NBER WORKING PAPER SERIES

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

Hamish Low
Costas Meghir
Luigi Pistaferri
Alessandra Voena

Working Paper 24356
http://www.nber.org/papers/w24356

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2018, Revised December 2025

We thank five anonymous referees and the editor, Magne Mogstad for detailed comments that substantially improved the paper. Abi Adams, Joe Altonji, Orazio Attanasio, Richard Blundell, Marianne Bruins, Tom Crossley, Mariacristina De Nardi, Felix Gray, John Kennan, Pat Kline, Robert Moffitt, Andreas Mueller, Brian Phelan, Ben Polak, Andrew Shephard, Melissa Tartari, Matt Wiswall, James Ziliak and participants in various seminars and conferences provided many helpful comments and suggestions. Bruno Escobar, Giacomo Fraccaroli, Davide Malacrino, Jorge Rodriguez Osorio, Matthew Schwartzman, Charlotte Woodacre, and Maria Alejandra Rodriguez provided excellent research assistance. Costas Meghir thanks the UK ESRC (grant number ES/K00624X/1), the ISPS at Yale, and the Cowles Foundation for financial assistance. The views expressed do not necessarily reflect the views of the Federal Reserve Bank of Chicago, the Federal Reserve System, or the National Bureau of Economic Research. All errors are our own.

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 December 2025
JEL No. H2, H31, H53, J08, J12, J18, J22

ABSTRACT

The 1996 US welfare reform introduced time limits on welfare receipt. We use quasi-experimental evidence and a rich life cycle model to understand the impact of time limits on different margins of behavior and well-being. We stress the impact of marital status and martial transitions on mitigating the cost and impact of time limits. Time limits cause women to defer claiming in anticipation of future needs and to work more, effects that depend on the probabilities of marriage and divorce. They also cause an increase in employment among single mothers and reduce divorce, but their introduction costs women 0.7% of lifetime consumption, gross of the redistribution of government savings.

Hamish Low
Federal Reserve Bank of Chicago
hamishlow1314@gmail.com

Costas Meghir
Yale University
Department of Economics
and IZA
and also NBER
c.meghir@yale.edu

Luigi Pistaferri
Stanford University
Department of Economics
and NBER
pista@stanford.edu

Alessandra Voena
Stanford University
Department of Economics
and NBER
avoena@stanford.edu

Marriage, Labor Supply and the Dynamics of the Social Safety Net *

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

November 26, 2025

Abstract

The 1996 US welfare reform introduced time limits on welfare receipt. We use quasi-experimental evidence and a rich life cycle model to understand the impact of time limits on different margins of behavior and well-being. We stress the impact of marital status and martial transitions on mitigating the cost and impact of time limits. Time limits cause women to defer claiming in anticipation of future needs and to work more, effects that depend on the probabilities of marriage and divorce. They also cause an increase in employment among single mothers and reduce divorce, but their introduction costs women 0.7% of lifetime consumption, gross of the redistribution of government savings.

Keywords: welfare reform, life-cycle, marriage and divorce, time limits, limited commitment, intrahousehold allocations

JEL codes: D91, H53, J12, J21

*We thank five anonymous referees and the editor, Magne Mogstad for detailed comments that substantially improved the paper. Abi Adams, Joe Altonji, Orazio Attanasio, Richard Blundell, Marianne Bruins, Tom Crossley, Mariacristina De Nardi, Felix Gray, John Kennan, Pat Kline, Robert Moffitt, Andreas Mueller, Brian Phelan, Ben Polak, Andrew Shephard, Melissa Tartari, Matt Wiswall, James Ziliak and participants in various seminars and conferences provided many helpful comments and suggestions. Bruno Escobar, Giacomo Fraccaroli, Davide Malacrino, Jorge Rodriguez Osorio, Matthew Schwartzman, Charlotte Woodacre, and Maria Alejandra Rodriguez provided excellent research assistance. Costas Meghir thanks the UK ESRC (grant number ES/K00624X/1), the ISPS at Yale, and the Cowles Foundation for financial assistance. The views expressed do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System. All errors are our own.

†Federal Reserve Bank of Chicago and IFS (email: hamish.low@chi.frb.org).

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

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

¶Stanford University, NBER, CEPR and SIEPR (email: avoena@stanford.edu).

Introduction

Welfare programs offer critical insurance for households who fall upon hard times. However, they can distort household decisions along several dimensions. In a number of countries, policymakers have expressed concern about the incentives that welfare programs generate for marriage and divorce, in addition to their impact on work choices. Eligibility for transfers and their generosity are often directly conditioned on the size and composition of household units. From an individual point of view, the value of insurance from welfare programs depends on marital status and expectations of changes in marital status.

The aim of this paper is to understand how time limits in eligibility for welfare programs affect labor supply, welfare participation, marital status, and ultimately individual well-being, by characterizing the dynamic trade-offs between incentives and insurance. We do so in the context of the major US welfare reform of 1996. A key element of this reform was to make access to welfare temporary. The pre-existing Aid to Families with Dependent Children (AFDC) provided coverage based on income and asset requirements, but without a time limit. This was replaced by the Temporary Assistance to Needy Families (TANF), which covers adult recipients with dependent children for up to 5 years over their lifetime. TANF introduced additional restrictions, such as work requirements, and was brought in at a time of positive shocks to the aggregate economy, and particularly to employment opportunities, which means that controlling for these features to isolate the effect of time limits is key.

Central to our analysis is allowing for marriage and divorce dynamics in a lifecycle framework. Family formation was an explicit target of the TANF reform. Further, the possibility of transitions between marital states is a crucial factor determining the importance of safety net programs and the impact of time limits on individuals. The importance of marriage arises, in part, because of the insurance it provides. The prospect of marriage offers the possibility of future insurance, while the prospect of divorce raises the need for insurance. In addition, marriage may limit access to benefits because of eligibility rules that depend on family income.

To motivate our framework, we start by empirically documenting the short-run effects of the time limits on labor supply, welfare participation, and marital status. In particular, we apply the quasi-experimental approach of Grogger and Michalopoulos (2003) on data from the Survey of Income and Program Participation (SIPP) and the Current Population Survey (CPS) to estimate the response of households to the prospect of losing benefit eligibility because of time limits. Identification relies on the fact that women whose youngest child

is close enough to age 18 at the time of the reform would have been unaffected by the introduction of time limits because eligibility was always restricted to those with a dependent child. Women with younger children saw their potential eligibility horizon shortened by the introduction of time limits. Other elements of the reform, such as work requirements, affect all mothers in a similar way, with the exception of mothers of very young children who were exempt from work requirements. Thus our quasi-experimental approach is designed to isolate the effects of time limits alone.

We find that welfare utilization declined dramatically and persistently right after the introduction of TANF, both for single mothers (see also Grogger and Michalopoulos, 2003, amongst others), and for married mothers (a group that has been subject to less empirical scrutiny). Employment of single mothers increased, partially offsetting the decline in benefit use (see also Grogger, 2003), but we find that it did not change for married ones. In terms of family structure, we find that the prevalence of divorce declined significantly.

These empirical results support the idea that the introduction of time limits had a substantial anticipatory effect on labor supply and household formation. However, these results are silent about the mechanisms that underlie the long-term consequences on behavior and about the effects of the reform on individual well-being. To make progress on these questions, we develop a lifecycle model where individuals make choices on being single, married, or divorced, alongside employment and savings and the decision of whether to claim welfare benefits if eligible. Married couples share resources and enjoy economies of scale in a limited commitment collective framework, where outside options affect both within-couple sharing and the prevalence of divorce. The economic environment is characterised by uncertainty and by the broad welfare system, including AFDC/TANF, the Earned Income Tax Credit (EITC), and Food Stamps. In the marriage process, uncertainty is over the probability of meeting a partner and over their characteristics (assets, earnings capacity, and match quality). When married, couples face shocks to their marriage match quality and the income of each partner. In the labor market, the primary source of uncertainty is persistent shocks to productivity, which can generate substantial lifetime risk. When shocks are persistent, people on welfare anticipate that their wages will continue being relatively low in the future, and hence time limits will be more likely to bind. With i.i.d. transitory shocks, as is common in the literature (e.g. Swann, 2005; Chan, 2013), this mechanism is absent.

We use the method of simulated moments (McFadden, 1989; Pakes and Pollard, 1989) to estimate our model on a sample of women and men without a college degree and of working

age, who are most likely to be affected by changes in the welfare system. We use pre-reform data because the post-reform system included other elements that our model does not focus on and varied substantially across states. With our approach, we are able to estimate the parameters based on the simpler welfare regime before the changes were implemented, and then focus on the impact of time limits.

We validate our approach by comparing the simulated short-term effects of time limits generated by the estimated model to our reduced form quasi-experimental results. The estimated model can account for the decline in welfare participation among married and single women, the increase in employment among single mothers, and the decline in divorce that we observe following the introduction of TANF. The baseline model also matches the time dynamics in the decline of welfare participation, which occurs shortly after the reform before welfare participants had hit the time limit. Alternative estimated versions of the model that do not allow for marriage transitions, divorce transitions, or asset accumulation, cannot explain these patterns. Similarly, a version of the model that includes work requirements alone can explain certain features of the data but misses some key impacts of the reform.

Our estimated model highlights that time limits reduce the scope of insurance provided by welfare programs and lead mothers to "bank" their benefit eligibility in anticipation of greater future needs. The imposition of time limits creates a trade-off between using benefits today vis-à-vis having them available for insurance in the future. Given this trade-off, some mothers decide to forgo benefits today, preserving the option for periods and states of the world where the marginal utility of consumption is higher.[1] Allowing for the possibility of divorce is an important element for explaining the "banking" of benefits by married women: they may prefer not to claim now, even if they are not working and eligible, so as to have the possibility of claiming if and when divorced. Indeed, we find that in a counterfactual scenario in which divorce is not an option, "banking" by married women is much reduced. On the other hand, for singles, the possibility of future marriage reduces the negative impact of time limits because marriage is associated with an income stream from the husband, reducing the risk of low consumption and the need for insurance. Consistent with this observation, we find that in a counterfactual scenario in which we eliminate the possibility of marriage, "banking" by single households increases substantially. Thus, allowing for marriage and divorce matters

[1]We use the term banking to refer to the deferral of claiming benefits. This is in contrast to claiming benefits and saving the money for future use. Such a decision is not practical because claiming means forgoing current job opportunities and the resulting need to finance current consumption would push individuals to spend current benefits rather than saving them. In principle, however, the availability of a savings vehicle may mitigate the desire to "bank" benefits. Allowing for this is one of the reasons our model includes savings.

for understanding the impact of imposing time limits on receipt of welfare benefits and in determining how safety net programs affect behavior.

To investigate further, we evaluate the welfare implications of the reform under different assumptions on marriage (and, hence, divorce) prospects, and on how revenue savings are redistributed. Time limits impose a utility loss on low-educated women of 0.7% of lifetime consumption on average, with negligible effects on men, before any redistribution of gains generated by government savings. Time limits are substantially more costly when the marriage option is suppressed.

Yet the welfare effects of time limits also depend on how the government savings are redistributed. If the redistribution of government savings targets single mothers through lump-sum payments without any stigma cost, women's well-being increases with time limits relative to AFDC, while men's well-being declines because of worsening marriage prospects. When savings are instead redistributed through lower payroll taxes, women's well-being declines, because they lose resources in states of high marginal utility of consumption and gain them in states of lower marginal utility of consumption (when they earn more labor income). Men's well-being, in contrast, improves: the loss in benefits during marriage is more than compensated by the tax decline.

Our paper builds on existing work on welfare reform and life cycle behavior as well as on earlier studies of the impact of the 1996 welfare reform. The empirical literature on the effects of this reform is large, with excellent overviews by Blank (2002) and Grogger and Karoly (2005). A number of empirical studies highlighted that time limits discouraged the use of welfare (Mazzolari and Ragusa, 2012) and led women to "bank" their future eligibility (Grogger and Michalopoulos, 2003). Other studies show that the 1996 reform increased the employment of eligible women (Meyer and Rosenbaum, 2001; Grogger, 2003; Fang and Keane, 2004; Kline and Tartari, 2016). All these studies focus on single mothers, the main recipients of cash benefits. However, married women represent 25% of welfare recipients.

A smaller empirical literature examined whether marriage and divorce are affected by time limits (e.g. Schoeni and Blank, 2000). Moffitt, Phelan and Winkler (2020) consider the impact of various components of the reform. They distinguish between households where the biological father is present as opposed to either no partner or one unrelated to the children. While they find that some work-related elements of the welfare reform increased single parenthood, they find no impact of time limits. Finally, Bitler et al. (2004) presents strong evidence that divorces decline but does not find that marriages increased. Our main

point is that marriage (and expected transitions into and out of it) matters not only as an outcome to study, but also as a conditioning state affecting behavior.

Our approach builds on an important literature that developed dynamic models of labor supply, some of which include choices of welfare use. Notable examples include Keane and Wolpin (2010) and Blundell et al. (2016), both of which specify rich life cycle models focusing on the role of welfare use while allowing for education choice and human capital accumulation. Both allow for the birth of children and transitions between being single and married, although only Keane and Wolpin (2010) treat these as endogenous choices. On the other hand, Blundell et al. (2016) allow for savings, which are important sources of self-insurance and need to be accounted for when evaluating the impact of benefit reform.

Our paper 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 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.[2]

Intrahousehold allocations are based on a dynamic collective model with limited commitment. The theoretical underpinnings for this are found in Chiappori (1988, 1992) and Blundell, Chiappori and Meghir (2005) and its dynamic extension by Mazzocco (2007). The risk-sharing framework with limited commitment draws from Ligon, Thomas and Worrall (2000) and Ligon, Thomas and Worrall (2002*b*) as extended to the life cycle marriage model by Mazzocco, Yamaguchi and Ruiz (2013) and Voena (2015).

Turning to the literature that has considered the dynamic aspects of the reform, Fang and Silverman (2009) estimate a model of work and welfare participation allowing for time inconsistency through present bias. Beyond this, there are two important papers with dynamic models that address explicitly the role of time limits imposed by the 1996 welfare reform and are therefore the closest antecedents to our paper.

The first is Swann (2005). His model considers the joint decisions of program participation, work, and marriage. However, it does not model savings, removing an important channel for self-insurance; moreover, it switches off demographic uncertainty by assuming perfect foresight for fertility. Welfare participation is limited to single mothers alone and wage shocks are assumed to be i.i.d., reducing the implied uncertainty that individuals face and hence understating the need for "banking".

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

The second is the paper by Chan (2013). It models the choices to work and to claim benefits and captures in much detail the institutional framework in place at every point in time. However, his model focuses only on continuously single mothers, which may overstate the role of welfare. As reported by Bane and Ellwood (1996), 29.4% of AFDC termination of claims occurred when the single mother "became a wife", which is similar to what we find in our data. Moreover, as in Swann (2005), savings are absent. By ignoring the family and ignoring the option to save, the paper eliminates important mechanisms providing insurance (marriage and assets), potentially overstating the effects of the reform on individual welfare. Further, divorce is an important channel of risk for married women and can affect their claims behavior. Ignoring family transitions leads to an incomplete picture of the effects of benefits and distorts the assessed impact even for single mothers, the group that is the focus of Chan (2013).

We provide detail of the 1996 welfare reform and the interaction between marriage and welfare in Section 1. The data and reduced form analysis of the effects of time limits are presented in Section 2. In section 3 we present our life-cycle model and discuss estimation of the model's parameters. The core results are presented in counter-factual analysis in section 4 and welfare analysis in Section 5. Section 6 concludes. Further details are given in the Online Appendix, where Appendices A to E correspond respectively to sections 1 to 5.

1 AFDC, TANF and the Welfare Reform

The Personal Responsibility Work Opportunity Reconciliation Act (PRWORA) was signed by President Clinton in August 1996. It gave states considerable latitude in setting parameters for welfare within broad federal guidelines (Ziliak, 2015). 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.

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, and states started counting such months at different times between 1995 and 1998.[3] Moreover, 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

[3]The size of the blocks grant was based on three-year average spending computed over the years preceding the reform and was independent of the business cycle.

funds. This flexibility meant that TANF varied significantly across states.[4] Information about the presence of time limits was systematically disseminated to beneficiaries through welfare offices, official notices and letters:

> ""The clock is ticking" is a common sight on welfare office walls and a catch-phrase used by many caseworkers. Images of ticking clocks and hourglasses have become familiar" (Bloom, Farrell and Fink, 2002).

These changes in eligibility requirements were not accompanied by appreciable changes in the generosity of benefits available conditioning on eligibility. During the same period, the EITC was expanded and this increased labor force participation among low-income individuals.

A further element of the welfare 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 childcare assistance. Mothers of children below the age of 1 are, in most states, exempt from work requirements. 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 may make measurement and enforcement problematic.[5] While work requirements and child care assistance may be important and also potential empirical confounds 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 (see Figure A.1 in the Appendix). We use this variation to isolate the effect of time limits from other elements of the reform and perform a series of robustness checks to ensure that the empirical analysis isolates the effect of time limits. In Section 4.3 we introduce work requirements in the model and study their implications.

An explicit goal of replacing AFDC in 1996 was to promote marriage (Ziliak, 2015). When AFDC was first introduced, benefits were limited to single parents. Partly in response to concerns that this was creating disincentives to forming dual-parent households, eligibility was later extended to include couples where the second adult was unemployed. As a consequence, welfare use among married couples became more common by the late 1980s, accounting for 20 to 30% of welfare recipients (Moffitt, Reville and Winkler, 1998). Similarly, in our SIPP

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

[5]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.

sample, as shown in Appendix Figure A.2, roughly a quarter of mothers on AFDC are married.[6] Besides maintaining eligibility for all households satisfying the income and asset tests, PRWORA also incentivized states to implement additional pro-family policies, such as the adoption of family caps, family planning provisions, step-parents' income disregards, etc. Moffitt, Phelan and Winkler (2020) provide a comprehensive discussion of how TANF rules changed incentives for marriage.

Figure 1: AFDC Benefit Amounts by Marital Status

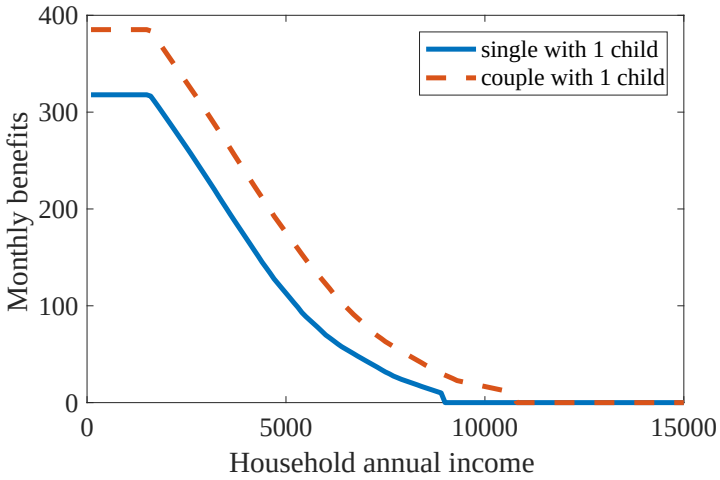

Notes: Monthly AFDC benefits by household annual income. AFDC benefits are a population-weighted state-level average. Amounts are in 1995 real U.S. dollars.

Marital status is strongly associated with changes in welfare eligibility. To provide an intuition for this, in Figure 1 we show how AFDC benefits varied 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). Most states condition AFDC eligibility on having gross income below the poverty line.[7] 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. If she were to marry a man whose earnings are half of the poverty line, her benefits would drop by more than 50%.[8]

[6]Our sample, which comprises women without a college degree, is described in detail in the next section.

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

[8]By contrast, for example, food stamps are available to all households, irrespective of the presence and of age of the children, and the amount paid is considerably larger for couple-headed households. 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 some variation across states.

1.1 Data and Descriptive Statistics

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 the years 1990 to 2007, focusing on the period up to 2002 in our main analysis.[9] 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.[10] The CPS is primarily a labor force survey that is conducted monthly. Its March supplement includes additional information on family characteristics and program participation and is nationally representative. The CPS Outgoing Rotation Group provides a limited longitudinal sample.

For the reduced form evidence in section 2, we restrict the sample to women between 18 and 60 who are not college graduates, and with at least one child under age 19.[11] We add men to this sample when we estimate the lifecycle model. 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 observation in a given wave.

Our main analysis includes 40,405 women in the SIPP who are heads or spouses of the head of their household, leading to a total of 254,627 quarterly observations. Our CPS sample includes 65,530 women and a total of 113,399 annual observations.[12] Table B.1 in the Appendix provides summary statistics for the SIPP and the CPS, pre- and post-reform.

[9]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.

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

[11]Since we are interested in the impacts of means-tested welfare benefits, such as TANF, we focus on lower-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.

[12]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).

2 Empirical Evidence on AFDC, Marriage and Time Limits

This section provides two types of evidence to show how welfare reform interacts with marital status. First, we show how current marital status and transitions in marital status are associated with welfare participation, and flows on and off welfare. Second, we show the relationship between the introduction of time limits and key outcome variables: welfare benefit utilization, women's employment, marriage, divorce, and fertility. In our reduced form analysis in this section, we consider these outcomes separately. We report results on employment and benefit utilization conditioning on marital status to highlight its importance. However, given that marital status is an evolving choice, the empirical evidence in this section serves to highlight the need for a model that considers these outcomes simultaneously.

2.1 Marriage, Divorce, and Welfare Use

Marital arrangements affect welfare generosity and the need for support, and so changes in marital status are likely to be closely tied to transitions in and out of welfare participation. Figure 2 shows that the rates of moving on and off welfare differ by marital transition. Figure 3 shows the marital circumstances of individuals that are moving on and off welfare. Figure 4 shows the time path of welfare use around marriage and divorce.

Panel (a) of Figure 2 shows that in SIPP data pre-reform, out of women who are getting divorced, the share of women who enter welfare is almost 10%. By contrast, out of women who marry or experience no change in marital status, the share entering welfare is less than 2%. Panel (b) shows that the share of women who exit welfare among those getting married is an order of magnitude larger than for women who remain single or become divorced. Further, these marital status transitions are relatively frequent: the annual probability of entering marriage for a single woman in our sample is over 5 percentage points up until age 40, while the probability of exiting marriage via separation or divorce ranges from 6pp to 1pp (Appendix Figure A.3).

