all states have at least five years of post-reform data regardless of when the time limits were implemented). We start by examining changes in welfare utilization. Using the SIPP, in the raw data, before the reform 10% of households claimed benefits, and among unmarried women, the rate was 30%.[19] These numbers fell to 4% and 11% respectively. Our reduced form analysis for welfare utilization shows how much of this fall is due to time limits. Exposed households have a 3 percentage point (pp) lower probability of claiming benefits after the introduction of time limits. Unmarried women have a 8.7pp 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. Longer-term effects (Panel B of Table 1) are qualitatively similar, and if anything larger. This may be because,

---

[19]Further details of the raw numbers are in the descriptive statistics in Table B.1

as time goes by, other program changes may interact with the welfare reform in similar directions. Similar effects are found in the CPS.

Table 1: Use of Benefits and Employment after the Reform

	Whole sample		Married women		Unmarried women	
Panel A: Up to 2002						
	SIPP	CPS	SIPP	CPS	SIPP	CPS
AFDC/TANF Utilization						
$Exposed_{dst}Post_{st}$	-0.030***	-0.016***	-0.011***	-0.003**	-0.087***	-0.084***
	(0.004)	(0.003)	(0.003)	(0.002)	(0.015)	(0.013)
Mean pre-reform	0.098	0.077	0.035	0.019	0.297	0.304
Obs	254,627	112,128	188,483	88,522	66,144	23,606
R^2	0.12	0.07	0.08	0.03	0.26	0.15
Employment						
$Exposed_{dst}Post_{st}$	0.014	-0.002	-0.001	-0.017	0.050***	0.054**
	(0.012)	(0.011)	(0.014)	(0.011)	(0.014)	(0.026)
Mean pre-reform	0.640	0.647	0.643	0.654	0.631	0.620
Obs	254,627	112,128	188,483	88,522	66,144	23,606
R^2	0.12	0.06	0.11	0.05	0.21	0.13
Panel B: Whole sample period						
	SIPP	CPS	SIPP	CPS	SIPP	CPS
AFDC/TANF Utilization						
$Exposed_{dst}Post_{st}$	-0.038***	-0.022***	-0.013***	-0.005***	-0.108***	-0.111***
	(0.004)	(0.002)	(0.002)	(0.001)	(0.012)	(0.010)
Mean pre-reform	0.098	0.077	0.035	0.019	0.293	0.298
Obs	336,129	153,498	242,825	119,905	93,304	33,593
R^2	0.11	0.07	0.07	0.03	0.27	0.15
Employment						
$Exposed_{dst}Post_{st}$	0.007	-0.014	-0.014	-0.031***	0.055***	0.053**
	(0.011)	(0.009)	(0.014)	(0.010)	(0.013)	(0.021)
Mean pre-reform	0.641	0.648	0.644	0.655	0.632	0.623
Obs	336,129	153,498	242,825	119,905	93,304	33,593
R^2	0.11	0.06	0.11	0.05	0.19	0.12

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. Data from the 1990-2004 SIPP panels and 1990-2007 March CPS. Sample of female heads of household who are not college graduates and have children aged 18 and below. Sampling weights used in the SIPP regressions. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-by-year fixed effects, state-by-demographics fixed effects, unemployment-rate-by-demographics fixed effects, and a dummy for two or more children interacted with dummies for post-1993 and post-1996.

Next, consider the the main impacts of the reform on employment levels. The introduc-

tion of time limits does not significantly change employment among married women, while the impact on single mothers is to increase employment by about 5pp. The evidence is qualitatively similar if we look at effects in the longer term (Panel B of Table 1).[20]

The increase in employment in the group of unmarried women is likely to be a direct consequence of the decline in welfare utilization, suggesting that labor supply is the main mechanism of substitution. However, we show in Table C.2 in Appendix C that there is an increase of between 2 and 5pp in the fraction of single mothers who are neither employed nor on welfare. The question remains of what happens to the income of those who move off welfare but do not become employed. Table C.2 further shows a small decline in employed women on welfare, and this reinforces the view that work requirements are unlikely to be driving the decline in welfare use.

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 specification, 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 4 percentage points by 1999. It appears that households reduce their benefit utilization *before* anyone is likely to have run out of benefits eligibility. The decline is around 10 percentage points for unmarried women.
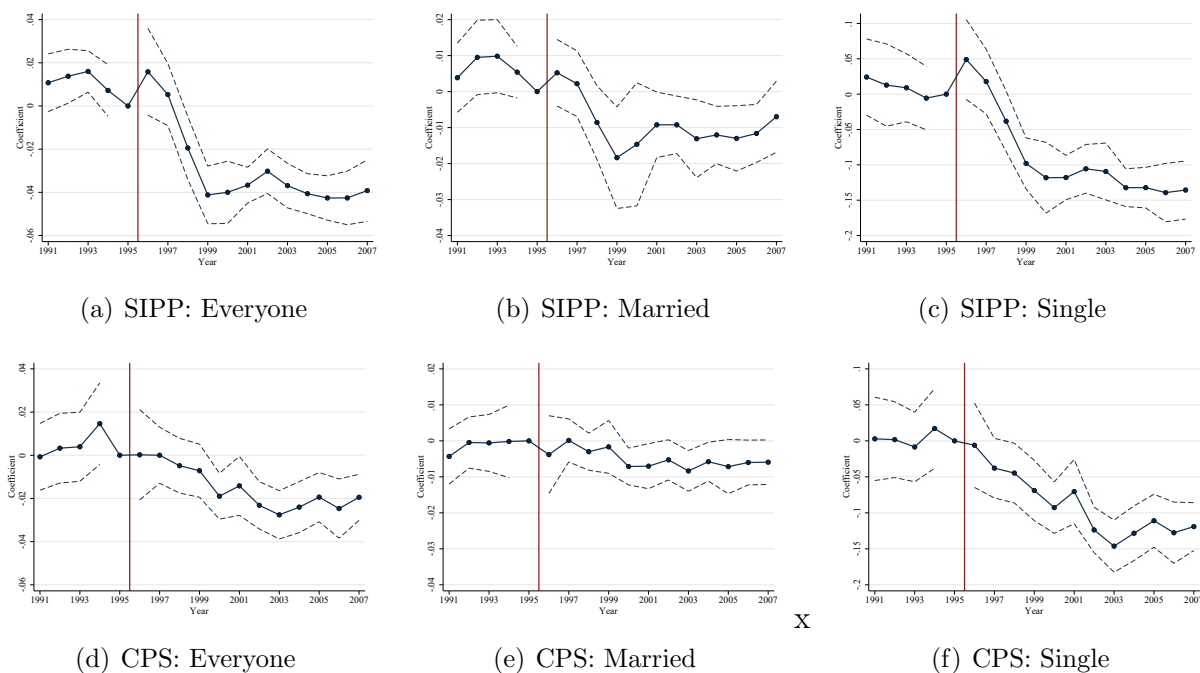
We provide additional evidence of forward looking behavior in two ways: first, we count the number of quarters since the introduction of time limits in each state and re-define the annual exposure dummy as $Exposed_{dst} \times \mathbf{1}\{\tau \, quarters \, since \, TL\}_{st}$. In Figure 4, we show the decline in welfare use on the sample that excludes states with shorter time limits (less than 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).

In Appendix C we perform a number of additional robustness checks. We interact the exposure coefficient with dummy variables for the age of the youngest child. We show in Figure C.1 in the Appendix that 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, consistent with the dynamic incentives introdced by time limits.

---

[20]The only notable deviation from the main effects is the significant decrease in employment among married women in the CPS sample. This highlight the importance of distinguishing between main effects and longer-term effects, since by the end of our sample period other welfare program changes may have interacted with the effects of time limits we focus on.

Figure 3: Program Participation Dynamics

(a) SIPP: Everyone   (b) SIPP: Married   (c) SIPP: Single

(d) CPS: Everyone   (e) CPS: Married   (f) CPS: Single

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. Data from the 1990-2004 SIPP panels and 1990-2007 March CPS. Sample of female heads of household who are not college graduates and have children aged 18 and below. Sampling weights used in the SIPP regressions. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-by-year fixed effects, state-by-demographics fixed effects, unemployment-rate- by-demographics fixed effects, and a dummy for two or more children interacted with dummies for post-1993 and post-1996.

We run three further robustness checks. First, we consider only the first wave of the SIPP for every household to address potential selective attrition (Table C.4). Second, we run a falsification exercise where we use a sample of college graduates to show that their behavior is little affected by time limits (Table C.5). Finally, we only consider mothers of children aged 6 and above (Table C.6). The similarity of 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. Our findings across these robustness checks are consistent with our baseline results.

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

Figure 4: Program Participation Following the Introduction of Time Limits

(a) Full sample

(b) Time limits above 24 months
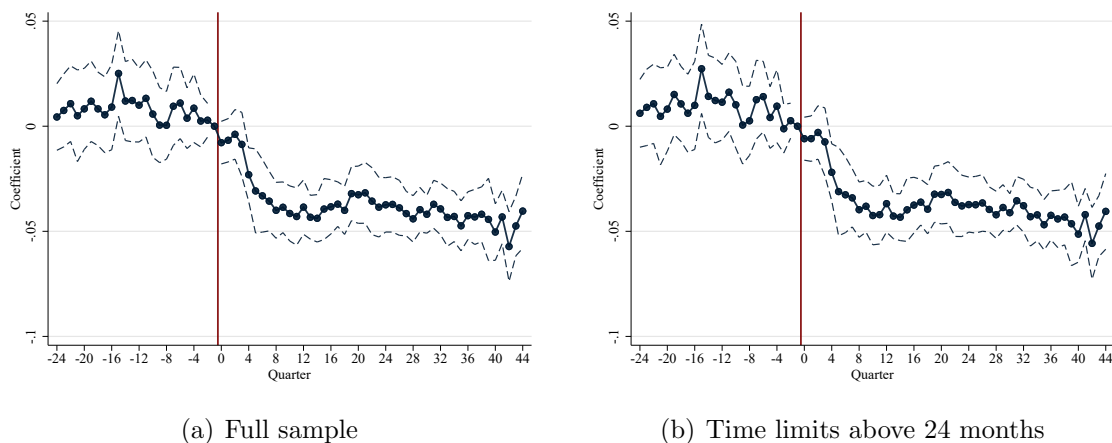
Notes: 95% confidence intervals based on standard errors clustered at the state level. Data from the 1990-2004 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. Sampling weights used in the SIPP regressions. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-by-year fixed effects, state-by-demographics fixed effects, unemployment-rate- by-demographics fixed effects, and a dummy for two or more children interacted with dummies for post-1993 and post-1996.

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. In this section, we consider whether there is direct evidence that time limits affected divorce and marriage.

Table 2 studies these outcomes, distinguishing between the main effects of the reform on marital status up to 2002 and effects in the longer term. The two are fairly similar, so here we comment on the former. Start with the impact of the welfare reform on the probability of being divorced or separated for women. We observe an imprecise decline in the probability of transitioning into divorce conditional on being married during the previous interview. In terms of stocks, women exposed to time limits are 1.5-2.7pp less likely to be divorced or separated, and this is statistically significantly different from zero in both the SIPP and CPS. Figure C.3 reports the dynamics of these outcomes. The last set of rows of Table 2 Panel A show that this decline in divorce was not associated with an increase in marriage.

These estimates suggest 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.[21] The changes in insurance will not only affect

---

[21]The literature has emphasized that, in theory, as discussed by Bitler et al. (2004), the effects of the

Table 2: Marital Status

	Panel A: Up to 2002			
	SIPP	CPS	SIPP	CPS
	Gets Divorced/separated		*Divorced/separated*	
$Exposed_{dst}Post_{st}$	0.000	0.003	-0.027***	-0.015*
	(0.001)	(0.005)	(0.007)	(0.008)
Mean pre-reform	0.009	0.014	0.150	0.126
Obs	160,210	37,617	254,627	112,128
R^2	0.01	0.02	0.03	0.01
	Gets Married		*Married*	
$Exposed_{dst}Post_{st}$	-0.000	-0.016	0.004	-0.007
	(0.003)	(0.015)	(0.007)	(0.010)
Mean pre-reform	0.025	0.047	0.758	0.796
Obs	54,441	9,727	254,627	112,128
R^2	0.04	0.08	0.05	0.05
	Panel B: Whole sample period			
	SIPP	CPS	SIPP	CPS
	Gets Divorced/separated		*Divorced/separated*	
$Exposed_{dst}Post_{st}$	-0.001	-0.002	-0.033***	-0.013*
	(0.001)	(0.003)	(0.009)	(0.006)
Mean pre-reform	0.009	0.014	0.151	0.126
Obs	207,562	52,528	336,129	153,498
R^2	0.01	0.02	0.03	0.01
	Gets Married		*Married*	
$Exposed_{dst}Post_{st}$	-0.001	-0.019*	-0.002	-0.014*
	(0.003)	(0.011)	(0.011)	(0.008)
Mean pre-reform	0.025	0.045	0.756	0.793
Obs	77,489	14,157	336,129	153,498
R^2	0.04	0.07	0.05	0.05

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. Data from the 1990-2004 SIPP panels and 1990-2007 March CPS. Sample of female heads of household who are not college graduates and have children aged 18 and below. Sampling weights used in the SIPP regressions. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-by-year fixed effects, state-by-demographics fixed effects, unemployment-rate-by-demographics fixed effects, and a dummy for two or more children interacted with dummies for post-1993 and post-1996.

