

TANF, Childcare and Child Well-Being in Sole-Parent Families

Marianne Bruins*

24 August 2017

Abstract

I estimate a structural model of sole-parent families to analyse the impact of TANF. Its benefits fail to reach the poorest mothers, who typically prefer not to meet the 30 hr/wk work requirement. Using the model to measure the intra-household allocation of resources, I find that poverty rates among children of sole mothers have risen by four percentage points since 1996. I compare TANF to alternative policies, such as free childcare, which promote labour supply by increasing the returns to work. Such policies are more than twice as effective at targeting household resources to children, per dollar spent.

*Email: marianne.bruins@economics.ox.ac.uk. Nuffield College and Department of Economics, University of Oxford. This research has greatly benefited from discussions with R. Blundell, S. Bond, M. Brewer, M. Browning, J. Duffy, J. Fenske, C. Meghir, S. Quinn, J. Ziliak and comments from seminar participants at Oxford and ISER (Essex). The author thanks the British Academy for funding (Award no. SG152880). This work was supported by the HPC facilities operated by, and the staff of, the Yale Center for Research Computing.

Contents

1	Introduction	1
2	Background	4
2.1	Welfare programmes and reform	4
2.2	Trends in household behaviour since 1996	7
3	Structural model	9
3.1	Preferences	10
3.2	Constraints	13
3.3	Solution and two-stage budgeting	13
4	Estimation procedure	15
4.1	Data sources	16
4.2	Simulated method of moments	19
4.3	Imputation (via simulation) of missing variables	21
4.4	Sample moments: selection and construction	23
4.5	Identification of structural parameters	26
4.6	Identifying the TANF enrolment disutilities	27
5	Estimates and model fit	28
5.1	Parameter estimates	29
5.2	Model fit: in-sample performance	31
5.3	Model fit: external validation	35
5.4	Identification: a numerical analysis	36
6	Child welfare, TANF and alternative policies	38
6.1	Child welfare in the welfare-to-work era	40
6.2	PROWRA, child welfare and maternal labour supply	43
6.3	Alternatives to welfare-to-work programmes	46
6.4	Implications for welfare policies	50
7	Conclusion	51
8	References	52

Appendices (for online publication)

A	Data sources and construction	A1
B	Estimation procedure: further details	A3
C	Parameter estimates	A7
D	Calculation of taxes and welfare payments	A12

1 Introduction

The introduction of welfare-to-work programmes, in many OECD countries, has been the most significant change to cash welfare in the postwar era. These programmes are designed to reduce welfare dependency by tying eligibility to work requirements. They were introduced with a limited understanding of their direct impact on child welfare, in spite of the welfare and economic opportunities of disadvantaged children being the focus of other policy initiatives. In this paper I open the black box of household behaviour to examine how welfare-to-work programmes influence the allocation of household resources to children. That this remains poorly understood is perhaps not surprising, given the taxing data requirements and sophisticated modelling of households' preferences and the tax and welfare system necessary to answer this question. But it is of pressing concern in view of the extensive evidence suggesting that what goes on within the household is crucial for children's immediate welfare and outcomes later in life.

In the US, the major policy reform of this kind was effected with the introduction of Temporary Assistance for Needy Families (TANF) in 1996, which replaced the existing cash welfare programme designed for sole mothers. This paper provides a detailed analysis of how TANF affected the allocation of household resources to children over 1996–2008, and to the best of my knowledge is the first paper to do so. My analysis is based on a structural model of household decision-making – estimated on a sample of sole mothers (without a college degree) – in which children benefit from consumption, time with their mother, and a domestically produced public good. With the estimated model, I am able to perform a comprehensive analysis of the impact of the introduction of TANF on child welfare. I also consider how welfare policies might be designed so as to better target household resources to children, while still meeting TANF's original objective of promoting female labour supply.

Using the estimated model, I demonstrate that, at 30 hours per week (hr/wk), the work requirements for TANF are set so high that many low-wage mothers would be better off spending their time in activities other than market work, than meeting the work requirements and receiving TANF payments -- even though their weekly earnings put them well below the poverty line. For this group, the main consequence of welfare-to-work reforms has been simply a withdrawal of the benefits to which they had previously been entitled. This loss of income has been mostly offset by an increase in their labour supply of around 5 hr/wk, primarily at the expense of time spent in housework.

Accordingly, poverty measures based on household-level income and consumption have indicated no increase in child poverty rates since 1996 (Blank, 2002; Meyer and Sullivan, 2004). But such measures fail to account for the decline in mothers' time available for

activities outside market work, and the subsequent decline in home production which has negatively affected children. I therefore use the model to construct a measure of child poverty which is based on the resources children actually receive from the household, and thus appropriately values the time mothers spend with their children, and fully accounts for home production. I find that, once the intra-household allocation of resources is accounted for in this manner, child poverty has in fact increased by four percentage points since the introduction of TANF.

I also use the model to perform a number of counterfactual experiments, so as to evaluate the consequences of both actual and hypothetical welfare reforms. As discussed above, few sole mothers enrol in TANF and meet the work requirements, so TANF only provides meaningful work incentives for a very small group of mothers, less than five per cent of my sample. In contrast, the Earned Income Tax Credit (EITC), an in-work subsidy, raises the net returns to work for all households with sufficiently low incomes. Thus, it is not surprising that I find the EITC to be more effective than TANF at promoting maternal labour supply. I also find the EITC to be twice as efficient at targeting resources to children, measured in terms of the (money-metric) improvement in children's welfare per (net) dollar spent on the policy. (This calculation of 'net spending' on the EITC – and on other programmes – fully accounts for its secondary effects on government tax revenues, due to households' behavioural responses.)

These results motivate me to consider a number of alternative policies that similarly promote labour supply by increasing the returns from working. These alternative policies take the form of childcare subsidies and wage subsidies. Their main point of difference from the EITC lies in that they pay benefits that are proportional to hours worked, rather than to total earnings. One drawback of a policy whose benefits are contingent on total earnings, instead of hours worked, is that it fails to distinguish between e.g. two mothers, one of who works part time and the other full time, but who have the same total earnings. Clearly, in such a case, the mother working part time will be able to devote more time to housework and children, both of which contribute to children's welfare, and thus her children may be considerably better off than those of the mother working full time.

My analysis of these alternative policies reveals that they are indeed more effective than TANF and the EITC at increasing maternal labour supply, and twice as efficient as TANF at targeting household resources to children. More specifically, among the policies I consider are the provision of free childcare, and two wage subsidies: a flat-rate wage subsidy of \$1/hr, and a targeted wage subsidy which brings all wages up to \$11.60/hr (both in 2000 dollars).¹

¹This is equivalent to \$15/hr in 2016 dollars, which compares to the \$15/hr minimum wage scheduled for New York in 2018.

For both the free childcare and targeted wage subsidies, the extent to which expenditure on these policies passes through to children is comparable to the EITC, at \$0.41 and \$0.50 of every (net) dollar spent, respectively. The \$1/hr wage subsidy has a pass-through rate almost twice that of the EITC, at \$0.76 per dollar spent. To the best of my knowledge, the favourable effects of such policies on children have not previously been identified, most likely because their impact on the allocation of resources within the household has not been properly understood accounted for.

In recent decades, many OECD countries have increased childcare subsidies, and in a few cases have provided universal free childcare. Recent literature examining the effects of these policies on the cognitive and non-cognitive development of young children has focused largely on the relative quality of care available at home and at centres. My analysis identifies a secondary channel through which such policies may improve child welfare: by making additional household resources available to children. In quantitative terms, I find that for each (net) dollar spent on such policies, \$0.50 gets through to children. Depending on the household's income level, either channel – the relative quality of care, or the allocation of household resources to children – may predominate: and my results suggest that the household allocation channel is particularly relevant for children in low-income households.

To account for the benefits children derive from maternal time spent in activities aside from paid work, the structural model developed here contains a detailed treatment of mothers' time use. This is facilitated by the use of time diary data, which allows me to construct a measure of parents' time with children that aggregates those activities that the literature has identified as being important for child development. To keep the complexity of the modelling and the estimation procedure within reasonable bounds, while allowing for such a detailed treatment of the household's allocation and budget constraint, I abstract as far as possible from the intertemporal aspects of the household's decision making. The structural model is therefore static – but I am careful to ensure that the solution to the model is consistent with that of the household's intertemporal choice problem. This consistency permits the parameters of the structural model to be identified from data on households' actual choices, even though those choices arise from intertemporally optimising behaviour.

Methodologically, this paper demonstrates how to combine separate time-use and consumption datasets to estimate a model of household resource allocation, yielding much larger sample sizes and more precise estimates.² (Estimation by simulated method of moments is necessary here due to the use of these multiple datasets.) Only a few datasets provide both

²The datasets drawn used are: American Time Use Survey (ATUS, 2003–08), the American Heritage Time Use Survey (AHTUS, 1995), the Consumer Expenditure Survey (CE, 1993–2008), the Current Population Survey (CPS, 1993–2008) and the Survey of Income and Program Participation (SIPP, 1996, 2001, 2004, and 2008).

types of data, and those that do have only a small number of observations, which is particularly problematic in this context due to the measurement error problems associated with time diaries and consumption data. On the other hand, estimating the model with only consumption data – as has been the practice in much of the previous empirical work involving these models – means that the additional information on the household allocation provided by time-use data is ignored. By examining how the (estimated) asymptotic variances of the parameter estimates would be altered if certain groups of moments were ignored, I obtain empirical evidence that the use of moments derived from time-use data significantly improves the precision of the estimates.

Outline The remainder of this paper is organised as follows. Section 2 provides more detailed background information on changes in welfare policy over the sample period (1993–2008), and observed trends in the behaviour of sole-mother households over this period. The structural model, and the procedure used to estimate it, are elaborated in Sections 3 and 4. Section 5 presents parameter estimates and analyses the model’s fit. Section 6 presents the main findings of the paper: here I develop money-metric measures of child welfare, provide estimates of the effect of TANF on child poverty rates, and use the model to perform a series of counterfactual experiments, to evaluate a range of alternatives to TANF. Section 7 concludes.

2 Background

In this section I discuss the introduction of TANF, and other welfare programme reforms that occurred over 1993–2008. I then describe some of the trends in sole mothers’ behaviour during this period: the most striking of which is a reallocation of around 5hr/wk from market work to home production. These trends suggest that standard measures of child welfare, such as the official poverty measure, that fail to account for the value of mothers’ time spent in activities aside from market work, would give a misleading impression of the impact of TANF on child welfare. (I later return to this issue in Section 6, where I use the model to evaluate the totality of resources allocated to children, including time with parents and the provision of a public good, and consider how this has changed since 1996.)