Panel (a) of Figure 3 reports the composition of those moving onto welfare, and panel (b) the composition of those flowing off, decomposing by major life events associated with marriage, employment, and fertility. Pre-reform, 9% of women who enter welfare simultaneously experience a divorce, similar to the share of entrants who experience a job loss (8%). Among those exiting welfare, almost 10% get married, while 24% move into work. Fertility

Figure 2: Welfare Transition Probabilities for Women Changing Marital Status

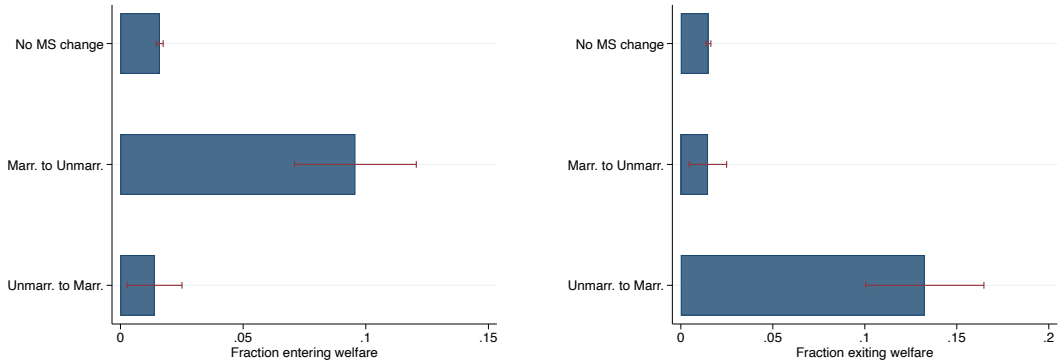

(a) Probability of entry (b) Probability of exit

Notes: We use pre-1996 data from SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below. In panel (a) we plot the shares entering welfare among women experiencing no change in marital status and among those who do. In panel (b) we plot the corresponding shares leaving welfare from one year to the next. We plot the associated 95% confidence intervals.

Figure 3: Life Events Associated with Welfare Entry and Exit

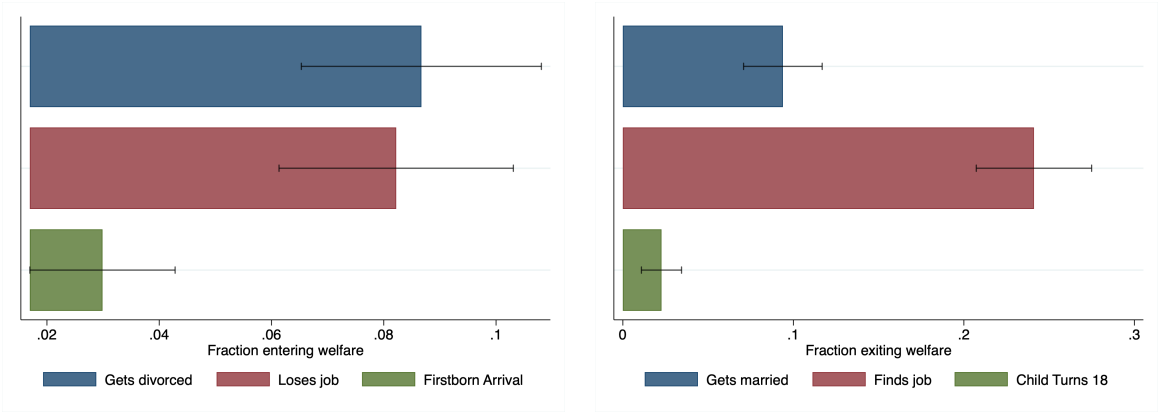

(a) Probability of entry (b) Probability of exit

Notes: We use pre-1996 data from SIPP panels. Sample of female heads of household who are not college graduates. In panel (a) we plot the fractions of welfare-entrant mothers who, at the same time as entry, (i) get divorced/separated, (ii) lose their job, and (iii) have a newborn while previously their youngest child (if any) was over 18. In panel (b) we plot the corresponding shares of welfare exiters from one year to the next who, at the same time as leaving, (i) get married, (ii) find a job, and (iii) whose youngest child turns 18 years old. We plot the associated 95% confidence intervals.

events (changes in the number of dependent children) have much smaller effects. This ev-

idence corroborates the earlier findings of Bane and Ellwood (1996) from the Panel Study of Income Dynamics on the importance of marriage and divorce for understanding welfare transitions. They show that, in the 1980s, divorce or separation were the most common factors associated with the beginning of a welfare spell and marriage or reconciliation with the end of a welfare spell.[13]

This suggestive evidence is reinforced by Figure 4, where we plot the rate of welfare participation by month before and after a change in marital status. Getting divorced or separated is associated with an increase in welfare participation of 10pp (pre-1996, panel a) and 4pp (post-reform, panel b). Getting married is associated with a reduction in welfare participation of 5pp (pre-1996, panel c) and 8pp (post-reform, panel d). This shows that marital transitions continued to be associated with transitions into and out of welfare even after the reform. Appendix Figure A.4 shows corresponding transitions into unemployment.

2.2 Reduced Form Evidence of the Impact of Time Limits

We present reduced-form evidence on the impact of limiting access to welfare via time limits on welfare use, employment, and marital status. We consider these outcomes separately, and we show results conditioning on whether mothers are married or unmarried. However, a core part of the argument of this paper is that there are important linkages between outcomes, and further, that conditioning on marital status can be problematic because marital status is not constant regardless of whether there is welfare reform.

2.2.1 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, 2007; Mazzolari and Ragusa, 2012) to examine the behavior of married couples and the role of household

[13]In particular, "Wife became female head" was the *most* common reason behind the beginning of a welfare spell in the 1980s (42.1% of spells in the PSID sample between 1980 and 1988). Symmetrically, "Female household head becom[ing] a wife" was one of the most common reasons behind the end of a welfare spell, accounting for 29.4% of all spells. The book clearly shows us that, in terms of correlations, changes in marital status were strong determinants of changes in welfare participation and a much stronger one than earnings shocks.

Figure 4: Welfare Participation Before and After Changes in Marital Status

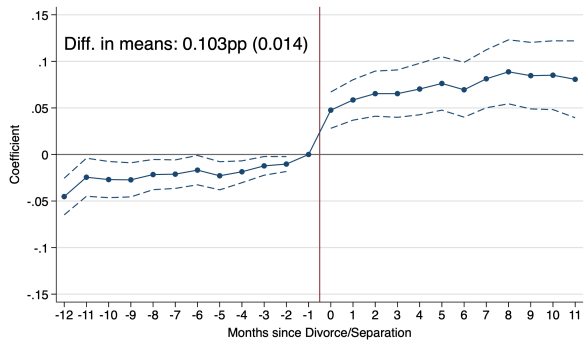

(a) Pre-1996 Sample

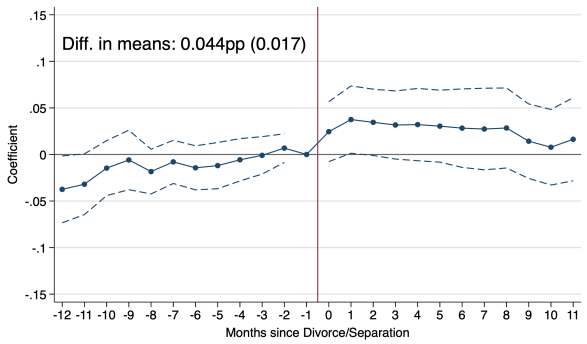

(b) Post-1996 Sample

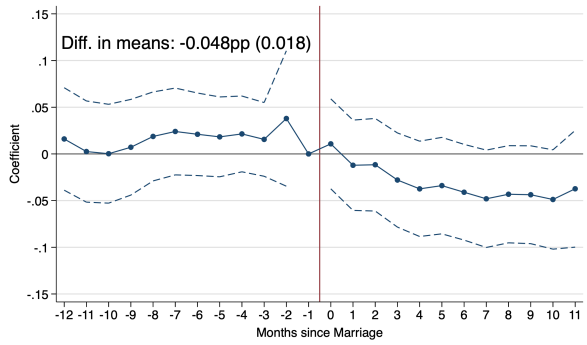

(c) Pre-1996 Sample

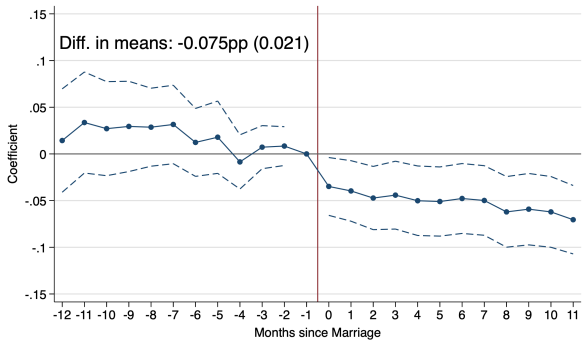

(d) Post-1996 Sample

Notes: 95% confidence intervals based on standard errors clustered at the person level. Data from the 1990-1996 SIPP panels; thus, until the year 2000. Sample of female heads of household who are not college graduates, have children aged 18 and below, and got a divorce over the years at stake. In panel (a), we plot the coefficients for AFDC participation controlling for the months from the divorce for the pre-reform period and for individual FEs (Obs. = 23,453; Unique women = 1,150). In panel (b), we plot the coefficients for TANF participation controlling for the months from the divorce for the post-reform period and for individual FEs (N = 9,373; Unique women = 359). In panel (c), we plot the coefficients for AFDC participation controlling for the months from the marriage for the pre-reform period and for individual FEs (Obs. = 16,561; Unique women = 998). In panel (d), we plot the coefficients for the TANF participation controlling for the months from the marriage for the post-reform period and for individual FEs (Obs. = 11,844; Unique women = 528).

structure.[14]

[14]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 payments of about $250 on average. They find no evidence that such loss is offset by increases in other income sources and conclude that time limit enforcement resulted in an increase in the rates of deep poverty for households hitting the time limits.

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.[15] 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. Mothers who are exposed to the reform varies by state depending on the time limit imposed. For example, in a state where the limit is 5 years, women whose youngest child is younger than 13 are exposed. The threshold is 16 in states with a two-year limit. Hence, *Exposed* is a function of the demographic characteristics of a household and the rules of the state in which the household resides. Appendix Figure B.1 reports the potential variation for the year 2000.

In most states, the reform imposed stricter work requirements, so that a level effect on employment may be expected across both treated and control groups. However, the work requirements are independent of the age of the children, and consequently comparing the impact of the reforms within states and across families with children of different ages identifies the effect of time limits alone. The exception is families with very young children, who are not subject to work requirements (and may received childcare assistance); we thus experiment with removing them from the sample.

The value of *Exposed* may change over time because some states change their statutory time limits during the sample period and because states differ in the date when the time limit clock starts to tick. Therefore, we construct a variable $Post_{st}$ indicating the post-reform period, based on the timing of the introduction of time limits.[16] We look at the impacts after the reform up to 2002, as well as the longer-term effects for further validation (Appendix B). We do not observe how much of their benefits people have used, and so we focus mostly on effects in the period shortly after the reform, where the risk of treatment/control contamination is minimal. We interact the $Post_{st}$ variable with $Exposed_{ds}$ and estimate:

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

where α is the average effect of the reform, identified from changes in the behavior of households with a youngest child of different ages, $\mathbf{X}$ is a vector of individual and state-level

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

[16]For example, Arizona had the clock starting in November 1995, while California started the clock in July 1997. Years of time limits are reported in Mazzolari and Ragusa (2012).

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.[17]

2.2.2 Empirical Results

We start by showing the relationship between time limits and welfare use and employment. This provides the context for thinking about household structure.

Benefits Utilization and Employment Table 1 reports the main effects of time limits within the first six years after 1996. 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 approximately 30%.[18] These numbers fell to 4% and 11% respectively. Our reduced form analysis of welfare utilization shows how much of this fall is due to time limits because we compare within state and between demographic groups that are both subject to the same work requirements. Exposed households have a 3 percentage point (pp) lower probability of claiming benefits after the introduction of time limits. Unmarried women have an 8.7pp lower probability of claiming welfare, a large effect consistent with those reported by Grogger and Michalopoulos (2003) and Grogger (2003).[19] We also document a 1.1pp decline in the welfare use of married women, who have low rates of welfare use pre-reform. CPS data shows similar declines across these outcomes. Since the aggregate decline in the share of unmarried (married) women on welfare is approximately 20pp (2pp), one interpretation is that the imposition of time limits explain roughly half of the decline for the relevant group. There are other important forces that were also driving down aggregate welfare use post-reform, in particular the economy-wide boom post-1996 (Ziliak et al., 2000), and the introduction of work requirements, which we discuss further below.

The second panel of Table 1 shows the impact on employment levels. The introduction of time limits increases the employment probability of single mothers by about 5pp, in line

[17]To study the dynamics of the outcome variables, we also interact $Exposed_{ds}$ with dummies for each calendar year between 1990 and 2002 (excluding 1995 for scaling). We estimate pre-reform interactions with year dummies to rule out pre-reform trends across demographic groups:

$$y_{idst} = \sum_{\tau=1990}^{2002} \alpha_\tau \, Exposed_{dst} \times \mathbf{1}\{t=\tau\}_{st} + \mathbf{X}_{idst}\boldsymbol{\beta} + f_{st} + f_{ds} + f_{mt} + \epsilon_{idst}. \quad (2)$$

[18]A full set of descriptive statistics are provided in Table B.1.

[19]Longer-term effects, using data until 2007 (Appendix Table B.3) are qualitatively similar.

Table 1: Use of Benefits, Employment, Marital Status and Fertility after the Reform

	AFDC/TANF Utilization					
	Whole sample		Married women		Unmarried women	
	SIPP	CPS	SIPP	CPS	SIPP	CPS
$Exposed_{dst}Post_{st}$	-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	0.101	0.077	0.034	0.019	0.306	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					
	Whole sample		Married women		Unmarried women	
	SIPP	CPS	SIPP	CPS	SIPP	CPS
$Exposed_{dst}Post_{st}$	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	0.642	0.647	0.647	0.654	0.624	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
	Divorce/Separated		*Married*		$Newborn_{t+1}$ *(Annual)*	
	Whole sample		Whole sample		Whole sample	
	SIPP	CPS	SIPP	CPS	SIPP	CPS
$Exposed_{dst}Post_{st}$	-0.027***	-0.015*	0.004	-0.007	-0.001	-0.002
	(0.007)	(0.008)	(0.007)	(0.010)	(0.005)	(0.004)
Mean pre-reform	0.151	0.126	0.753	0.796	0.060	0.049
Obs	254,627	112,128	254,627	112,128	55,142	47,344
R^2	0.03	0.01	0.05	0.05	0.08	0.04

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. We use quarterly data from 1990-2001 SIPP panels, until 2002 and annual data from 1990-2002 March CPS. We annualized SIPP data for the newborn outcome. 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.

with Grogger (2003). 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 an important mechanism of mitigation. However, there is an increase of between 2 and 5pp in the fraction of single mothers who are neither employed nor on welfare and a small decline in the probability of being simultaneously employed and on welfare (see Table B.4 in Appendix C). On the other hand, time limits had no detectable effect on the employment of married women, despite the fall in benefit claiming.

Single women behave differently from married women pre-reform: only 7% of single women are neither on benefits nor employed, whereas the fraction for married women is 32%. Married women have income streams from their husbands, and so are more often ineligible for benefits, and yet may not need to work. This means that despite the reduction in benefit use, married women do not have to return to work to smooth consumption to the extent that unmarried women have to.[20]

Figure 5: Program Participation Dynamics Following the Introduction of Time Limits by Quarter

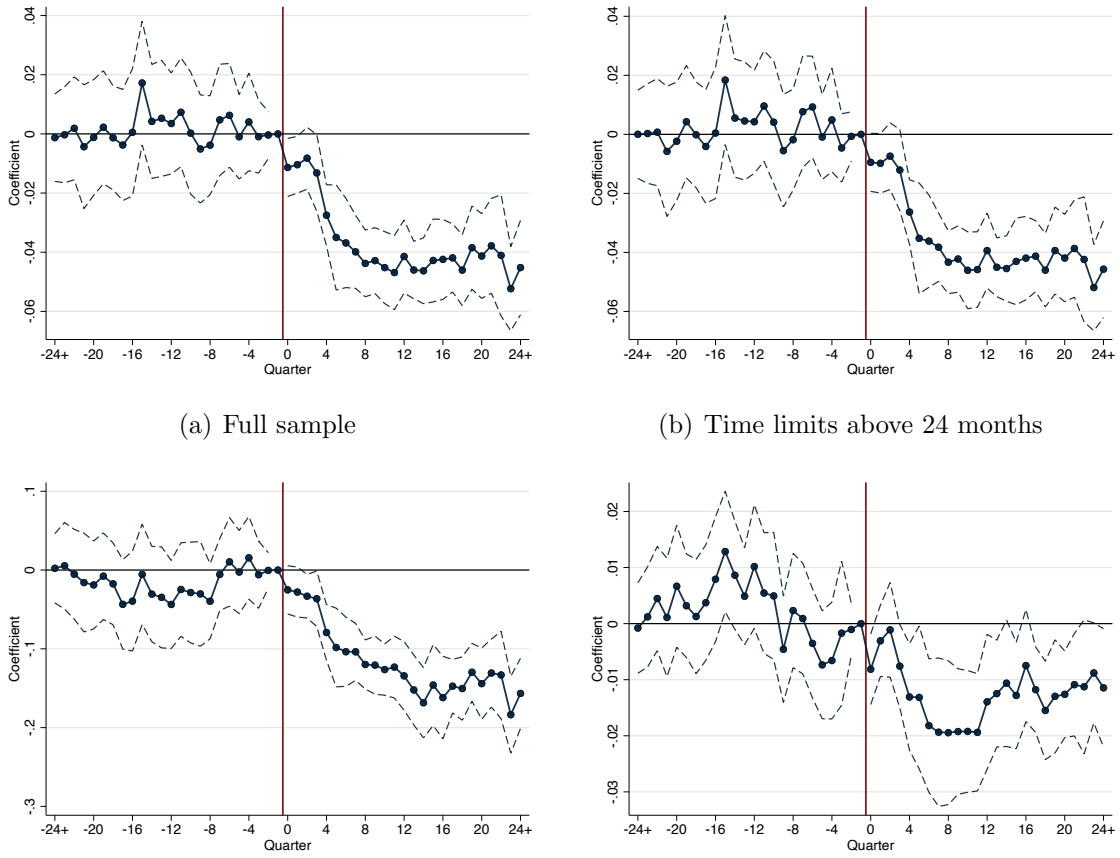

(a) Full sample

(b) Time limits above 24 months

(c) Single

(d) Married

Notes: 95% confidence intervals based on standard errors clustered at the state level. Data from the 1990-2001 SIPP panels, until 2002. 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.

[20]In Table B.2 in Appendix B, we re-estimate our reduced form welfare use and employment regressions for married women, splitting the sample into couples when the husband is employed compared to when unemployed. The impact of time limits on welfare use is greater for couples where the husband is unemployed, although from a larger base utilization. However, the impact on employment is not significantly different.

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 addresses this question. Figure 5 plots the coefficients of the variable *Exposed* interacted with quarterly dummies before and after the introduction of time limits in a given state on the overall sample (subfigure a). Even after excluding states with a limit below 2 years (subfigure b), the utilization rate begins to decline significantly within a year of the reform to a persistent drop of about 4 percentage points by year 2: households reduce their benefit utilization *before* anyone is likely to have run out of benefits eligibility. The decline is over 10 percentage points for unmarried women, and around 1.5 pp for married women. The pre-emptive decline suggests forward-looking behavior (in the form of "banking" of benefits).[21]

Our results focus on the average effect across mothers with children of different ages, but the intensity of treatment differs. To address this, we interact the exposure coefficient with dummy variables for the age of the youngest child: parents of younger children respond more strongly to the introduction of time limits compared to parents whose child is closer to age 13, consistent with the dynamic incentives introduced by time limits (Appendix B, Figure B.5).

Household Structure: Divorce, Marriage, and Births A central motivation, and indeed a stated goal, of the 1996 welfare reform was to encourage "the formation and maintenance of two-parent families". In this section, we consider whether there is direct evidence that time limits affected household structure, particularly marriage, divorce and fertility. The bottom panel in Table 1 reports these outcomes.

Marriage and divorce 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. This is a substantial decline: the pre-reform fraction of women who were divorced was between 12 and 15%.

The decline in divorce was not associated with an increase in the stock of married people. As discussed by Bitler et al. (2004), the effects of the welfare reform on household formation and dissolution are not obvious. The welfare reform, by curtailing the extent of public insurance available to low-income women, may have induced those who were already married to attach a higher value to marriage as a valuable risk-sharing tool (through male labor

[21] In the Appendix we estimate specification (2) for three outcomes: welfare use (see Figure B.2), employment (Figure B.3), and marital status (Figure B.4). The findings are similar to those of Table 1.

supply, for example), and reduced the option value of being single (and potentially claiming benefits). The reform increased the gains from marriage for those who would have been eligible for welfare benefits if single and this immediately translates into lower divorce rates. However, moving from being single to married takes time because of search costs and not every match will translate into a marriage that is beneficial. As a result, any tendency for an increased marriage rate is attenuated and becomes too small to detect in the data. Further, the changes in the implied insurance from marriage will not only affect the likelihood of marriage but also women's bargaining power and allocations of resources within marriage, as we capture in our model below. Figure B.4 reports the dynamics of marriage and divorce based on Equation (2). This highlights the time needed to adjust marital status.

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 birth 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), and use a dummy for having a newborn in $t+1$ as the dependent variable. We annualize the SIPP data for this purpose since we are examining the effect on birth in the next year. We find no effect of time limits on subsequent births. This finding partly justifies our choice to treat fertility choices in the model as an exogenous stochastic process, that is unaffected by changes in the environment.[22]

2.2.3 Other Aspects of the Reform

A key question is whether the approach above separates the effect of time limits from other features of welfare reform, such as work requirements, increased stigma, or childcare provisions, and from features of the aggregate economy. 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 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

[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.

by time limits because mothers of older children were under more pressure to work. By contrast, childcare provisions, 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 ages 6 and 18, who should be facing similar regulations and claiming conditions but be differentially affected by the likelihood of incurring binding time limits. See Appendix Table B.5. Our results on the impact of time limits are unchanged.

To address whether other aspects of the reform could drive our estimates we have collected further information about additional features of the reform. In Appendix B, we describe and report the findings from controlling for variation induced by other policies, such as work requirements (Appendix Table B.7), child support enforcement (Appendix Table B.8), and waivers (Appendix Table B.9). Our strategy for controlling for other aspects of the reform accounts for the direct effect of these features in the estimation of the impact of time limits, but does not capture interactions between time limits and these features.

2.2.4 Robustness Checks

In Appendix B, we describe and report the findings from additional robustness checks: (a) To address attrition in the SIPP sample, we only consider the first wave of each panel (Appendix Table B.10); (b) We examine the behavior of college graduates, as a sample that is less likely to be affected by time limits (Appendix Table B.11). The findings are consistent with the interpretation of our baseline results.

3 Life-Cycle Model

3.1 Overview

We specify a life-cycle model without bequests that introduces the channels through which the economic and policy environment affect labor supply behavior, savings, and family structure. The estimated model allows us to evaluate how individuals make choices to mitigate the expected increased exposure to risk resulting from imposing time limits to welfare benefits, or other policy counterfactuals. With it, we can assess the longer-run changes in labor supply and family structure and quantify the effects on well-being.

As a brief overview of the model, in each period single women may meet a man, drawn

from a distribution of single men of similar age and characterized by their wealth and labor productivity. They mutually decide whether to marry or not, depending on the joint value of marriage relative to the value of being single and on a random marital match quality effect. Married couples, on the other hand, decide whether to continue in their relationship or to divorce, a decision that depends both on the economic environment defining outside options and on the shock to a marriage match. Children are not a choice, but they arrive randomly with a probability that is a function of the mother's marital status, her age and the age of any existing child.