---

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 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

the likelihood of marriage, but also women's bargaining power 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. We estimate equation 1 but use a dummy for having a newborn in $t+1$ as the dependent variable. 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 the SIPP and in the CPS. In no specification do we find that exposure to time limits influences the probability of future births, irrespective of marital status. This finding partly justifies our choice to treat fertility choices in the model as stochastic.[22]

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. Women also 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 a distribution of singles' types, 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 or consume. We allow for savings, both because these households do hold assets in the data and because any analysis of the welfare effects of the reform would be incomplete without taking explicitly

---

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.

[22]A separate question, that cannot be examined with this strategy, is whether welfare reform may have affected the onset of fertility, and particularly teenage fertility. We leave this important topic for future research.

into account self-insurance. Modeling savings is also important to prevent overstating the benefits of banking welfare benefits as one of the main forms of intertemporal choice. 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 consider the three main decision problems once marital status for a given period has been decided: the problem of a single woman, the problem of a single man, and the problem of a couple. The subscript i indexes the woman, the subscript j indexes the man, and ij indexes the corresponding couple.

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,[23] 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 ($c_{i,t}^{F_s}$) 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 $b_{i,t}$. 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 any stigma associated with receiving welfare (Moffitt, 1983). In addition, she makes a choice to marry, which will also depend on meeting a man and whether he will agree to marry her. 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

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

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

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,

---

[23]Since 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.

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 obtained 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 FS), EITC and AFDC/TANF, where the latter subject to time limits. The value of all programs depends on demographics and income.

The state space for a single woman is $\boldsymbol{\Omega_{i,t}^{F_s}} = \{A_{i,t}^{F_s}, w_{i,t}^F, k_{i,t}^a, TB_{i,t}\}$, 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. EITC is a function of the vector $\{k_{i,t}^a, w_{i,t}^F P_{i,t}^{F_s}\}$, food stamps are a function of $\{k_{i,t}^a, w_{i,t}^F P_{i,t}^{F_s}, A_{i,t}^F\}$, while AFDC/TANF are a function of the vector $\{k_{i,t}^a, w_{i,t}^F P_{i,t}^{F_s}, TB_{i,t}, A_{i,t}^F\}$. The dependence of AFDC/TANF and food stamps on assets is due to the presence of an asset test.

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 $m_{ij,t} = 1$, and 0 otherwise. 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.[24]

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

$$
V_t^{F_s}\left(\boldsymbol{\Omega_{i,t}^{F_s}}\right) = \max_{\boldsymbol{q_{i,t}^{F_s}}} \left\{
\begin{array}{c}
u^{F_s}\left(c_{i,t}^{F_s}, P_{i,t}^{F_s}, B_{i,t}^{F_s}\right) \\[2mm]
+\beta E_t \left[\lambda_{t+1} \left[\begin{array}{c} (1 - m_{ij,t+1}) V_{t+1}^{F_s}\left(\Omega_{i,t+1}^{F_s}\right) \\[2mm] + m_{ij,t+1} V_{t+1}^{F_m}\left(\Omega_{i,t+1}^m\right) \end{array} \right] \right. \\[4mm]
\left. + (1 - \lambda_{t+1}) V_{t+1}^{F_s}\left(\Omega_{i,t+1}^{F_s}\right) \right]
\end{array}
\right\}
$$

subject to the intertemporal budget constraint (3).

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

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

---

[24]In 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. Solving for the equilibrium distribution in two dimensions is likely to be very complicated computationally.

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.[25]

4.2 Problem of the Single Man

Single men are subject to an exogenous employment and wage process, and do not have an option to receive welfare benefits (other than food stamps). Children affect the man's problem only when he is married to the child's mother. The state space for a single man is defined by his asset holdings and his wage: $\boldsymbol{\Omega_{j,t}^{M_s}} = \{A_{j,t}^M, y_{j,t}^M\}$. The state space when married, $\boldsymbol{\Omega_{ij,t}^m}$, contains both the husband and wife's economic state variables. These assumptions determine $V_t^{M_s}\left(\boldsymbol{\Omega_{j,t}^{M_s}}\right)$, the man's value function when he is single, and $V_t^{Mm}\left(\boldsymbol{\Omega_{ij,t}^m}\right)$, the value accruing to a married man. The budget constraint of the single man is given by:[26]

$$\frac{A_{j,t+1}^{M_s}}{1+r} = A_{j,t}^{M_s} - c_{j,t}^{M_s} + y_{j,t}^M + FS_{j,t} \tag{5}$$

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

The problem for the single man is thus defined by

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

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:

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

---

[25]We 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.

[26]We 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 guarantee symmetry with the women's utility function. Note that P^M is not a choice for the man, but the result of an exogenous employment process.

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). Household choices 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 $\boldsymbol{\Omega_{ij,t}^m}$.[27] 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 $\left(d_{ij,t+1}\left(\boldsymbol{\Omega_{ij,t+1}^m}\right) \in \{0,1\}\right)$, and so the joint problem that the couple solves is

$$
V_t^m\left(\boldsymbol{\Omega_{ij,t}}^m\right) = \max_{\boldsymbol{q_{ij,t}^m}} \left\{ \begin{array}{c} \theta_{ij,t}^F u^{Fm}\left(c_{ij,t}^F, P_{ij,t}^F, B_{ij,t}^F\right) + \theta_{ij,t}^M u^{Mm}\left(c_{ij,t}^M, P_{j,t}^M\right) + L_{ij,t} \\ +\beta E_t \left[\begin{array}{c} (1 - d_{ij,t+1}) V_{t+1}^m\left(\boldsymbol{\Omega_{ij,t+1}^m}\right) \\ +d_{ij,t+1}\left(\theta_{ij,t}^F V_{t+1}^{F_s}\left(\boldsymbol{\Omega_{i,t+1}^{F_s}}\right) + \theta_{ij,t}^M V_{t+1}^{M_s}\left(\boldsymbol{\Omega_{j,t+1}^{M_s}}\right)\right) \end{array} \right] \end{array} \right\} \tag{6}
$$

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

$$
L_{ij,t} = L_{ij,t-1} + \xi_{ij,t} \tag{7}
$$

where $\xi_{ij,t}$ can be interpreted as a "love shock" to the marriage. The optimization above is subject to the budget constraint:

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

To capture economies of scale in marriage (including public goods), we assume that the

---

[27]The state variables for the couple (represented by $\boldsymbol{\Omega_{ij,t}}^m$), 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 $\theta_{ij,t}^M, \theta_{ij,t}^F$.

individual consumptions $c_{ij,t}^F$ and $c_{ij,t}^M$ and the equivalence scale $e(k_t^a)$ imply an aggregate household expenditure of

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

The extent of economies of scale is controlled by ρ and $e(k_t^a)$. If $\rho > 1$, consumption is partially public, and the sum of spouses' consumption exceed what they would consume if single and spending the same amount. If the couple divorce at the start of $t+1$, then $A_{ij,t+1}$ is divided equally.[28]

The value for each spouse is defined recursively so that in the last period, it is equal to the flow utility for each spouse evaluated at the optimum

$$V_T^{Fm} = u^{Fm}(c_{ij,T}^{*F}, P_{ij,T}^{*F}, B_{ij,T}^{*F}) + L_{ij,T} \quad \text{and} \quad V_T^{Mm} = u^{Mm}\left(c_{ij,T}^{*M}\right) + L_{ij,T}$$

and in other periods

$$V_t^{Fm}\left(\boldsymbol{\Omega_{ij,t}}^m\right) = u^{Fm}(c_{ij,t}^{*F}, P_{ij,t}^{*F}, B_{ij,t}^{*F}) + L_{ij,t} + \beta E_t\left[(1 - d_{ij,t+1})V_{t+1}^{Fm}\left(\boldsymbol{\Omega_{ij,t+1}^m}\right) + d_{ij,t+1}V_{t+1}^{Fs}\left(\boldsymbol{\Omega_{i,t+1}^{Fs}}\right)\right]$$

and

$$V_t^{Mm}\left(\boldsymbol{\Omega_{ij,t}}^m\right) = u^{Mm}\left(c_{ij,t}^{*M}\right) + L_{ij,t} + \beta E_t\left[(1 - d_{ij,t+1})V_{t+1}^{Mm}\left(\boldsymbol{\Omega_{ij,t+1}^m}\right) + d_{ij,t+1}V_{t+1}^{Ms}\left(\boldsymbol{\Omega_{i,t+1}^{Fs}}\right)\right]$$

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:

$$V_{t+1}^{Fm}\left(\boldsymbol{\Omega_{ij,t+1}^m}\right) \geq V_{t+1}^{Fs}\left(\boldsymbol{\Omega_{i,t+1}^{Fs}}\right)$$

$$V_{t+1}^{Mm}\left(\boldsymbol{\Omega_{ij,t+1}^m}\right) \geq V_{t+1}^{Ms}\left(\boldsymbol{\Omega_{j,t+1}^{Ms}}\right) \tag{8}$$

---

[28]This 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.

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 (2002a). If it is not feasible to satisfy both spouses' participation constraints and the intertemporal budget constraint for any allocation of resources, then divorce follows.[29] 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 (L), and possibly also because features of the welfare system promote it. And indeed, marriage can break down efficiently if no Pareto weights exist that would imply positive gains from marriage for each partner. 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 life as a single person 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 a 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, $\boldsymbol{\theta}_{ij,t_0}$, distribute these gains and we assume they are are chosen as the solutions to a symmetric Nash bargaining game between spouses, as described in Appendix D.[30] Ours is a context of

---

[29]We 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.

[30]In the estimation below, we experiment with varying the share of the surplus that goes to each person

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[31]

4.4 Sources of Uncertainty

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

Female wages and male earnings

Male earnings and female wages are respectively specified as

$$\log(w_{j,t}^M) = a_0^M + a_1^M t + a_2^M t^2 + z_{jt}^M + \varepsilon_{jt}^M$$

$$\log(w_{it}^F) = a_0^F + a_1^F t + a_2^F t^2 + z_{it}^F + \varepsilon_{it}^F$$

$$z_{jt}^M = z_{j,t-1}^M + \zeta_{jt}^M$$

$$z_{it}^F = z_{i,t-1}^F + \zeta_{it}^F.$$

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.[32] We allow for a stochastic employment shock that draws men's income to zero, while women can chose to work or not. Men's employment status evolves stochastically following a Markov process.

---

away from the one determined by symmetric Nash bargaining. 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.

[31]Characterizations 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.

[32]One 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.

Fertility

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

$$Pr\left(k_{t+1}^a | k_t^a, m_{t-1}, t\right). \tag{9}$$

Each woman can only have one child at a time, with age of the child treated at a state variable that can re-start if a newborn arrives. 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 marital status, fertility is partially endogenized through the marital decision.

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:

$$\lambda_t = min\{max\{\lambda_0 + \lambda_1 \cdot (t-1) + \lambda_2 \cdot (t-1)^2, 0\}, 1\}. \tag{10}$$

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 drawn 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.

5.1 Externally Set Parameters

Panel A of Table 3 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 the calibration in 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 along with household assets determine eligibility for AFDC. We calculate AFDC benefit for different household types by taking a population-weighted average value of benefits across states for different income levels, as reported in Figure 1.

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) as the relevant child care cost.

The Fertility Process

Each woman has at most one child at a time, and the arrival rate of a newborn is a function of a woman's age, marital status and the age of youngest child, which are state variables in the model. We compute the Markov process for fertility by examining transition probabilities in the SIPP data, as in equation (9).

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 $\{A_t, y_t, k_t\}$, where k_t indicates whether or not she has

had a child.[33] We allow for mass for the cases in which $A_t^j = 0$ and $y_t^H = 0$ and model non-negative income and assets to follow a bivariate log-normal distribution. 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.[34] Since we do not model the intensive margin, we assume that participating women work 1,530 hours per year (the median hours worked in the 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 employment transition matrix directly from the SIPP data. 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 3, 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.

---

[33]The 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.

[34]Female 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" × the number of weeks at the job within the month × 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%.