2.1 Welfare programmes and reform

TANF and AFDC The Personal Responsibility and Work Opportunity Reconciliation Act of 1996 (PRWORA) effected the greatest change in the US welfare system since the

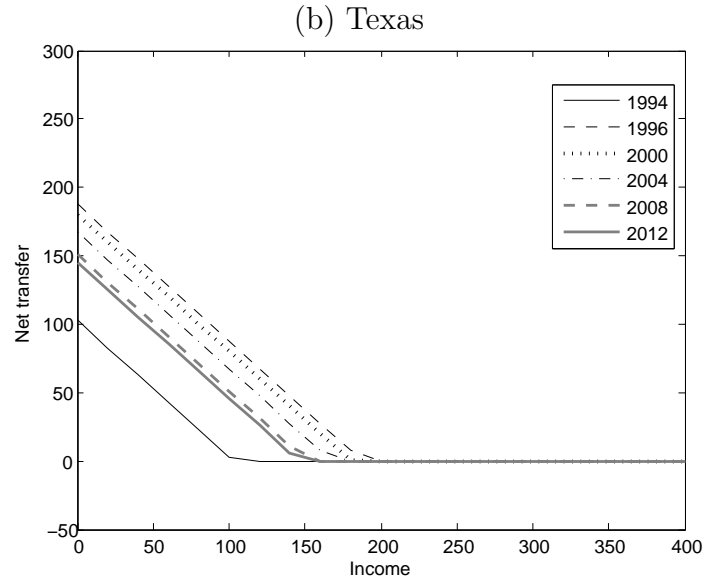
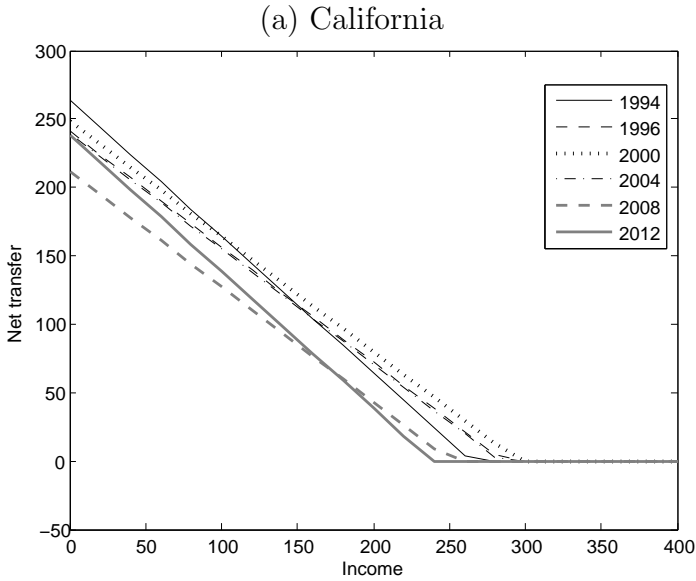
New Deal, by replacing Aid to Families with Dependent Children (AFDC) – the main cash welfare programme available to sole mothers – with Temporary Assistance for Needy Families (TANF). Unlike AFDC, TANF imposes: (a) lifetime limits on receipt, allowing individuals to claim benefits for a maximum of five years; and (b) work requirements, which make eligibility contingent on working a certain number of hours (usually 30) per week. (Although AFDC did impose some work requirements, these were much more limited in extent.) In practice, these work requirements have been unevenly enforced, with certain activities permitted to substitute for paid employment, such as actively searching for work or participation in a job training programme. Indeed only a third of TANF recipients worked in 2008, the final year of the sample period.

PRWORA granted states greater discretion over the provision of TANF than they had been allowed over AFDC: in the setting of eligibility rules, benefit levels, and the extent to which non-compliance with work requirements are penalised. The result is considerable heterogeneity in the way TANF is administered across states, as illustrated for example by the marked differences in the generosity of benefits provided in California and Texas, displayed in panels (a) and (b) of Figure 1. Some states, such as Alabama, in addition to the five-year lifetime limit, prohibit individuals from receiving TANF for more than two consecutive years, whereas others, such as Michigan, draw upon their own funds to extend payments to individuals who have exceeded the five-year limit. States have also exercised their autonomy to redirect federal funding for TANF away from cash welfare, towards such services as marriage counselling, family planning, training programmes, and childcare subsidies. Thus in aggregate, only 28 per cent of the federal block grant is spent on by cash welfare payments; a further 16 per cent is spent on childcare subsidies, and 17 per cent on administration (Schott, Pavetti, and Finch, 2012).

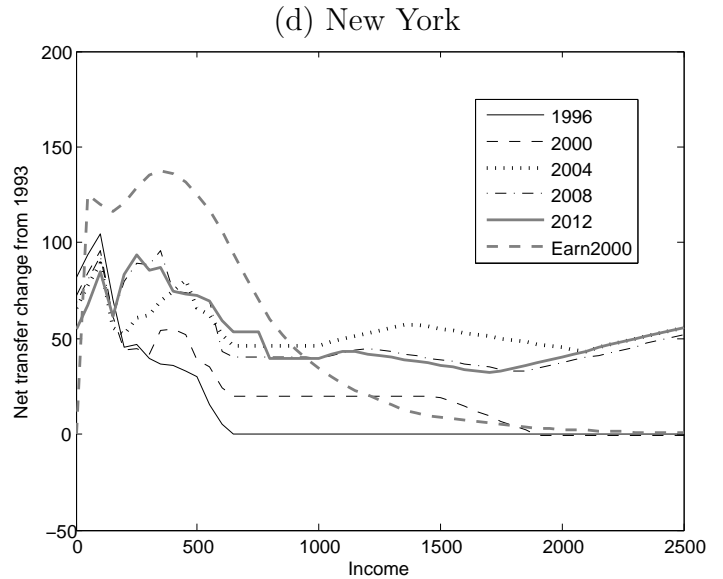

Other programmes While this paper is principally concerned with the impact of TANF on sole-parent households, the estimation of my structural model requires a comprehensive modelling of the entire tax and welfare system. In addition to AFDC and TANF, the model used in this paper accounts for all state and federal taxes on earned incomes, and the most significant transfer expenditures and welfare programmes targeted at children, excluding those related to healthcare. The four largest are the Child Tax Credit (CTC), the Supplemental Nutrition Assistance Program (SNAP), the Dependent Exemption, and the Earned Income Tax Credit (EITC) (see Isaacs, Toran, Hahn, Fortuny, and Steuerle, 2012); to these I add the Child and Dependent Care Tax Credit (CDCTC). Any state-level variation in these programmes (where applicable) is accounted for in the model; see Appendix D for further details.

Figure 1: Tax and welfare programmes
 Transfers and incomes per week (year 2000 dollars)

Maximum benefit payable by AFDC (1994) or TANF (1996–2012)



*Change in the net transfer received by households, relative to 1993**



• All transfers computed for a one-parent family with two children aged under 16.

* Assuming no rebate is paid on childcare expenses. *Earn2000* is the density of earnings for one-parent families in 2000 (computed from the CPS).

Panels (c) and (d) of Figure 1 plot the change in the net transfer received, under the foregoing taxes and welfare programmes, by a one-parent household with two children in California and New York, relative to 1993 (adjusted for inflation). These illustrate that, aside from the introduction of TANF in 1996, there have been substantial changes to the tax and welfare system at both the state and federal level over the sample period (1993–2008), something which aids the identification of the structural model’s parameters.