Single women decide whether to work or not and how much to save or consume, and whether to claim AFDC/TANF, if they have children. The decisions of married women/couples are governed by intrahousehold bargaining where their relative power is defined by Pareto weights. These weights, originally set at the point of marriage, evolve based on a limited commitment framework where updates take place only when one of the household members has a credible threat to leave the marriage.[23] Within this framework, married women decide whether to work or not and those with children decide whether to claim AFDC/TANF benefits or not. Both members of the household save in a common asset, which upon divorce is split equally among them. Whether men work or not is exogenous and is governed by a random process.

Savings are an important source of self-insurance, yet typically ignored in most of the literature studying the impact of welfare reform. Ignoring savings would distort the impact of welfare benefits. First, it would lead us to overstate the impacts of reducing or eliminating benefits. Second, in the presence of time limits, it would distort the decision of whether to defer claiming (bank) benefits. We thus allow for savings, both because these households hold assets in the data and because any analysis of the welfare effects of the reform that changes insurance options would be incomplete without explicitly taking into account self-insurance.

In Subsection 4.3, we consider how time limits interact with the other key element of the reform, the imposition of work requirements. Because the model has no job search, we assume that work requirements act as a utility cost imposed on mothers who claim benefits but do not work. This is similar to Mullins (2022).

[23]See Chiappori et al. (2020) and Lise and Yamada (2019) for evidence in support of limited commitment.

3.2 Problem of the Single Woman

We focus on women who have completed their schooling choices and have not obtained a college degree. The subscript s indicates that individual i is single, and F indicates a woman. The subscript j denotes a man, and ij the corresponding couple. At each age t, women decide 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 $b_{i,t}$. 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 is implemented.[24] Employment, savings, and program participation decisions will be conditional on the marriage decision that occurs at the beginning of the period.

If a woman remains single, her budget constraint is given by

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

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

where $A_{i,t}^{F_s}$ are assets and $CC_{i,t}^a$ is the financial cost of childcare paid only if the woman works. The woman's wage rate $w_{i,t}^F$ follows a persistent stochastic process, detailed below. The term $e(k_{i,t}^a)$, the amount of consumption obtained by spending \$1, is an equivalence scale to account for the presence of a child of age a ($k_{i,t}^a$). The arrival of children is stochastic and exogenous, albeit varying with the woman's marital status, own age, and age of an existing child. Children affect consumption, benefit eligibility, and the opportunity cost of women's time in the labor market. We account for three important social safety net programs: food stamps (denoted FS), EITC and AFDC/TANF. The value of all programs depends on demographics and income.

The state space for a single woman is $\mathbf{\Omega_{i,t}^{F_s}} = \{A_{i,t}^{F_s}, w_{i,t}^F, k_{i,t}^a, TB_{i,t}\}$, where $TB_{i,t}$ is the number of years of benefit use. 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

[24] Positive assortative matching arises endogenously because a match occurs only if *both* agents prefer to match rather than continue searching.

$\{k^a_{i,t}, w^F_{i,t}P^{F_s}_{i,t}, TB_{i,t}, A^F_{i,t}\}$. 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 $\{A^M_{j,t}, y^M_{j,t}\}$ (assets and exogenous earnings) and together they draw an initial match quality $L^0_{ij,t}$. 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 3.3 and in Appendix C.[25]

We denote by $V^{F_s}_t\left(\boldsymbol{\Omega^{F_s}_{i,t}}\right)$ the value function for a single woman at age t and $V^{F_m}_t\left(\boldsymbol{\Omega^{m}_{i,t}}\right)$ the value function for a married woman at age t. A single woman has the following value function:

$$V^{F_s}_t\left(\boldsymbol{\Omega^{F_s}_{i,t}}\right) = \max_{\boldsymbol{q}^{F_s}_{i,t}} \left\{ \begin{array}{c} u^{F_s}\left(c^{F_s}_{i,t}, P^{F_s}_{i,t}, B^{F_s}_{i,t}\right) \\ +\beta E_t \left[\lambda_{t+1} \left[\begin{array}{c} (1-m_{ij,t+1}) V^{F_s}_{t+1}\left(\Omega^{F_s}_{i,t+1}\right) \\ +m_{ij,t+1} V^{F_m}_{t+1}\left(\Omega^{m}_{i,t+1}\right) \end{array} \right] \right. \\ \left. + (1-\lambda_{t+1}) V^{F_s}_{t+1}\left(\Omega^{F_s}_{i,t+1}\right) \right] \end{array} \right\}$$

subject to the intertemporal budget constraint (3), and where $\boldsymbol{q^{F_s}_{i,t}} = \{c^{F_s}_{i,t}, P^{F_s}_{i,t}, B^{F_s}_{i,t}\}$ are the choices of the single woman.

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

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

When a woman works ($P = 1$), her marginal utility of consumption changes according to ψ, which depends on whether she has a child and on her marital status. The parameter η represents the utility cost of claiming AFDC/TANF benefits.[26] The model has finite horizon and is solved by backward recursion.

The problem of a single man, which is similar to the one of a single woman but with no AFDC/TANF option, is described in Appendix C.

[25] In principle, this distribution is endogenous and as economic conditions change, the associated marriage market will change, as supply and demand change. 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.

[26] We assume there is no utility cost from claiming Food Stamps or EITC benefits because for these two programs we do not endogenize the claiming decision. When married, we assume that both partners experience a utility cost from receiving AFDC or TANF (η for the woman and η^S for the man).

3.3 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 power of each member, which in turn depends on the outside option of being single and possibly remarrying. The bargaining power of each household member is reflected in the Pareto weights $(\theta^M_{ij,t}, \theta^F_{ij,t})$, which evolve endogenously. We focus in this section on transitions between marital states and the evolution of the Pareto weights. In Appendix C, we define the full optimization problem for the married couple, given the overall state space $\boldsymbol{\Omega^m_{ij,t}}$.[27]

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 a single woman results in marriage.

For the marriage to continue, individual participation constraints need to be satisfied. The value of marriage must be larger than the value of being single for both spouses:[28]

$$\begin{aligned} V^{F_m}_{t+1}\left(\boldsymbol{\Omega^m_{ij,t+1}}\right) &\geq V^{F_s}_{t+1}\left(\boldsymbol{\Omega^{F_s}_{i,t+1}}\right) \\ V^{M_m}_{t+1}\left(\boldsymbol{\Omega^m_{ij,t+1}}\right) &\geq V^{M_s}_{t+1}\left(\boldsymbol{\Omega^{M_s}_{j,t+1}}\right). \end{aligned} \tag{5}$$

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 the value of remaining married. If one person's participation constraint is not satisfied, the Pareto weight moves the minimal amount needed to satisfy it. This is based on the dynamic contracting literature with limited commitment, such as Kocherlakota (1996) and Ligon, Thomas and Worrall (2002*a*). If there is no Pareto weight that satisfies both the spouses' participation constraints and the intertemporal budget constraint, then divorce follows.[29] Divorce can

[27] The state variables for the couple (represented by $\boldsymbol{\Omega^m_{ij,t}}$), are: the combined assets, each spouse's productivity, whether the woman worked in the previous year, the number of periods of welfare benefits utilization, age of the youngest child present $(k^a_{i,t})$, and the Pareto weights $\theta^M_{ij,t}, \theta^F_{ij,t}$.

[28] Assets are divided equally between spouses upon divorce, and there is no loss of wealth or direct financial cost of divorce in the model. This implies that divorce is an option also for very low-income household, in agreement with the notion that, empirically, divorce is more common among such household than among higher-income ones (see, for example, Lundberg, Pollak and Stearns, 2016, on the negative relationship between educational attainment and divorce).

[29] We can rewrite equation 10 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

take place unilaterally, and if it does, it is efficient because there is no allocation such that each person can have a non-negative 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 (5).[30]

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 (5)). 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 being chosen as the solutions to a symmetric Nash bargaining game between spouses, as described in Appendix C.[31] Ours is a context of imperfectly transferable utility, which implies that the Pareto weight affects the size of the gains to be shared. Since the outside option depends on the possibility of future marriages, the anticipated share of future gains will affect the probability of marriage in this indirect way.[32]

3.4 Sources of Uncertainty

Underlying the decisions of single women, single men, and couples, there are two main sources of uncertainty: wages and the marriage process. We describe those below. Appendix C describes the fertility process.

---

the women only optimize over their own utility, as shown by single men and single women having their own budget constraints.

[30] 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. When a 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.

[31] In the estimation below, we experiment with varying degrees of commitment in the marriage decision. This has some effect on parameter estimates. However, our conclusions on the effects of time limits on behavior are unchanged, as we show in Appendix D.

[32] 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.

Female Wages and Male Earnings

The wage process for men and women takes the form

$$\log(w_{j,t}^M) = a_0^M + a_1^M t + a_2^M t^2 + z_{jt}^M \tag{6}$$

$$\log(w_{it}^F) = a_0^F + a_1^F t + a_2^F t^2 - \delta^F \mathbf{1}[P_{i,t-1}^F = 0] + z_{it}^F \tag{7}$$

$$z_{jt}^M = z_{j,t-1}^M + \zeta_{jt}^M, \qquad z_{j,0}^M \sim N(0, \sigma_{0M}^2)$$

$$z_{it}^F = z_{i,t-1}^F + \zeta_{it}^F, \qquad z_{i,0}^F \sim N(0, \sigma_{0F}^2).$$

The term z_t^G $(G = F, M)$ represents the permanent income realization at time t, which evolves as a random walk following innovation $\zeta_t^G \sim N(0, \sigma_{\zeta G}^2)$. Initial dispersion in wages, reflecting heterogeneity, is distributed normally with mean zero and variance σ_{0G}^2. We do not impose correlation in the shocks to wages between men and women, but correlation in wages will arise endogenously through the marriage decision.[33] The parameter δ^F captures the depreciation of a woman's productivity when she did not work in the previous year.

While women's employment is a choice, men's employment status evolves stochastically following a Markov process. We allow for a stochastic employment shock that drives men's income to zero.[34]

Marriage

For those who are single, there is uncertainty about the marriage process. We parameterize the rate of arrival of meetings λ_t to vary with the woman's age. When two individuals meet, at time t_0 they draw an initial match quality L_{ij,t_0} from a distribution $N(0, \sigma_{\xi 0}^2)$.

For those who are married there is uncertainty about the evolution of the match value $L_{ij,t}$, which evolves according to a random walk process:

$$L_{ij,t} = L_{ij,t-1} + \xi_{ij,t} \tag{8}$$

[33] One question 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 substitutability, we do not expect important general equilibrium effects.

[34] In estimation, we also allow for transitory variation in men's earnings and women's wages by adding an i.i.d. error term. This term is not relevant to the economic decision process of the agents, but is included to account for measurement errors.

where $\xi_{ij,t}$ can be interpreted as a "love shock" to the marriage. The innovations to match quality $\xi_{ij,t}$ are drawn from a distribution $N(0, \sigma_\xi^2)$.

3.5 Claiming Benefits

In the absence of time limits, the decision to claim benefits in our model is driven by a direct financial incentive to claim offset by the utility cost of claiming. Further, there is an associated disincentive to work because AFDC benefits decline as earnings increase, which may be compensated by the depreciation of wages that occurs from not working. The decline in amount of benefits as earnings increase may discourage smaller claims because of stigma costs. These factors, together with the work disincentives of food stamps and the work incentives of EITC, determine the decision to work as well as the decision to claim AFDC. These trade-offs are particularly acute for unmarried mothers. For married mothers, the decision to claim will depend additionally on their husband's income and on the Pareto weight.

Time limits imply that claiming now reduces the availability of benefits in the future. This in turn means there will be less insurance against future fluctuations induced by wage and other shocks. Hence, claiming now may mean having to return to work at a future time of low productivity.

The strength of these incentives to claim benefits and the likely impact of time limits depend on marital status and potential transitions in marital status. For unmarried mothers, the possibility of marrying in the future, and the income stream of a potential future husband, may partially mitigate the risk of being left without benefits, and may therefore reduce their incentive to bank benefits. For married mothers, the possibility of divorce increases the risk associated with claiming benefits early, and this channel will increase the extent that benefits are banked. On the other hand, time limits worsen the outside option of married women compared to their husbands, potentially leading to greater early claiming since the utility cost of claiming benefits is primarily incurred by women.

The important point is that time limits will induce deferral of claims for both unmarried and married mothers, but the extent to which this happens depends on the marriage and divorce probabilities, respectively, and on the relative bargaining power of the spouses. The importance of marital status and marital transitions is explored quantitatively through our model.

3.6 Estimation of Model Parameters

We select the parameters of the model in two steps. First, some parameters are set using standard values in the literature or estimated directly from the data without imposing the model's structure. Second, the remaining parameters are estimated using the method of simulated moments (MSM), matching data and model-based simulated moments. We use moments based on pre-reform data and use post-reform data to validate the model.

Our choice of estimating on pre-reform data only is to avoid our parameter estimates being impacted by the other aspects of TANF and being impacted by the growth in the aggregate economy post-1996. However, this implies that when we simulate the introduction of time limits into our model, we will not explain all of the decline in welfare use or increase in employment post-1996. Instead, our interpretation of our results is that we show the impact of time limits and their interaction with marital status if there had been no improvement in employment opportunities.

3.6.1 Externally Set Parameters

Panel A of Table C.1 reports parameters taken from external sources or estimated outside of the model.

Welfare Programs Design 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.

Risk Aversion, Discounting, Interest Rate, and Economies of Scale 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 perform sensitivity analyses for different choices of each of these 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 daycare and babysitting for working women in the data (by child age) as the relevant childcare cost.

The Distribution of Characteristics of Single Men and Women Individuals in the model use the age-dependent distribution of characteristics of those currently single to form expectations about potential matches. Single men are characterized by the age-dependent distribution of $\{A_t, y_t\}$. Single women are characterized by the age-dependent distribution for $\{A_t, y_t, k_t^a, P_{t-1}^W\}$. We allow for additional mass for the cases in which $A_t^j = 0$ and $y_t^H = 0$ and model strictly-positive income and assets as a bivariate log-normal distribution.[35]

Men's Earnings Process Men's employment is an exogenous random process in the model, not involving choice, so the labor earnings process can be estimated outside of the model. We estimate the employment transition matrix directly from the SIPP data and the variance of the permanent component of log annual earnings ($\sigma^2_{\zeta^M}$) using GMM, as described in Appendix C. In Panel B of Appendix Table C.1, we report variances in the wage process. Earnings are subject to relatively high variance of permanent shocks, and initial heterogeneity is large, implying large initial dispersion in productivity. This also reflects differences in schooling among our (non-college graduate) group.

3.6.2 Method of Simulated Moments

We estimate the remaining parameters of the model by MSM (McFadden, 1989; Pakes and Pollard, 1989). The 17 unknown parameters are: the disutility from working for unmarried women without children (ψ^{0s}), married women without children (ψ^{0m}), married women with a child (ψ^{1m}), and unmarried women with a child (ψ^{1s}); the variance of match quality at marriage ($\sigma^2_{\xi_0}$); the variance of innovations to match quality (σ^2_{ξ}); the parameters characterizing the probability of meeting a partner over the life cycle ($\lambda_0, \lambda_1, \lambda_2$); the utility cost of being on welfare (η and η^S); and the parameters that govern women's lifetime wages ($a_0^W, a_1^W, a_2^W, \delta, \sigma^2_{0F}, \sigma^2_{\zeta F}$).

[35]The normality assumption may not be appropriate 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 education sample with no college degree who have lower income and assets.

Empirical moments $\boldsymbol{\phi}_{data}$ are calculated in the pre-reform period (1990-95) on women that are between age 19 and 38. 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.[36] We calculate the standard errors of our parameter estimates using the standard asymptotic formula (McFadden, 1989).[37]

Target Moments and Local Identification Our target moments for estimation can be classified into three sets: employment, marital status transitions and wages. The model fit is shown in Appendix Figures C.1 and C.2. While all moments jointly contribute to the estimation of the structural parameters, some moments are more closely tied to a particular parameter.[38] In Appendix Figures C.3 and C.4 we hold all parameters constant at the estimated values, except one parameter which varies along the horizontal axis. The vertical axis reports the values taken by the target moment that is best linked to the structural parameter of interest.[39]

The first set of moments includes the fraction of women who are employed, by marital and fertility status (shown in Appendix Figure C.1, panel a). These moments are primarily responsible for pinning down the utility cost of working ψ (which varies by marital status and presence of an eligible child). We show the impact of varying utility cost parameters in Appendix Figure C.3 panels (a) to (d). The fraction employed is monotonically increasing as the cost of working declines for each demographic subgroup. Appendix Figure C.1, panel (b) shows 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 leads to lower welfare use, keeping the other parameters constant, as shown in Appendix Figure C.4 (panels b and c).

The second set of moments includes the life cycle variation in the fraction of women who marry having been single the previous year (Appendix Figure C.1, panel c), and the fraction getting divorced for women who were married in the previous year, varying the age of their

[36] We use the full variance-covariance matrix because of the correlation between different moments.

[37] 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).

[38] The dynamics of marriage and divorce are intertwined with the savings, employment and claiming decisions, making formal identification challenging.

[39] This is a similar approach in Voena (2015) and Autor et al. (2019), for related parameters.

youngest child (Appendix Figure C.2, panel a). Both sets of moments are averaged over two years to ensure large enough sample sizes. This life cycle variation in the fraction getting married contributes to the identification of the parameters characterizing the probability of meeting a potential partner λ_t at different ages (Appendix Figure C.3 panel g and h, and C.4, panel a). A faster decline in λ_t by age increases the probability of accepting an offered match to avoid the risk of not encountering a partner in the future. The moments on marriage and divorce transitions also contribute to the variance of the initial match quality draw $\sigma^2_{\xi_0}$ and of the innovations to match quality σ^2_{ξ}: a higher initial variance increases the probability of accepting marriage and reduces the probability of early divorce while a higher variance in innovations increases the probability of divorce over time. Appendix Figure C.3 (panel e and f) show the simulated mapping from these variances to divorce probabilities at different ages of the youngest child.

The third set of moments are based on wage data. We target the age profile of women's wages from age 19 to age 53 (Appendix Figure C.2, panel b), the difference in wage growth between women who worked at time $t-2$ and women who did not (the left bars in panel c of Appendix Figure C.2), the variance of wages at the point of entry in the labor market (mid bars), and the variance of wage growth residuals (right bars).[40] These moments map directly into the profile of women's lifetime offer wages, variance of wage shocks, and wage depreciation (Appendix Figure C.4, panels d-i).

Estimated Model Parameters Table 2 reports the estimates of the structural parameters. We estimate a large disutility of work for single women (similarly to Blundell et al. (2016)). This is driven by low employment rates in the data for unmarried women, despite the absence of income from a spouse. For married women, we estimate a lower disutility of work because the presence of spousal income itself discourages women's employment.

The estimated utility cost of claiming welfare benefits is high, driven by women who are not claiming benefits despite being eligible. In the pre-reform period, there was no intertemporal tradeoff to claiming benefits, and so the model attributes the decision not to claim to the utility cost of claiming. This highlights the advantage – from an identification point of view – of using pre-reform data to estimate preference parameters. In the counterfactual simulations, for the post reform period with time limits in place, we hold constant the utility cost of claiming, but the overall cost of claiming is higher because of the loss of future

[40]The residuals are obtained from regressions of log hourly wage on a quadratic in age, and dummies for cohort, state, and race.

benefits.

The variance of innovations to match quality is two thirds of the variance of initial match love quality. The arrival rate of marital offers is relatively flat with age.

Finally, the estimates for women's hourly wage profile predict that wages increase with age but at a diminishing rate; we also estimate a small depreciation when women have not worked in the previous year, equal to 1.8%.

Table 2: Estimated Parameters

Parameter		**Estimate**	**(s.e.)**
Cost of work			
Unmarried, no children	$\exp\{\psi^{s0}\}$	0.320	(0.006)
Married, no children	$\exp\{\psi^{m0}\}$	0.352	(0.010)
Unmarried, with child	$\exp\{\psi^{s1}\}$	0.352	(0.007)
Married, with child	$\exp\{\psi^{m1}\}$	0.362	(0.005)
Cost of being on AFDC			
For direct recipient	η	0.008	(0.000)
For spouse	η^S	0.001	(0.000)
Match quality			
Variance at marriage	σ_0^2	0.006	(0.000)
Variance of innovations	σ_ξ^2	0.004	(0.000)
Probability of meeting partner by age	λ_0	0.105	(0.002)
	λ_1	-0.003	(0.000)
	λ_2	0.000	(0.000)
Woman's life cycle earnings profile	a_0^W	1.879	(0.005)
	a_1^W	0.004	(0.000)
	a_2^W	-0.000	(0.000)
Permanent wage process			
Variance at entry	σ_{0F}^2	0.125	(0.004)
Variance of innovations	$\sigma_{\zeta F}^2$	0.021	(0.000)
Human capital depreciation	δ	0.018	(0.000)

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 Equation (4) this is $\exp\{\psi\}$. The underlying parameter vector for ψ is given by $\{\psi^{s0}, \psi^{m0}, \psi^{s1}, \psi^{m1}\} = \{-1.139, -1.044, -1.044, -1.016\}$. The standard errors reported are for the underlying ψ.

4 The Effect of Time Limits on Behavior and the Role of Marriage and Divorce

In this section, we use our model to evaluate the effects of introducing time limits, focusing first on the short-run responses to the reform, to mimic what happened when the welfare reforms were enacted in 1996, as captured by our reduced-form estimates. In subsection 4.1, we compare simulated responses to the difference-in-difference estimates and show to what extent the response that we observe in the data is due to forward-looking optimizing behavior. In subsection 4.2, we examine the long-term consequences of the reform under several counterfactual scenarios, to show the role that marriage and divorce play in mitigating the effects of the reform.

4.1 Response to the Introduction of Time Limits

We simulate the introduction of time limits for women of different ages at the time when the reform took place. We match the age distribution at the introduction of time limits to the 1996 age distribution in our data, and the distribution of time limits stringency from 2 to 5 years across states. We use the simulated data to estimate short-run effects equivalent to those estimated by the difference-in-difference specification in section 2.

Table 3 reports the actual estimates in the SIPP and the estimated coefficients from the simulations.[41] The simulated difference-in-differences estimates are very close to the empirical ones. Our baseline matches the large decline in welfare utilization and the partial offsetting increase in employment among unmarried women, the decline in welfare utilization among married women, and the reduction in the fraction of women who are divorced. There is no significant effect in the data of time limits on the fraction who are married (and the sign differs between the SIPP and CPS, Table 1). In the model simulations, there is a small positive effect on marriage. Our model allows for three key sources of heterogeneity: wages, marriage market, and match quality, and these appear to replicate data features well. However, there may be other forms of heterogeneity, such as in costs of work, stigma, etc., that could play a role in explaining selection into employment or welfare. See, for example, Chan (2013).

[41] Note that the SIPP estimates differ from those reported in Table 1 because, to be consistent with the model, we estimate the reduced form difference-in-differences using annualized data and a smaller set of controls (age, marital status, and time since the reform).