Table 3: Externally Set and Directly Estimated Parameters of the Model

Parameter	Value/source
Panel A: Externally Set Parameters	
Relative risk aversion (γ)	1.5
Discount factor (β)	0.98
Economies of scale in marriage (ρ)	1.23
Welfare program parameters	Statutory rules
Panel B: Directly Estimated Parameters	
Childcare costs (CC^a)	CEX (see text)
Fertility process	SIPP
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 annual earnings (a_0^M, a_1^M, a_2^M)	9.76, 0.043, -0.001
Life cycle profile of log female hourly wages (a_0^F, a_1^F, a_2^F)	1.96, 0.022, -0.0003

This also reflects differences in schooling among our (non-college graduate) group. Male and female wages have a concave lifecycle profile as expected.

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). The remaining unkown parameters are: 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 $\boldsymbol{\phi}_{sim}$ are computed using the full numerical solution of the model. We use the inverse of the variance-covariance matrix of the empirical moments as the weighting matrix, computed using the bootstrap method.[35]

We calculate the standard errors of our parameter estimates using the standard asymptotic formula (McFadden, 1989).[36]

5.3.1 Target Moments

Our estimation target moments can be classified into three sets. While all moments jointly contribute to the estimation of the structural parameters, each set is more closely tied to a particular corresponding set of parameters.

The first set of moments includes the fraction of women who are employed by marital and fertility status (Figure 5, panel a). These moments are primarily responsible for pinning down the utility cost of working $\psi(M, k^A)$: the higher such parameter, holding other parameters fixed, the lower the probability of employment.

The second set of moments includes the proportion of single mothers and of married mothers who are on AFDC. These moments play an important role in identifying the utility cost of being on welfare η: a higher utility cost naturally leads to lower welfare use, keeping the other parameters constant (Figure 5, panel b). Note that we can fit both of these moments closely without requiring η to vary by marital status.

The third set of moments includes the life cycle profile of the probability of being married (evaluated between the ages of 19 and 33) and the life cycle profile of the probability of being divorced (evaluated between the ages of 21 and 33), see Figure 5, panel c. These life cycle profiles contribute to determining the parameters characterizing the probability of meeting a potential partner λ_t (i.e., the parameters λ_0, λ_1 and λ_2 in equation (10)). *Ceteris paribus*, a higher λ_t corresponds to a higher chance of marriage at time t. A faster decline in λ_t increases the probability of marrying early given the heightened risk of not encountering a partner in the future. Note that these moments also contribute to determining the variance of the initial match quality draw σ_0^2 and of the innovations σ_ξ^2: a higher initial variance increases the probability of marriage and reduces the probability of early divorce while higher variance in innovations increases the probability of divorce over time.

---

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

[36]We calculate the derivatives of each moment with respect to each parameter, the 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 (Judd, 1998).

Table E.2 in Appendix E reports the complete set of target and simulated moments.

Figure 5: Target Moments of SIPP and Simulated Data

(a) Employment moments

(b) Welfare moments

(c) Marriage market moments
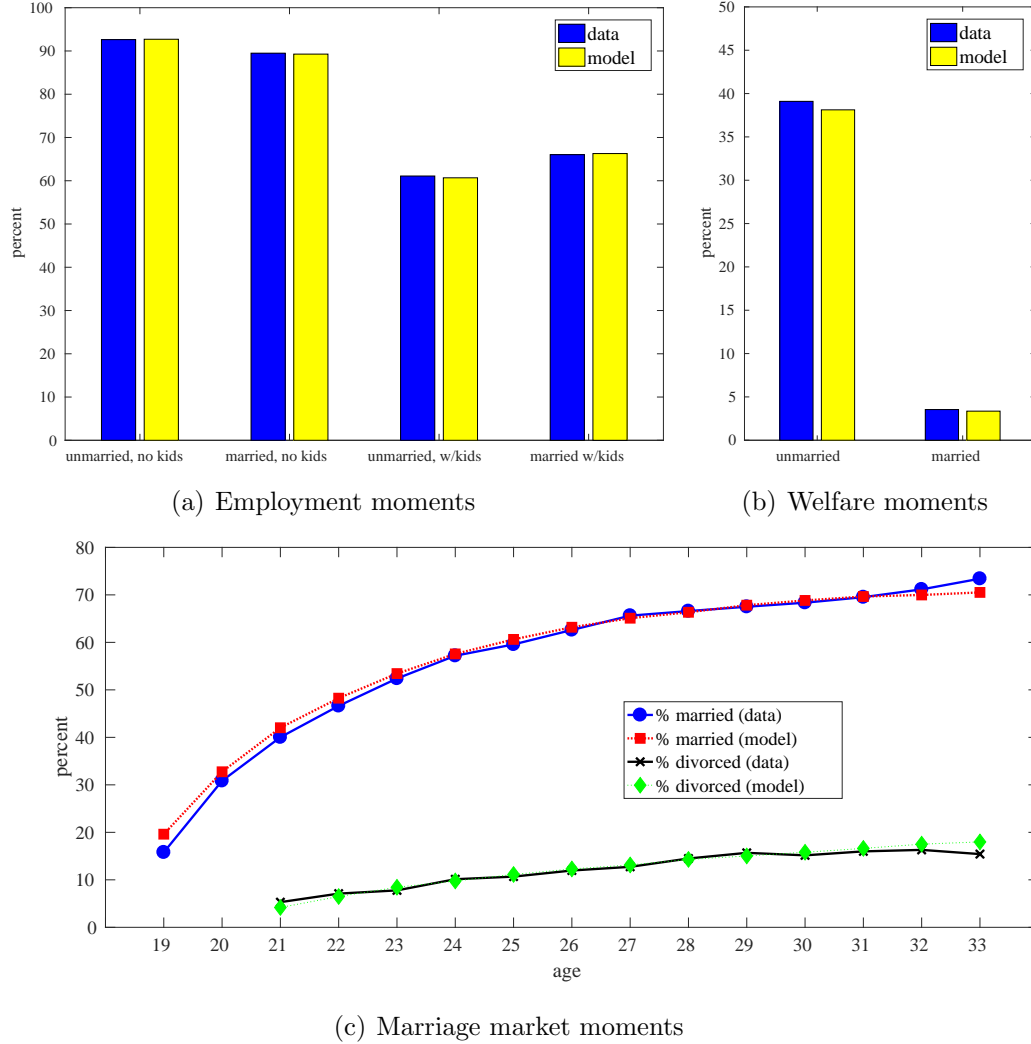
Notes: Target moments in the pre-reform SIPP data and in estimated model.

5.3.2 Estimated Model Parameters

Table 4 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 Figure 5, the estimation requires a large disutility of work for single women (similarly to Blundell et al. (2016)). For married women, the employment rate can be matched with a lower disutility of work because of the explicit presence of spousal income which discourages women's employment.

The estimated utility cost of claiming welfare benefits is high, and is identified by the women who are not claiming benefits despite being eligible given their income. In the pre-reform period, there was no intertemporal tradeoff to claiming benefits, and so we attribute the decision not to claim to utility or other costs of claiming. Our estimate of η implies that a single woman with annual baseline consumption of $6000 requires minimum monthly benefits of $78 to overcome the utility cost. In the counterfactual simulations, for the post reform period with time limits in place, the cost of claiming now is higher because of the loss of future benefits, and so we identify the direct utility cost from the pre-reform period.

The variance of innovations to match quality is one third of the variance 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.

Table 4: Parameters Estimated by Method of Simulated Moments

Parameter		Estimate	(s.e.)
Cost of work			
Unmarried, no children	$\exp\{\psi^{s0}\}$	0.33	(0.018)
Married, no children	$\exp\{\psi^{m0}\}$	0.59	(0.012)
Unmarried, with child	$\exp\{\psi^{s1}\}$	0.43	(0.031)
Married, with child	$\exp\{\psi^{m1}\}$	0.43	(0.010)
Cost of being on AFDC	η	0.0018	(0.0002)
Match quality			
Variance at marriage	σ_0^2	0.097	(0.027)
Variance of innovations	σ_ξ^2	0.031	(0.009)
Probability of meeting partner by age	λ_0	0.426	(0.007)
	λ_1	-0.034	(0.002)
	λ_2	0.001	(0.0001)

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 $\{\psi^{s0}, \psi^{m0}, \psi^{s1}, \psi^{m1}\} = \{-1.1215, -0.5208, -0.8335, -0.8426\}$

At the heart of the model is a within-couple allocation process which is affected by outside options and preferences. Across our simulated individuals, the average Pareto weight

for women is about two thirds of the weight for men.[37] The average of a wife's share of consumption is highly correlated with her share of potential earnings, both at the time of marriage and over time (Appendix Figure E.2): a 1 percentage point increase in potential earnings leads to a 0.25 percentage points increase in the consumption share. However, the consumption share captures only one aspect of the utility because changes in potential earnings will also change labor force participation which directly changes utility for women. 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 they increase the value of marriage. Shocks to the match quality can also cause the participation constraints to bind. However, in practice the constraints may bind only rarely following shocks.

6   The Effect of Time Limits on Behavior

In this section, we use our model to evaluate the effects of introducing time limits, focusing first on the short-run responses of the transition into the reform, to mimic what happened when the welfare reforms were enacted in 1996. We can use these simulations to validate our model against the difference-in-difference estimates and to show to what extent the response that we observe in the data is due to forward-looking optimizing behavior.

6.1   Response to the Introduction of Time Limits

We simulate the introduction of time limits for women of 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 our data. We then use the simulated data to estimate short run effects equivalent to those estimated by the difference-in-difference specification in section 3.

Table 5 reports the estimated coefficients from the simulations and also the actual estimates in the CPS, focusing on the sample of women aged 21 to 53 in the first 6 years post

---

[37]This is in line with estimates from the literature on collective household models for the Unites States, the United Kingdom, and Japan, which tend to find larger sharing rules for men than for women (Lise and Seitz, 2011; Mazzocco, Yamaguchi and Ruiz, 2013; Voena, 2015; Lise and Yamada, 2014).

reform.[38] 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 partial offsetting increase in employment among unmarried women. The model also replicates the small decline in welfare utilization among married women. We also find that time limits lead to a reduction in the number of divorces, although the effects is small.[39]

Table 5: Difference-in-Differences Estimates: Data and Model

	data	95% c. i.	model
Welfare use, unmarried	-0.108	[-0.131, -0.085]	-0.090
Employed, unmarried	0.085	[0.049, 0.121]	0.081
Welfare use, married	-0.004	[-0.007, -0.001]	-0.009
Employed, married	-0.007	[-0.029, 0.015]	-0.009
Divorced	-0.013	[-0.025, -0.001]	-0.004

Notes: Estimates from the CPS data between 1990 and 2002 (first 6 years after the reform). Sample of women without college degrees, age 21 to 53. Controls in the data include age fixed effects, education, number of children, state-by-year fixed effects, year fixed effects, age of youngest child-by-year fixed effects, dummies for having 2 or more children post-1993 and post-1996.

To what extent do these responses reflect forward looking behavior and the banking of benefits? Over time, the difference-in-differences estimates compound anticipatory effects with the effect of differential welfare use by age of youngest child both in the data and in the model. 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 6 compares welfare utilization by the number of years since the reform took place, comparing the data with the full model and with the myopic model. Myopic individuals do not cut their use of welfare when the reform takes place. By contrast, in the baseline model and in the data, there is a decline directly after the reform.[40]

---

[38]We choose the CPS for validation to avoid the complexity of the post-reform sampling of the SIPP, which was no longer nationally representative starting with the 1996 panel.

[39]Married women whose husbands are working are unlikely to qualify for AFDC even without time limits because of their total household income. On the other hand, married women whose husbands do not work respond in a similar way to single women.

[40]For visual clarity, we do not include the 95% confidence interval of the data. Given the low stan-

Figure 6: Dynamic Response of Welfare Utilization to Time Limits for Mothers

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.

We also examine the effect of the introduction of time limits on intra-household allocation, as summarized by the women's Pareto weight θ_t^W. The effect of a 5-year time limit on married mothers is reported in Appendix table F.1, where we compare overall changes (column 1) with changes *within* a marriage spell (column 2) and changes for new marriages (column 3), i.e., we compare changes due to re-negotiation vs. changes in the initial Pareto weight. The change in women's weight is small, since for many women on the verge of marriage time limits will become binding only after having children. The estimates of the effect of a 5-year time limit in table F.1 correspond to a 0.1 percentage point reduction in women's expenditure share overall at the mean Pareto weight, a 0.06 percentage point reduction for existing marriages and a 0.3 percentage points reduction for new marriages.

---

dard errors, the baseline model predictions fall well within the confidence interval, while the myopic model predictions lie outside of it.

6.2 Lifetime Effects of Time Limits

In this section, we compare the simulated lifetime behavior of women who always live without time limits to the simulated lifetime behavior of the same women who always live under time limits.