2.2 Trends in household behaviour since 1996

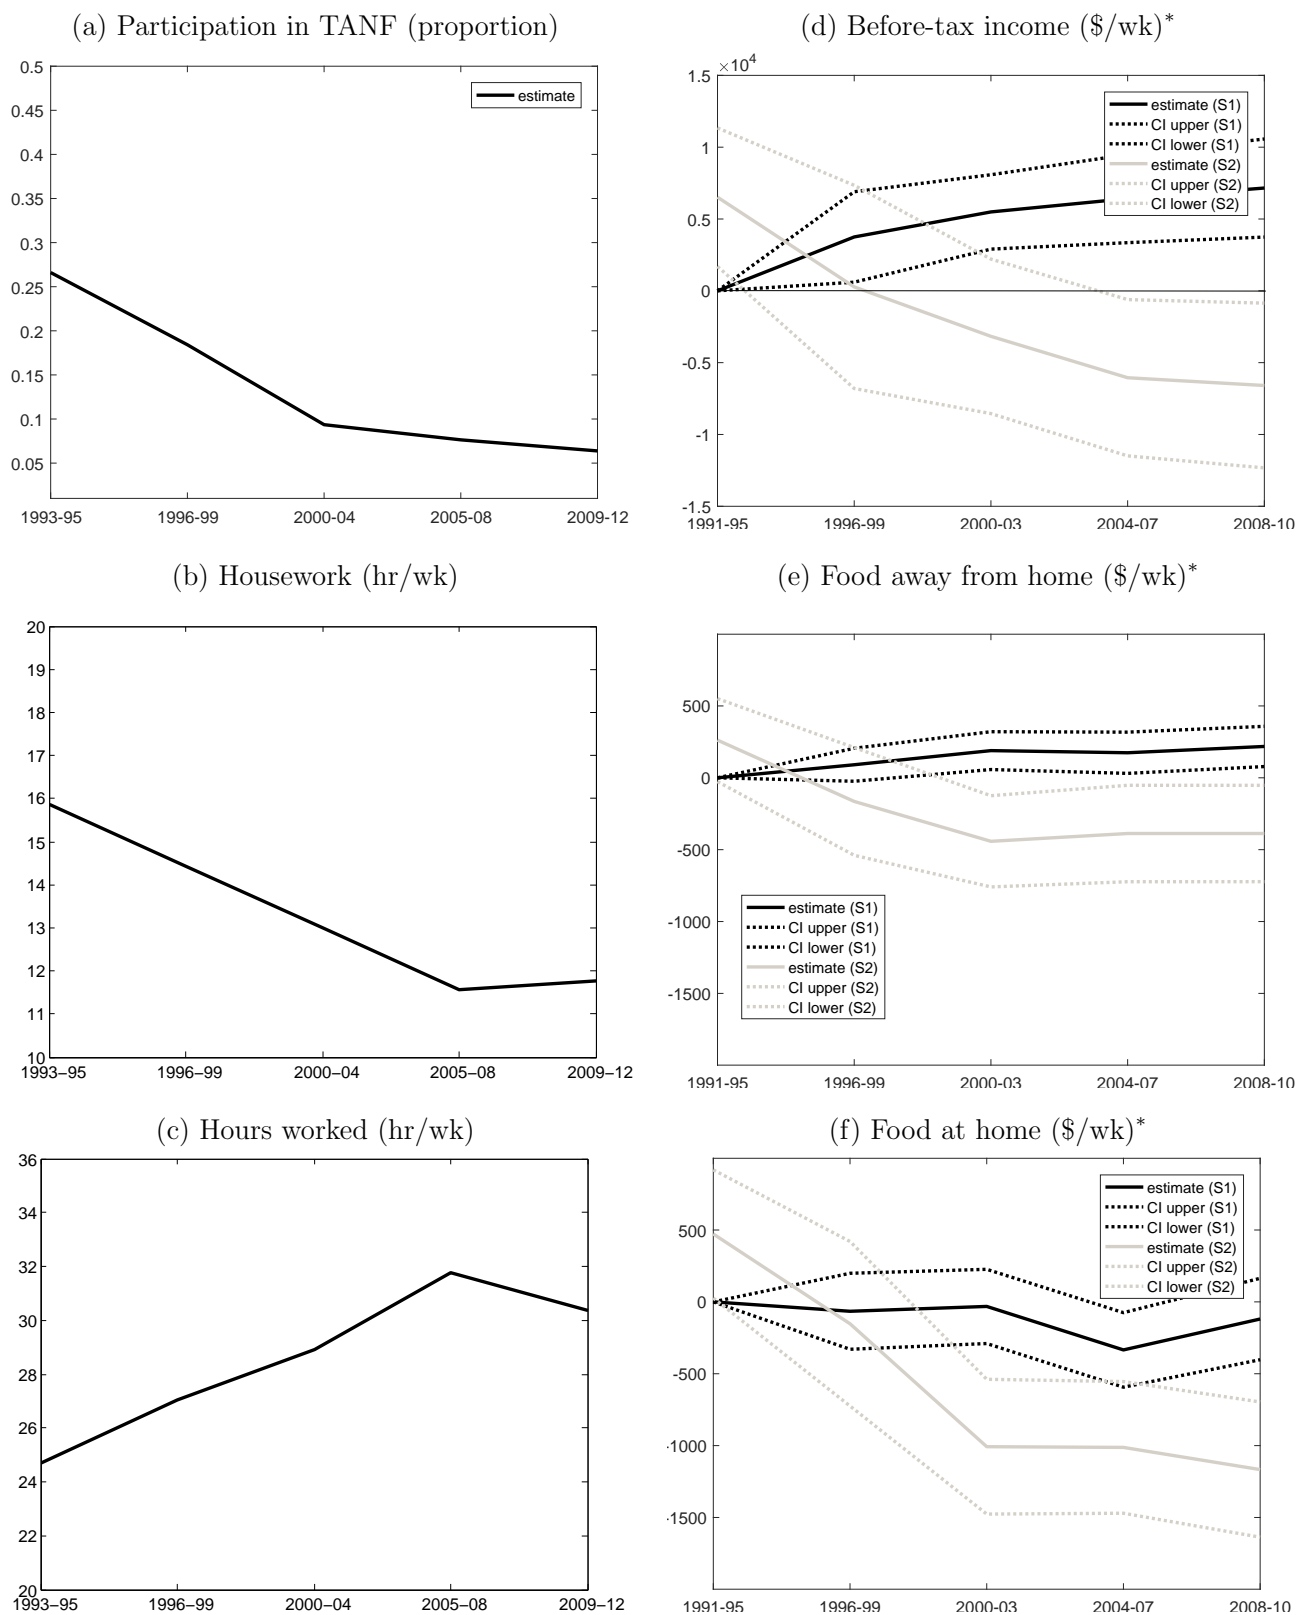
One of the major objectives of TANF was to reduce welfare dependency, and since 1996 there has been a marked decline in caseloads: by 2008, the number of TANF enrollees had fallen by two-thirds (see panel (a) of Figure 2). The other major aim of TANF was to promote employment: panel (b) shows that hours worked by single mothers (without a college degree) indeed increased over the same epoch.

This paper focuses on the impact of TANF on the allocation of resources within the household, particularly on the resources allocated to children. Panels (d)–(f) display trends in household income and food expenditures separately for single women with children (S1) and without (S2), controlling for the number of children, age, education, race, and state (see the notes to the figure for further details). This extends Meyer and Sullivan’s (2004) analysis out to 2010: and as in their work, the comparison allows me to control for any shocks that would have affected both groups similarly, such as a change in the relative price of food away from home. Any remaining differences in these trends may therefore be ascribed to something that has affected single women *with* children differently from those without, the introduction of TANF being a plausible candidate.

Meyer and Sullivan (2008) find that sole mothers have increased their hours worked since 1996, and that this has come at the expense of non-market work. Consistent with Meyer and Sullivan (2008), I find a decline in housework of around 5 hr/wk between 1993 and 2008 (see panel (c)). I also find that sole mothers have increased their expenditure on food away from home, a trend not observed for single women without children (panel (e)). In contrast, expenditure on food at home for sole mothers has declined, if anything (panel (f)). Before-tax income also increased significantly for single women with children, but not for those without, reflecting the increase in their labour supply (panel (d)). Overall, these figures illustrate a decline in home production and an increase in market work following the introduction of TANF.

Standard measures used to evaluate how children welfare has changed since the introduction of TANF are based on income or consumption at the household level, and do not take

Figure 2: Trends in sole mothers' behaviour



• Sole mothers without a college degree. Sources: CPS (panels (a)–(b)); A(H)TUS (panel (c)); CE (panels (d)–(f)). Dollar values are in year 2000 dollars.

* Panels (d)–(f) extends the work of Meyer and Sullivan (2004) to 2010. Figures report coefficients on time dummies (1993–95, 1996–99, 2000–03, 2004–07 and 2008–10), interacted with whether or not the single woman has a child (S1: has a child; S2: does not); base case is a single woman with children in 1993–95. Coefficients are estimated by an OLS regression of the reported outcome on the preceding dummies, and additionally controls for: the number of children, the woman's age, education (high school, some college dummies), race (black, white dummies), and region (nine) fixed effects.

into account home production, the value of maternal time in activities other than market work, and the distribution of resources within the household. The structural model estimated in this paper will subsequently be shown to closely match the trends described above (see Section 5.2), and will also allow me to uncover the underlying changes in child welfare (see Section 6.2).

3 Structural model

The model elaborated in Section 3.1–3.2 (henceforth ‘the structural model’) provides a detailed treatment of maternal time use: both to explain the trends described above, and because the allocation of maternal time not spent in market work – i.e. between leisure, time with children and housework – is likely to be important for child welfare. It differs from other household models estimated in the literature along two dimensions (see e.g. Apps and Rees, 1996; Cherchye, De Rock, and Vermeulen, 2012; Blundell and Shephard, 2011). Firstly, child utility is modelled separately from mothers’ utility and the home production of a public good. Secondly, children’s utility depends on this public good, which is produced from mother’s time in housework and public consumption expenditures. As I only observe data on the inputs into the children’s utility function, its shape is determined by the mother’s choices regarding the allocation of household resources. (Thus a key assumption underlying the welfare analysis of Section 6 below is that the mother is altruistic towards her children.)

To keep the complexity of the modelling and the estimation procedure within reasonable bounds, I abstract as far as possible from the intertemporal aspects of the household’s decision making. The structural model developed below is therefore static: it describes the household’s resource allocation problem within a given period. I discuss in Section 3.3 how the solution to the structural model may be rendered consistent with that of the household’s intertemporal choice problem. This consistency permits the parameters of the structural model to be identified from data on households’ actual choices, even though those choices arise from the households’ intertemporal optimisations. My arguments here build upon the two-stage budgeting approach of Blundell and Walker (1986), extended so as to take account of the intertemporal aspects of TANF enrolment (due to lifetime limits). I also discuss how the estimated model may be used to reliably evaluate counterfactual changes to the tax and welfare system, subject to the qualifications discussed at the end of Section 3.3.

3.1 Preferences

The mother's utility is a weighted sum of: her private utility, u ; the children's utility, K ; a disutility τ from participating in the labour force; and a disutility, ψ , from being enrolled in AFDC/TANF,

$$U = u + \delta_k K - \tau \cdot \mathbf{1}\{h_m > 0\} - \psi \cdot d, \quad (3.1)$$

where h_m denotes (weekly) hours in paid employment, and $d = 1$ if the mother enrolls in TANF (and is zero otherwise). Her private utility has the CES form

$$u(c_m, l_m, q) = \log(\gamma_{m,c}^{1-\eta_m} c_m^{\eta_m} + \gamma_{m,l}^{1-\eta_m} l_m^{\eta_m} + \gamma_{q,l}^{1-\eta_m} q^{\eta_m})^{1/\eta_m} \quad (3.2)$$

where c_m denotes mother's consumption, l_m leisure, and q the public good. ($\gamma_{m,c} + \gamma_{m,l} + \gamma_{m,q} = 1$; and $\gamma_{m,c}, \gamma_{m,l}, \gamma_{m,q} \in [0, 1]$, $\eta_m \leq 1$.)

As noted above, I model children's utility separately from their mother's, to facilitate the measurement of child welfare. Children's utility takes as inputs: mother's time with children t_m ; children's private consumption, c_k ; and the (home-produced) public good, q ; it is assumed to have the CES form

$$K(c_k, t_m, q) = \log(\gamma_{k,c}^{1-\eta_k} c_k^{\eta_k} + \gamma_{k,t}^{1-\eta_k} t_m^{\eta_k} + \gamma_{k,q}^{1-\eta_k} q^{\eta_k})^{1/\eta_k} \quad (3.3)$$

($\gamma_{k,c} + \gamma_{k,t} + \gamma_{k,q} = 1$; and $\gamma_{k,c}, \gamma_{k,t}, \gamma_{k,q} \in [0, 1]$, $\eta_k \leq 1$). I include mothers' time with children in the model, because it has been demonstrated to be important for children's cognitive and non-cognitive development (see e.g. Brooks-Gunn and Markman 2005; ?; Phillips 2011). The constant elasticity of substitution (CES) functional form adopted here nests (at $\eta_k = 0$) the Cobb-Douglas specification that has been used in some previous work (see e.g. Del Boca, Flinn, and Wiswall, 2014).