In the rest of Table 3 we consider alternative specifications of our model: eliminating marriage and, hence, divorce transitions, eliminating divorce transitions only, and eliminating the possibility of accumulating assets to self-insure against shocks. For each of these specifications, we re-estimate the parameters of the model and use the new parameter estimates to simulate the introduction of time limits. We estimate our difference-in-difference regressions on the simulated data from each counterfactual.[42]

Modeling marital status transitions and asset accumulation are key to allowing our model to match the reduced-form effects. When marriage transitions are not possible, time limits induce excessively large declines in welfare use among single women because the absence of marriage shuts down that insurance channel and so benefits are banked even more. Further, time limits do not induce sufficient declines in welfare use among married women when divorce transitions are not possible. Finally, in the absence of asset accumulation, time limits induce excessive labor supply responses among single women: those coming off benefits all move into employment, in contrast to the data.

Table 3: Difference-in-Difference Estimates in the SIPP Data and in the Model

	(1) Benefits Unmarr.	(2) Benefits Married	(3) Employed Unmarr.	(4) Employed Married	(5) Divorced	(6) Married
Coeff. - SIPP data	-0.093	-0.011	0.066	0.004	-0.021	-0.003
	[-0.131,-0.055]	[-0.017,-0.004]	[0.039,0.093]	[-0.023,0.030]	[-0.041,-0.001]	[-0.023,0.017]
Time limits - baseline model	-0.148	-0.012	0.062	0.027	-0.008	0.005
Time limits in no marriage and no divorce model	-0.185		0.104			
Time limits in no divorce model	-0.195	-0.005	0.121	0.010		0.009
Time limits in no assets model	-0.146	-0.002	0.147	-0.006	-0.001	0.003

Notes: Estimates from the annualized SIPP 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.

4.1.1 Intra-household Allocations

An important question is how the reform changed the balance of power within a household. We examine the effect of the introduction of time limits on intrahousehold allocation, as summarized by the women's Pareto weight θ_t^W. We find that women's weight declines by

[42]For the "no marriage option" specification, we suppress all marriage parameters and only target non-marriage moments. Similarly, for the "no divorce option" specification, we suppress all divorce-related parameters and moments in the estimation. The "no assets" specification uses the same parameters and moments as the baseline model.

0.25 percentage points. This is driven by new marriages: weights decline by 0.14 percentage points for new marriages formed after the reform, while changes among existing couples are negligible (Appendix Table C.2). These estimates suggest that the reform had only limited implications for intrahousehold allocations on average.

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

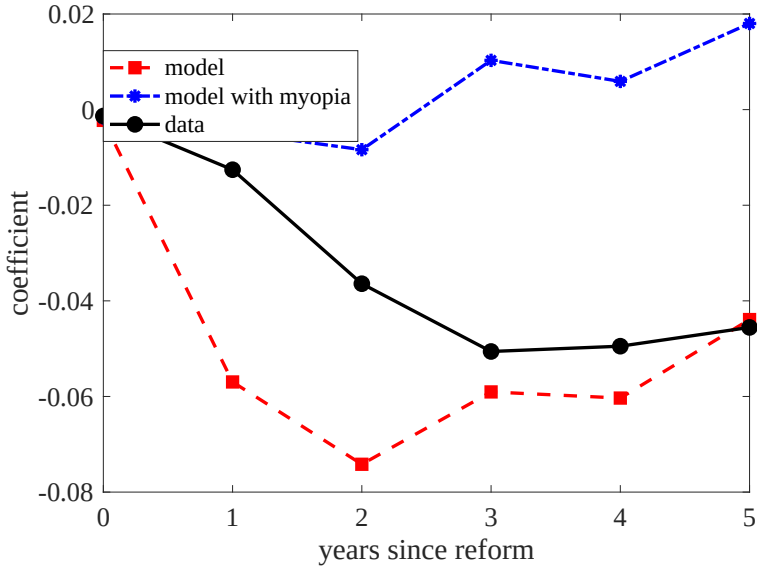

Notes: Effects estimated using data 7 years before the reform and 14 years afterwards. By *Model with myopia* we mean individuals who behave as if the introduction of time limits had not occurred (until they actually run out of benefits), but are forward-looking in terms of other behavior.

4.1.2 Myopia and Awareness of Time Limits

To address whether the documented responses reflect forward-looking behavior and the banking of benefits, 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.

4.2 Lifetime Effects of Time Limits by Productivity

The reduced-form results in the data and in the simulations show only the transitional impact of the reform. In this section, we use the model to examine the effect of time limits in the long run, comparing the simulated lifetime behavior of women who always live without time limits (the AFDC world) to the simulated lifetime behavior of the same women who always live under time limits (the TANF world). Since there might be substantial heterogeneity even in this low-income group, we break the sample down into quintiles of productivity z_{it}^F.

Panels (a) and (b) of Figure 7 show the effect of introducing a 5 year time limit on the fraction using welfare and the fraction employed. These effects are concentrated in the bottom two quintiles: the bottom quintile experiences the greatest decline in welfare use and the largest increase in employment, while the second quintile experiences effects of the same sign, but smaller magnitudes.

To mimic this split by productivity in the data, we run our reduced-form regressions separately by education, reported in Table 4. The lower the education, the greater the impact of time limits.

In the Appendix, we break down the simulated effects on welfare use and employment separately for unmarried and married mothers, holding marital composition fixed (Appendix Figures D.1 and D.2, respectively).[43] Welfare use is highly concentrated among unmarried mothers, especially in the bottom income quintile, where over 90% are using welfare in the absence of time limits. For these mothers, the introduction of time limits halves welfare use and increases employment from close to zero to around 40%. For married mothers, welfare use in the bottom quintile is lower but not negligible, with 15% using welfare. The introduction of time limits for this group halves their use of welfare but leads to only a small increase in employment.

The reduction in welfare use under time limits in Panels a and b in Figure 7 happens partly because some individuals hit the limit and are refused further assistance (a mechanical effect) and partly because some anticipate exhausting their eligibility and decide to defer claiming their benefits for a time of greater need.[44] To quantify these anticipatory effects, we examine welfare use and employment among women who have not yet used 5 years of

[43] We hold composition fixed by taking a woman's marital status under the regime with time limits and examine her behavior under both regimes.

[44] Even if the children are about to age out of eligibility the mother may still not claim the remaining benefits because to do so she may have to give up work and a higher income stream.

Figure 7: Welfare Use and Employment of Mothers by Productivity

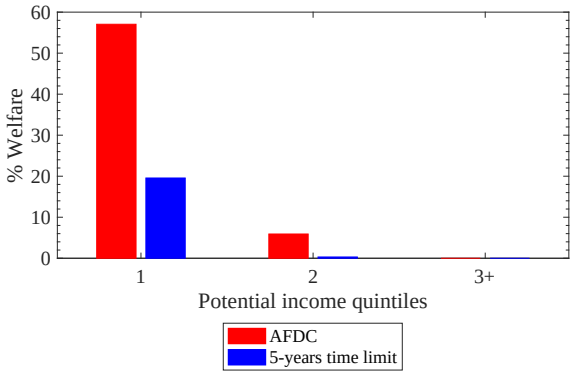

(a) Welfare use

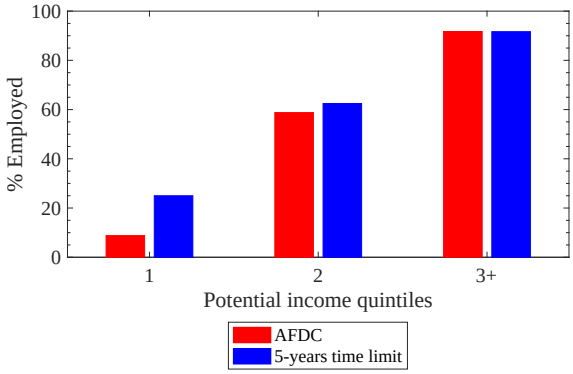

(b) Employment

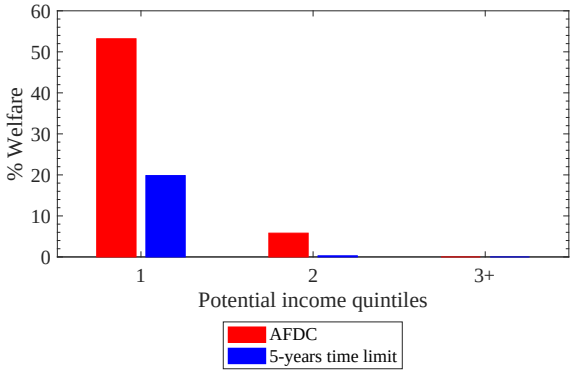

(c) Welfare use - not run out of benefits

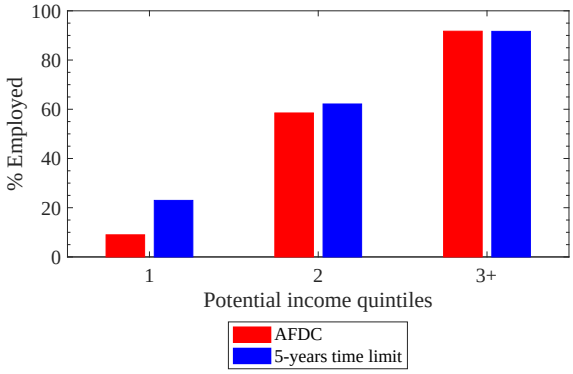

(d) Employment - not run out of benefits

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. Panels c and d are for the subset of women who, in the counterfactual with time limits, have not yet used 5 years of benefits.

Table 4: Effects of Time Limits for Women by Education Level

	AFDC/TANF Utilization			*Employment*		
Sample (Education):	Less than HS	HS degree	Some college	Less than HS	HS degree	Some college
$Exposed_{dst}Post_{st}$	-0.049***	-0.027***	-0.016***	0.048*	0.014	-0.002
	(0.014)	(0.005)	(0.006)	(0.025)	(0.016)	(0.015)
Mean pre-reform	0.244	0.088	0.048	0.420	0.658	0.730
Obs	44,797	124,712	85,118	44,797	124,712	85,118
R^2	0.18	0.09	0.07	0.13	0.09	0.10

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. We use data from the 1990-2001 SIPP panels, until 2002. Sample of female heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-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.

benefits. Changes in behavior between AFDC and time limits in this subgroup are only due to anticipatory effects. The decline in claiming and the rise in employment happens before 5 years of benefits are claimed. These anticipation effects are sizable in both the first and second quintiles of productivity (Figure 7, panels c and d). Within the second quintile, the risk of dropping further down the productivity distribution and needing benefits at that stage provides a strong incentive for banking benefits. These anticipatory effects are present for married and single mothers, shown in Appendix Figure D.3. We return in Section 4.4 below to the importance of marriage dynamics to understanding these effects and the extent of the anticipatory motives.

4.3 The Role of Work Requirements

As discussed in Sections 1 and 2.2, the welfare reform introduced not only time limits, but also work requirements. In this section we extend the model to study the implications of this feature of the reform, focusing first on the short-run effects and next on the lifetime effects.

4.3.1 Short-term Effects

Similar to the approach in Mullins (2022), we characterize work requirements as a utility cost imposed on women who claim benefits but do not work. The utility function of mothers

is rewritten as follows:

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

We consider whether the introduction of a work requirement into our model could have generated difference-in-differences estimates similar to those we document in the empirical section. We calibrate the parameter η^{WR} to match the decline in welfare utilization among single mothers following the reform, as reported in Appendix Table B.1. To make the starkest case for work requirements, we assume in this calibration that there are no time limits, so that any decline in welfare utilization can only be attributed to work requirements. As discussed above, work requirements have exemptions for mothers of very young children, which could generate heterogeneous effects among mothers with children of different ages. We therefore focus on the sample of mothers of children over age 5.

In Appendix Table D.1, we verify first in the SIPP that following the introduction of time limits, mothers of younger children (aged 6-13) experienced a larger decline in welfare use (-7%) and a larger increase in employment (+6%) relative to mothers of older children (aged 14-18), similarly to the evidence for the whole sample. Our baseline model, with time limits but no work requirements, is capable of explaining these differences (second row of Table D.1). In contrast, the model with work requirements and no time limits (third row of Table D.1) can explain the decline in welfare utilization but not the employment increase: the model predicts that mothers of younger children increase their employment *less* than mothers of older children. There is no economic mechanism in the "work requirement but no time limits" version of the model that is capable of explaining the differential impact of the reform on mothers of younger (above age 5) vs. older children, since they are all subject to work requirements.

4.3.2 Lifetime Effects

In Figure 8 we compare the lifetime effects of time limits and work requirements, analogously to Figure 7. We consider four different scenarios: a pre-1996 world (red bar), a scenario with time limits alone (blue bar), one with work requirements alone (grey bar), and one that has both time limits and work requirements (green bar).[45]

[45]As discussed above, we characterize the introduction of work requirements as an increase in the utility cost of claiming welfare only for mothers who do not work while on benefits (Mullins, 2022). We consider a substantial penalty, tripling the utility cost of being on welfare for these mothers compared to those who

The effect of introducing work requirements alone is shown by comparing the red and grey (first and third) bars within each quintile of Figure 8. Work requirements have a large effect when introduced in the absence of time limits: a 27pp fall in welfare use and a 15pp increase in employment for the bottom quintile. However, comparing the blue and green (second and fourth) bars within each quintile, shows that the effect of work requirements is much less for mothers already exposed to time limits: an additional 5pp decline in welfare use and an additional 6pp increase in employment for the bottom quintile. Comparing the grey and green (third and fourth) bars show the impact of introducing time limits if work requirements are in place.

We can draw two broad lessons from this exercise and from the comparison of empirical diff-in-diff and the simulated diff-in-diff (Table D.1). First, the introduction of time limits or the introduction of work requirements, in isolation, each have large effects on claiming and employment behavior. Second, work requirements appear to have a smaller marginal impact for explaining welfare use or employment responses relative to time limits. Third, while both features of the reform can explain, on their own, the decline in claiming and the increase in labor supply, only the model with time limits can explain the variation across demographic groups, and in particular the larger response of mothers of younger children relative to mothers of older children.

4.4 The Role of Marriage and Divorce Dynamics on the Effect of Time Limits

Our reduced form and simulation results show that single mothers are the most affected by welfare reform. Further, much of the literature estimates the effects of welfare reform focusing on the sample of single mothers (Chan, 2013). However, this section shows the important role that marriage and divorce transitions, and the expectation of these transitions, play in shaping the responses to time limits of unmarried and married women.

We perform counterfactual simulations where the marriage or the divorce option is suppressed and examine the behavior of households. In particular, we examine how the behavior of unmarried women (in the baseline model) changes if they do not expect a marriage transition, and how the behavior of married women changes if they do not expect a divorce transition.

claim benefits and work, close to what we calibrated in the previous section.

Figure 8: Welfare Use and Employment of Mothers by Productivity - Time Limits and Work Requirements

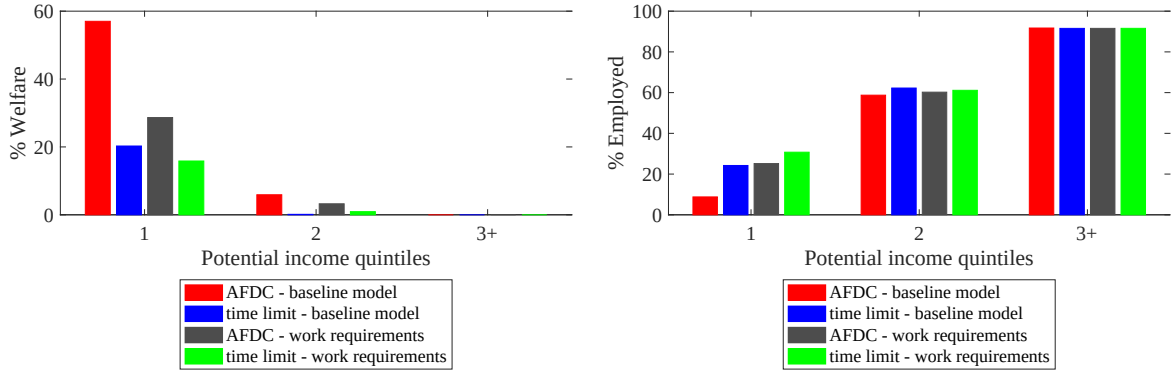

(a) Welfare use

(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}^{jk} and $j \in \{0,1\}$ and $k \in \{0,1\}$ where 0 is without any time limit or work requirements and 1 is with a 5-year time limit imposed or work requirements imposed, the bars represent $P(y_{it}^1 = 1)$ and $P(y_{it}^0 = 1)$ of mothers, pooling across i and t.

4.4.1 Eliminating Marriage Transitions

Figure 9 shows how marriage affects the likelihood that time limits will bind. Figure 9 (panel a) shows the change in the fraction that has used up that number of years of benefits over their lifetime (see Appendix Figure D.5 for the underlying levels). With time limits in place, the fraction who previously used more than 5 years of benefits (over 10% of women) would bunch at 5 years or reduce welfare use even further. An increased fraction never uses benefits, and an increased fraction use 1 to 4 years of benefits. Individuals anticipate the cut-off of benefits at 5 years and therefore do not reach the limit. Panel (a) shows that about 15% of women will face binding time limits (they would have used 6+ years of AFDC benefits had AFDC not been reformed). Absent behavioral responses, these women will bunch at the statutory 5 year time limit. However, the figure shows that almost half of the decline is absorbed by "never takers". These impacts of time limits are amplified if there is no marriage option. Almost 25% of women would now find the 5-year time limit binding. Approximately half of the decline in the share of women who face binding time limits now use less than 5 years of welfare.

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

Figure 9: Changes in Welfare Use and Time Limits with and Without a Marriage Transition Option

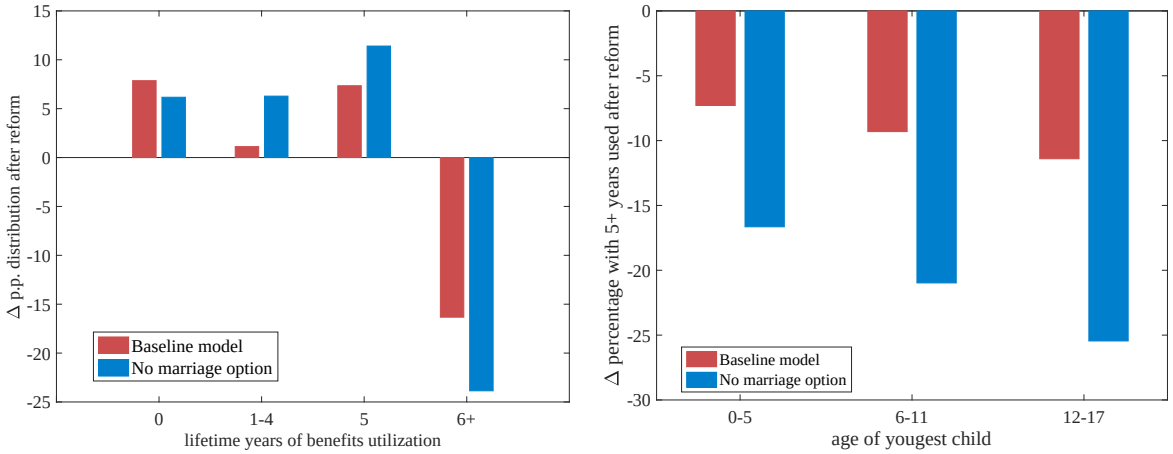

(a) Distribution of lifetime welfare utilization (b) Benefits Exhaustion by Age of Youngest Child

Notes: Panel a: Model simulations of the changes in distribution of lifetime welfare use under time limits compared to AFDC in simulated data with and without a marriage transition option. Panel b: Model simulations of the changes in the fraction of mothers who would exhaust benefits by the age of the youngest child under time limits compared to AFDC.

the maximum. To distinguish the former mechanical effect from the latter optimal-timing effect, we show first in Figure 9 (panel b) the change in how quickly mothers use up 5 years of benefits (see Appendix Figure D.5 for the underlying levels). The decline in usage occurs right from the early years of the child. 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 B.5 in Appendix B, showing that mothers of younger children cut welfare use more, which is the premise of our difference-in-difference strategy. Eliminating the prospect of marriage again leads to even more dramatic banking of benefits.

Further, Figure 10 plots changes in welfare use and employment level by single mother-ss between time limits and AFDC, with and without the option of marriage, focusing on mothers who have not yet run out of benefits (anticipatory effect).[46] The key takeaway of Figure 10 is that the declines in welfare use and the rise in employment caused by time limits are greater if there is no possibility of marriage. As can be observed from Appendix Figure D.6, a substantial decline in welfare use that arises when there is no marriage option occurs because of anticipatory effects: when there is no prospect of getting married in the future,

[46]This figure controls for 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.

individuals have to bank their benefits and return to work. The difference in this effect is strongest in the second-lowest quintile of the productivity distribution: anticipatory effects for this subgroup are approximately twice as large when there is no marriage option.

Figure 10: Anticiptory Effect on Welfare Use and Employment of Unmarried Mothers with and without a Marriage Transition Option

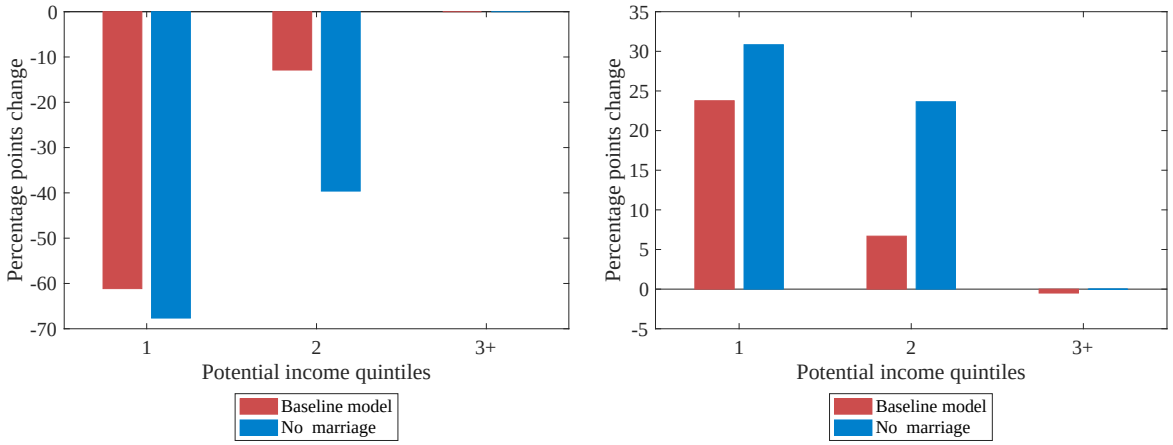

(a) Welfare (b) Employment

Notes: Simulated means of lifetime welfare use and employment marital status by policy regime and marriage option for unmarried women. 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 black 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 report 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 a marriage transition, but only average the choices of women i who, at time t, would be unmarried if marriage were an option, eliminating selection effects. Levels reported in Appendix Figure D.6.