6.2.1 Effects of Time Limits by Productivity

We compare the welfare utilization and employment behavior of women who do not face a time limit with those of women subject to a 5-year time limit, breaking the sample down into quintiles of productivity z_{it}^F for each age. Figure 7 shows the impact on welfare use (Panel a) and employment (Panel b) for all mothers, by quintile. The effects of time limits are concentrated in the bottom two quintiles. The bottom quintile experiences the greatest decline in welfare use (from 34% to 18%, Panel a) and the largest increase in employment (from 10% to 22%, Panel b). The second quintile experiences effects of the same sign, but smaller magnitudes.[41]

Figure 7: Welfare Use and Employment of Mothers by Productivity
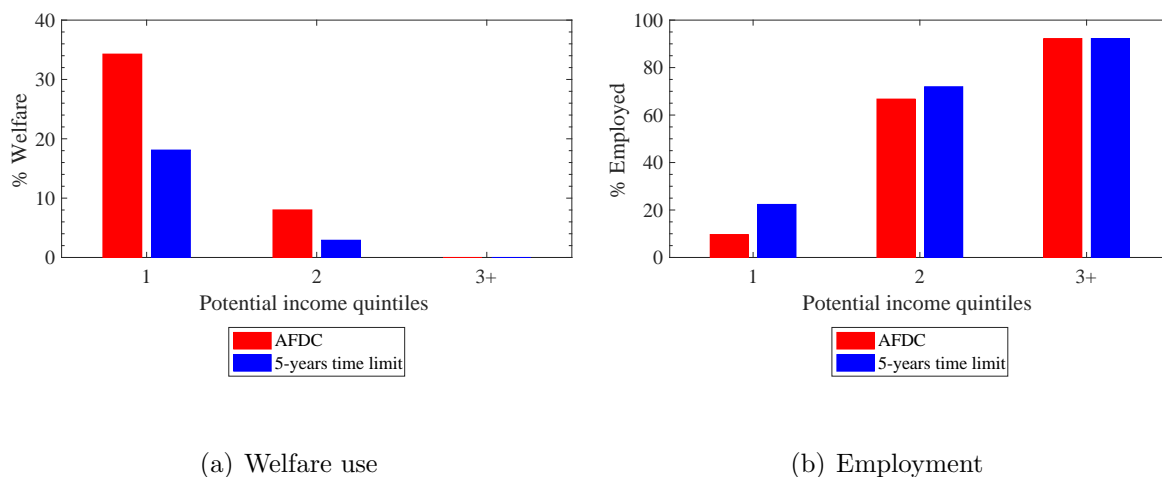
(a) Welfare use   (b) Employment

Notes: Simulated means of welfare use and employment by policy regime and by age-specific productivity z_{it}^F quintile. For outcome y_{it}^j and $j \in \{0,1\}$ where 0 is without any time limit and 1 is with a 5-year time limit imposed, the bars represent $P(y_{it}^1 = 1)$ and $P(y_{it}^0 = 1)$ of mothers, pooling across i and t.

---

[41]In addition to the impacts on welfare use and employment, we simulate how households expenditure differs across quintiles in the two regimes (figure F.2). Within quintile, the average level of household expenditure does not change with the introduction of time limits. However, this lack of impact on the average masks heterogeneity on winners and losers within a quintile: those who move off benefits and into employment (the majority) have higher income and spend more; those who move off benefits but are not employed (a small portion of the population), end up spending less.

In Figure 8, we break down the effects on welfare use and employment separately for unmarried and married mothers, holding compositional effects fixed. We do so by considering a woman's marital status under time limits and examine her behavior under both regimes. Without time limits, welfare use is highly concentrated among unmarried mothers, and especially among unmarried mothers in the bottom income quintile, where over 95% are using welfare in the absence of time limits. For these mothers, the introduction of time limits halves welfare utilization, and increases employment from close to zero to around 40%. For married mothers, welfare use even in the bottom quintile is much lower, with only 10% using welfare. The introduction of time limits for this group also halves their use of welfare but leads to only a negligible increase in employment.

Figure 8: Welfare Use and Employment of Married and Unmarried Mothers

(a) Welfare use, unmarried

(b) Welfare use, married

(c) Employment, unmarried

(d) Employment, married

Notes: Simulated means of lifetime welfare use and employment by policy regime, by age-specific productivity z_{it}^F quintile and by (counterfactual) marital status. For outcome y_{it}^r and $r \in \{0, 1\}$ where 0 is AFDC and 1 is a 5-year time limits, the bars represent $P(y_{it}^1 = 1 | married_{it}^1 = s)$ and $P(y_{it}^0 = 1 | married_{it}^1 = s)$ of mothers, pooling across i and t.

6.2.2 The Anticipatory Effect of Time Limits

The reduction in welfare use under time limits raises again the question of to what extent are the reductions happening because individuals anticipate reaching the time limit, and so preemptively reduce welfare use. To address this question, we show in Figure 9 the distribution of total years of lifetime welfare use. Panel (a) reports the distribution for all mothers, Panel (b) reports the distribution for women in the bottom productivity quintile

at age 19. Each bar reports the fraction of the relevant population that have used up that number of years of benefits over their lifetime. With time limits in place, the fraction who previously used more than 5 years of benefits would bunch at 5 years or reduce welfare use even further. There is evidence of this bunching both for the sample of all mothers and for the sample of the bottom quintile. There is also evidence of an increased fraction never using benefits and of increased fractions using only 2, 3 or 4 years of benefits. This is indicative of individuals anticipating the cut-off of benefits at 5 years and then not reaching the limit.

Figure 9: Distribution of lifetime welfare utilization

(a) All Mothers   (b) Bottom quintile of productivity at $t = 1$
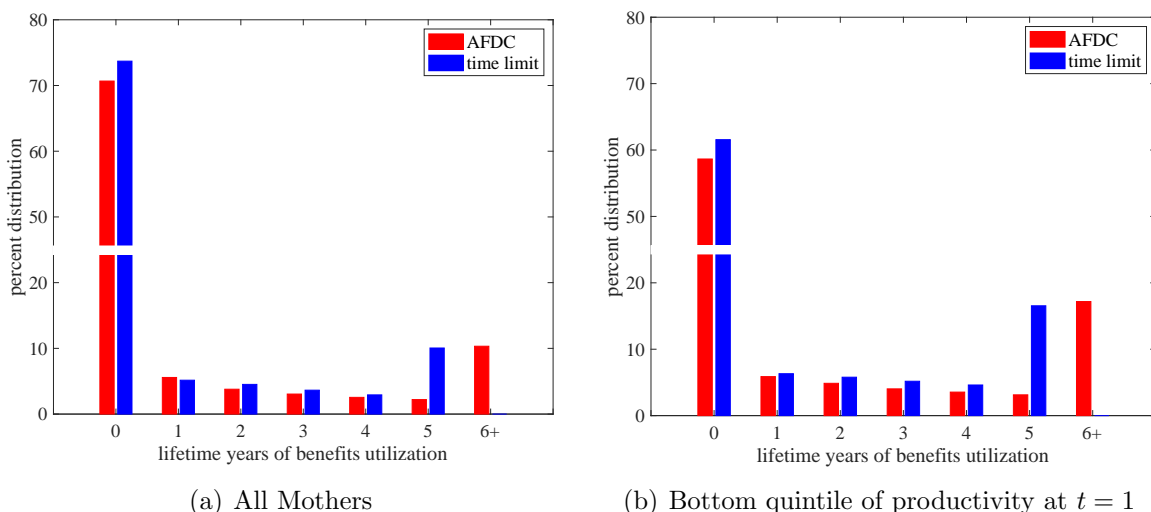
Notes: Distribution of lifetime welfare use in simulated data. Panel (a) plots the overall distribution, while Panel (b) limits the simulated sample to women belonging to the bottom quintile of the productivity distribution at age 19 z_{i1}^F.

Among those who reach the 5 year limit, we cannot tell whether the restriction was purely mechanical and so individuals acted as if there were no time limit until their benefits actually ran out, or whether individuals adjusted when they were claiming but still ended up claiming the maximum. To distinguish the former mechanical effect from the latter optimal-timing effect, we show in Figure 10 how quickly mothers use up 5 years of benefits when they are unrestricted and when they are restricted. The Figure shows that the presence of time limits leads individuals to delay claiming: under time limits, the fraction who reach 5 years of benefits at age 17 is the same as the fraction who reach 5 years at age 7 under AFDC. While this may be the optimal response of individuals to the presence of time limits, it does mean a substantial decline in resources from welfare for families with young children. This is consistent with the reduced form evidence in Figure C.1 in Appendix C showing that

Figure 10: Benefits use by age of youngest child

Notes: Percentage of women who have used five or more years of benefits, by age of the youngest child.

mothers of younger children cut welfare use more.

6.3 The Importance of Marriage and Divorce in Determining the Response to Welfare Reform

The reduced form results in Section 3 and the simulations in Section 6.2 highlight that single mothers are the most affected by welfare reform. Further, much of the literature estimates the effects of welfare reform focusing on single women (Chan, 2013). On the other hand, we showed in our reduced form evidence that the introduction of time limits reduced the divorce rate: the stock of divorcees decline by at least 1.3 percentage points. This motivated the inclusion of the marriage and divorce decision in our model: the possibility of divorce in the future and the possibility of marriage changes the behavior of single and married women. Further, this behavior will change in the face of time limits because the introduction of time limits changes outside options, and in particular, worsens the outside option associated with getting divorced.

6.3.1   The Marriage Option

To show how marriage and the option of marriage changes behavior, we solve and simulate the model when there is no possibility of marriage (that is, where $\lambda_t = 0 \;\; \forall\, t$). We then introduce time limits, and explore how welfare use and employment choices in the presence of time limits change when there is no option of marriage. We show in Figure 11 the impact on welfare use. We report the percentage of women who have used up to 5 years of benefits by age of child, grouped in 6 year age bands, similarly to Figure 10. The first two bars in each band correspond to the baseline AFDC, and the addition of time limits to the baseline; the second two bars report usage when there is no marriage. The level of take-up of benefits is substantially higher when there is no marriage market, irrespective of time limits. However, when time limits are introduced without the option of marriage, the use of benefits decreases by over 50%.

Figure 11: Benefits Use by Age of Youngest Child
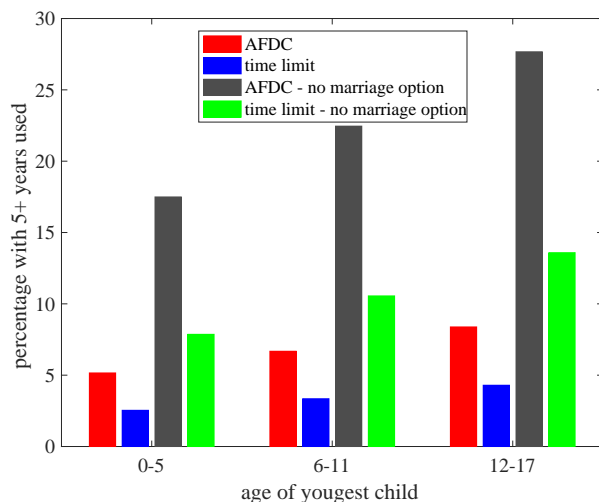
Notes: Percentage of mothers who, in model simulatotions, have used five or more years of benefits, by age of the youngest child. In the "no marriage option" counterfactual, we solve and simulate our model eliminating the possibility of marriage.

As discussed above, these changes in claiming may be mechanical, or there may be a reduction in anticipation of benefits running out. In Figure 12 we plot the change in welfare use and the change in employment with and without the option of marriage, and we further separate out the overall effect from the anticipatory effect. Panel (a) shows the effect on welfare use, separately by quintiles of potential income. For each quintile, we report four

Figure 12: The Effect of Time Limits on Unmarried Mothers With or Without a Marriage Option

(a) Effect of time limit on welfare

(b) Effect of time limit on employment

Notes: Differences in simulated means of lifetime welfare use and employment between policy regimes, by age-specific productivity z_{it}^F quintile. For outcome y_{it}^r and $r \in \{0,1\}$ where 0 is AFDC and 1 is a 5-year time limits, the bars represent $P(y_{it}^1 = 1|married_{it}^1 = 0) - P(y_{it}^0 = 1|married_{it}^1 = 0)$ of mothers, pooling across i and t. Anticipatory effects are computed by taking the sample of mothers i who, at time t, have not yet used 5 years of benefits under the time limit regime $r = 1$, and calculating the difference in average welfare use and employment for this group with and without time limits. In the "no marriage option" counterfactual, we solve and simulate our model eliminating the possibility of marriage, but only average the choices of women i who, at time t, would be unmarried if marriage were an option. See Appendix Figure F.4 for the underlying means of lifetime welfare use and employment and a detailed explanation of the construction of the simulated sample.

numbers: the overall change in welfare use when time limits are introduced (with and without the marriage option); and the change in welfare use in anticipation of a binding time limit (again with and without the marriage option). Panel (b) shows analogous effects on employment.