The quantity of the public good produced by the household is given by the (constant-returns-to-scale) CES production function

$$q(c_q, q_m) = (\gamma_{q,c}^{1-\eta_q} c_q^{\eta_q} + \gamma_{q,m}^{1-\eta_q} q_m^{\eta_q})^{1/\eta_q},$$

where q_m denotes the mother's time devoted to housework, and c_q the expenditure on consumption goods which are used as inputs in the production of the public good, which I refer to as 'public consumption'. ($\gamma_{q,c} + \gamma_{q,m} = 1$; and $\gamma_{q,c}, \gamma_{q,m} \in [0, 1]$, $\eta_q \leq 1$.) This CES specification has been used in prior literature to model home production (see e.g. Aguiar and Hurst, 2007).

AFDC/TANF enrolment The AFDC/TANF enrolment disutility ψ decomposes as

$$\psi = \psi_{\text{st}} + \psi_{\text{it}}. \quad (3.4)$$

ψ_{st} represents the ‘static’ non-pecuniary costs incurred when enrolling in AFDC/TANF. For example, to be eligible for TANF payments without meeting the work requirements, a mother would usually need to be engaged in certain other activities, such as actively searching for work or participating in a job training programme. (The extent of federal funding that states receive for TANF is related to the proportion of their caseload meeting the work requirements, and thus states generally make it more difficult for individuals to participate in TANF without meeting these requirements.) Those who do meet the work requirement still face non-negligible administrative costs of applying for benefits. ψ_{st} is accordingly allowed to vary, depending on whether or not work requirements are met (or in the case of AFDC, whether or not the mother works).

ψ_{it} is intended to capture the ‘intertemporal’ cost of enrolling in TANF in the current period, due to foregoing the option of enrolling in (a year’s worth of) TANF in the future. (ψ_{it} is therefore set to zero for AFDC.) As discussed in Section 3.3 below, ψ_{it} is needed to make the structural model consistent with the mother’s intertemporal optimisation problem. (See Section 4.6 below for the moments used to separately identify ψ_{st} and ψ_{it} .)

Receipt of other transfers and benefits Participation in the other transfer programmes considered in this paper – the EITC, SNAP, CTC and the CDCTC – is not modelled. Rather, when calculating the household’s after-tax income (see Section 3.2 below) it is assumed that eligible households always receive the payments (or tax credits) to which they are entitled. Relative to TANF, which had a take-up rate of around 30.7 per cent in 2013 (Crouse and Macartney, 2016), the take-up rates for these programmes are much higher, so this seems an acceptable simplifying assumption. (According to Dahl and Lochner (2011), the take-up rate for the EITC was over 80 per cent in every year of the sample period, while that for SNAP always exceeded 50 per cent, reaching a maximum of 87 per cent in 2007). Moreover, TANF’s job search and training requirements impose burdens on enrollees that have no counterpart in these other programmes.

Preference heterogeneity and parametrisation To permit heterogeneity in preferences, a number of parameters are allowed to vary with both observables and unobservables. In the following paragraphs, ϵ signifies an i.i.d. standard normal disturbance (all of which are mutually uncorrelated), and x a vector of observed household characteristics.

Regarding the mother's preferences, her weight on the child's utility is parametrised as

$$\delta_k = \exp(x'_\delta \beta_\delta + \sigma_\delta \epsilon_\delta), \quad (3.5)$$

where x_δ includes a constant, and dummies for: the presence of two or more children, and a child aged five or under. To ensure that the mother's utility weights $\gamma_{m,\cdot}$ satisfy the the stipulated adding-up and range constraints, I first specify

$$\tilde{\gamma}_{m,c} = \exp(x'_{m,c} \beta_{m,c} + \sigma_{m,c} \epsilon_{m,c}) \quad \tilde{\gamma}_{m,l} = \exp(x'_{m,l} \beta_{m,l} + \sigma_{m,l} \epsilon_{m,l}) \quad \tilde{\gamma}_{m,q} = 1$$

where $x_{m,l}$ and $x_{m,q}$ both include a constant and an education dummy (indicating whether the mother has some tertiary education), and $x_{m,l}$ additionally includes the residual from the mother's wage equation (see Section 4.3 below). The utility weights are then constructed from the $\tilde{\gamma}_{m,\cdot}$ as per $\gamma_{m,i} = \tilde{\gamma}_{m,i} / \sum_{j \in \{c,l,q\}} \tilde{\gamma}_{m,i}$ for $i \in \{c, l, q\}$. The parameter governing the elasticity of substitution is specified as

$$\eta_m = 1 - \exp(x'_{m,\eta} \beta_{m,\eta})$$

where $x_{m,\eta} = x_\delta$ above.

The children's utility weights are specified as $\gamma_{k,i} = \tilde{\gamma}_{k,i} / \sum_{j \in \{c,t,q\}} \tilde{\gamma}_{k,i}$, where

$$\tilde{\gamma}_{k,c} = \exp(x'_{k,c} \beta_{k,c} + \sigma_{k,c} \epsilon_{k,c}) \quad \tilde{\gamma}_{k,t} = \exp(x'_{k,t} \beta_{k,t} + \sigma_{k,t} \epsilon_{k,t}) \quad \tilde{\gamma}_{k,q} = 1$$

and $x_{k,c}$ and $x_{k,t}$ both include a constant and a dummy for the presence of a child aged five or under. The elasticity is parametrised as $\eta_k = 1 - \exp(x'_{k,\eta} \beta_{k,\eta})$, where $x_{k,\eta}$ additionally includes a dummy for the presence of two or more children. The public good production weights $\tilde{\gamma}_{q,\cdot}$ are constructed analogously from

$$\tilde{\gamma}_{q,c} = \exp(x'_{q,c} \beta_{q,c} + \sigma_{q,c} \epsilon_{q,c}) \quad \tilde{\gamma}_{q,m} = 1,$$

while $\eta_q = 1 - \exp(x'_{q,\eta} \beta_{q,\eta})$, where both $x_{q,c}$ and $x_{q,\eta}$ include a constant and a dummy for the presence of a child aged five or under.

The enrolment disutilities ψ_{st} and ψ_{it} are parametrised as $\exp(x' \beta + \sigma \epsilon)$, where ϵ is a standard Gaussian disturbance. For ψ_{st} , x consists of dummies for the year groups 1993–95, 1996–99, 2000–04, and 2005–08; both β and σ are allowed to vary depending on whether the mother meets the mandated work requirements (for TANF; for AFDC, it varies according to whether or not she works). For ψ_{it} , x collects the same year group dummies, and additionally dummies for the presence of: two or more children, and a child aged five or under.

3.2 Constraints

The mother’s utility in (3.1) will be maximised subject to her time and budget constraints. I also discretise hours choices so that sole mothers choose their hours of work h_m from $\mathbb{H} = \{0, 10, 20, 30, 40, 50, 60\}$. This may be justified by the presence of frictions in the labour market (see Hoynes, 1996; Blundell and Shephard, 2011) and simplifies the computation of the household’s optimal choices, which would otherwise be greatly complicated by the non-convexity of the budget set.

The mother’s (weekly) time constraint is

$$l_m + h_m + t_m + q_m \leq 105, \quad (3.6)$$

(nine hours per day are excluded for sleep and personal care). The budget constraint is

$$\underbrace{c_m + c_k + c_q}_{\text{non-durables consumption}} + \underbrace{n_{o5}p_{o5}[h_m - a_{o5} - 30]_+}_{\text{cost of childcare for children aged 6–13}} + \underbrace{n_{u5}p_{u5}[h_m - a_{u5}]_+}_{\text{cost of childcare for children aged 0–5}} \leq e(h_m, w_m; d) + y - s \quad (3.7)$$

where: $[x]_+ = \max\{x, 0\}$; and for children aged 6–13, p_{o5} , a_{o5} , and n_{o5} respectively denote the price of childcare, the hours of informal free childcare available (such as would be provided by a neighbour or relative), and number of children (and p_{u5} , a_{u5} and n_{u5} similarly for children aged 5 and under). Note that the ‘30’ in equation (3.7) reflects the 30 hr/wk that children aged 6–13 would spend in school. y denotes nonlabour income, s savings, w_m the mother’s wage, h_m her hours worked, and $e(h_m, w_m; d)$ her implied after-tax earnings. Recall that $d = 1$ if the mother elects to receive TANF payments; and zero otherwise.

Calculation of after-tax income $e(\cdot)$ includes: taxes at the federal and state levels; tax credits (EITC, CTC, CDCTC); SNAP (‘food stamps’); and AFDC/TANF (see Section 2.1 above and Appendix D for further details). As noted above, for all programmes *except* TANF it is assumed that the household always receives the full transfer for which it is eligible. I also assume that SNAP receipts are always less than the household’s desired expenditure on food, and so can be treated as a pure cash transfer.

3.3 Solution and two-stage budgeting

Solving the static model Let $\mathcal{C} = (c_m, c_k, c_q, l_m, h_m, t_m, q_m)$ denote the mother’s choices over time and expenditure, and

$$\mathcal{U}(\mathcal{C}) = u[c_m, l_m, q(q_m, c_q)] + \delta_k K[c_k, t_m, q(q_m, c_q)] \quad (3.8)$$

her total utility, abstracting from TANF enrolment costs. Given $y - s$ and her TANF enrolment status $d \in \{0, 1\}$, her feasible set of choices $\mathbb{C}(y - s, d)$ is determined by her weekly time and budget constraints ((3.6) and (3.7)); recall that I also require $h_m \in \mathbb{H}$. Thus according to the model above (‘the structural model’) her optimal choices for \mathcal{C} and d may be determined by solving

$$\max_{d \in \{0, 1\}} \max_{\mathcal{C} \in \mathbb{C}(y-s, d)} \{\mathcal{U}(\mathcal{C}) - d \cdot \psi\} = \max_{d \in \{0, 1\}} \left\{ \max_{\mathcal{C} \in \mathbb{C}(y-s, d)} \mathcal{U}(\mathcal{C}) - d(\psi_{\text{st}} + \psi_{\text{it}}) \right\}. \quad (3.9)$$

I now turn the question of how the structural model may be reconciled with the mother’s intertemporal decision problem.

Two-stage budgeting A sole mother may be regarded as solving an intertemporal decision problem, the solution to which involves choices over: (i) savings (or equivalently, asset holdings), (ii) enrolment in TANF, and (iii) the allocation of time and income within each period. My structural model is concerned with the second and third of these choices. The modelling of TANF enrolment is complicated by the presence of lifetime limits, which give TANF enrolment an intertemporal aspect that is not shared by other welfare programmes. Together with the mother’s (unmodelled) savings decision, this raises two issues. The first concerns whether the parameters of the structural model can be identified from data on household’s choices, given that those observed choices arise from the solution to those household’s intertemporal optimisation problems. This will be the case if the solutions to the structural model and intertemporal choice problem are consistent with each other; such consistency is established immediately below. The second issue, which is discussed at the end of this section, concerns the extent to which the model can be reliably used to conduct counterfactual experiments.