4.4.2 Eliminating Divorce Transitions

Figure 11 shows the anticipatory effects of time limits on welfare use and employment by married women when there is no possibility of divorce. The first bar for each quintile shows that married mothers cut welfare use and increase employment in anticipation of hitting time limits. However, the second bar shows that these effects are mitigated in the absence of divorce, and there is less banking of benefits by married women, who no longer hold back the

use of benefits against the possibility of a future divorce.[47] This discussion highlights that conclusions about the effects of welfare reform depend in essential ways on whether marriage and divorce choices are part of the forward-looking decision-making process.

Figure 11: Welfare Use and Employment of Married Mothers With or Without a Divorce Transition Option

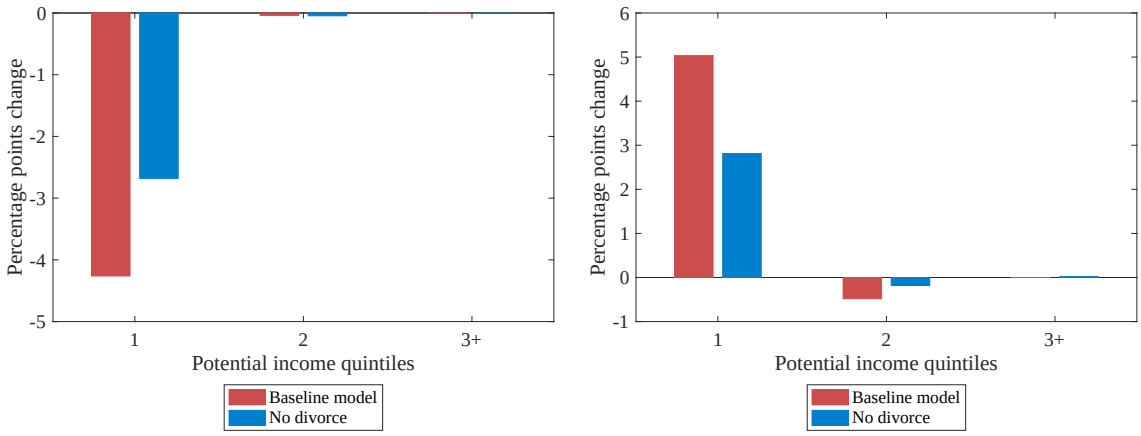

(a) Welfare use

(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 a divorce transition, but only average the choices of women i who, at time t, would be married if divorce were an option.

4.5 Robustness to modeling assumption about the degree of commitment in marriage

We examine the robustness of the model implications to alternative specifications of the degree of commitment among married couples. In particular, we re-estimate the model under the assumptions of full commitment and of no commitment (Pollak, 2019). Under full commitment, spouses choose a Pareto weight by symmetric Nash bargaining at the time of marriage and cannot change it while they remain married. Under no commitment, couples set the Pareto weight *in each period* by symmetric Nash bargaining. Figure D.4 shows that

[47] See Appendix Figure D.7 for the underlying levels of welfare use and employment that give rise to these effects. This figure abstracts from composition effects because the differences are computed of the simulated sample of women who are married in both the baseline model and the "no divorce" counterfactual, even when the divorce option is removed.

the commitment assumptions have a very small impact on the estimated effects of time limits on the behavior of married women.

5 Implications of Welfare Reform for Lifetime Utility

We now evaluate the consequences of time limits for lifetime utility. We first introduce time limits without adjusting taxes or benefits, effectively reducing resources available to the population of non-college graduates. As described in Appendix E, we compute a consumption equivalent measure, separately for men and women. This represents the consumption compensation that agents would need to receive, under time limits, to achieve the same utility level they achieve without time limits.

The introduction of time limits reduces government spending. We need to consider redistributing this saving to make the time limits revenue neutral. We therefore compute the consumption equivalent measures under two alternative methods of redistributing the government surplus. One is a lump-sum cash transfers to single mothers, and the other a negative payroll tax to all workers, male and female. These welfare calculations likely underestimate the cost on well-being (utility) of time limits because they do not take into account the effects of limiting benefits on the children (see Mullins, 2022; Bruins, 2017).

An important issue in our analysis is understanding whether the value of TANF is coming from redistribution or from insurance. The conceptual difficulty of contrasting the level and insurance effects of marriage and self-insurance is that level differences (at least in our model) arise because of the realization of negative shocks from the beginning of working life. This means that the distinction between redistribution towards those with low income and providing insurance to those who have bad shocks is blurred. We deal with this by looking at expected utility from the perspective of the beginning of the first period, and we calculate the consumption equivalent of shutting down time limits using this perspective.

5.1 Uncompensated Lifetime Utility Cost of Time Limits

Imposing time limits leads to a per-capita expected financial loss of benefits equal to 0.45% of the lifetime aggregate consumption of women in our model. Since benefits are accrued in states of high marginal utility of consumption, the consumption equivalent is equal to 0.72% (Figure 12, columns a and b).

When we eliminate the possibility of marriage from the model, we find that the financial

loss in benefits is equal to 0.75% of the lifetime aggregate consumption of women. Correspondingly, we find that women under time limits, would require a compensation equal to 0.9% of lifetime consumption.

Men are basically unaffected by time limits when the benefits loss is uncompensated: the direct loss of benefits hurts them in marriage, but the gains in intrahousehold allocations benefit them, leading to a negligible negative consumption equivalent in the baseline mode, and 0 when marriage is not an option.

5.2 Revenue Neutral Lifetime Utility Cost of Time Limits

5.2.1 Lump-Sum Payment to Single Mothers

One way to redistribute the government savings that arise from imposing time limits is through a lump-sum transfer to single mothers, independent of employment status. We consider the welfare implications of such a redistribution scheme in Figure 12, columns (c) and (d). The negative sign means women prefer a lump-sum annual transfer to all single mothers alongside the introduction of time limits, rather than maintaining the AFDC regime with no time limits but no lump sum transfer. On the other hand, men prefer the AFDC because the lump-sum payment improves women's outside option to marriage. Women would need to give up 0.83% of their lifetime consumption to achieve the same utility as without time limits and without the cash transfer, while men would need to receive a 0.29% consumption compensation for the loss of marriage prospects and intrahousehold allocation. With no marriage option, men are indifferent with respect to time limits, while women experience even larger utility gains from experiencing time limits and receiving the government saving via a cash transfer for single mothers, with a consumption equivalent equal to -1.25%.

Note that we have assumed no utility cost associated with the receipt of this cash grant, which is reasonable for universal benefits. Thus, part of the welfare improvement is due to the removal of any stigma effects from receiving payments in this way.

5.2.2 Negative Payroll Tax

The second redistribution scheme we consider is through the labor income tax system. In the baseline model, the payroll tax rate on labor is set to 0, and hence we let the government run a deficit to fund AFDC and EITC for eligible mothers, and food stamps for all eligible

Figure 12: Lifetime Utility Costs Of Time Limits on Women and Men by Government Savings Redistribution Method

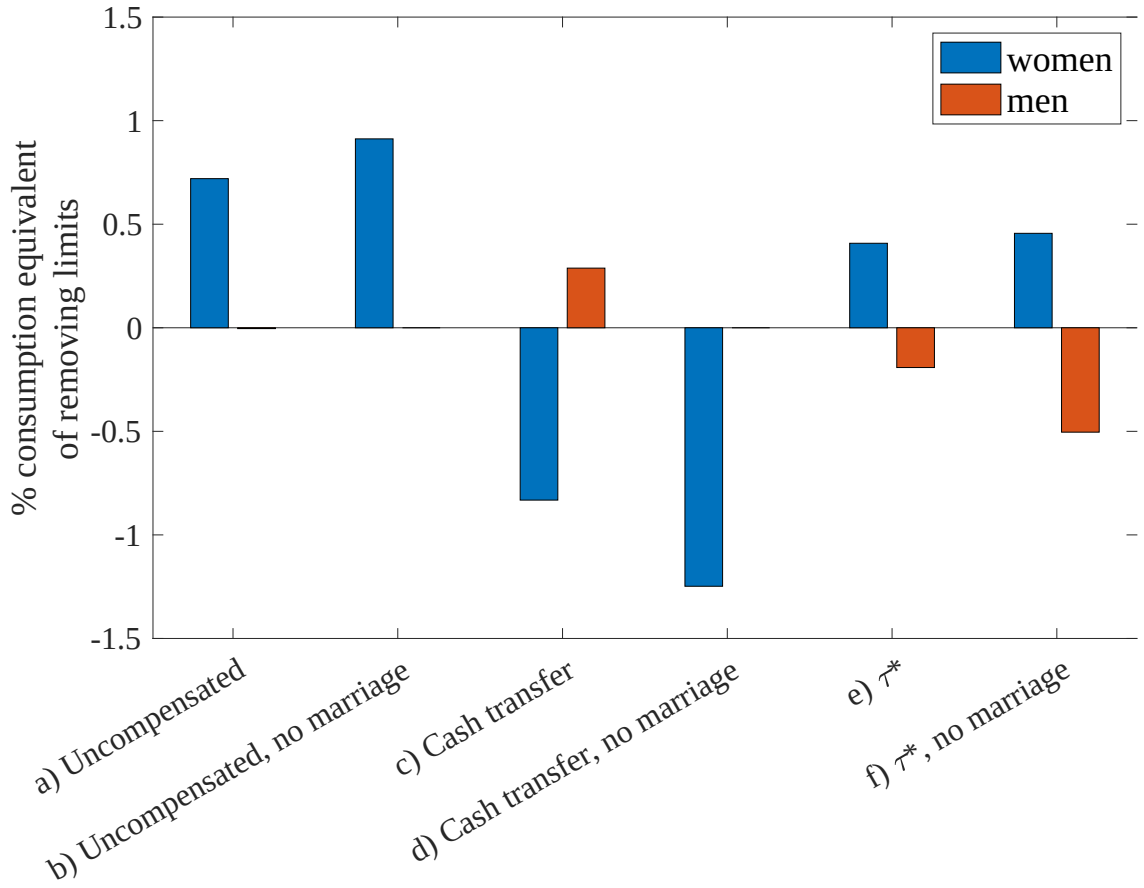

Notes: Simulated consumption equivalents of AFDC without time limits relative to AFDC with a 5-year time limit, under different assumptions for revenue neutrality achieved through a cash transfer to single mothers or through a negative payroll tax on women's and men's earnings. The consumption equivalent is percentage of annual consumption over a lifetime that an agent requires under time limits to achieve the same utility as without time limits. See Appendix E for a detailed description of how consumption equivalents are computed. The "no marriage" case is computed by eliminating the marriage option and re-estimating all relevant parameters accordingly. The simulations are based on estimates from a dataset of non-college graduates.

agents. When introducing time limits, to hold the government budget deficit constant, we adjust a negative proportional payroll tax on labor earnings, τ^*, such that the present discounted value of net revenue flows remains constant with and without time limits.

Once government savings are rebated to the sample population (women and men without a college degree) through a negative payroll tax, women still require a compensation equal to 0.4% their lifetime consumption, in the baseline model, to achieve the same utility as pre-reform, and around 0.5% when the marriage option is suppressed. Despite the loss of benefits of their partner, men are better off with time limits and lower taxes. However, considering men and women, the overall willingness to pay to remove time limits is positive in our baseline model: time limits reduce aggregate well-being.

Our conclusion is that while the 1996 reform may have achieved its stated goals of reducing welfare dependence, emphasizing self-sufficiency through work and reducing divorce, the effects on lifetime expected utility depends very much on how the government savings have been spent. Tax reductions transfer resources from the state of the world in which the marginal utility of consumption is low to states where it is high, and generally lead to a utility loss for women, and gain for men. Redistributing through universal benefits for single mothers leads to welfare gain for women, but a loss for men. The post-1996 era was characterized by negligible support for UBI-style programs, an expansion of EITC (including state-options), and a reduction in marginal tax rates on wages, including at the bottom.[48] While these programs were not necessarily funded out of savings generated by the welfare reform, the policy landscape of the post-1996 period has been closer to the tax reduction experiment than to the conditional lump-sum experiment considered in this section, suggesting that for low-skill women the overall impact was likely negative.

Our simulations also indicate that the utility cost of time limits on women become larger as marriage becomes less likely. With marriage rates declining in recent decades (Stevenson and Wolfers, 2007), declines in the generosity of welfare programs for low-income mothers become more costly in terms of lifetime utility.

5.3 Sources of Insurance

To benchmark the importance of marriage as an insurance mechanism, we examine counterfactuals in which another key source of consumption insurance, the possibility of self-

[48]For individuals earning an amount equal to the poverty line, unadjusted gross tax liability after deductions but before any credits fell, in real terms, from $412 to $199 between 1992 and 2003.

insurance through saving, is shut down in the model. We contrast the consumption equivalent of having (a) no access to marriage, with (b) no access to self-insurance, and with (c) no access to either.[49] If there were no restrictions on borrowing and no frictions in the marriage market, we would expect to estimate a lower benefit for government welfare programs. There will still be value if the government provided insurance against permanent shocks (a key component of risk in our model).

Figure 13 reports the percentage of consumption that women gain upon the removal of time limits in an uncompensated scenario (as in the first bar of Figure 12 above), under different assumptions about agents' ability to save. Figure 12 showed that when both marriage and saving are possible, the consumption equivalent to remove time limits is 0.72%, rising to 0.91% with no marriage. Figure 13 shows, similarly, that the consumption equivalent of eliminating time limits also increases to 0.86% when individuals cannot self-insure through saving but when marriage is an option. When neither self-insurance through saving nor marriage is an option, the consumption equivalent rises to 1%. These results show the value of marriage as mitigating the cost of time limits is similar to the value of being able to self-insure through saving. Nonetheless, even with self-insurance and marriage, there are substantial welfare losses from introducing time limits.

6 Conclusions

We provide evidence on the impact of time limits on the various margins of behavior targeted by the 1996 welfare reform, and offer a life-cycle framework to understand what theoretical forces explain the evidence, focusing in particular on marriage and divorce transitions. Analyzing the impact of time limits focusing only on single women can exaggerate the impact and cost of time limits. This is because women in our sample who are continuously single over a prolonged period are often heavily reliant on welfare and are a highly selected group.

Our empirical analysis shows that the introduction of time limits led to immediate and sizable reductions in welfare participation, to increased employment, and to a decline in divorce rates. These responses emerge well before households become constrained by time limits, consistent with forward-looking "banking" of benefits in anticipation of future periods

[49] We focus on shutting down self-insurance (and in particular savings) rather than changing the ability to borrow because borrowing in our model is already "naturally constrained" by the infeasibility of dying in debt (Aiyagari, 1994).

Figure 13: Lifetime Utility Costs Of Time Limits on Women by the Availability of a Savings Technology

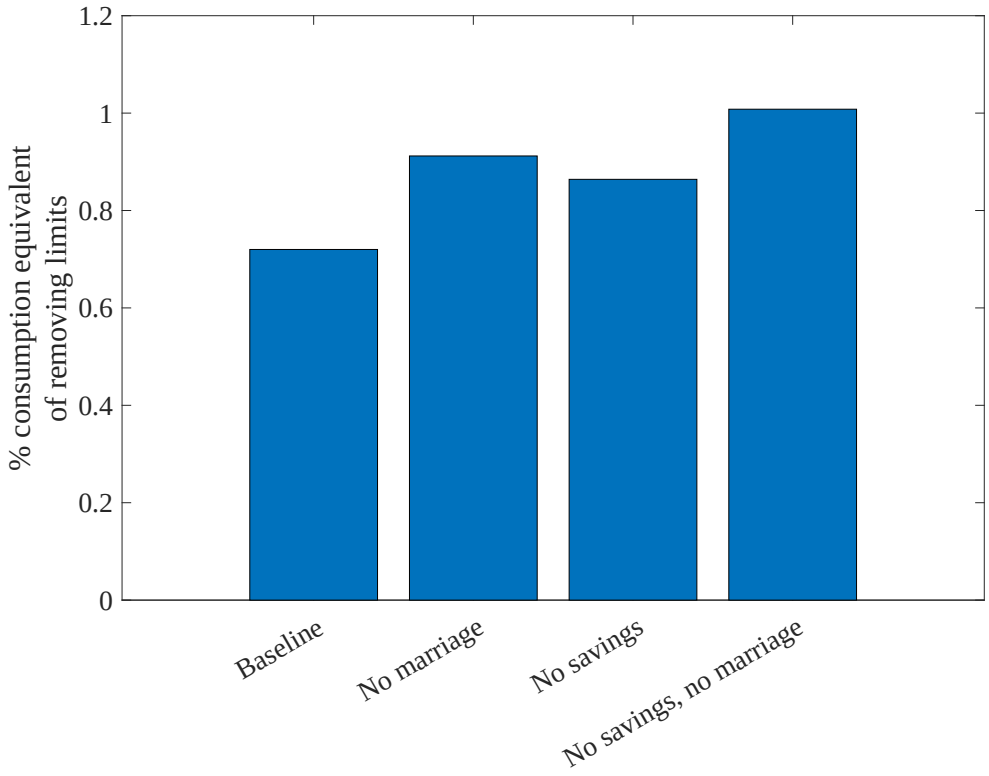

Notes: Simulated consumption equivalents of AFDC without time limits relative to AFDC with a 5-year time limit, under different assumptions about savings. The consumption equivalent is percentage of annual consumption over a lifetime that an agent requires under time limits to achieve the same utility as without time limits. See Appendix E for a detailed description of how consumption equivalents are computed. The "no marriage" and "no savings" cases are computed by eliminating the marriage option and re-estimating all relevant parameters accordingly. The simulations are based on estimates from a dataset of non-college graduates.

of greater need. These outcomes are clearly interconnected and this underpins building a model to capture the complex interactions between marital status (and changes thereof) and welfare use.

We thus estimate a rich life-cycle model that jointly incorporates choices over labor supply, savings, welfare use, and marital transitions. The model replicates the short-run empirical responses. Further, time limits reduce welfare utilization throughout the life cycle, increase employment among unmarried mothers, and decreases divorce by lowering the value of being single relative to marriage.

Allowing for marital transitions permits a more complete accounting for risk and insurance. Focusing on single mothers ignores the fact that some women are only temporarily single and that most of them transition into marriage, mitigating the adverse effects of negative income shocks and thus reducing the need for welfare benefits. On the other hand, some married women become single due to divorce, facing an environment with less support and greater need for welfare which justifies deferral of benefits. We show that changing the expectations of these transitions has important behavioral consequences: in a world where single women had no marriage opportunities, there would be much more banking (and labor participation); symmetrically, in a world where married women had no exposure to divorce risk, the effects would move in the opposite direction.

These counterfactual analyses suggest that accounting for marriage and divorce is crucial even if divorce and marriage rates were not affected directly by the introduction of time limits: marital status matters not only as an outcome, but also as a conditioning state. Our conclusion is that these effects of marital status and marital transitions change the way we evaluate welfare reform: the utility cost of welfare reform is reduced when we account for these dimensions of insurance. This changes how we should think about the importance of the welfare system in societies that now see substantial shifts in marriage and divorce rates.

While the success of welfare reforms is often judged by the decline in welfare rolls or the increase in labor force participation they generate, the effects on well-being are less obvious and can be very heterogeneous. We show how the welfare consequences of time limits depend on how the resulting government savings are redistributed. When savings are returned through lower payroll taxes, women lose because resources shift away from states of the world with high marginal utility of consumption, while men tend to gain. In contrast, redistributing savings through lump-sum transfers to single mothers generates welfare gains for women but losses for men.

Our model does not explain the full decline in welfare utilization or the increase in employment that occurred after the 1996 reform. The two main additional contributors were the imposition of work requirements and the post-1996 boom (Ziliak et al., 2000). Instead, our interpretation of our results is that we show the impact of time limits and their interaction with marital status if there had been no sharp rise in employment opportunities or constraints.

The introduction of time limits induced large falls in payments to families, particularly those with young children (Grogger and Michalopoulos, 2003; Mullins, 2022; Bruins, 2017). However, in our welfare calculations, we do not capture the long run effects on children of these cuts. Nonetheless, the introduction of time limits reduced lifetime utility, despite the self-sufficiency brought about by increasing employment. The extent of the utility cost depends on individual productivity (driving potential earnings), on how the revenue that is saved is redistributed, and on expectations of future marital transitions.

References

Aiyagari, S. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *The Quarterly Journal of Economics*, 109(3): 659–684.

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.

Bane, Mary Jo, and David T Ellwood. 1996. *Welfare realities: From rhetoric to reform.* Harvard University Press.

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.

Bloom, Dan, Mary Farrell, and Barbara Fink. 2002. "Welfare Time Limits: State Policies, Implementation, and Effects on Families."

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-André, J. Ignacio Gimenez-Nadal, José Alberto Molina, Alexandros Theloudis, and Jorge Velilla. 2020. "Intrahousehold Commitment and Intertemporal Labor Supply." *IZA Discussion Papers 13545.*

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 Dan Silverman. 2009. "Time-Inconsistency and Welfare Program Participation: Evidence from the NLSY." *International Economic Review*, 50(4): 1043–77.

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.

Huang, Chien-Chung, Irwin Garfinkel, and Jane Waldfogel. 2004. "Child support enforcement and welfare caseloads." *Journal of Human Resources*, 39(1): 108–134.

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. 2019. "Household sharing and commitment: Evidence from panel data on individual expenditures and time use." *The Review of Economic Studies*, 86(5): 2184–2219.

Lundberg, Shelly, Robert A Pollak, and Jenna Stearns. 2016. "Family inequality: Diverging patterns in marriage, cohabitation, and childbearing." *Journal of Economic perspectives*, 30(2): 79–102.

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. 2007. "Welfare use when approaching the time limit." *Journal of Human Resources*, 42(3): 596–618.

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.

Meyer, Bruce D, and Dan T Rosenbaum. 2001. "Welfare, the earned income tax credit, and the labor supply of single mothers." *The quarterly journal of economics*, 116(3): 1063–1114.

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. 2020. "Welfare rules, incentives, and family structure." *Journal of Human Resources*, 55(1): 1–42.

Moffitt, Robert A, Robert Reville, and Anne E Winkler. 1998. "Beyond single mothers: Cohabitation and marriage in the AFDC program." *Demography*, 35(3): 259–278.

Mullins, Joseph. 2022. "Designing Cash Transfers in the Presence of Children's Human Capital Formation."

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*.

Pollak, Robert A. 2019. "How bargaining in marriage drives marriage market equilibrium." *Journal of Labor Economics*, 37(1): 297–321.

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.

Schoeni, Robert F, and Rebecca M Blank. 2000. "What has welfare reform accomplished? Impacts on welfare participation, employment, income, poverty, and family structure."

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. 2015. "Temporary assistance for needy families." In *Economics of Means-Tested Transfer Programs in the United States, Volume 1*. 303–393. University of Chicago Press.

Ziliak, James P, David N Figlio, Elizabeth E Davis, and Laura S Connolly. 2000. "Accounting for the decline in AFDC caseloads: Welfare reform or the economy?" *Journal of Human Resources*, 570–586.