To calculate the effect of anticipating a binding time limit, we calculate first the fraction of those who choose to use welfare out of those that have used less than 5 years of benefits. If there were no anticipation effects of time limits binding, then this fraction would not change when a five year time limit was introduced. The larger the change in this fraction, the more the anticipation of the effect is binding.[42]

The key takeaway of Figure 12 is that the decline in welfare use and the rise in employment associated with time limits are significantly greater if there is no possibility of marriage. Further, a large part of the decline in welfare use that arises when there is no marriage option occurs because of the anticipatory effect: when there is no prospect of getting married in the future, individuals have to bank their benefits and return to work. This effect is strongest in the lowest quintile of the productivity distribution.[43]

6.3.2 The Divorce Option

Analogously to Figure 12, we show in Figure 13 the effects on choices about welfare use and employment by married women when the possibility of divorce is not present. As we discussed above, the effects of time limits on married mothers are substantially lower than the effects on unmarried mothers: in the baseline, there are effects only in the first quintile. Ruling out divorce also changes the extent that women change behavior in anticipation of hitting the time limit: there is less of a decline in welfare use in anticipation of the time limit being hit when divorce is not possible. There is still a decline that arises from the direct effect, but there is less banking of benefits. This is because women will not be holding back using benefits in case they divorce.[44]

This discussion highlights that estimates of the effects of welfare reform depend on whether the marriage and divorce choices are part of the model. Modeling the effect of policy changes

---

[42]This figure abstracts from composition effects because the differences are computed from the simulated sample of women who are unmarried in the baseline model, even when the marriage option is removed.

[43]The percentages of welfare use and employment by regime are reported in the Appendix in figure F.4.

[44]The percentages of welfare use and employment by regime are reported in the Appendix in figure F.5, alongside a note which explains the construction of the variables. Note that this figure, again, abstracts from composition effects because the differences are computed of the simulated sample of women who are unmarried in both the baseline model and the "no divorce" counterfactual, even when the marriage option is removed.

Figure 13: Welfare Use and Employment of Married Mothers With or Without a Divorce
Option

(a) Welfare use

(b) Employment

Notes: Differences in simulated means of lifetime welfare use and employment between policy regimes, by age-specific productivity z_{it}^F quintile. For outcome y_{it}^r and $r \in \{0,1\}$ where 0 is AFDC and 1 is a 5-year time limits, the bars represent $P(y_{it}^1 = 1 | married_{it}^1 = 1) - P(y_{it}^0 = 1 | married_{it}^1 = 1)$ of mothers, pooling across i and t. Anticipatory effects are computed by taking the sample of mothers i who, at time t, have not yet used 5 years of benefits under the time limit regime $r = 1$, and calculating the difference in average welfare use and employment for this group with and without time limits. In the "no divorce option" counterfactual, we solve and simulate our model eliminating the possibility of divorce, but only average the choices of women i who, at time t, would be married if marriage were an option. See Appendix Figure F.5 for the underlying means of lifetime welfare use and employment and a detailed explanation of the construction of the simulated sample. Note we use a different scale in this Figure compared to Figure 12.

ignoring these choices is likely to give misleading answers, particularly when expectations about the future are an important element of the response.

7 Implications of Welfare Reform for Lifetime Utility

Our results so far document a rich array of dynamic behavioral responses to the introduction of time limits. Time limits reduce the use of benefits and increase employment; and further, behavior changes in anticipation of time limits binding. On the other hand, we find that a substantial fraction of those coming off welfare do not move into employment and so their consumption declines. There is a fall in divorce rates because outside options get worse, and this can reduce lifetime utility for married women with a weakened bargaining position within marriage. In this section we use our model to calculate the net effect on lifetime utility of these offsetting consequences of the reform.

In computing the consequences for litetime utility, we introduce a negative payroll tax to hold constant government net spending, defined as total spending on welfare programs minus tax receipts from the payroll tax. Holding revenue constant is important because the fall in welfare use induced by the introduction of time limits leads to savings on government spending on transfers while the rise in employment leads to an increase in tax revenues. Hence, revenue neutrality means that the tax rate needs to be cut once the time limit is introduced. We implement this by considering two possible scenarios:

a) a payroll tax/subsidy, denoted as τ^F, falling only on women;

b) a payroll tax/subsidy, denoted as τ, falling on both women and men.

In both cases, the focus is on women without a college degree (and their partners). Hence, we transfer the revenue gains to them in terms of lower taxes and abstract from redistribution across groups or from further up the income distribution. We describe the procedure in full in Appendix G.

To define the welfare cost or benefit of introducing time limits, we begin by focusing on women and calculate the proportion of consumption that an individual is willing to pay *ex-ante*, denoted as π, to be indifferent between a new scenario (for example a 5 year time limit on welfare benefits) and the baseline (AFDC). We consider which demographic groups are affected the most, by conditioning the sample over which the the expectation is computed based on ex-post outcomes, such as marital status and fertility at age 25. Finally, we compute

a further counterfactual calculation, which includes the lifetime utility of men, that achieves revenue-neutrality by changing the tax rate paid by both men and women (τ).

7.1 The Lifetime Utility Cost of Time Limits on Women

Table 6 reports values of π with and without revenue neutrality. The table shows that women are willing to pay 0.54% 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 self-sufficiency through work, might have induced a net welfare loss by reducing insurance available to low-income women. In contrast, without the revenue-neutrality adjustment increasing income ($\Delta\tau^F = 0$), the willingness to pay to avoid time limits for these women is equal to 0.74% of lifetime consumption. In the final three columns, we consider how this willingness to pay changes when we condition on the *ex post* presence of children and on marital status at age 25. The willingness to pay increases to 2.2% for unmarried mothers even allowing for the tax rebate.

Table 6: Lifetime Utility Costs Of Time Limits: Women only

marriage option	revenue neutrality	$\Delta\tau^F$	Cons. Equiv. (π)	Cons. Equiv. unmarried mothers at 25	Cons. Equiv. married mothers at 25	Cons. Equiv. non-mothers at 25
yes	yes	-0.375%	0.54%	2.19%	1.56%	0.06%
yes	no	0	0.74%	2.33%	1.76%	0.28%
no	yes	-0.975%	0.80%	2.91%	-	-0.19%
no	no	0	1.58%	3.63%	-	0.61%
limited	yes	-0.5%	0.67%	2.68%	0.93%	0.09%

Notes: Simulated consumption equivalents of AFDC without time limits relative to AFDC with a 5-year time limit, under different assumptions for revenue neutrality achieved through a negative payroll tax on women's earnings. See Appendix G for a detailed description of how tax rates and consumption equivalents are computed. The "no marriage" case is computed by eliminating the marriage option. The "limited marriage" marriage refers to simulations with restricted marriage opportunities: we reduce the arrival rate of offers so that the model generates a 10 percentage points decline in marriage rate, similar to what was experienced by the 1980s birth cohort relative to the 1960s one. In this scenario, it is harder to marry and this means there is less divorce.

We consider further how the consequences of time limits for lifetime utility change when we restrict the option of marriage. When there is no option at all, and time limits are introduced, there is a substantial financial saving for the government because of the extent that benefits are scaled back as shown in Figure 11, and the tax rate is cut substantially ($\Delta \tau^F = -0.975\%$). Nonetheless, this rebate is not enough to offset the lost insurance and consumption willingness to pay falls by 0.8%. This result highlights the extent of interaction between different mechanisms of insurance.

In a final counterfactual, we limit, but do not eliminate, the marriage prospects. We do so by reducing the probability of marriage by age by a fixed factor that leads the marriage rate to drop by 10 percentage point, from approximately 70% to 60%, which is a reduction that mimics what more recent generations have been experiencing relative to our estimation cohort. We find that even such a relatively small reduction in marriage prospects leads women to experience a drop in lifetime well-being equal to 0.67% of lifetime consumption, relative to 0.54% in our baseline.

7.2   The Lifetime Utility Cost of Time Limits on Men and Women

Table 7 reports consumption willingness to pay when we also include the utility of men, so that tax changes affect both genders. When revenue-neutrality is imposed, men are willing to pay to move to a world with time limits for two reasons: the lower tax rate increases take-home pay and they realize modest increases in intra-household allocations despite the loss of benefits of their partner. The last column calculates an overall welfare measure by imposing the same consumption equivalent for men and women, and shows that overall welfare would increase if time limits were removed.

One important caveat of these welfare calculations is that in our model, parents do not internalize the effect of time limits on their children, a topic we leave for future work.

7.3   The Lifetime Utility Cost of Cutting the Generosity of AFDC

Table 6 and Table 7 show the extent that the introduction of time limits worsened lifetime utility. This loss is despite the revenue saving incurred by introducing time limits. We can use our model to consider whether alternative reforms to AFDC that saved the same amount of revenue could have incurred less of a lifetime utility loss. We focus on a revenue-neutral policy of reducing the generosity of benefits but without imposing any limit on the number

45

Table 7: Lifetime Utility Costs Of Time Limits: Men and Women

marriage option	revenue neutrality	$\Delta\tau$	Cons. Equiv. for women	Cons. Equiv. for men	Cons. Equiv. in expectation
yes	yes	-0.114%	0.68%	-0.07%	0.31%
yes	no	0	0.74%	$\sim 0\%$	0.35%
no	yes	-0.325%	1.32%	-0.30%	0.61%
no	no	0	1.58%	0%	0.88%
limited	yes	-0.178%	0.82%	-0.21%	0.33%

Notes: Simulated consumption equivalents of AFDC without time limits relative to AFDC with a 5-year time limit, under different assumptions for revenue neutrality achieved through a negative payroll tax on women's and men's earnings. See Appendix G for a detailed description of how tax rates and consumption equivalents are computed. The "no marriage" case is computed by eliminating the marriage option. The "limited marriage" marriage refers to simulations with restricted marriage opportunities: we reduce the arrival rate of offers so that the model generates a 10 percentage points decline in marriage rate, similar to what was experienced by the 1980s birth cohort relative to the 1960s one. In this scenario, it is harder to marry and this means there is less divorce.

of years that a person can receive benefits. Spending the same revenue as TANF on reducing generosity would be welfare increasing. Generosity falls to 73.5% of its baseline value, and yet women are better off because because they have access to the (reduced) benefits for as many years as needed. The lifetime utility cost of such a regime compared to time limits is equal to -0.11%, where the negative number indicates an increase.

8   The Role of EITC and Food Stamps

An important aspect of the model developed and estimated in Sections 4 and 5 is that it allows for the interaction between different aspects of government support. In particular, in addition to AFDC, we model the earned income tax credit (EITC) and food stamps, which provide different sorts of support to low income individuals. Allowing for these alternative programs means allowing for potentially important substitution between welfare programs, as shown by Keane and Moffitt (1998). In this section, we show the importance of allowing for these alternative programs when we consider the introduction of time limits to the AFDC programs. We focus this section on showing how the presence of EITC and Food Stamps change the anticipatory effects of time limits, i.e. behavior in anticipation of the time limit

being reached.

In Panels (a) and (b) of Figure 14 we plot welfare use (for AFDC and under time limits) and employment when we include and exclude an EITC option. This exercise is done only for those who have not hit the time limit constraint. In the first quintile of potential income, there is little difference in behavior in the absence of EITC, since EITC requires that people work, and in this group labor force attachment is very low. Greater differences are observed in the second quintile, where employment is lower in the absence of EITC, and welfare use is higher. The introduction of time limits has a bigger impact in the absence of EITC because that quintile is more likely to be using welfare. Finally, in the top 60% the effects are absent because EITC is phased out.

In Panels (c) and (d) of Figure 14 we plot welfare use and employment, respectively, when we include and exclude Food Stamps. The elimination of Food Stamps, a key element of the safety net in the US, leads to more employment and consequently reduces welfare use. The impact of time limits is similar for the first quintile whether or not Food Stamps are present. However, as with EITC, the effects on the second quintile are more marked: without Food Stamps, there is very little use of AFDC in the second quintile and so time limits make no difference.

These counterfactual exercises emphasize the role that other welfare programs play in complementing or substituting for the benefits lost due to time limits. Without EITC, welfare use is affected much more by time limits; whereas without food stamps, welfare use is affected much less by time limits.

9   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 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

Figure 14: Welfare Use and Employment With and Without EITC or Food Stamps


(a) Welfare use, with and without EITC

(b) Employment, with and without EITC

(c) Welfare use, with and without FS

(d) Employment, with and without FS

Notes: The figure shows simulated fractions of welfare use and employment for women who have not used 5 years of benefits under time limits ($E_{it}^1 = 1$ represents having used fewer than five years of welfare under time limits.). For outcome y^j and $j \in \{0, 1\}$ where 0 is without any time limit and 1 is with a 5-year time limit imposed, the bars represent $P(y_{it}^0 = 1 | E_{it}^1 = 1)$ and $P(y_{it}^1 = 1 | E_{it}^1 = 1)$ of mothers by age-specific productivity z_{it}^F quintile, pooling across i and t. Bars of different colors represent different regimes: (i) AFDC without time limits, (ii) AFDC with a 5-year time limit (iii) AFDC without time limits and without EITC (top) or food stamps (bottom), (iv) AFDC with a 5-year time limit and without EITC (top) or food stamps (bottom).