In view of (3.8) above, and recalling that ψ_{st} embodies the ‘static’ cost of enrolling in TANF (see Section 3.1 above), assuming additive separability the mother’s intertemporal problem may be put into recursive form as

$$\mathcal{V}(A, D) = \max_{s \in \mathbb{R}} \max_{d \in \{0, 1\}} \left\{ \max_{\mathcal{C} \in \mathbb{C}(y-s, d)} \mathcal{U}(\mathcal{C}) - d \cdot \psi_{\text{st}} + \rho \mathbb{E} \mathcal{V}[A'(s), D + d] \right\}, \quad (3.10)$$

where A denotes her asset holdings, $D \leq 5$ the number of times that she has enrolled in TANF prior to the current period, \mathcal{V} is the current value of the problem, ρ is the mother’s discount rate, and \mathbb{E} is computed conditional on the information available in the current period. Asset holdings A' in the next period evolve according to $A'(s) = (1 + r')(A + s)$, where r' is the (unknown) rate of return on the household’s portfolio between the current

period and the next; non-labour income in the next period is given by $y' = r'(A + s)$.

A key implication of (3.10) is the following. Conditional on her (optimal) choices of s and d , the mother's optimal choice of \mathcal{C} may be computed simply by maximising $\mathcal{U}(\mathcal{C})$ over $\mathbb{C}(y - s, d)$. This corresponds exactly to the maximisation problem (3.9) solved in the structural model, with respect to the \mathcal{C} variables. So far as TANF enrolment (d) is concerned, it is also evident from (3.10) that the mother chooses to enrol ($d = 1$) if and only if

$$\max_{\mathcal{C} \in \mathbb{C}(y-s,1)} \mathcal{U}(\mathcal{C}) - \max_{\mathcal{C} \in \mathbb{C}(y-s,0)} \mathcal{U}(\mathcal{C}) \geq \psi_{\text{st}} + \rho \mathbb{E} \{ \mathcal{V}[A'(s), D] - \mathcal{V}[A'(s), D + 1] \} = \psi_{\text{st}} + \psi_{\text{it}}.$$

I interpret the second term on the r.h.s. of the inequality as the ‘intertemporal cost’ of enrolling in TANF, due to foregoing the option of (a year’s worth of) TANF in the future; this term is subsumed in ψ_{it} in the structural model. Although ψ_{it} should depend on the household’s current levels of A and D , I do not observe these variables, and thus any variation in these across the sample is treated as a form of unobservable heterogeneity, i.e. it is absorbed into the stochastic component of ψ_{it} . (Regarding the separate identification of ψ_{st} and ψ_{it} , see Section 4.6 below.) Thus by including ψ_{it} in the structural model, the solution to that model (for \mathcal{C} and d) can be made consistent with that of the household’s intertemporal optimisation, thereby permitting the identification (and thence estimation) of the structural model’s parameters.

Regarding the counterfactual exercises conducted in Section 6 below, in most of these TANF is eliminated and replaced by an alternative policy – the parameters of TANF itself are not adjusted. In such cases, the question of how the (unmodelled) intertemporal trade-offs related to TANF enrolment might be altered by the counterfactual policy simply does not arise. On the other hand, I must assume that savings decisions – also omitted from the structural model – are not materially different under the counterfactual. I justify this on the grounds that sole mothers – particularly those with lower earnings, and who are therefore more likely to be eligible for cash welfare – have low asset holdings (Del Boca, Flinn, and Wiswall, 2014) and behave essentially as hand-to-mouth consumers. Moreover, the hypothetical policy changes that I consider are intended to be permanent: and as such would induce less substitution in consumption across time periods than a temporary policy change.

4 Estimation procedure

I estimate the structural model using simulated method of moments (SMM), drawing upon multiple datasets to build up an accurate picture of the allocation of household resources. To

estimate the model using only one dataset would require that dataset to contain both time use and consumption data. Such datasets are particularly rare: the most notable instances of previous work estimating household models using only one dataset have involved either the Panel Study of Income Dynamics (Child Development Supplement) or the Longitudinal Internet Studies for the Social Sciences (a Dutch panel), as used by Del Boca, Flinn, and Wiswall (2014) and Cherchye, De Rock, and Vermeulen (2012) respectively. The small size of these datasets, particularly for the subgroup of sole mothers considered here (considerably less than 500 observations in either survey) is problematic due to the measurement errors associated with time diaries and expenditure surveys.

Moreover, because I will subsequently consider counterfactual policy interventions that differentially affect households according to their income level, the model needs to accurately describe the full distribution of household behaviours. Accordingly, as discussed in Section 3.1 above, the model’s parametrisation allows for much behavioural heterogeneity across households. The estimation of such a richly parametrised model by SMM requires a collection of sample moments that is highly informative as to the distribution of the allocation of household resources. This makes accurate data on the full range of household choices highly valuable, and for this reason I draw upon five datasets to estimate the model, as explained in Section 4.1.

The mechanics of the estimation procedure are described in Sections 4.2–4.3. The choice of sample moments, and the identification of the model parameters, are discussed in Sections 4.4–4.5, following which I address the problem of how the static and intertemporal costs of TANF enrolment (ψ_{st} and ψ_{it}) may be separately identified and estimated (Section 4.6).

4.1 Data sources

To estimate the model I draw upon five datasets: the American Time Use Survey (ATUS, 2003–08), the American Heritage Time Use Survey (AHTUS, 1995), the Consumer Expenditure Survey (CE, 1993–2008), the Current Population Survey (CPS, 1993–2008) and the Survey of Income and Program Participation (SIPP, 1996, 2001, 2004, and 2008). Sample sizes – not counting those households that are excluded on the grounds noted below – range from more than 3500 in the CE to almost 50,000 in the CPS. As described in more detail subsequently, the CE provides the ‘base sample’ of households in the structural model. That model is estimated by matching moments simulated from the model against sample moments constructed from the A(H)TUS, the CE and the CPS. The SIPP is used to estimate equations for the price of childcare and the availability of informal care, which I use to impute

Table 1: Descriptive statistics

		ATUS (03-08)	CE (93-08)	CPS (93-08)	SIPP (96-08)*
<i>Household demographics</i>					
1	High school diploma (% of sample)	41	35	41	40
	Some tertiary education (%)	39	45	40	35
3	At least one child under 5 (%)	43	20	31	26
	More than one child (%)	42	44	41	29
	Mother's age (mean)	36	38	35	37
<i>Labour supply</i>					
6	Paid work (hr/wk)		32	29	26
	Participation rate (%)		83	78	68
	Wages (mean for participants; \$/hr)		11.3		
<i>Weekly earnings (\$/wk)</i>					
9	Mean		375		
	25th percentile		180		
	50th percentile		326		
	75th percentile		530		
<i>Intra-household allocation</i>					
13	Public consumption (\$/wk)		251		
	Private consumption (\$/wk)		302		
	Non-labour income [†] (\$/wk)		173		
16	Time with children (hr/wk)	6.0			
	Housework (hr/wk)	11.7			
	<i>n</i>	8484	3624	49676	2986

• Sample: as described in Section 4.1 (see also Appendix A). Sample period in parenthesis: 93-08 denotes 1993-2008, etc. All dollar values are deflated to the year 2000. Statistics computed using inverse population weights.

* SIPP only records data for the years 1996, 2001, 2004 and 2008.

† Total non-durables consumption expenditure less after-tax earnings (constructed as per Section 4.4).

these missing childcare variables to households in the structural model (see the final part of Section 4.2 below).

The A(H)TUS and the CE record time use and consumption expenditures at a much finer level of detail than is needed for the broad categories (leisure, time with children, etc.) to which the model refers. This has the considerable advantage of permitting me to exercise discretion as to how different uses of time and expenditure, should be classified. For example, I construct a measure of mother’s time with children that is focussed on those activities that have been found to be beneficial for child development, as discussed in Section 4.4 below. Likewise, my measure of ‘public consumption’ includes expenditure on goods that are not ‘public’ in the strict sense, such as food consumed at home, because it is intended to capture all forms of expenditure that, together the mother’s housework, are inputs to home production.

For all datasets, I restrict the sample by retaining only households in which: the mother is aged 18–58 (or 25–58, for the CE); the mother is not self-employed or in the armed forces; the household contains no more than one adult (i.e. one individual over the age of 18; this restriction is not applied to the SIPP); the household’s residence is occupied by only one family; the mother does not have a college degree (only 16 percent of sole mothers have a college degree); and the mother’s hourly wage (if observed) lies in the range \$2.50–\$250 (in 2000 dollars). (See Appendix A for further details, including an explanation as to why a different age range is used for the CE.)

Table 1 presents summary statistics for the ATUS, CE, CPS and SIPP samples. Here, as is the case throughout this paper, all statistics are computed using the inverse population weights provided by these surveys (normalised so that each year receives equal weight). The results suggest that the samples are reasonably similar, particularly when allowance is made for the different epochs to which they refer. The first two rows give the proportion of the sample with only a high school diploma, and the proportion with a higher level of educational attainment (i.e. a vocational qualification; recall that college-educated mothers are excluded from the sample). Across all four datasets these figures differ by a maximum of ten percentage points. Similarly, the mother’s average age does not vary much over the three datasets, differing by only three years (see row 5). Female labour supply patterns are also similar across the CE and CPS (see rows 6 and 7, cols 2 and 3, Table 1).

Table 1 also displays weekly (before-tax) earnings for sole mothers, drawn from the CE. In this sample, 25 per cent of households earn under \$180/wk, or \$9,380/yr (in 2000 dollars), while 50 per cent of households earn less than \$16,900/yr. This compares to an official poverty line of \$11,250/yr for a one-child, sole-parent family, and \$14,150/yr for a two-child, sole-parent family.