Marriage, Labor Supply and the Dynamics of the Social Safety Net

Online Appendix

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

Appendix A: AFDC, TANF and the Welfare Reform

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

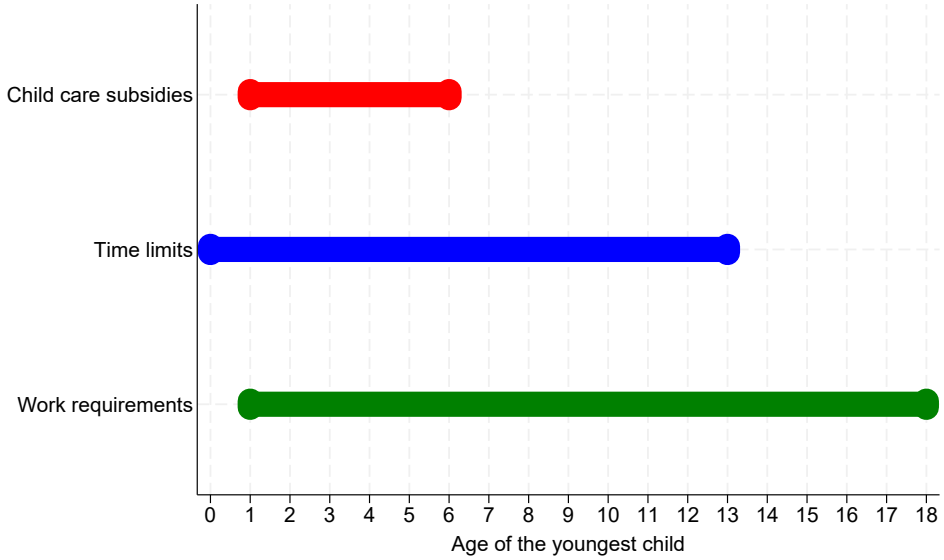

Figure A.1: Elements of the reform, by age of the youngest child

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. Figure A.1 shows which elements of the reform applied to women as a function of the age of their youngest child (assuming the statutory 5-year time limits).

Welfare and Marital Status

Welfare-part-Unemp

Figure A.2: Distribution of Welfare Participants by Marital status

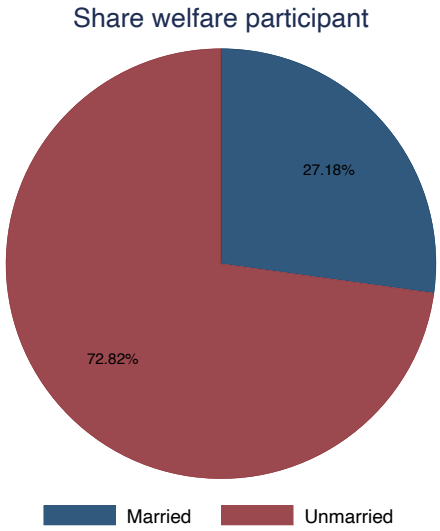

Notes: We use pre-1996 data from SIPP panels. Sample of female heads of household who are not college graduates and have children aged 18 and below.

Figure A.3: Marital Status Transitions by Women's Age

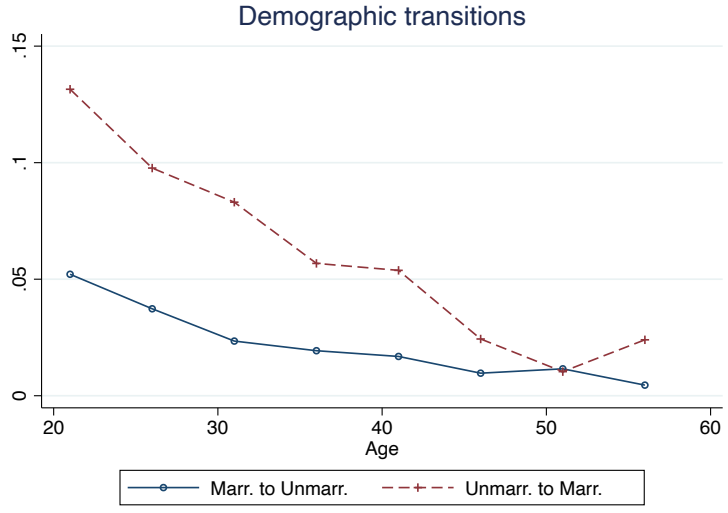

Notes: We use pre-1996 data from the SIPP panel. Sample of female heads of household who are not college graduates and have children aged 18 and below. We plot the probability of transitioning into different marital status from one year to the next for women in the same 5-year age window.

Figure A.4: Welfare Participation Before and After Unemployment Shock

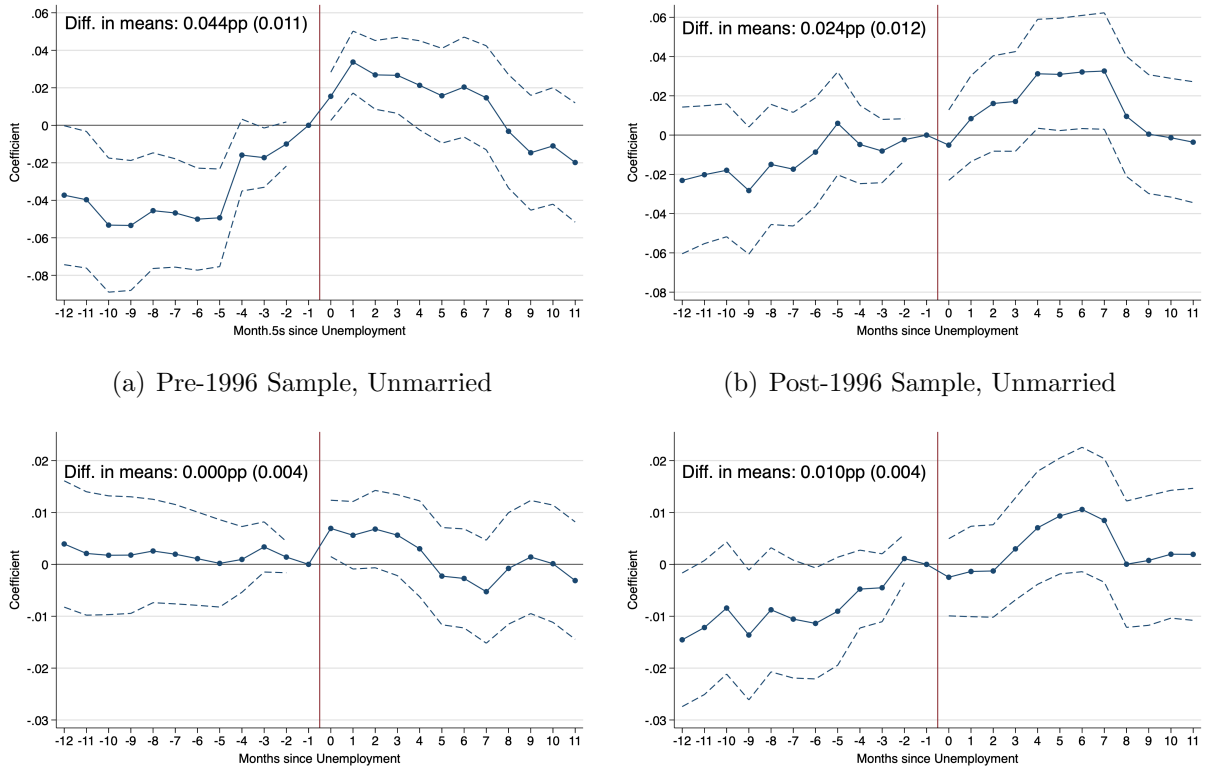

(a) Pre-1996 Sample, Unmarried

(b) Post-1996 Sample, Unmarried

(c) Pre-1996 Sample, Married

(d) Post-1996 Sample, Married

Notes: 95% confidence intervals based on standard errors clustered at the person level. Data from the 1990-2001 SIPP panels; thus, until the year 2003. Sample of female heads of household who are not college graduates, have children aged 18 and below, and are married. In panel (a), we plot the coefficients for AFDC participation controlling for the months from the start of unemployment for the pre-reform period and with individual FEs (**N = 51,419; Unique women = 3,504**). In panel (b), we plot the coefficients for TANF participation controlling for the months from the start of unemployment for the post-reform period and with individual FEs (**N = 28,486; Unique women = 1,968**).

Appendix B: Reduced-Form Evidence on Time Limits

Data and Descriptive Statistics

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.[50]

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: We use data from the 1990-2001 SIPP panels, until 2002, and the 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.

[50]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.

Empirical Strategy

Figure B.1: Variation Exploited in Time Limits Across States (2000)

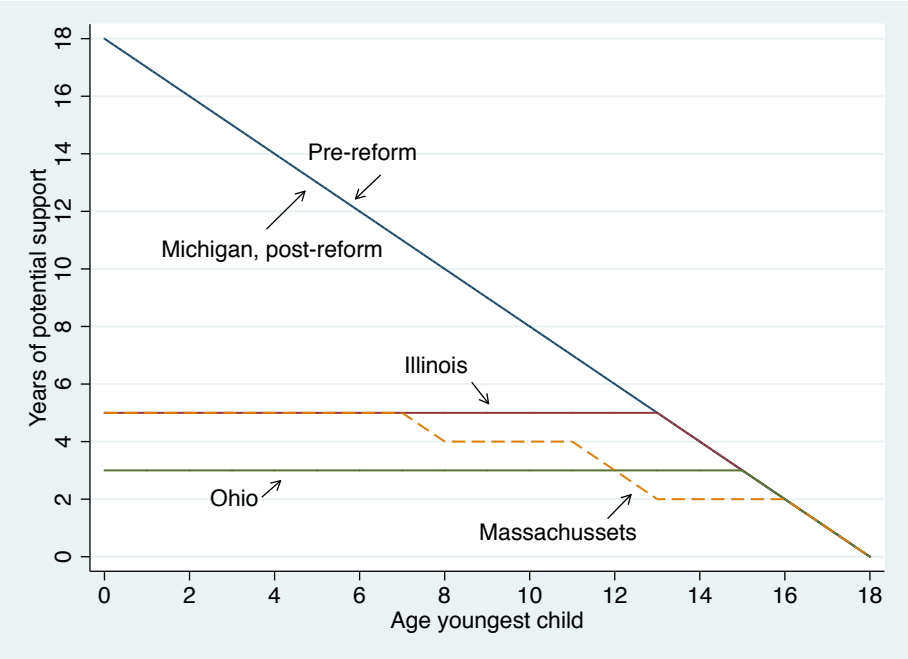

Additional Empirical Analysis

Controlling for Husband's Unemployment

Table B.2: Effects of Time Limits for Married Women by Spouse Employment Status

	AFDC/TANF		*Employment*	
Sample:	Spouse Employed	Spouse Unemployed	Spouse Employed	Spouse Unemployed
$Exposed_{dst}Post_{st}$	-0.008***	-0.040**	-0.006	0.022
	(0.002)	(0.016)	(0.015)	(0.041)
Mean pre-reform	0.015	0.190	0.660	0.526
Obs	166,001	16,374	166,001	16,374
R^2	0.05	0.28	0.10	0.27

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. We use data from the 1990-2001 SIPP panels, until 2002. Sample of female married heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-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.

Long-run Results

In Table B.3, we estimate the baseline results for all outcomes on a longer-run pool of SIPP and CPS dataset; we use data until the 2004 SIPP panels (until year 2007) and 2007 CPS.

Table B.3: Long-run Effects of Time Limits on Benefits, Employment, Marital Status, and Fertility

	AFDC/TANF Utilization					
	Whole sample		Married women		Unmarried women	
	SIPP	CPS	SIPP	CPS	SIPP	CPS
$Exposed_{dst}Post_{st}$	-0.038***	-0.022	-0.013***	-0.005	-0.108***	-0.111***
	(0.004)	(.)	(0.002)	(.)	(0.012)	(0.010)
Mean pre-reform	0.101	0.077	0.034	0.019	0.303	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					
	Whole sample		Married women		Unmarried women	
	SIPP	CPS	SIPP	CPS	SIPP	CPS
$Exposed_{dst}Post_{st}$	0.007	-0.014	-0.014	-0.031	0.055***	0.053**
	(0.011)	(.)	(0.014)	(.)	(0.013)	(0.021)
Mean pre-reform	0.642	0.648	0.647	0.655	0.625	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
	Divorce/Separated		*Married*		$Newborn_{t+1}$ (*Annual*)	
	Whole sample		Whole sample		Whole sample	
	SIPP	CPS	SIPP	CPS	SIPP	CPS
$Exposed_{dst}Post_{st}$	-0.033***	-0.013	-0.002	-0.014	-0.004	-0.001
	(0.009)	(.)	(0.011)	(.)	(0.005)	(0.003)
Mean pre-reform	0.151	0.126	0.752	0.793	0.058	0.049
Obs	336,129	153,498	336,129	153,498	74,215	66,685
R^2	0.03	0.01	0.05	0.05	0.08	0.04

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. We use data from the 1990-2004 SIPP panels, until 2007, and the 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.

Dynamic Effects of Time Limits

Figure B.2: 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. We use data from the 1990-2001 SIPP panels, until 2002, and the 1990-2002 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 B.3: Employment Dynamics

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

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

Notes: We use data from the 1990-2001 SIPP panel, until 2002, and the 1990-2002 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 B.4: Marital Status Dynamics

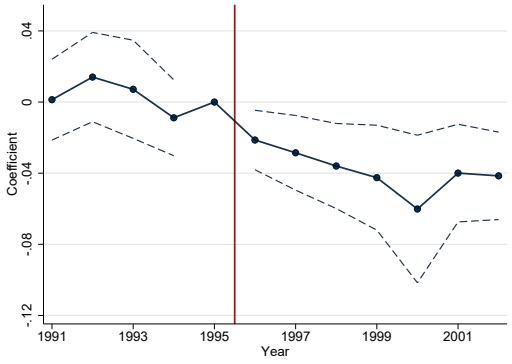

(a) SIPP: Divorce/Separation

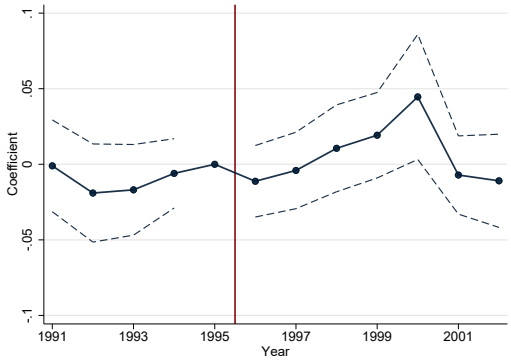

(b) SIPP: Marriage

(c) CPS: Divorce/Separation

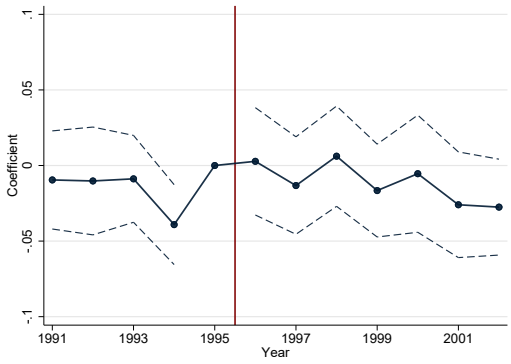

(d) CPS: Marriage

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. We use data from the 1990-2001 SIPP panels, until 2002, and 1990-2002 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.

Effects By Age of Youngest Child

Figure B.5: Program Participation and Employment Dynamics by Child Age

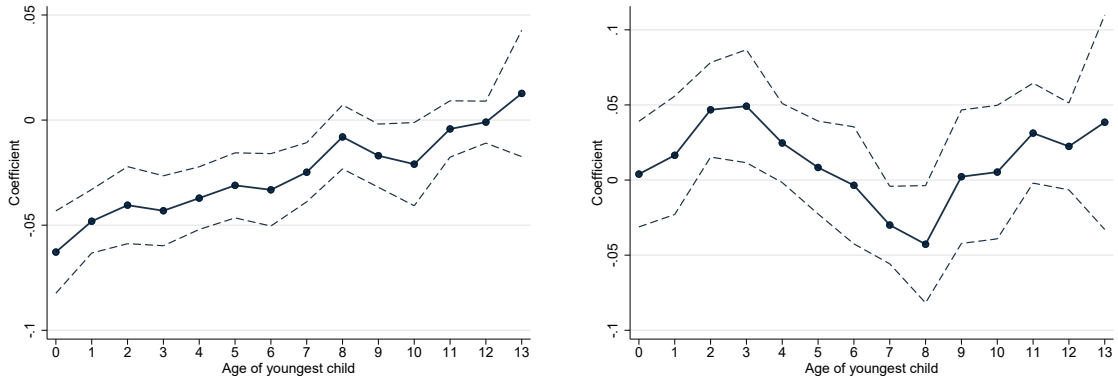

(a) Welfare participation

(b) Employment

Notes: We use data from the 1990-2001 SIPP panels, until 2002. 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.

Joint Employment and Welfare Use Status

Table B.4: 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
$Exposed_{dst}Post_{st}$	-0.022***	0.072***	-0.064***	0.015
	(0.007)	(0.015)	(0.014)	(0.013)
Mean pre-reform	0.062	0.562	0.244	0.132
Obs	66,144	66,144	66,144	66,144
R^2	0.07	0.23	0.24	0.10
			CPS	
Dependent Var:	Employed On Welfare	Employed Not on Welfare	Not Employed On Welfare	Not Employed Not on Welfare
$Exposed_{dst}Post_{st}$	-0.005	0.059**	-0.080***	0.025
	(0.009)	(0.023)	(0.012)	(0.028)
Mean pre-reform	0.067	0.554	0.237	0.142
Obs	23,606	23,606	23,606	23,606
R^2	0.04	0.14	0.15	0.05

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. We use data from the 1990-2001 SIPP panels, until 2002, and the 1990-2002 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.

Robustness Checks

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

Excluding Mothers of Young Children

A potential concern is that our results are driven by changes in the behavior of households with small children after welfare reform as a result of the more generous childcare provisions in the PRWORA.[51] Appendix Table B.5 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. However, this is unlikely to represent a significant bias given the size of the population exempted.

Additional Features of Welfare Reform: Work Requirements, Child Support, and Waivers

There were additional aspects to the welfare reform of 1996, particularly the introduction of work requirements and the strengthening of child support enforcement. The state and time fixed effects, and their interaction will capture the common effect of these additional aspects on all mothers irrespective of the age of the youngest child. Our identification of the effect of time limits depends on differences in behaviour between mothers of young and old children. Nonetheless, the additional aspects of reform could confound our core estimates if the impact differs by age of the youngest child.

Work Requirements To study potential age-specific effects of work requirements that are not captured by the state-by-year fixed effects, we consider an specification that includes

[51] 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.

Table B.5: Women with Children above Age 5

			SIPP			
Dependent Var:	**AFDC/TANF**		**Employed**		**Div/Sep**	**Married**
Sample:	Whole	Unmarried	Whole	Unmarried	Whole	Whole
$Exposed_{dst}Post_{st}$	-0.022***	-0.057***	0.007	0.039**	-0.029***	0.011
	(0.005)	(0.017)	(0.013)	(0.016)	(0.010)	(0.009)
Mean pre-reform	0.067	0.198	0.717	0.713	0.180	0.745
Obs	141,741	38,137	141,741	38,137	141,741	141,741
R^2	0.09	0.21	0.10	0.18	0.03	0.07
			CPS			
Dependent Var:	**AFDC/TANF**		**Employed**		**Div/Sep**	**Married**
Sample:	Whole	Unmarried	Whole	Unmarried	Whole	Whole
$Exposed_{dst}Post_{st}$	-0.007**	-0.046***	0.003	0.044	-0.009	-0.011
	(0.003)	(0.012)	(0.011)	(0.029)	(0.007)	(0.010)
Mean pre-reform	0.055	0.205	0.714	0.705	0.156	0.777
Obs	66,129	15,088	66,129	15,088	66,129	66,129
R^2	0.05	0.14	0.07	0.14	0.02	0.04

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. We use data from the 1990-2001 SIPP panels, until 2002, and 1990-2002 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.

controls for stringency of work requirements policies for each state interacted with the after-policy time variable $Post_{st}$ and dummies for whether age of youngest child is less than 1 and 5 years old.

We classify each state into two of six categories as described in Table B.6, which encompass both sanctions (lenient, intermediate, and stringent), and age exemptions (lenient, intermediate, and stringent). We then group states into three categories that encompass a similar number of states: overall lenient, overall intermediate, and overall stringent. Overall lenient states are lenient on both dimensions or intermediate on one dimension and lenient on the other. Overall intermediate states are intermediate on both dimensions or stringent on one and lenient on the other. Overall stringent states are stringent on both dimensions or stringent on one and intermediate on the other.

Table B.6: Description of State-level Work Requirements Index

	Sanctions	**Age exemption**
Lenient	Partial or no sanction	Over 12 months
Intermediate	Full sanction after repeated offences	4-12 months
Stringent	Full sanction at first offence	0-3 months (incl. state with no exemptions)

In Table B.7, we consider the interaction terms between age of the youngest child groups, stringency and time as additional controls for AFDC utilization and employment. Controlling for potential age-of-youngest-child-specific effects of having stronger work requirements does not affect the estimates of the effects of time limits.

Child Support Enforcement We coded a state-level dataset that captures whether a given state had a child support enforcement rule already in place by 1990 and, hence, at the time of the reform (Huang, Garfinkel and Waldfogel, 2004). In Table B.8 we interact the treatment variable with this child support enforcement variable. Of particular interest is the coefficient on these interactions: we find that the reform had the same effect irrespective of whether child support enforcement provisions had already been long established in the state or not, suggesting that these rules are unlikely to be confounds for the effect of time limits.

Table B.7: AFDC Utilization Regressions with Work Requirements Interactions

	AFDC/TANF Utilization			*Employment*		
Sample:	Whole	Married	Unmarried	Whole	Married	Unmarried
$Exposed_{dst}Post_{st}$	-0.024***	-0.009***	-0.067***	0.011	-0.001	0.045***
	(0.005)	(0.003)	(0.017)	(0.012)	(0.014)	(0.015)
$[Child \leq 1_{dt}]Interm_sPost_{st}$	-0.024*	-0.009	-0.088***	-0.005	-0.020	0.030
	(0.013)	(0.013)	(0.031)	(0.020)	(0.027)	(0.027)
$[Child \leq 1_{dt}]Stringent_sPost_{st}$	-0.033***	-0.024***	-0.089***	-0.010	-0.025	0.037
	(0.009)	(0.008)	(0.020)	(0.016)	(0.019)	(0.026)
$[Child \leq 5_{dt}]Interm_sPost_{st}$	-0.007	-0.003	-0.010	0.030**	0.026	0.008
	(0.011)	(0.008)	(0.030)	(0.013)	(0.019)	(0.024)
$[Child \leq 5_{dt}]Stringend_sPost_{st}$	-0.015*	0.002	-0.057**	-0.003	-0.011	-0.000
	(0.008)	(0.005)	(0.025)	(0.014)	(0.015)	(0.027)
Mean pre-reform	0.101	0.034	0.306	0.642	0.647	0.624
Obs	254,627	188,483	66,144	254,627	188,483	66,144
R^2	0.12	0.08	0.26	0.12	0.11	0.21

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. We use data from the 1990-2001 SIPP panels, until 2002. Sample of female heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-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. Additionally, we interact work requirement indices with the age of youngest child and the calendar year.