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. In detail, we document that welfare use falls substantially for single women: pre-reform, 30% 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 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.

These vast effects have important welfare consequences. We use our model to consider welfare implications, alternative benefit reforms, and interactions with other safety net programs. Welfare analyses reveal that the groups mostly targeted by the reform (i.e., low educated women) suffered a net welfare loss, despite the increase in self-sufficiency brought about by increasing employment. Time limits induce larger welfare losses than alternative reforms (such as cutting benefits) that produce the same government savings.[45] Finally, because of substitution and income effects, changes in one program induce spillovers and reallocation into and out of other anti-poverty programs, as we show when considering the impact of the reform on EITC and food stamps. This is important for a more comprehensive evaluation of the reform.

While our model simulations suggest that the reform caused a utility loss for women and only a negligible gain for men, it is important to note several caveats and avenues for future research. First, we ignore welfare issues associated with the change in time and resources allocated to children. Second, we have assumed that the costs of children are borne entirely by women after marriage ends, and so benefit eligibility affects women more than men. Third, ours is not an equilibrium model of the labor market or of the marriage market, and equilibrium effects may dampen - or exacerbate - some of the effects we document.

---

[45] Although cost savings may not have been the main rationale for the reform.

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.

Autor, David, Andreas Kostøl, Magne Mogstad, and Bradley Setzler. 2019. "Disability Benefits, Consumption Insurance, and Household Labor Supply." *American Economic Review*, 109(7): 2613–54.

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.

Fang, Hanming, and Michael P Keane. 2004. "Assessing the impact of welfare reform on single mothers." *Brookings Papers on Economic Activity*, 2004(1): 1–116.

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. 2003. "The effects of time limits, the EITC, and other policy changes on welfare use, work, and income among female-headed families." *Review of Economics and statistics*, 85(2): 394–408.

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.

Judd, Kenneth L. 1998. *Numerical methods in economics.* MIT press.

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. 2002*a*. "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. 2002*b*. "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. 1983. "An Economic Model of Welfare Stigma." *The American Economic Review*, 73(5): 1023–1035.

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.

Stevenson, Betsey, and Justin Wolfers. 2007. "Marriage and divorce: Changes and their driving forces." *Journal of Economic perspectives*, 21(2): 27–52.

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

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://wrd.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 reports summary statistics for the two data sets. In the first two columns, we report statistics for our regression sample in the SIPP and the CPS. In the following four columns, we break these samples 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 estimation uses only pre-reform data (with the added restriction of focusing only on the 1960's birth cohorts).[46]

---

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

Table B.1: Summary statistics

	Regression Sample		Pre-reform		Post-reform	
	SIPP	CPS	SIPP	CPS	SIPP	CPS
On Welfare	0.086	0.063	0.101	0.077	0.040	0.040
On Welfare (married)	0.029	0.015	0.034	0.019	0.014	0.010
On Welfare (unmarr.)	0.253	0.240	0.306	0.304	0.108	0.148
Employed	0.650	0.666	0.642	0.647	0.674	0.694
Employed (married)	0.650	0.666	0.647	0.654	0.659	0.685
Employed (unmarr.)	0.648	0.664	0.624	0.620	0.712	0.726
Divorced or separated	0.152	0.125	0.151	0.126	0.154	0.124
Div/sep if $m_{t-1} = 1$	0.009	0.013	0.009	0.014	0.009	0.013
Married	0.745	0.789	0.753	0.796	0.720	0.780
Married if $m_{t-1} = 0$	0.025	0.047	0.025	0.047	0.025	0.047
Less than high school	0.172	0.245	0.171	0.311	0.174	0.142
High school	0.491	0.377	0.493	0.337	0.484	0.438
Some college	0.337	0.378	0.336	0.352	0.343	0.420
White	0.805	0.833	0.810	0.835	0.790	0.831
Age	36.035	36.256	35.837	35.921	36.653	36.780
Number of children	1.991	2.113	1.992	2.111	1.988	2.118
Age of youngest	7.248	7.562	7.163	7.436	7.515	7.759
$Exposed_{dst}Post_{st}$	0.181	0.287	0.000	0.000	0.745	0.736
N. of Obs.	254,627	112,128	171,062	68,353	83,565	43,775

Notes: Data from the 1990-2002 SIPP panels and 1990-2002 March CPS. 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.

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

	SIPP		
Dependent Var:	**Food Stamps Utilization**		
Sample:	Whole	Married	Unmarried
$Exposed_{dst}Post_{st}$	-0.011	-0.007	-0.041***
	(0.008)	(0.006)	(0.014)
Mean pre-reform	0.156	0.077	0.401
Obs	336,129	242,825	93,304
R^2	0.14	0.12	0.23

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. Data from the 1990-2004 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. Sampling weights used in the SIPP regressions. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-by-year fixed effects, state-by-demographics fixed effects, unemployment-rate- by-demographics fixed effects, and a dummy for two or more children interacted with dummies for post-1993 and post-1996. Standard errors in parentheses, clustered at the state level.

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

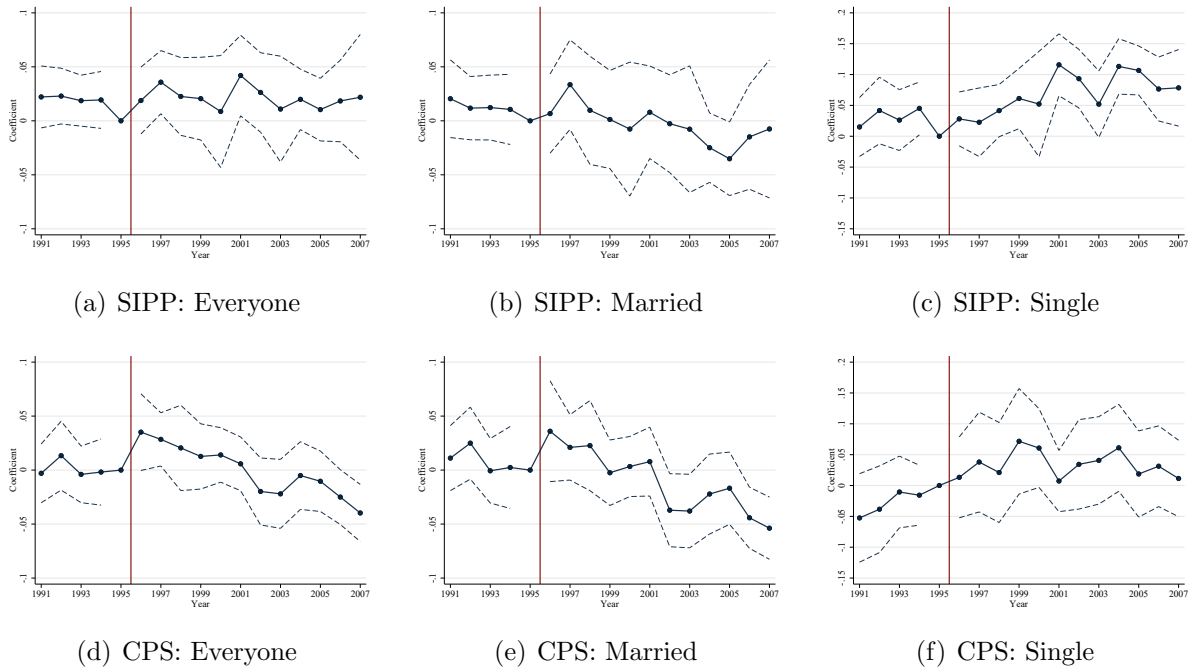|	SIPP			
Dependent Var:	Employed On Welfare	Employed Not on Welfare	Not Employed On Welfare	Not Employed Not on Welfare
$Exposed_{dst}Post_{st}$	-0.029***	0.084***	-0.080***	0.024*
	(0.005)	(0.015)	(0.011)	(0.014)
Mean pre-reform	0.062	0.570	0.231	0.137
Obs	93,304	93,304	93,304	93,304
R^2	0.07	0.21	0.24	0.10
	CPS			
Dependent Var:	Employed On Welfare	Employed Not on Welfare	Not Employed On Welfare	Not Employed Not on Welfare
$Exposed_{dst}Post_{st}$	-0.011	0.064***	-0.100***	0.047**
	(0.008)	(0.019)	(0.010)	(0.021)
Mean pre-reform	0.066	0.557	0.232	0.145
Obs	33,593	33,593	33,593	33,593
R^2	0.04	0.13	0.15	0.06

Notes: Standard errors in parentheses clustered at the state level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$. Data from the 1990-2004 SIPP panels and 1990-2007 March CPS. Sample of female heads of household who are not college graduates and have children aged 18 and below. Sampling weights used in the SIPP regressions. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-by-year fixed effects, state-by-demographics fixed effects, unemployment-rate-by-demographics fixed effects, and a dummy for two or more children interacted with dummies for post-1993 and post-1996.

Figure C.1: Program Participation and Employment Dynamics by Child Age
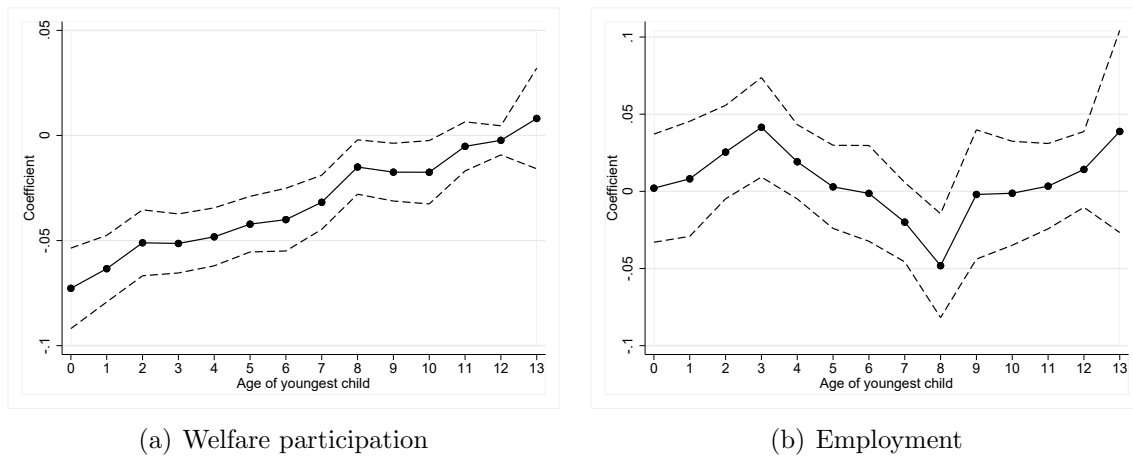
(a) Welfare participation

(b) Employment

Notes: Data from the 1990-2004 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. Sampling weights used in the SIPP regressions. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-by-year fixed effects, state-by-demographics fixed effects, unemployment-rate-by-demographics fixed effects, and a dummy for two or more children interacted with dummies for post-1993 and post-1996. Child age is defined as the age of the youngest child.

Figure C.2: Employment Dynamics

(a) SIPP: Everyone (b) SIPP: Married (c) SIPP: Single

(d) CPS: Everyone (e) CPS: Married (f) CPS: Single

Notes: Data from the 1990-2004 SIPP panels and 1990-2007 March CPS. Sample of female heads of household who are not college graduates and have children aged 18 and below. Sampling weights used in the SIPP regressions. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-by-year fixed effects, state-by-demographics fixed effects, unemployment-rate- by-demographics fixed effects, and a dummy for two or more children interacted with dummies for post-1993 and post-1996. Standard errors in parentheses, clustered at the state level.

Figure C.3: Marital Status Dynamics

(a) SIPP: Divorce/Separation

(b) SIPP: Marriage

(c) CPS: Divorce/Separation

(d) CPS: Marriage

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. Data from the 1990-2004 SIPP panels and 1990-2007 March CPS. Sample of female heads of household who are not college graduates and have children aged 18 and below. Sampling weights used in the SIPP regressions. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-by-year fixed effects, state-by-demographics fixed effects, unemployment-rate- by-demographics fixed effects, and a dummy for two or more children interacted with dummies for post-1993 and post-1996.

Table C.3: Marital Status Transitions

Dependent Var:	Gets married		Gets married	
Sample:	$m_{t-1}=0,\ d_{t-1}=0$		$m_{t-1}=1,\ d_{t-1}=1$	
	SIPP	CPS	SIPP	CPS
$Exposed_{dst}Post_{st}$	0.001	-0.000	0.001	-0.020
	(0.005)	(0.001)	(0.004)	(0.019)
Mean pre-reform	0.017	0.003	0.030	0.052
Obs	33,102	58,540	44,387	8,145
R^2	0.10	0.02	0.06	0.13

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. Data from the 1990-2004 SIPP panels and 1990-2007 March CPS. Sample of female heads of household who are not college graduates and have children aged 18 and below. Sampling weights used in the SIPP regressions. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-by-year fixed effects, state-by-demographics fixed effects, unemployment-rate- by-demographics fixed effects, and a dummy for two or more children interacted with dummies for post-1993 and post-1996.