4.2 Simulated method of moments

Marshallian demands Let $\mathcal{Y} = (y - s, w_m, p_{u5}, p_{o5}, a_{u5}, a_{o5})$ collect all the variables relevant to the household's budget constraint, and ϕ the household's preference parameters. (Though it would be tedious to enumerate ϕ in its entirety, it records the values of the parameters δ_k , τ , ψ_{st} and ψ_{it} in (3.1), $\gamma_{m,c}$, $\gamma_{m,l}$, $\gamma_{q,l}$ and η_m in (3.2), etc.) Recall that to permit preference heterogeneity, some elements of ϕ are allowed to vary parametrically across households. I accordingly write $\phi = \phi(x, \xi; \theta)$, where x denotes the household demographic variables (age, education, race etc.), which I regard as exogenous, ξ the random disturbances capturing the unobservable preference heterogeneity in the model, and θ the parameters indexing $\phi(\cdot)$. Thus, with $\delta_k = \exp(x'_\delta \beta_\delta + \sigma_\delta \epsilon_\delta)$ as per equation (3.5) above, ϵ_δ is an element of ξ , and β_δ a subvector of θ .

Since the feasible choice set $\mathcal{C}(y - s, d)$ depends on the whole of \mathcal{Y} (not merely $y - s$), while each of $u(\cdot)$, $K(\cdot)$ and $q(\cdot)$ are parametrised by ϕ , the solution to (3.9) yields the Marshallian demands

$$\mathcal{C} = f(\mathcal{Y}; \phi) = f[\mathcal{Y}; \phi(x, \xi; \theta)] = g(\mathcal{Y}, x, \xi; \theta) \quad (4.1)$$

In principle, (4.1) should allow one to derive the density of \mathcal{C} conditional on (\mathcal{Y}, x) , and thence to estimate θ by maximum likelihood. In practice, there are two obstacles to this route: (i) the evaluation of f requires solving the household's problem numerically (no analytical solution for the household's problem is available), which makes calculation of the likelihood challenging; and (ii) more significantly, as discussed in the introduction to this section, the joint distribution of \mathcal{C} is available in only a few small and – for our purposes – unrepresentative datasets.

Construction of the moments For these reasons I estimate the model by simulated method of moments. Since this does not require the recovery of the entire joint (conditional) distribution of \mathcal{C} , I am able to combine several datasets to build up an accurate picture of the observed allocation of household resources, as described in more detail in Section 4.4 below. Regarding the selection of suitable moments, these could ordinarily be drawn from the household's first-order conditions: but with the present model, most of these would involve non-separable functions of variables appearing in different datasets. I therefore instead select the means, standard deviations and correlations of household choices, computed both unconditionally and conditional on the household characteristics noted in Section 4.4 below.

More formally, the simulated moments are constructed by averaging across households

$i \in \{1, \dots, n\}$ and simulation draws $s \in \{1, \dots, S\}$, as per

$$\hat{m}_n(\theta) := \frac{1}{n} \sum_{i=1}^n \frac{\pi_i}{S} \sum_{s=1}^S \mathbf{m}[g(\mathcal{Y}_{(i)}, x_{(i)}, \xi_{(is)}; \theta); x_{(i)}], \quad (4.2)$$

where the π_i 's are the inverse population weights from the CE, normalised so that $\sum_{i=1}^n \pi_i = n$; \mathbf{m} is a vector-valued transformation, chosen so as to deliver a vector of sample means of the levels, squares and cross products of elements of \mathcal{Y} ; and S is the number of simulation draws. The final estimates are computed with $S = 10$. A further transformation puts the moments into their required form (as means, standard deviations, and correlations); I write this more compactly as $\hat{\mu}_n(\theta) = \varphi[\hat{m}_n(\theta)]$.

An estimate of θ is then obtained by minimising the following criterion,

$$Q_n(\theta) = \sum_{k=1}^K \frac{(\hat{\mu}_{nk}(\theta) - \bar{\mu}_{nk})^2}{\hat{\omega}_k^2} = \|\hat{\mu}_n(\theta) - \bar{\mu}_n\|_{\hat{W}}^2 \quad (4.3)$$

where $\bar{\mu}_n = \varphi(\bar{m}_n)$ denotes the corresponding vector of sample moments, constructed from data on household's actual choices; \hat{W} is a diagonal matrix with k th diagonal element $\hat{\omega}_k^{-2}$; and $\|x\|_A := x'Ax$. The inverse weights $\hat{\omega}_k^2$ are computed from the estimated asymptotic variance of the sample moment $\bar{\mu}_{nk}$; that is, as an estimate of ω_k^2 in $n^{1/2}(\bar{\mu}_{nk} - \mu_k) \xrightarrow{d} N[0, \omega_k^2]$.³ Although this does not correspond to the theoretically optimal weighting, it at least ensures that sample moments from the same dataset are weighted proportionally to the precision with which they are estimated, and renders Q_n invariant to the units in which the households' choices are measured.

The estimation procedure described by (4.2) and (4.3) has to be modified somewhat, to take account of: (i) the imputation (via simulation) of those elements of \mathcal{Y} that are missing from the CE; and (ii) the introduction of additional moments, whose simulated counterparts are computed on a slightly different basis from $\hat{\mu}_n(\theta)$, to ensure the identification of ψ_{it} . The former is dealt with in Section 4.3; the latter is deferred to Section 4.6.

Computation of $\hat{\theta}$ The minimisation of Q_n is undertaken as follows. I first draw a large number (20,000) of candidate values of θ , uniformly from a large, bounded subset of the parameter space (chosen conservatively, so as to span all the economically plausible values of the parameters). I evaluate Q_n at each of these points, and retain $\theta^{(k)}$'s corresponding

³Note that ω_k^2 does not depend on the size of the sample used to estimate $\bar{\mu}_{nk}$: were I to instead weight the moments by their approximate *finite-sample* variances $\hat{\omega}_k^2/n$, I would grossly over-weight moments constructed from the CPS, which has a much larger sample size than either the CE or the A(H)TUS (see Table 1).

to the 1000 smallest values of $Q_n(\theta^{(k)})$ thus obtained. From each of these points, I run 300 iterations of a Gauss-Newton optimisation routine, and then retain the 500 best-performing optimisations, following which I run a further 600 iterations of Gauss-Newton. (I use the implementation of the Gauss-Newton routine provided by version 10 of the Artelys Knitro software package.) I then take the best 100 points remaining, and iterate Gauss-Newton from these until convergence. The estimates reported in this paper correspond to the parameters delivering the minimum value of Q_n achieved by this algorithm. (Four optimisers converged to this point of the parameter space; others terminated at higher values of Q_n .) As a check on my results, I also ran ten independent simulated annealing chains, each of 1 million draws in length, starting from the ten best randomly-sampled values (as drawn at the first stage of the preceding algorithm), followed by a Gauss-Newton optimisation, iterated to convergence: but in no case did this beat the estimates obtained by the preceding algorithm.

Inference Under the assumption that the model is correctly specified (with parameters θ_0), and the number of simulation draws S is fixed as $n \rightarrow \infty$, standard results on minimum-distance-type estimators (see e.g. Gourieroux, Monfort, and Renault, 1993) imply that $n^{1/2}(\hat{\theta} - \theta_0) \xrightarrow{d} N[0, V]$, where

$$V = (D'WD)^{-1}D'W\Sigma WD(D'WD)^{-1} \quad (4.4)$$

for $D = \partial_{\theta}\mu(\theta_0)$ the Jacobian of $\mu(\cdot) = \text{plim } \hat{\mu}_n(\cdot)$ at θ_0 , Σ the limiting variance of $n^{1/2}[\hat{\mu}_n(\theta_0) - \bar{\mu}_n]$, and $W = \text{plim } \hat{W}$. (Note that V depends on S through Σ .) The estimation of V requires estimators for D and Σ . D can be estimated by numerically differentiating $\hat{\mu}_n$ with respect to θ at $\hat{\theta}$. The estimation of Σ is somewhat complicated by the use of multiple datasets, and so a discussion of this is deferred to Appendix B.3. (Note that inferences based on V ignore the additional variability introduced by the imputation procedure described in the following section.)

4.3 Imputation (via simulation) of missing variables

Recall that I use the CE (1993–2008) as to construct the ‘base sample’ of households in the model: this provides me with data on x , w_m (for most households), and $y - s$; this last being computed as the excess of total consumption expenditures (including childcare) over after-tax earnings (see (3.7) above). However, w_m is not observed for working mothers, while childcare prices and the availability of informal care, $(p_{u5}, p_{o5}, a_{u5}, a_{o5})$, are not observed for any household in the CE. I therefore estimate a collection of models that allow values of these missing variables to be simulated, prior to estimating the structural model. To describe these,

let each of x_m , x_y , x_p , and x_a denote subvectors of the household demographic variables x ; and each of ϵ_y , ϵ_w , $\epsilon_{p,u5}$, $\epsilon_{p,o5}$, $\epsilon_{a,u5}$, and $\epsilon_{a,o5}$ i.i.d. standard Gaussian disturbances.

Wages and non-labour income (less savings) These are modelled according to

$$\begin{bmatrix} y - s \\ \log w_m^* \end{bmatrix} = \begin{bmatrix} x_y' \beta_y \\ x_w' \beta_w \end{bmatrix} + L \begin{bmatrix} \epsilon_y \\ \epsilon_w \end{bmatrix} \quad (4.5)$$

where L is a lower triangular matrix, and w_m^* is the market wage (observed only if the mother participates in the labour force). x_w consists of: education dummies (high school, some tertiary education), two race dummies (black and white), mother's age and age-squared, and time and (nine) region dummies (the regions are as defined in Appendix A.6). x_y additionally include the Case-Shiller house price index, interacted with an indicator for whether the household owns their residence (similarly to Lise and Seitz, 2011).

(4.5) is estimated using the CE (1993–2008). The equation for $y - s$ is first estimated by OLS. The residual from this equation is added as a regressor to the mother's wage equation (to permit estimation of L), which is itself then estimated by OLS, using a Heckman selection correction to account for the mother's labour force participation decision. The first-stage probit for participation includes an urban dummy (and the $y - s$ residual), in addition to all the variables in x_w .

Childcare Equations for the missing childcare variables are estimated using the SIPP (1996, 2001, 2004, and 2008), which provides a detailed record of the cost of childcare, hours spent in childcare and the availability of free informal childcare. The price of childcare for children aged 6–13 is modelled as

$$p_{o5} = p_{\min} + \max\{x_p' \beta_{p,o5} + \sigma_{p,o5} \epsilon_{p,o5}, 0\}, \quad (4.6)$$

where $p_{\min} = 2$, and x_p includes: age of the mother; two race dummies (white and black); a dummy for tertiary education; and (nine) region dummies. (4.6) is estimated via a censored (Tobit) regression. To account for the fact that p_{o5} is only observed for households that purchase childcare, when estimating (4.6) I augment the r.h.s. by the residual from a censored regression for the hours of childcare purchased (see Appendix B.1 for further details). To aid identification, the r.h.s. of that equation includes, in addition to x_p , a Hispanic dummy and the age of the youngest child. The Hispanic dummy is significant at the 1 per cent level, and is included because Hispanic extended families tend to be larger and live closer together (Kimmel and Connelly (2007) use the presence of additional adults in a household

for a similar purpose).