Time Limits Waivers As a robustness check, we coded statewide time-limit waivers (see Table 2.5 in Grogger and Karoly (2005)) as an alternative way to identify the introduction of the time limits. In Table B.9 we present results for this alternative timeline, which leads to very similar conclusions to the one using the official timeline.

Table B.8: Results with Child Support Interactions

	AFDC/TANF Utilization		Employment	
$Exposed_{dst}Post_{st}$	-0.087***	-0.084***	0.050***	0.053***
	(0.015)	(0.020)	(0.014)	(0.017)
$Exposed_{dst}Post_{st}\times$ Pre-reform CS		-0.012		-0.013
		(0.029)		(0.025)
Mean pre-reform	0.306	0.306	0.624	0.624
Obs	66,144	66,144	66,144	66,144
R^2	0.26	0.26	0.21	0.21

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. We use data from the 1990-2001 SIPP panels, until 2002. Sample of female single heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-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. Additionally, we interact child support state dummy with the age of youngest child and the calendar year.

Table B.9: Effects of Time Limits under Different Policy Timelines

	AFDC/TANF Utilization			Employment		
Sample:	Whole	Married	Unmarried	Whole	Married	Unmarried
$Exposed_{dst}Post_{st}$ (Alter.)	-0.028***	-0.011***	-0.083***	0.015	-0.000	0.048***
	(0.004)	(0.003)	(0.015)	(0.012)	(0.013)	(0.015)
Mean pre-reform	0.101	0.034	0.306	0.642	0.647	0.624
Obs	254,627	188,483	66,144	254,627	188,483	66,144
R^2	0.12	0.08	0.26	0.12	0.11	0.21

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. We use data from the 1990-2001 SIPP panels, until 2002. Sample of households where both spouses are present, neither are college graduates, and have children aged 18 and below. We use use statewide time limit waivers from Grogger and Karoly (2005) to code this alternative timeline for the roll-out of the time limits policy. 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 B.10 we show that this adjustment leads to larger estimates in absolute value, but to qualitatively similar results.

Table B.10: OLS Regressions with First Wave of Each SIPP panel

Dependent Var:	**AFDC/TANF**		SIPP **Employed**		**Div/Sep**	**Married**
Sample:	Whole	Unmarried	Whole	Unmarried	Whole	Whole
$Exposed_{dst}Post_{st}$	-0.045***	-0.124***	0.060***	0.142***	-0.037*	-0.010
	(0.007)	(0.022)	(0.019)	(0.040)	(0.019)	(0.016)
Mean pre-reform	0.095	0.298	0.639	0.626	0.149	0.756
Obs	32,928	8,775	32,928	8,775	32,928	32,928
R^2	0.14	0.33	0.15	0.30	0.06	0.08

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. We use data from the 1990-2001 SIPP panels, until 2002. 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 sample 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 B.11).

Table B.11: Effects of Time Limits on College Graduates

			SIPP			
Dependent Var:	**AFDC/TANF**		**Employed**		**Div/Sep**	**Married**
Sample:	Whole	Unmarried	Whole	Unmarried	Whole	Whole
$Exposed_{dst}Post_{st}$	-0.008***	-0.055***	-0.019	-0.008	-0.015	0.006
	(0.002)	(0.015)	(0.012)	(0.027)	(0.016)	(0.017)
Mean pre-reform	0.010	0.051	0.768	0.891	0.091	0.878
Obs	93,233	12,384	93,233	12,384	93,233	93,233
R^2	0.04	0.28	0.09	0.23	0.09	0.10
			CPS			
Dependent Var:	**AFDC/TANF**		**Employed**		**Div/Sep**	**Married**
Sample:	Whole	Unmarried	Whole	Unmarried	Whole	Whole
$Exposed_{dst}Post_{st}$	-0.001	-0.033**	-0.018	-0.012	-0.005	-0.002
	(0.002)	(0.015)	(0.012)	(0.030)	(0.013)	(0.016)
Mean pre-reform	0.005	0.037	0.774	0.909	0.077	0.901
Obs	36,089	3,676	36,089	3,676	36,089	36,089
R^2	0.02	0.20	0.04	0.21	0.03	0.04

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. We use data from the 1990-2001 SIPP panels, until 2002, and 1990-2002 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 C: Life-Cycle Model

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 $V_t^{M_m}\left(\boldsymbol{\Omega_{ij,t}^{m}}\right)$, the value accruing to a married man. The budget constraint of the single man is given by:[52]

$$\frac{A_{j,t+1}^{M_s}}{1+r} = A_{j,t}^{M_s} - c_{j,t}^{M_s} + y_{j,t}^M + FS_{j,t} \tag{9}$$

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

The problem for the single man is thus defined by

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

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:

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

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.

[52] 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).

Problem of the Couple

We define the optimization problem for the married couple, given the overall state space $\boldsymbol{\Omega}^{\boldsymbol{m}}_{\boldsymbol{ij,t}}$.[53]

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

$$V^m_t\left(\boldsymbol{\Omega_{ij,t}}^m\right) = \max_{\boldsymbol{q}^m_{ij,t}} \left\{ \begin{array}{c} \theta^F_{ij,t} u^{F_m}\left(c^F_{ij,t}, P^F_{ij,t}, B^F_{ij,t}\right) + \theta^M_{ij,t} u^{M_m}\left(c^M_{ij,t}, P^M_{j,t}\right) + L_{ij,t} \\ +\beta E_t \left[\begin{array}{c} \left(1 - d_{ij,t+1}\right) V^m_{t+1}\left(\boldsymbol{\Omega}^{\boldsymbol{m}}_{\boldsymbol{ij,t+1}}\right) \\ +d_{ij,t+1}\left(\theta^F_{ij,t} V^{F_s}_{t+1}\left(\boldsymbol{\Omega}^{\boldsymbol{F_s}}_{\boldsymbol{i,t+1}}\right) + \theta^M_{ij,t} V^{M_s}_{t+1}\left(\boldsymbol{\Omega}^{\boldsymbol{M_s}}_{\boldsymbol{j,t+1}}\right)\right) \end{array} \right] \end{array} \right\} \tag{10}$$

where $\boldsymbol{q}^{\boldsymbol{m}}_{\boldsymbol{ij,t}} = \{c^M_{ij,t}, c^F_{ij,t}, P^F_{ij,t}, B^F_{ij,t}\}$ 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:

$$L_{ij,t} = L_{ij,t-1} + \xi_{ij,t} \tag{11}$$

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

$$\frac{A_{ij,t+1}}{1+r} = A_{ij,t} - x(c^F_{ij,t}, c^M_{ij,t}, k^a_{i,t}) + (w^F_{i,t} - CC^a_t)P^F_{ij,t} + y^M_{j,t} + B_{ij,t}b_{ij,t} + FS_{ij,t} + EITC_{ij,t}$$

To capture economies of scale in marriage (including public goods), we assume that the individual consumptions $c^F_{ij,t}$ and $c^M_{ij,t}$ and the equivalence scale $e(k^a_t)$ imply an aggregate household expenditure of

$$x_{ij,t} = \frac{\left((c^F_{ij,t})^\rho + (c^M_{ij,t})^\rho\right)^{\frac{1}{\rho}}}{e(k^a_{i,t})}$$

The extent of economies of scale is controlled by ρ and $e(k^a_t)$. 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.[54]

[53] The state variables for the couple (represented by $\boldsymbol{\Omega_{ij,t}}^m$), are: the combined assets, each spouses' productivity, whether the woman worked last year, the number of periods of welfare benefits utilization, age of the youngest child present ($k^a_{i,t}$), and the Pareto weights $\theta^M_{ij,t}, \theta^F_{ij,t}$.

[54] This assumption is a good approximation of the legal position (see Voena (2015)). After marriage, spouses' assets merge into one value: $A_{ij,t} = A^{F_s}_{i,t} + A^{M_s}_{j,t}$ and so, in the computation, we do not keep track of individual assets going into the marriage.

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

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

and in other periods

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

and

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

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:

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

The 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.

Fertility

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

$$Pr\left(k_{t+1}^{a}|k_{t}^{a}, m_{t-1}, t\right). \tag{12}$$

We capture the arrival of additional children by treating the age of the child as 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.

Our assumptions about fertility mean that we will not capture any effects of benefits on children's welfare. Mullins (2022) shows that more generous welfare benefits have an additional value through increasing children's human capital development.

Marriage Process

We use the following process to allow marriage market opportunities to vary as people become older:

$$\lambda_t = min\{max\{\lambda_0 + \lambda_1 \cdot (t-1) + \lambda_2 \cdot (t-1)^2, 0\}, 1\}. \tag{13}$$

Externally Set and Directly Estimated Parameters of the Model

Table C.1: 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)	1.4023
Welfare program parameters	Statutory rules
Length of life	40
Panel B: Directly Estimated Parameters	
Childcare costs (CC^a)	CEX (see text)
Fertility process	SIPP
Distribution of single characteristics	SIPP (see text)
St. dev. of men's unexplained earnings in period 1 (σ^2_{0M})	0.464
St. dev. of men's earnings shocks ($\sigma^2_{\zeta M}$)	0.166
Life cycle profile of log male annual earnings (a^M_0, a^M_1, a^M_2)	9.43, 0.0641, -0.0013

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:

$$E[\Delta u_t^2] = \sigma^2_{\zeta^M} + 2\sigma^2_{\varepsilon^M} \tag{14}$$

$$E[\Delta u_t \Delta u_{t-1}] = -\sigma^2_{\varepsilon^M} \tag{15}$$

where u_t is the residual log earnings obtained after regressing earnings on dummies for age, disability status, and year. We use the SIPP sample of prime-aged men without a college degree.

Target Moments Fit

Sensitivity Analysis

Figure C.1: Target Moments of SIPP and Simulated Data

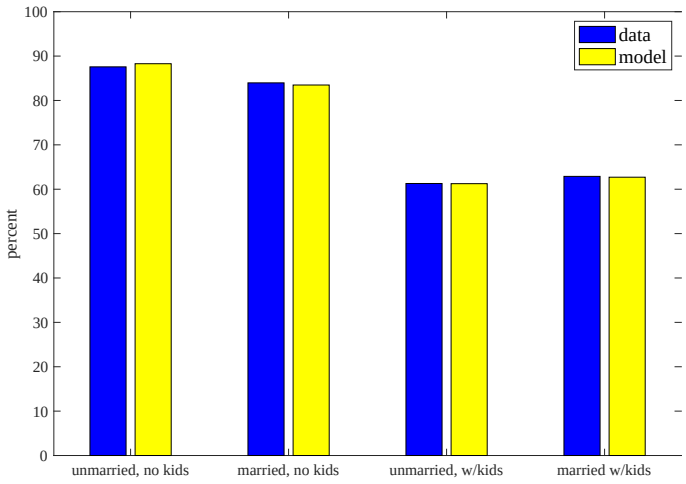

(a) Employment moments

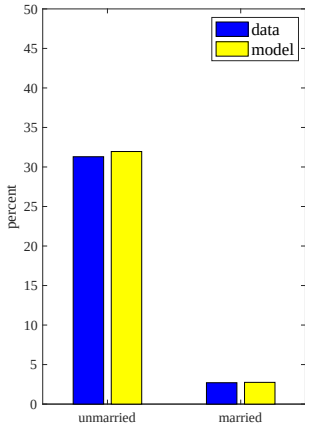

(b) Welfare moments

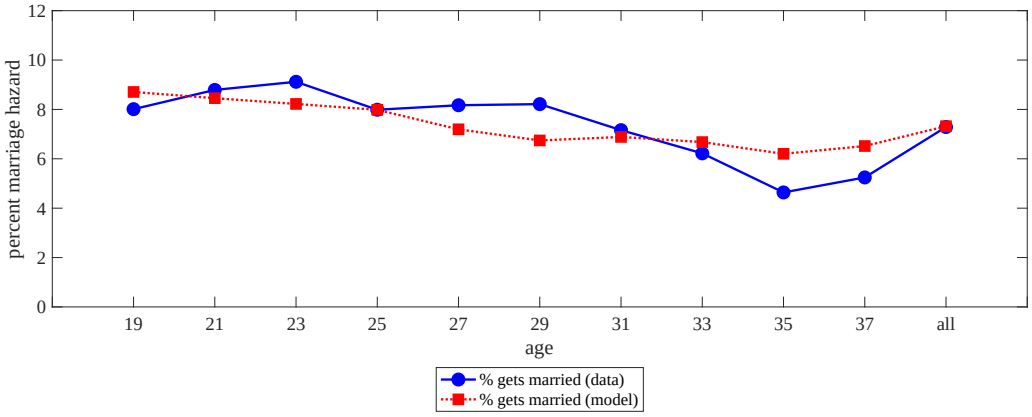

(c) Marriage transition moments

Figure C.2: Target Moments of SIPP and Simulated Data (cont.)

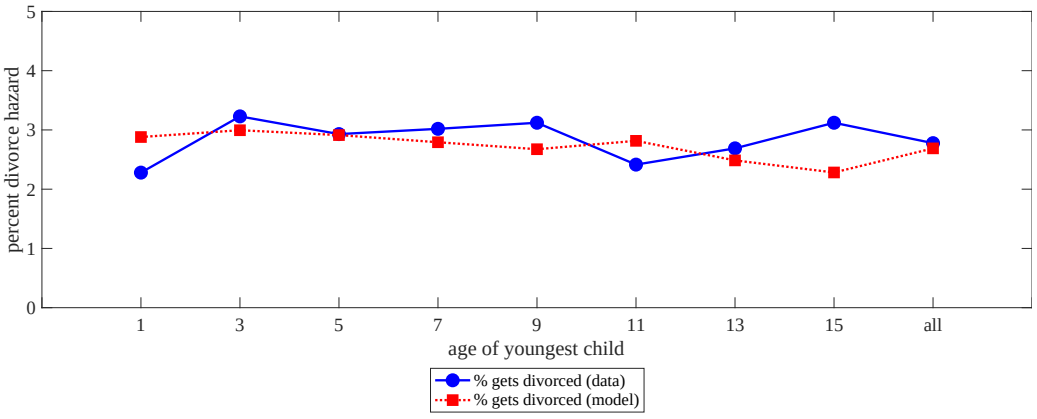

(a) Divorce transition moments

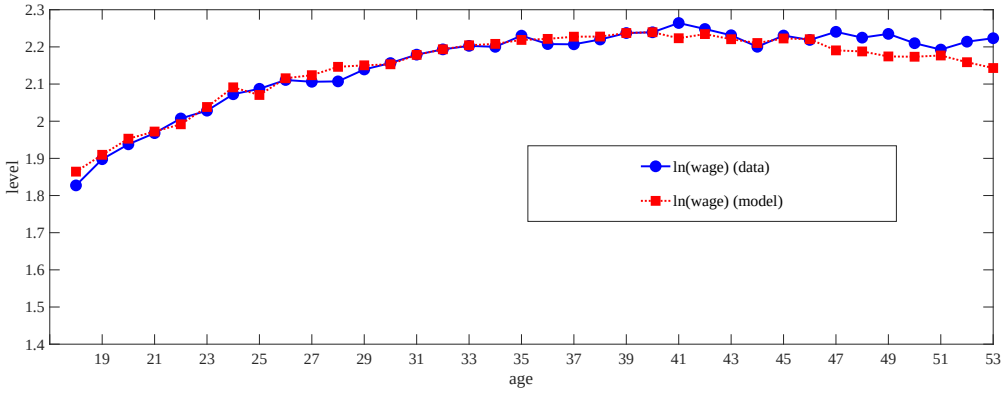

(b) ln(wage) profile moments

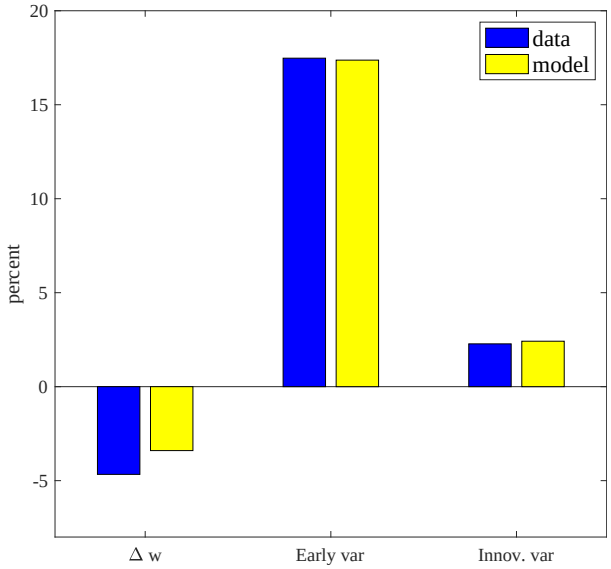

(c) Wage growth moments

Notes: Target moments in the pre-reform SIPP data and in the estimated model.

Figure C.3: Sensitivity Analysis

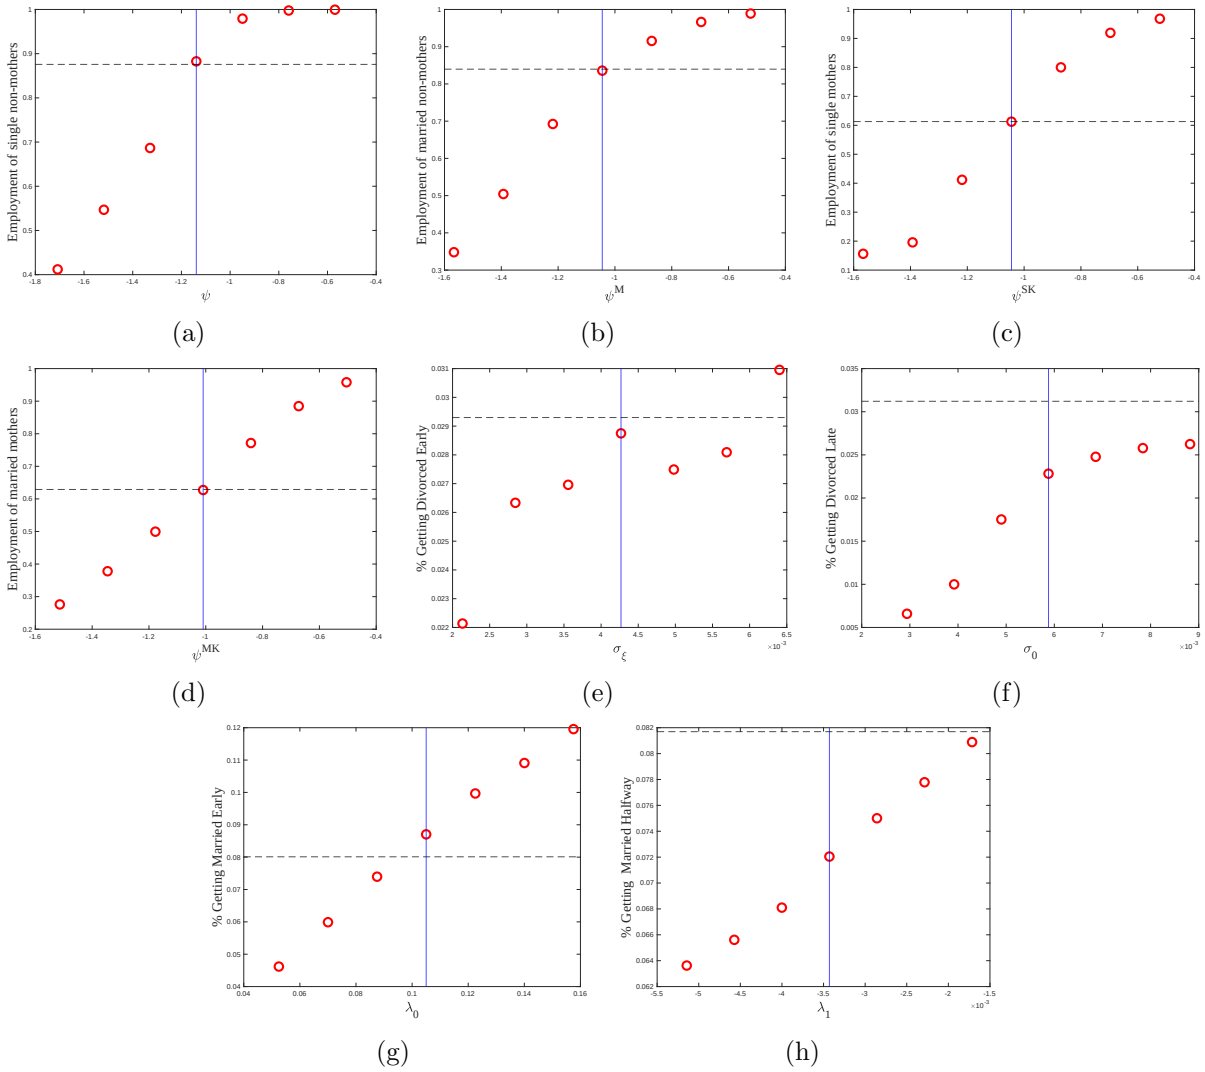

Notes: We hold all parameters constant at the estimated values, except one parameters which varies along the horizontal axis. The vertical axis reports the values taken by a given target moment.