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 Wave of Each SIPP panel

Dependent Var:	AFDC/TANF		Employed		Div/Sep	Married
Sample:	Whole	Unmarried	Whole	Unmarried	Whole	Whole
$Exposed_{dst}Post_{st}$	-0.049***	-0.153***	0.033**	0.114***	-0.031**	-0.027**
	(0.006)	(0.016)	(0.013)	(0.036)	(0.012)	(0.013)
Mean pre-reform	0.097	0.289	0.637	0.629	0.154	0.746
Obs	41,262	11,605	41,262	11,605	41,262	41,262
R^2	0.13	0.32	0.14	0.26	0.05	0.07

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. Data from the 1990-2004 SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. Sampling weights used in the SIPP regressions. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-by-year fixed effects, state-by-demographics fixed effects, unemployment-rate-by-demographics fixed effects, and a dummy for two or more children interacted with dummies for post-1993 and post-1996.

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 sameple of college graduates. We find very small effects on welfare utilization (-0.8pp 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).

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

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

Dependent Var:	SIPP					
	AFDC/TANF		Employed		Div/Sep	Married
Sample:	Whole	Unmarried	Whole	Unmarried	Whole	Whole
$Exposed_{dst}Post_{st}$	-0.008***	-0.067***	-0.010	0.005	-0.002	-0.013
	(0.002)	(0.015)	(0.012)	(0.020)	(0.014)	(0.015)
Mean pre-reform	0.010	0.050	0.772	0.895	0.094	0.875
Obs	141,336	19,348	141,336	19,348	141,336	141,336
R^2	0.03	0.23	0.08	0.19	0.08	0.09

Dependent Var:	CPS					
	AFDC/TANF		Employed		Div/Sep	Married
Sample:	Whole	Unmarried	Whole	Unmarried	Whole	Whole
$Exposed_{dst}Post_{st}$	-0.003**	-0.032***	-0.013	-0.004	0.001	-0.006
	(0.001)	(0.011)	(0.010)	(0.027)	(0.011)	(0.013)
Mean pre-reform	0.005	0.035	0.775	0.905	0.078	0.899
Obs	55,591	5,681	55,591	5,681	55,591	55,591
R^2	0.02	0.18	0.04	0.19	0.02	0.03

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. Data from the 1990-2004 SIPP panels and 1990-2007 March CPS. Sample of female heads of household who are not college graduates and have children aged 18 and below. Sampling weights used in the SIPP regressions. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-by-year fixed effects, state-by-demographics fixed effects, unemployment-rate-by-demographics fixed effects, and a dummy for two or more children interacted with dummies for post-1993 and post-1996.

in the PRWORA.[47] Appendix table C.6 shows that the results are robust to excluding 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.

---

[47]The 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 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.

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
			SIPP			
Sample:	Whole	Unmarried	Whole	Unmarried	Whole	Whole
$Exposed_{dst}Post_{st}$	-0.026***	-0.064***	-0.000	0.036**	-0.032***	0.008
	(0.004)	(0.015)	(0.012)	(0.015)	(0.009)	(0.011)
Mean pre-reform	0.065	0.189	0.713	0.716	0.180	0.746
Obs	189,950	54,296	189,950	54,296	189,950	189,950
R^2	0.08	0.20	0.09	0.16	0.03	0.07
			CPS			
Dependent Var:	AFDC/TANF		Employed		Div/Sep	Married
Sample:	Whole	Unmarried	Whole	Unmarried	Whole	Whole
$Exposed_{dst}Post_{st}$	-0.011***	-0.055***	-0.012	0.034	-0.003	-0.018**
	(0.002)	(0.009)	(0.010)	(0.023)	(0.006)	(0.007)
Mean pre-reform	0.055	0.202	0.715	0.706	0.157	0.774
Obs	91,826	21,643	91,826	21,643	91,826	91,826
R^2	0.05	0.14	0.06	0.13	0.02	0.05

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. Data from the 1990-2004 SIPP panels and 1990-2007 March CPS. Sample of female heads of household who are not college graduates and have children aged 18 and below. Sampling weights used in the SIPP regressions. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-by-year fixed effects, state-by-demographics fixed effects, unemployment-rate-by-demographics fixed effects, and a dummy for two or more children interacted with dummies for post-1993 and post-1996.

Fertility

Table C.7: Fertility

Dependent Var:	**Newborn in $t+1$**					
Sample:	Whole		Married		Unmarried	
	SIPP	CPS	SIPP	CPS	SIPP	CPS
$Exposed_{dst}Post_{st}$	-0.002	-0.001	-0.003	-0.000	0.004	0.006
	(0.002)	(0.003)	(0.002)	(0.003)	(0.003)	(0.007)
Mean pre-reform	0.072	0.049	0.075	0.051	0.065	0.039
Obs	198,657	66,685	145,256	52,528	53,401	14,157
R^2	0.42	0.04	0.42	0.06	0.43	0.09

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. Data from the 1990-2004 SIPP panels and 1990-2007 March CPS. Sample of female heads of household who are not college graduates and have children aged 18 and below. Sampling weights used in the SIPP regressions. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-by-year fixed effects, state-by-demographics fixed effects, unemployment-rate- by-demographics fixed effects, and a dummy for two or more children interacted with dummies for post-1993 and post-1996.

Appendix D: Life-Cycle Model

Determination of the Pareto Weight at the Time of Marriage

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:

$$\max_{\boldsymbol{\theta}_{ij,t_0}} \left(V_{t_0}^{Fm}\left(\boldsymbol{\theta}_{ij,t_0}\right) - V_{t_0}^{Fs} \right) \cdot \left(V_{t_0}^{Mm}\left(\boldsymbol{\theta}_{ij,t_0}\right) - V_{t_0}^{Ms} \right).$$

The solution to this maximization 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) $\frac{\partial V_t^{Fm}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = -\frac{\partial V_t^{Mm}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}}$. 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's utility will decline and so her consumption share will decline.

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

Earnings Process for Men

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

$$E[\Delta u_t^2] = \sigma^2_{\zeta M} + 2\sigma^2_{\varepsilon M} \tag{11}$$

$$E[\Delta u_t \Delta u_{t-1}] = -\sigma^2_{\varepsilon M} \tag{12}$$

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

$$\log w_{it} = \mathbf{X}_{it}\boldsymbol{\beta} + u_{it}. \tag{13}$$

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

Table E.1: Employment status Probit regressions - Women

VARIABLES	(1) coeff.	(2) marg. eff.
Average AFDC payment ($100)	-0.064***	-0.021***
	(0.007)	(0.003)
Average food stamps payment ($100)	-0.002	-0.008
	(0.095)	(0.031)
Average EITC payment ($100)	0.183***	0.060***
	(0.054)	(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.

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 $(\sigma^2_{\zeta F})$ is obtained by solving the following moment conditions:

$$E[\Delta u_t \mid P_t = 1, P_{t-1} = 1] = \sigma_{\nu \zeta F} \left[\frac{\phi(\alpha_t)}{1 - \Phi(\alpha_t)} \right]$$

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

$$E[\Delta u_t \Delta u_{t-1} \mid P_t = 1, P_{t-1} = 1, P_{t-2} = 1] = -\sigma^2_{\varepsilon F}$$

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.

Table E.2: Target moments

No.	Moment Description	Data Mean %	(s.e. in %)	Model Mean %
1	employed (married without children)	92.62	0.0073	92.71
2	employed (unmarried without children)	89.49	0.0068	89.28
3	employed (married with children)	61.10	0.0106	60.68
4	employed (unmarried with children)	66.03	0.006	66.28
5	on AFDC (unmarried with children)	39.10	0.0092	38.12
6	on AFDC (married with children)	3.54	0.0024	3.36
7	married at age 19	15.79	0.1169	19.61
8	married at age 20	30.85	0.0255	32.75
9	married at age 21	40.03	0.0227	42.02
10	married at age 22	46.62	0.0175	48.25
11	married at age 23	52.39	0.0129	53.43
12	married at age 24	57.18	0.0105	57.57
13	married at age 25	59.57	0.0101	60.63
14	married at age 26	62.61	0.0084	63.18
15	married at age 27	65.62	0.0074	65.09
16	married at age 28	66.59	0.0076	66.34
17	married at age 29	67.50	0.006	67.83
18	married at age 30	68.35	0.0067	68.83
19	married at age 31	69.52	0.0077	69.66
20	married at age 32	71.16	0.0082	69.98
21	married at age 33	73.41	0.009	70.52
22	divorced at age 21	5.30	0.0105	4.18
23	divorced at age 22	7.10	0.0082	6.48
24	divorced at age 23	7.78	0.0067	8.44
25	divorced at age 24	10.13	0.0069	9.76
26	divorced at age 25	10.68	0.0064	11.13
27	divorced at age 26	11.96	0.0064	12.29
28	divorced at age 27	12.74	0.0067	13.13
29	divorced at age 28	14.49	0.0068	14.30
30	divorced at age 29	15.68	0.0065	15.03
31	divorced at age 30	15.15	0.0058	15.80
32	divorced at age 31	15.99	0.008	16.62
33	divorced at age 32	16.30	0.0099	17.51
34	divorced at age 33	15.42	0.0122	17.98

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.

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

Figure E.1 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 4), 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.

Figure E.1: Life-cycle profiles of log-wages for women in the data and in the model


(a) All women   (b) Unmarried women
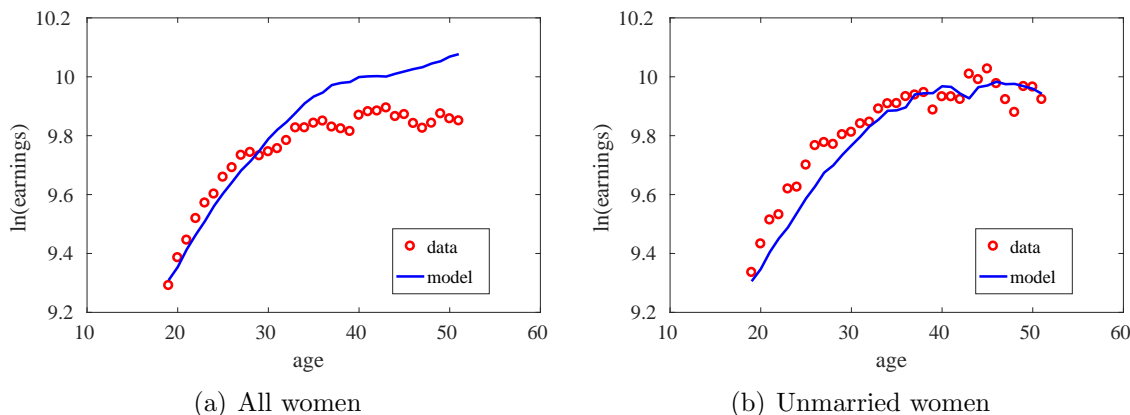
Notes: Data from 1990-2008 panels of SIPP. Data from the 1960-69 birth cohorts, pre-refom.

Counterfactual distributions of singles' characteristics in the marriage market

After time limits are introduced in a counterfactual exercise, unmarried 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.

Intrahousehold allocations

18

Figure E.2: Consumption Allocation in the Household

(a) Offer wages at marriage

(b) Offer wages
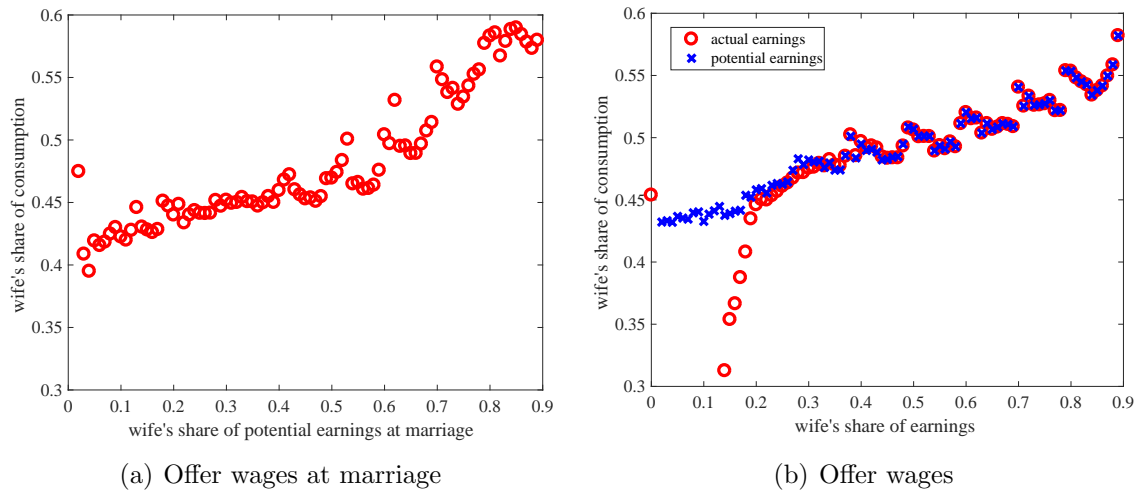
Notes: Simulations of consumption shares $\frac{c_{ij,t}^F}{x_{ij,t}}$ from the estimated model.