The hours of free informal childcare available to the household (i.e. care provided by relatives) for children aged 6–13 are given by

$$a_{o5} = \max\{x'_a \beta_{a,o5} + \sigma_{a,o5} \epsilon_{a,o5}, 0\}, \quad (4.7)$$

where x_a includes the same variables as x_p , and additionally an indicator for whether the family owns their residence. a_{o5} is subjected to further censoring and sample selection, since only $\min\{a_{o5}, h_m\}$ is observed, and that only if $h_m > 0$. Accordingly, I estimate (4.7) by a doubly-censored regression regression, with 0 as the left and h_m as the right censor point (see Appendix B.1). To control for the right censoring and sample selection, I augment the r.h.s. of (4.7) with the residual from an estimated censored regression for hours worked (the r.h.s. of that equation includes x_a , a Hispanic dummy and the age of the youngest child). Analogous equations to (4.6) and (4.7) hold for children aged 0–5, and are estimated in the same way (see Appendix B.1).

Imputation via simulation Estimates of the parameters of (4.5)–(4.7) are reported in Appendix C.1. Given these estimates, and data on $(y - s)_{(i)}$ and $x_{(i)}$ for the i th household, I can use these equations to simulate values for those elements of \mathcal{Y} that are missing, by drawing values for the disturbances $(\epsilon_w, \epsilon_{p,u5}, \epsilon_{p,o5}, \epsilon_{a,u5}, \epsilon_{a,o5})$ from a multivariate standard normal. (Since $y - s$ is always observed, the residual $\hat{\epsilon}_y$ from equation (4.5) is used to construct a simulated value for w_m .) Let $\mathcal{Y}_{(is)}$ denote the simulated values obtained for the s th simulation applied of the i th household, where $\mathcal{Y}_{(is)}$ records the observed values of $(y - s)_{(i)}$ and – when these are available – $w_{m,(i)}$. Then the simulated moments at θ can be written as

$$\hat{m}_n(\theta) := \frac{1}{n} \sum_{i=1}^n \frac{\pi_i}{S} \sum_{s=1}^S \mathbf{m}[g(\mathcal{Y}_{(is)}, x_{(i)}, \xi_{(is)}; \theta); x_{(i)}], \quad (4.8)$$

so that the only difference between the preceding and equation (4.2) is that the (partially unobserved) $\mathcal{Y}_{(i)}$ has been replaced by the (partially simulated) $\mathcal{Y}_{(is)}$.

4.4 Sample moments: selection and construction

As noted above, the ‘base sample’ of households in the structural model is drawn from the CE (1993–2008); this is used to construct simulated moments for a given value of θ . To construct the sample moments against which these are to be matched, I draw upon four datasets: the ATUS (2003–08) and the AHTUS (1995) for data on mother’s time use; the CE (1993–2008) for household consumption expenditures, wages, and non-labour income;

Table 2: Moments used to estimate the model

	t_m	q_m	h_m *	TANF [†]	c_m	c_k	c_q	c_{pr}
Source [‡]	A	A	C	C	E	E	E	E
Means	★	★	★		Est. [§]	Est. [§]	★	★
Std. dev.	★	★	★				★	★
Corr. with	h_m, w_m	h_m	w_m				h_m, w_m, c_{pr}	h_m, c_q
Year groups	★	★	★	★				

• For the means and standard deviations marked ★, and for all correlations, I match both the unconditional moment, and moments conditioned on each of the following: educational attainment (groups: at least a high school diploma; some tertiary), number of children (more than one child; at least one child under 6; at least one child under 3), and mother’s (predicted) wage lying in the lowest quintile.

* I also match the (conditional) proportions of mothers working at least one hour per week, and working sixty or more hours per week.

† Proportion of sample enrolled in TANF and meeting work requirements, and proportion enrolled and not meeting these requirements (or in 1993–95, for AFDC, the proportions working and not working). See Section 4.6 for the additional TANF-related moments that are included to ensure the separate identification of ψ_{st} and ψ_{it} .

‡ Dataset used to compute moments listed below: A = A(H)TUS, C = CPS, E = CE.

§ Estimates of means for these variables constructed using the approach of Dunbar, Lewbel, and Pendakur (2013): see Appendix B.2. Only the unconditional mean is matched.

|| I match means conditional on the year groupings: 1993–95, 1996–99, 2000–02, 2003–05, and 2006–08. t_m is only observed from 2003 onwards, while q_m is not observed in 1996–2002.

and the CPS (1993–2008) for labour supply. (Although labour supply data could also be sourced from the ATUS or CE, wherever possible it is drawn from the CPS, due to its larger sample size.) These allow me to construct moments that capture the key features of the distribution of the allocation of household resources. The associated large sample sizes ensure that these sample moments are very precisely estimated.

As summarised in Table 2, the selected moments consist primarily of means, standard deviations, and correlations involving household choices. I also match the proportion of sole mothers enrolled in AFDC/TANF, separately for those meeting and not meeting the work requirements (or merely working, in the case of AFDC). The sample correlations between household choices that may be computed is limited by their availability in a common dataset; since hours worked appears in all datasets, I match its correlations with all other household choices. I also match correlations between the mother’s wage (imputed for non-participants using the model (4.5)) and hours worked, time with children and public consumption.

Both conditional and unconditional moments are fitted. The conditioning is based on: mothers’ education (high school, some tertiary), the age of the youngest child (whether a child aged 0–5 years or 0–2 years is present), the number of children in the household (more than one child), and whether the mother’s wage is below the 20th percentile (imputed

wage for non-participants). Conditioning on this final group is motivated by low-income households being the most likely to receive government assistance, and as such being of particular relevance to this paper. For those variables indicated in the table, I also condition on the following year groups: 1993–95, 1996–99, 2000–02, 2003–05, and 2006–08.

Mothers’ time use Moments for mothers’ time use include the (conditional) means and standard deviations of hours worked, housework, and time with children (leisure is excluded, since all time uses must sum to 105 hr/wk).

The ATUS (2003–08) provides a single 24-hour time diary for each mother on a randomly selected day of the week. These time diaries are used to construct measures of parents’ time with children, and time spent in housework. (The ATUS separately records the number of hours worked per week.) My measure of time with children is an aggregate of time spent in activities that have been found beneficial for child development, including: reading to children (Scarborough and Dobrich, 1994), talking with parents (Tamis-LeMonda, Bornstein, and Baumwell, 2001), mealtime conversation (Snow and Beals, 2006), and novel experiences and places (Phillips, 2011). The AHTUS (1995) provides similar time diary data, but the significant differences in the classification of activities across the two surveys prevents the construction of a comparable measure of parents’ time with children. I therefore only draw data on time spent in housework from the AHTUS. (Further details are provided in Appendix A.1.)

A shortcoming of this data is that the time diaries only record activities during a 24-hour window, whereas I am interested in average behaviour over the course of a week. First moments (and certain cross-correlations) involving time use may still be consistently estimated from such time-diary data, but measures of variability will be upwardly biased. I accordingly adjust the sample standard deviations computed for these variables, using estimates of this bias computed from another dataset (from the Netherlands) which provides time diaries for an entire week for each respondent (see Appendix A.1 for details).

Consumption expenditures I match the (conditional) means and standard deviations of public and total private consumption expenditures, measures of which are constructed from the CE. Public consumption includes such categories as food, utilities, mortgage interest, and rent (see Appendix A.3). Private consumption is computed as total household expenditure less public consumption and childcare costs.

Two additional moments help to calibrate the breakdown of private consumption between household members. Since it is only possible to unambiguously assign expenditures on a few commodities recorded in the CE (e.g. clothing) to either the mother or her children, I use

the procedure developed by Dunbar, Lewbel, and Pendakur (2013) to infer the average breakdown of total private consumption expenditures, between mothers and children (see Appendix B.2 for further details). This provides me with estimates of the unconditional means of mother’s and children’s private consumption, which are matched by the estimation procedure; note that no conditional means are matched, so this only contributes two additional moments (out of a total of 173).

4.5 Identification of structural parameters

Two questions pertinent to the identification of the structural parameters are: (i) what exogenous variation is present in the sample that might help to identify those parameters; and, since those parameters are to be estimated by SMM, (ii) do the moments provide a sufficiently rich description of the data to capture the principal trade-offs faced by households? I give a brief discussion of both these questions here, before providing a further numerical analysis of parameter identification in Section 5.4 below.

The evolution of tax and welfare policies over the sample epoch provides the main source of exogenous variation in households’ budgets. Changes to these policies have clearly affected household behaviour: for instance, the reduced generosity of TANF relative to AFDC has likely contributed to the marked increase in the time that sole mothers have spent in paid work, instead of in housework, since 1996 (recall Section 2.2 above). The moments matched by the estimation procedure include the means of certain household choices conditional on each of five subperiods, thereby capturing these aggregate trends. As shown in Section 5.2 below, the model achieves a good fit to these trends, suggesting that it accurately predicts households’ responses to (exogenous) variation in the tax and welfare system.

Even stronger evidence for this comes from the external validation exercise performed in Section 5.3 below. TANF greatly widened regional disparities in the generosity of welfare payments, by devolving much of the responsibility for welfare provision to the states (see Section 2.2 above). Although none of the moments used to estimate the model condition on location – and thus omit any region-specific trends in household resource allocation that might have resulted from this policy variation – I show below that the estimated model is able to match region-specific trends in household time use and consumption over 1993–2008 with a high degree of accuracy. I interpret the model’s success in capturing how households respond to changes to welfare policies as implying that variation in these policies must be highly influential in pinning down the values of the estimated parameters.

With respect to the second aspect to identification, it should first be noted that in contrast to much of the prior empirical literature that estimates intra-household models

using only consumption data (Apps and Rees, 1996; Lise and Seitz, 2011; Dunbar, Lewbel, and Pendakur, 2013), I have at my disposal both time use and consumption data, which taken together provide a far richer description of the allocation of household resources. The correlations among household choices, and between those choices and wages, ought to be particularly informative as to the trade-offs households face when deciding between alternative uses of those resources. Although it is not possible for me to include all such correlations – because not all variables appear in a common dataset – I am able to include correlations between all household choices and hours worked.