Figure C.4: Sensitivity Analysis (part 2)

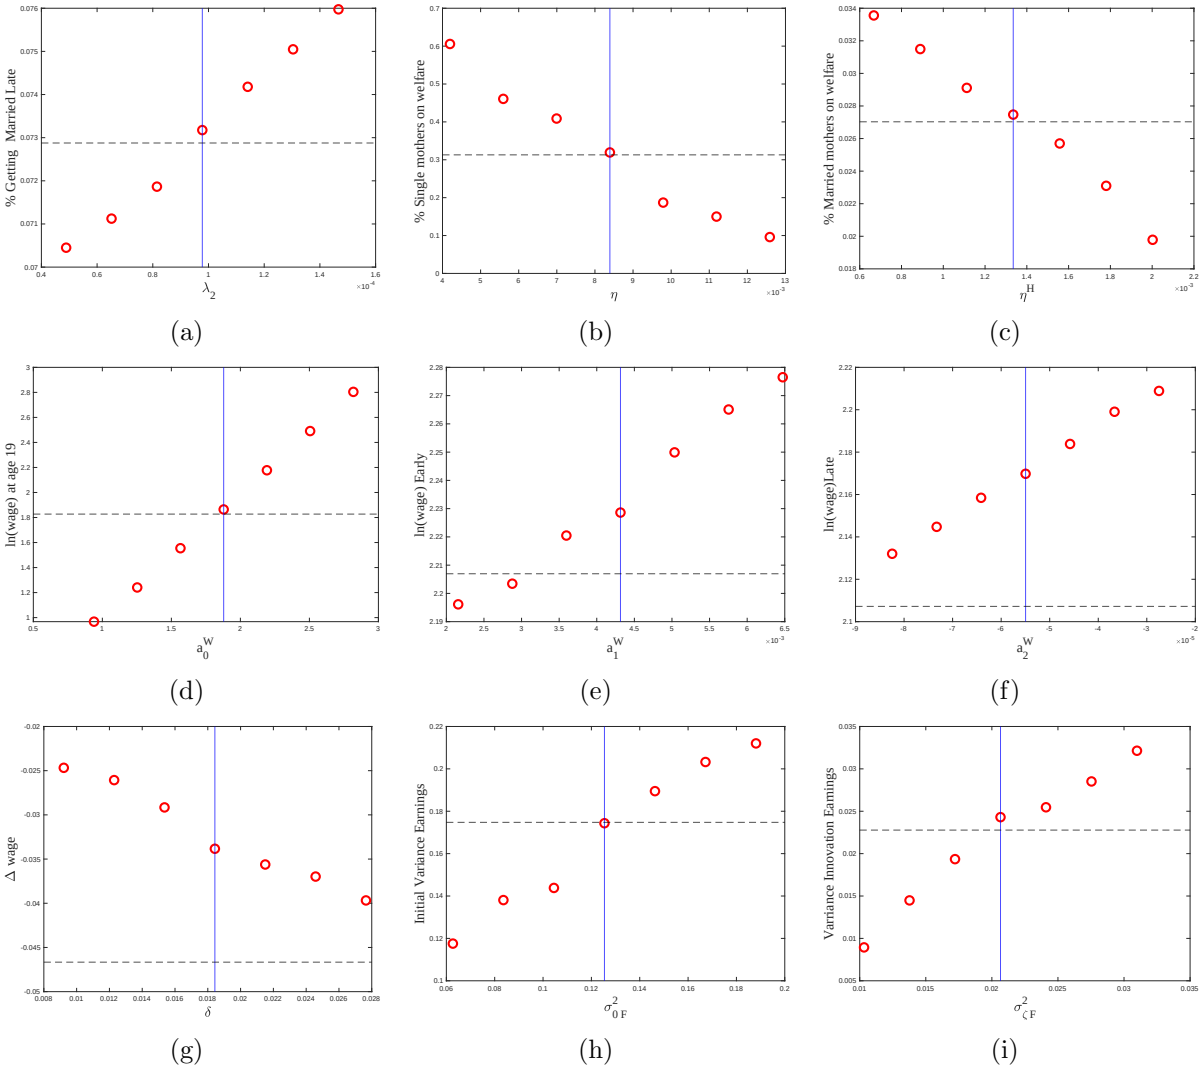

Notes: We hold all parameters constant at the estimated values, except one parameters which varies along the horizontal axis. The vertical axis reports the values taken by a given target moment.

Intrahousehold Allocations

Table C.2: Simulated Estimates of Effect of Time Limits on Women's Pareto Weight

	All marriage	Existing marriages	New Marriages
Effect of Time Limits	-0.0025	-0.0005	-0.0014

Notes: Model estimates of the effect of the introduction of time limits on women's Pareto weight.

Appendix D: The Effect of Time Limits on Behavior and the Role of Marriage and Divorce

Work Requirements

Table D.1: Difference-in-Difference Estimates in the SIPP Data and in the Work Requirement Counterfactual

	(1) Benefits Unmarr.	(2) Employment Unmarried
Coeff. - SIPP data	-0.067	0.055
	[-0.105,-0.029]	[0.027,0.084]
Time limits - No Work Requirements	-0.119	0.032
Work requirements - No Time Limits	-0.034	-0.029

Notes: Estimates from the annualized SIPP data between 1990 and 2002 (first 6 years after the reform) for children aged 6 and above. 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.

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.

The Effects of Time Limits by Productivity For Different Subgroups

Figure D.1: Welfare Use and Employment of Single Mothers

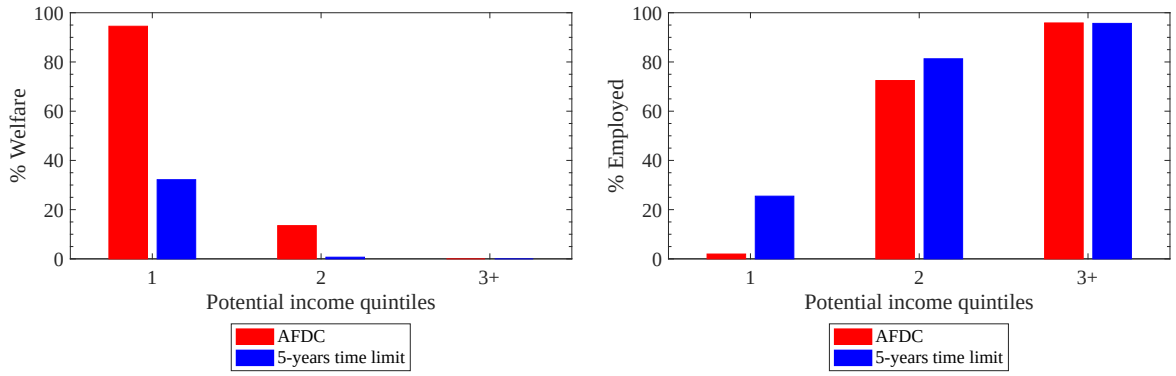

(a) Welfare use (b) Employment

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, for the subset of women who are unmarried under time limits.

Figure D.2: Welfare Use and Employment of Married Mothers

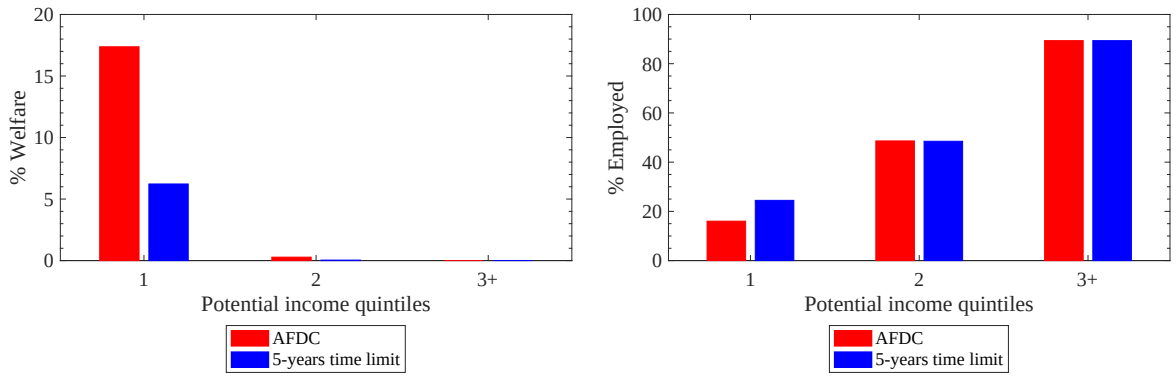

(a) Welfare use (b) Employment

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, for the subset of women who are unmarried under time limits.

Figure D.3: Welfare Use and Employment of Single and Married Mothers Who Have not Yet Run Out of Benefits, by Productivity

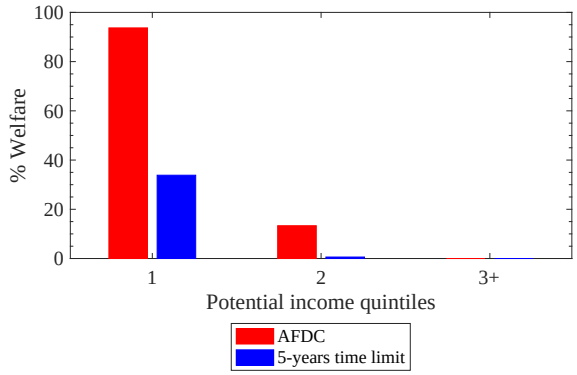

(a) Welfare use - single

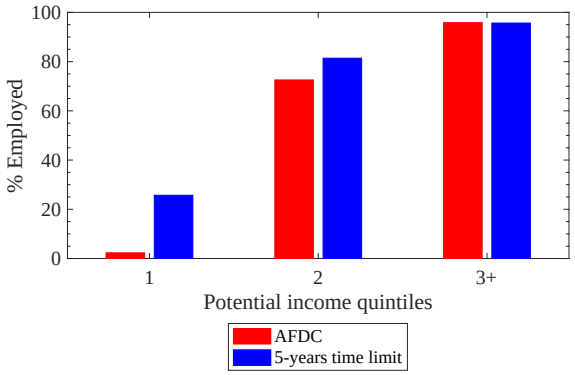

(b) Employment - single

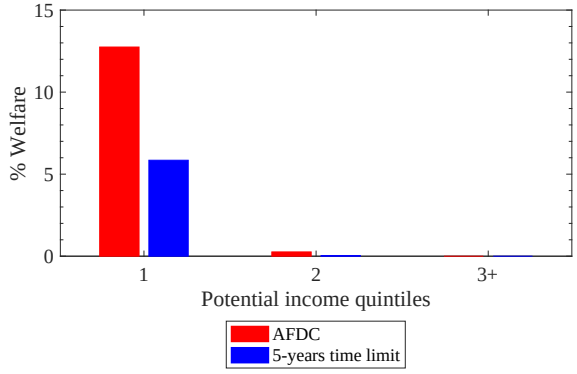

(c) Welfare use - married

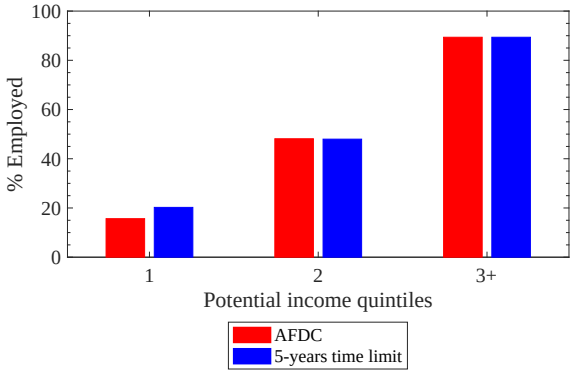

(d) Employment - married

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. All figures focus the subset of women who, in the counterfactual with time limits, have not yet used 5 years of benefits.

Figure D.4: Welfare Use and Employment of Married Mothers by Productivity: Robustness Analysis to Limited Commitment Assumption

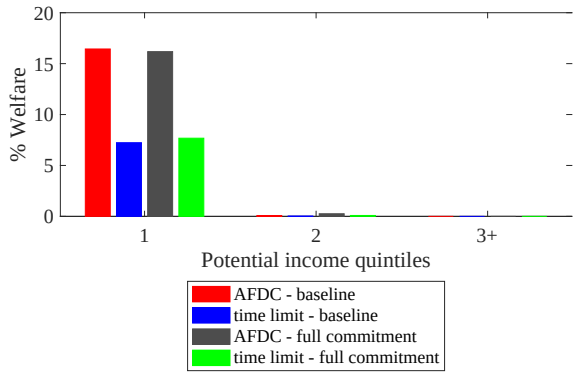

(a) Welfare use, full commitment

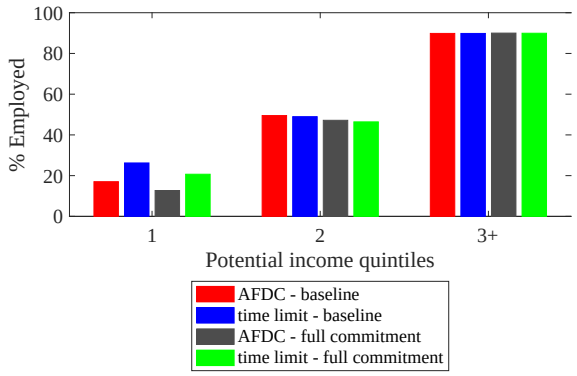

(b) Employment, full commitment

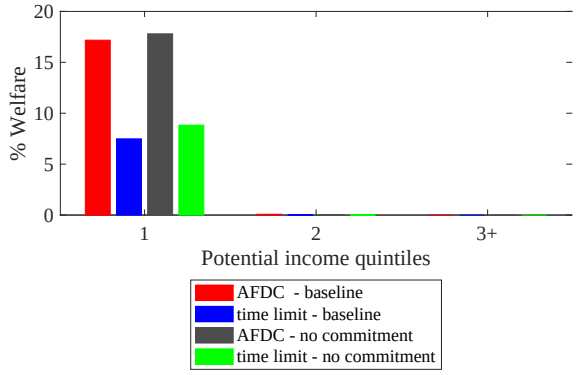

(c) Welfare use, no commitment

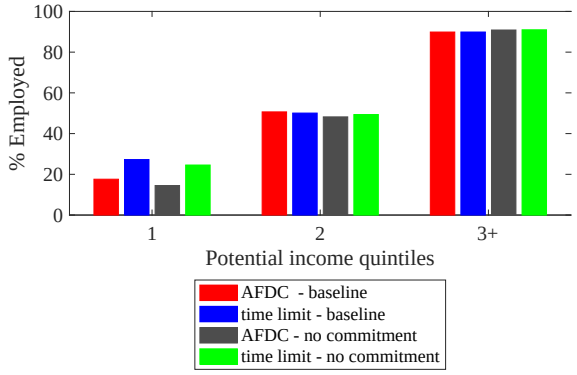

(d) Employment, no commitment

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 red and blue bars are baseline model counterfactuals. The grey and the green bars are model counterfactuals re-estimated under no commitment. In a full commitment model, spouses never renegotiate Pareto weights after marriage. In a no-commitment model, spouses renegotiate their Pareto weight each period by symmetric Nash bargaining.

Role of Household Structure

Figure D.5: Welfare Use and Time Limits with and Without a Marriage Transition Option

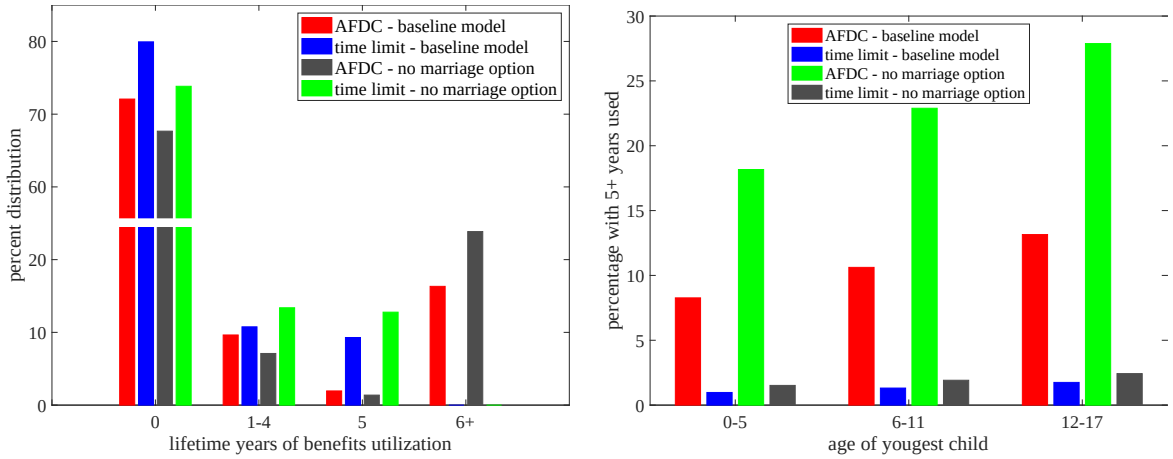

(a) Distribution of lifetime welfare utilization (b) Benefits Exhaustion by Age of Youngest Child

Notes: Panel a: Model simulations of the distribution of lifetime welfare use in simulated data with and without a marriage transition option. Panel b: Model simulations of the fraction of mothers who would exhaust benefits by the age of the youngest child.

Table D.2: Effects of Time Limits for Single Women by Future Marital Status

	AFDC/TANF Utilization		*Employment*	
	Future Marriage	No Future Marriage	Future Marriage	No Future Marriage
$Exposed_{dst}Post_{st}$	-0.088***	-0.117***	0.078*	0.047***
	(0.027)	(0.011)	(0.044)	(0.015)
Mean pre-reform	0.204	0.292	0.682	0.626
Obs	6,994	109,396	6,994	109,396
R^2	0.49	0.26	0.42	0.17

Notes: Standard errors in parentheses clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1. We use data from the 1990-2001 SIPP panels, until 2002. Sample of female single heads of household who are not college graduates and have children aged 18 and below. The full set of controls includes age dummies, education dummies, number of children dummies, year-by-month fixed effects, state-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. "Future marriage" sample includes currently unmarried women that marriage in some future observation in the panel. "No future marriage" includes currently unmarried women who never marry afterwards in the panel.

Figure D.6: Welfare Use and Employment of Unmarried Mothers with and without a Marriage Option

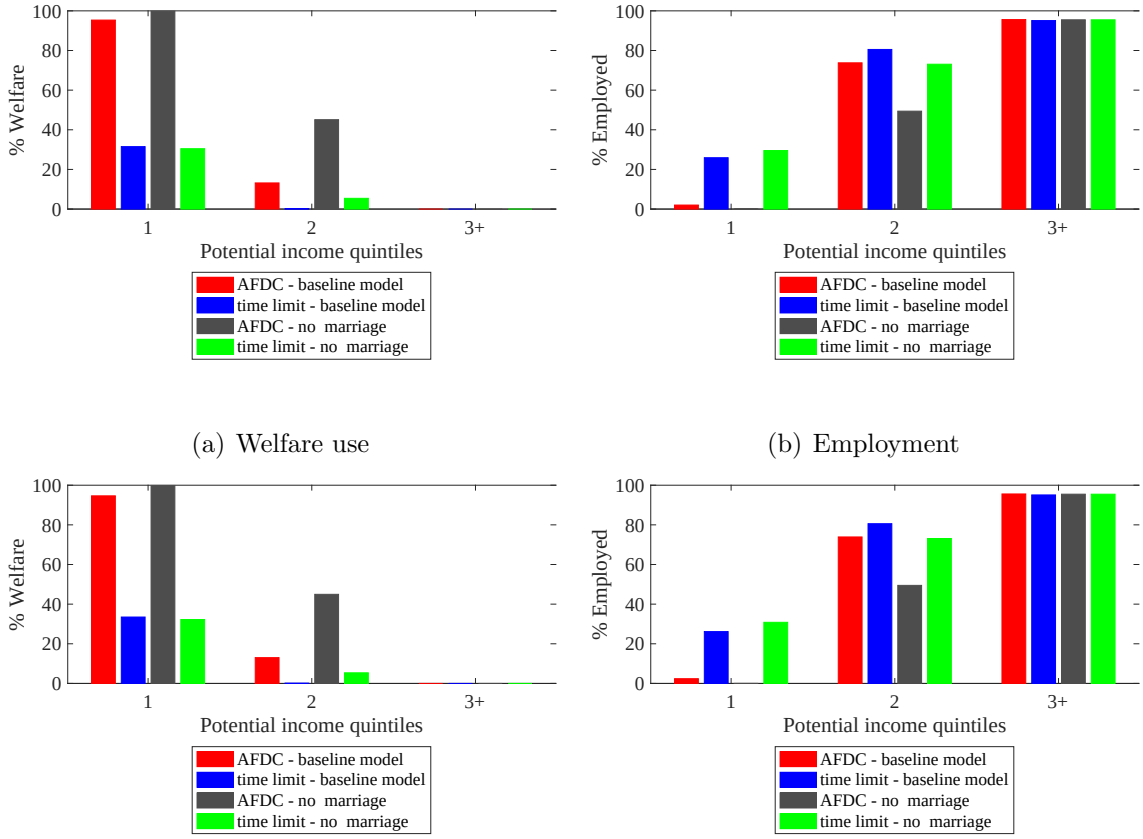

(a) Welfare use

(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 black 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 black 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 D.7: Welfare Use and Employment of Married Mothers with and without a Divorce Option

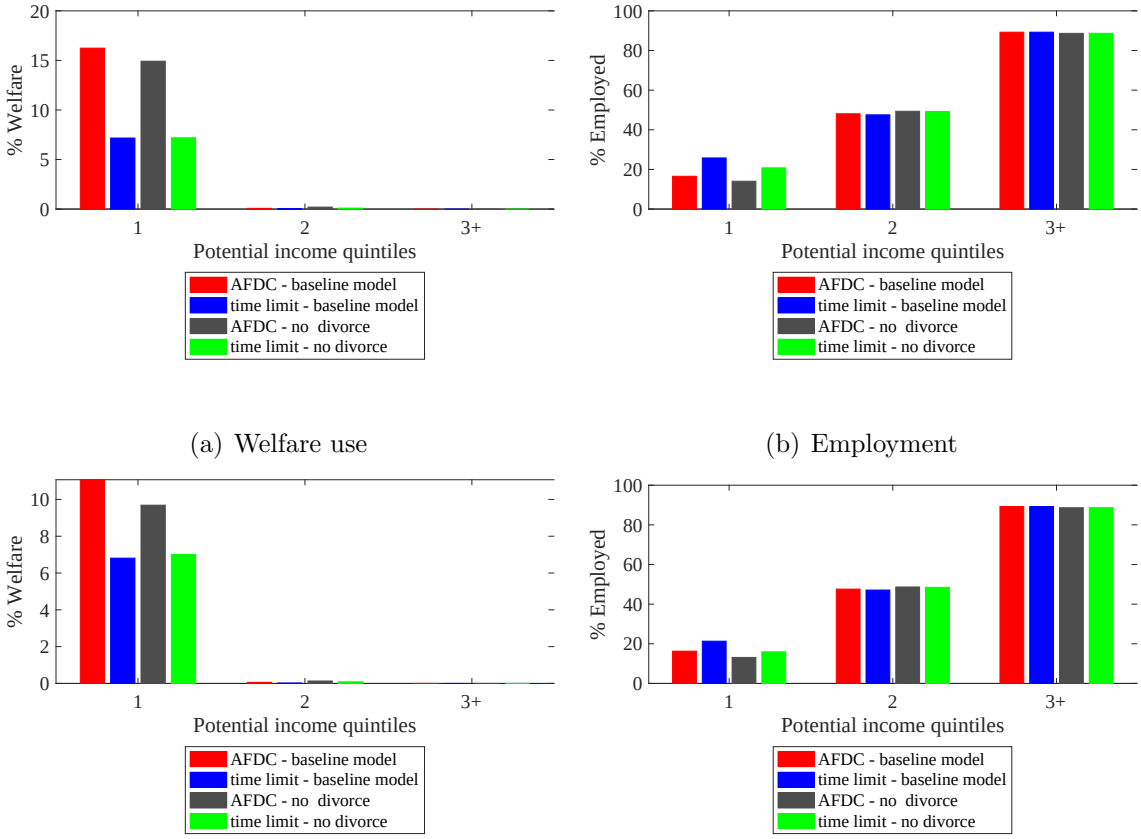

(a) Welfare use

(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 black 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 black 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 E: Implications of Welfare Reform for Lifetime Utility

Welfare Calculation

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

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

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 $s = AFDC$.[55]

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

We solve for π such that

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

where π can be interpreted as the consumption cost of going from AFDC to a 5-year limit.[56]

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

[55] 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.

[56] In varying π we do not reoptimize.

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

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

$$E_0\,U^M\left(5TL,\tau\right)\big|_\pi + E_0\,U^F\left(5TL,\tau\right)\big|_\pi = E_0 U^M\left(AFDC,0\right) + E_0\,U^F\left(AFDC,0\right)\big|_\pi. \quad (18)$$