Appendix F: The Effect of Time Limits on Behavior

Figure F.1: Distribution of lifetime welfare utilization

Notes: Model simulations of the distribution of lifetime welfare use in simulated data with and without a marriage option.

Table F.1: Simulated Data: Effect of Time Limits on Women's Pareto Weight

	(1)	(2)	(3)
	All marriages	Married pre-reform	At the time of marriage
5-year time limits	-0.0013	-0.0009	-0.0046
Age fixed effects	Y	Y	Y
Marriage duration fixed effects	Y	Y	N
Marriage spell fixed effects	N	Y	N

Notes: Regression estimates of model simulations where a 5-year time limit is introduced in an unanticipated manner. *5-year time limits* is a dummy that takes value 1 once the household is exposed to time limits, and 0 otherwise. Column 1 includes all marriages. Column 2 includes marriages that formed before the reform. In both cases, the regression features one observation per year of marriage. Column 3 only includes an observation for the first year of marriage.

Figure F.2: Consumption of Mothers by Productivity

Notes: Simulated fractions of welfare use and employment. For outcome y^j and $j \in \{0,1\}$ where 0 is without any time limit and 1 is with a 5-year time limit imposed, the bars represent $P(y_{it}^1 = 1)$ and $P(y_{it}^0 = 1)$ of mothers by age-specific productivity z_{it}^F quintile, pooling across i and t.

Figure F.3: Anticipatory Effects: Welfare Use and Employment

(a) Welfare use, < 5 years of past welfare use   (b) Employment, < 5 years of past welfare use

Notes: Simulated means of lifetime welfare use and employment. For outcome y_{it}^r and $r \in \{0,1\}$ where 0 is AFDC and 1 is a 5-year time limits, the bars represent $P(y_{it}^1 = 1 | E_{it}^1 = 1)$ and $P(y_{it}^0 = 1 | E_{it}^1 = 1)$ of mothers by age-specific productivity z_{it}^F quintile. Here, $E_{it}^1 = 1$ represents having used fewer than five years of welfare under time limits.

Figure F.4: Welfare Use and Employment of Unmarried Mothers with and without a Marriage option

(a) Welfare use

(b) Employment

(c) Welfare use, < 5 years of past welfare use

(d) Employment, < 5 years of past welfare use

Notes: Simulated means of lifetime welfare use and employment marital status by policy regime and marriage option for unmarried women. Subfigures (a) and (b): In the baseline model B (red and blue bars), outcome $y_{it}^{r,B}$ and $r \in \{0,1\}$ where 0 is AFDC and 1 is a 5-year time limits, the bars represent $P(y_{it}^{1,B} = 1|married_{it}^{1,B} = 0)$ and $P(y_{it}^{0,B} = 1|married_{it}^{1,B} = 0)$ of mothers by age-specific productivity z_{it}^F quintile. In the counterfactual model C (grey and green bars), there exists no marriage option. Outcome $y_{it}^{r,C}$ and $r \in \{0,1\}$ where 0 is AFDC and 1 is a 5-year time limits, the bars represent $P(y_{it}^{1,C} = 1|married_{it}^{1,B} = 0)$ and $P(y_{it}^{0,C} = 1|married_{it}^{1,B} = 0)$ of mothers by age-specific productivity z_{it}^F quintile. Subfigures (c) and (d) focus on women who have not hit their time limit. In the baseline model B (red and blue bars), outcome $y_{it}^{r,B}$ and $r \in \{0,1\}$ where 0 is AFDC and 1 is a 5-year time limits, the bars represent $P(y_{it}^{1,B} = 1|married_{it}^{1,B} = 0, E_{it}^{1,B})$ and $P(y_{it}^{0,B} = 1|married_{it}^{1,B} = 0, E_{it}^{1,B})$ of mothers by age-specific productivity z_{it}^F quintile. In the counterfactual model C (grey and green bars), there exists no marriage option. Outcome $y_{it}^{r,C}$ and $r \in \{0,1\}$ where 0 is AFDC and 1 is a 5-year time limits, the bars represent $P(y_{it}^{1,C} = 1|married_{it}^{1,B} = 0, E_{it}^{1,C} = 1)$ and $P(y_{it}^{0,C} = 1|married_{it}^{1,B} = 0, E_{it}^{1,C} = 1)$ of mothers by age-specific productivity z_{it}^F. Here, $E_{it}^{1,B} = 1$ represents having used fewer than five years of welfare under time limits in the baseline model and $E_{it}^{1,C} = 1$ r in the counterfactual model.

Figure F.5: Welfare Use and Employment of Married Mothers with and without a Divorce option

(a) Welfare use

(b) Employment

(c) Welfare use, < 5 years of past welfare use

(d) Employment, < 5 years of past welfare use

Notes: Simulated means of lifetime welfare use and employment marital status by policy regime and marriage option for married women. Subfigures (a) and (b): In the baseline model B (red and blue bars), outcome $y_{it}^{r,B}$ and $r \in \{0,1\}$ where 0 is AFDC and 1 is a 5-year time limits, the bars represent $P(y_{it}^{1,B} = 1|married_{it}^{1,B} = 0)$ and $P(y_{it}^{0,B} = 1|married_{it}^{1,B} = 1)$ of mothers by age-specific productivity z_{it}^F quintile. In the counterfactual model C (grey and green bars), there exists no divorce option. Outcome $y_{it}^{r,C}$ and $r \in \{0,1\}$ where 0 is AFDC and 1 is a 5-year time limits, the bars represent $P(y_{it}^{1,C} = 1|married_{it}^{1,B} = 1)$ and $P(y_{it}^{0,C} = 1|married_{it}^{1,B} = 1)$ of mothers by age-specific productivity z_{it}^F quintile. Subfigures (c) and (d) focus on women who have not hit their time limit. In the baseline model B (red and blue bars), outcome $y_{it}^{r,B}$ and $r \in \{0,1\}$ where 0 is AFDC and 1 is a 5-year time limits, the bars represent $P(y_{it}^{1,B} = 1|married_{it}^{1,B} = 1, E_{it}^{1,B})$ and $P(y_{it}^{0,B} = 1|married_{it}^{1,B} = 1, E_{it}^{1,B})$ of mothers by age-specific productivity z_{it}^F quintile. In the counterfactual model C (grey and green bars), there exists no marriage option. Outcome $y_{it}^{r,C}$ and $r \in \{0,1\}$ where 0 is AFDC and 1 is a 5-year time limits, the bars represent $P(y_{it}^{1,C} = 1|married_{it}^{1,B} = 1, E_{it}^{1,C} = 1)$ and $P(y_{it}^{0,C} = 1|married_{it}^{1,B} = 1, E_{it}^{1,C} = 1)$ of mothers by age-specific productivity z_{it}^F. Here, $E_{it}^{1,B} = 1$ represents having used fewer than five years of welfare under time limits in the baseline model and $E_{it}^{1,C} = 1$ r in the counterfactual model.

Appendix G: Implications of Welfare Reform for Lifetime Utility

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^F, such that the present discounted value of net revenue flows remains constant:

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

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.[48]

In practice, we first calculate the left hand side in the baseline. This gives the size of $\bar{D}$ per woman/year when $\tau^F = 0$. 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 ($\bar{D}'$ per woman/year). Third, we iterate on τ^F so that the deficit under the new policy is equal to $\bar{D}$. As τ^F 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 $\tau^F = -0.375\%$ for the case of a 5 year time limit.

---

[48]Note 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.

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

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

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 $s = AFDC$.[49]

$$E_0 U\left(5TL,\tau_w\right)|_\pi = \sum_{i=1}^{N}\sum_{t=0}^{T} \beta^t \left(\frac{\left((1-\pi)\, c^s \cdot e^{\psi(M,k^a)\cdot P^s}\right)^{1-\gamma}}{1-\gamma} - \eta B_t^s + L^t m_t^s \right) \qquad (14)$$

We solve for π such that

$$E_0\, U\left(5TL,\tau_w\right)|_\pi = E_0 U\left(b\right), \qquad (15)$$

where π can be interpreted as the consumption cost of going from AFDC to a 5-year limit.[50]

We also consider which demographic groups are affected the most, by conditioning the sample over which the the expectation is computed based on ex-post outcomes, such as marital status and fertility at age 25. Finally, we also compute a further counterfactual utility calculation, which includes the lifetime utility of men as

$$E_0 \left[U^M\left(\varsigma,\tau\right) \right]|_\pi = \frac{1}{N} \sum_{i=1}^{N}\sum_{t=0}^{T-R} \beta^t \left(\frac{\left((1-\pi^\varsigma)\, c_{i,t}^{M\varsigma} \cdot e^{\psi^{00}\cdot P_{i,t}^{M\varsigma}}\right)^{1-\gamma}}{1-\gamma} + L_{i,t} m_{i,t}^{M\varsigma} \right)$$

---

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

[50]In varying π we do not reoptimize.

where $\{c_t^{M\varsigma}, m_t^{M\varsigma}\}$ refer to the implied consumption (with $P_{i,t}^{M\varsigma}$ being, in the men's case, an exogenous employment process) and marital status in economy ς, and τ is the revenue neutral tax rate paid by both men and women. E_0 represents the expectation at the beginning of working life, before initial conditions are known. We solve for π such that

$$E_0\, U^M\, (5TL, \tau)\big|_\pi + E_0\, U^F\, (5TL, \tau)\big|_\pi = E_0 U^M\, (AFDC, 0) + E_0\, U^F\, (AFDC, 0)\big|_\pi. \quad (16)$$

27

Table G.1: Welfare Use Under Revenue-Neutral Taxes, With and Without Marriage

Marriage option	Time limits	τ	Benefits use	Benefits use by mothers	Benefits by married mothers	Benefits use by single mothers
yes	no	0%	4.96%	11.47%	2.71%	34.84%
no	no	0%	9.94%	34.53%		34.53%
yes	yes	0%	2.49%	5.75%	1.29%	17.65%
no	yes	0%	3.37%	11.71%		11.71%
yes	yes	-0.375%	2.48%	5.72%	1.29%	17.55%
no	yes	-0.975%	3.33%	11.55%		11.55%

Figure G.1: Welfare Use and Employment with reduced generosity
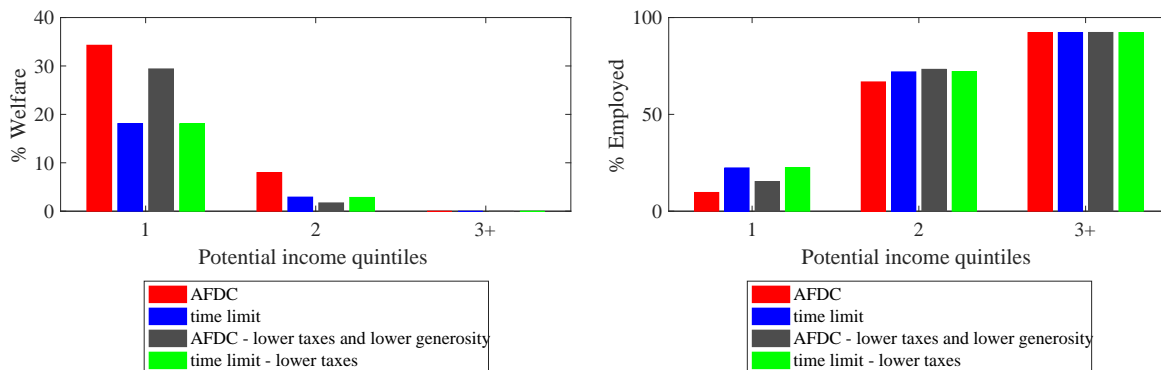
(a) Welfare use

(b) Employment

Notes: Simulated means of welfare use and employment by policy regime and by age-specific productivity z_{it}^F quintile. For outcome y_{it}^j and $j \in \{0,1\}$ where 0 is without any time limit and 1 is with a 5-year time limit imposed, the bars represent $P(y_{it}^1 = 1)$ and $P(y_{it}^0 = 1)$ of mothers, pooling across i and t. The four policy regimes are (i) AFDC with no time limits, (ii) AFDC with a 5-year time limit and no reduction in taxes, (iii) less-generous AFDC (reduced to 73.5% of case i) without time limits and a reduction in taxes by 0.375 percentage points, AFDC with a 5-year limit and a reduction in taxes by 0.375 percentage points.