4.6 Identifying the TANF enrolment disutilities

While the total disutility $\psi = \psi_{st} + \psi_{it}$ associated with TANF enrolment is identified from the observed enrolment rates (conditional on the year groups indicated), this does permit the separate recovery of its components, ψ_{st} and ψ_{it} . Recall that ψ_{it} captures the ‘intertemporal’ cost of enrolling in TANF, due to lifetime limits. Solving the model with $\psi_{it} = 0$ thus yields the enrolment rate, implied by the model’s other parameters, that would counterfactually obtain if lifetime limits were absent. I propose to match this to a sample estimate of the proportion of households would enrol in TANF, in the absence of lifetime limits. ψ_{it} (and thus also $\psi_{st} = \psi - \psi_{it}$) may then be identified from the estimated effect of lifetime limits on TANF enrolment.

Several papers have previously estimated this effect: here I follow the approach of Grogger and Michalopoulos (2003). They exploit the fact that when TANF was introduced, lifetime limits were not binding on households in which all children were aged over 12 – whereas all other households were to some extent constrained by these limits. More formally, their approach involves estimating the following probit regression for TANF enrolment

$$d = \mathbf{1}\{x'_d\beta_d + \beta_1z_1 + \beta_2z_2 + \epsilon > 0\}, \quad (4.9)$$

where: $d = 1$ if the mother enrolls in TANF; $\epsilon \sim N[0, 1]$; x_d is a vector of controls; and z_1 and z_2 are dummy variables, with $z_1 = 1$ for a mother who was never subject to time limits (because her youngest child was over 12 when time limits were introduced to her state), and $z_2 = 1$ for a mother who was only partially exposed to time limits (because her youngest child was born before time limits were introduced). I control for the age composition of the children in the household, by including dummies for the presence of children aged 0–2, 3–5, and 6–12 in x_d . I also include dummies for the periods 1996–99, 2000–02, 2003–05, 2006–08, to allow the disutility from enrolling in TANF to change over time. x_d additionally includes: the mother’s age and age-squared; race dummies (white and black); a dummy for whether

there is only one child in the household; a dummy for whether the mother has some tertiary education; and nine region effects (for the regions given in Appendix A.6).

Across 1996 to 2008, an average of 18 per cent of sole mothers in the CPS were enrolled in TANF. Using the procedure outlined above, this figure would rise to 22 per cent if time-limits were eliminated.

The parameters of (4.9) are estimated using the CPS (see Table 10 in Appendix C.2). Conditional on x_d , the estimated probability that a household that would enrol in TANF in the *absence* of time limits is given by $\Phi(x'_d \hat{\beta}_d + \hat{\beta}_1)$. Averaging this over the entire CPS sample (1996–2008) gives an implied TANF enrolment probability in the absence of lifetime limits of 22 per cent, as compared with an observed enrolment rate of 18 per cent. To construct the additional sample moments \bar{v}_n that are matched by the estimation procedure, I compute sample averages of $\Phi(x'_d \hat{\beta}_d + \hat{\beta}_1)$, both unconditionally and conditional on the presence of: a child aged 0–5, and two or more children. Doing this for each of time periods 1996–99, 2000–02, 2003–05 and 2006–08 yields a total of 12 additional moments.

To construct the simulated moments that correspond to \bar{v}_n , I partition the parameters of the structural model as $\theta = (\vartheta, \psi_{it})$, so that solving the model with $\psi_{it} = 0$ yields the enrolment rates that would obtain under ϑ , in the absence of lifetime limits. Let $\hat{v}_n(\vartheta)$ denote the vector of simulated moments thus computed. The overall criterion that I minimise to estimate the structural parameters is (4.3) augmented by the squared weighted distance between \bar{v}_n and $\hat{v}_n(\vartheta)$, i.e.

$$Q'_n(\vartheta, \psi_{it}) = \|\hat{\mu}_n(\vartheta, \psi_{it}) - \bar{\mu}_n\|_{\hat{W}}^2 + \|\hat{v}_n(\vartheta) - \bar{v}_n\|_{\hat{V}}^2$$

where \hat{V} is a diagonal weight matrix whose elements are computed on the same basis as those of \hat{W} (see Section 4.2 above).

5 Estimates and model fit

In this section I demonstrate that the model performs well according to three different criteria. Firstly, the mother’s labour supply elasticities are within the range of those obtained in previous studies (Section 5.1). Secondly, the model achieves a good (in-sample) fit to the moments used in estimation (Section 5.2). In particular, the model is able to replicate those aggregate trends in household behaviour that were discussed in Section 2.2 above.

Thirdly, the model is able to match ‘out-of-sample’ variation in household choices: that is, variation observed during the sample period, but which is not directly targeted by the estimation procedure. To this end, I show that observed region-specific trends in household

choices over the sample period correlate closely with those predicted by the model. Such geographical variation is of particular interest, as it has doubtless been influenced by the manner in which TANF extended considerably greater freedom to states in setting the parameters of their welfare programmes (recall Section 2.1 above). The good performance of the model in this respect lends greater credibility to the counterfactual exercises subsequently conducted in Section 6, which involve changes to welfare policy different from those observed historically.

Finally, I also investigate empirically how different subsets of the selected moments contribute to the identification of the model’s parameters, to complement the heuristic discussion of parameter identification in Section 4.4 above (see Section 5.4). I quantify this by computing how the precision of the parameter estimates, in terms of their estimated standard errors, would be affected if certain moments were ignored by the estimation procedure. Consistent with the heuristic discussion of identification (see Section 4.4 above), time-use moments and correlations among household choices significantly strengthen the identification of the structural parameters.

5.1 Parameter estimates

The average values of the behavioural parameters implied by the estimates, conditional on certain household characteristics, are displayed in Table 3. For the estimated CES utility and production functions, I report the (more readily interpretable) elasticities of substitution ς_i rather than the corresponding curvature parameters η_i ; the two are related via $\varsigma_i = 1/(1 - \eta_i)$. As discussed in Section 3.1 above, the model allows for considerable heterogeneity (both observable and unobservable) in household preferences. The number of underlying structural parameters is accordingly rather large (48), and so a complete listing of their estimates is deferred to Appendix C.3. (Estimates of the parameters for the models used to impute missing variables are given in Appendix C.1.)

The estimates indicate that the children’s utility function places a relatively large weight on the public good, with $\gamma_{k,q} = 0.42$ as compared with $\gamma_{m,q} = 0.18$ for mothers (see row 9). Thus children significantly benefit from expenditure on housing (c_q includes imputed rents), home cooked meals, and having a clean and tidy house. In households with children aged 0–5, the children’s utility places a higher weight on the mother’s time with children, and the degree of substitutability between inputs is lower than it is on average (see rows 8 and 12).

In home production, the elasticity of substitution between maternal time and consumption is estimated to be 1.07 (row 13, col 1). This is slightly higher than the average elasticity of substitution estimated for inputs into the child utility function, at 0.83 (row 11, col 2).

Table 3: Estimated behavioural parameters (group means)

Param.	Subgroup	Estimate	Param.	Subgroup	Estimate
<i>Mother's weight on children's utility</i>			<i>Utility cost of programme enrolment</i>		
1	δ_k	Full sample 1.00* —	TANF, not meeting work req.		1.00* —
2		Child aged 0–5 1.13 (0.27)	TANF, meets work req.		0.05 (0.56)
3		More than 1 child 1.17 (0.11)	AFDC, not working		0.27 (2.48)
4		Some tertiary educ. 0.99 (0.01)	AFDC, working		0.23 (1.98)
<i>Mother</i>			<i>Children</i>		
5	$\gamma_{m,c}$	Full sample 0.42 (0.11)	$\gamma_{k,c}$	Full sample	0.47 (0.07)
6	$\gamma_{m,l}$	Full sample 0.40 (0.07)		Child aged 0–5	0.33 (0.06)
7	$\gamma_{m,q}$	Full sample 0.18 (0.11)	$\gamma_{k,t}$	Full sample	0.11 (0.01)
8	ς_m^\dagger	Full sample 3.52 (1.03)		Child aged 0–5	0.19 (0.02)
9			$\gamma_{k,q}$	Full sample	0.42 (0.08)
10				Child aged 0–5	0.48 (0.05)
<i>Public good production</i>					
11	$\gamma_{q,c}$	Full sample 0.75 (0.02)	ς_k^\dagger	Full sample	0.83 (0.29)
12	$\gamma_{q,m}$	Full sample 0.25 (0.02)		Child aged 0–5	0.21 (0.25)
13	ς_q^\dagger	Full sample 1.07 (0.33)		More than 1 child	0.55 (0.22)

• Subgroup means of for estimated values of selected behavioural parameters; standard errors computed via the delta method. (The underlying parameter estimates are given, along with standard errors, in Appendix C.3.) Subgroups: *child aged 0–5* collects households with at least one child aged 0–5; *more than 1 child* households with two or more children; *some tertiary educ.* households in which the mother has some tertiary education (but less than a college degree).

* Mean values of δ_k within subgroups are expressed *relative* to the average value of δ_k obtained across the full sample, which has been normalised to unity (ex post). Similarly, utility costs of enrolment are expressed relative to the cost of enrolling in TANF, when not meeting the work requirement.

† $\varsigma_i = 1/(1 - \eta_i)$ denotes the implied elasticity of substitution, for $i \in \{m, k, q\}$.

It is somewhat lower than the estimates of 1.25, 2.44, and 1.80 obtained respectively by McGrattan, Rogerson, and Wright (1997), Chang and Schorfheide (2003) and Aguiar and Hurst (2007); however, it should be noted that those estimates are not exactly comparable with my own, since they are based on a sample of households that includes both sole mothers and married couples.

The disutility from TANF enrolment is estimated to be significantly higher when work requirements are met, than when they are not (see rows 1–2, col 2). Indeed, the estimated disutility from enrolment for those meeting the work requirements is so low as to have apparently little effect on enrolment: the model predicts that if this disutility were eliminated, the proportion of the sample enrolled in TANF *and* meeting work requirements, averaged over 1996–2008, would rise by only 0.9 percentage points to 3.4 per cent. For AFDC, on the other hand, the estimated disutility suffered by enrollees differs little between those working (0.23) and those not (0.27; see rows 3–4, col 2). These differences between the estimates for two programmes are entirely consistent with TANF more heavily sanctioning those who fail to meet its work requirements.

Maternal labour supply elasticity I compute the aggregate wage elasticity implied by these estimates by: (i) incrementing the wage of every household (including non-participants) in the structural model by \$1/hr (in 2000 dollars); (ii) solving each household’s optimisation problem under this counterfactually higher wage, to obtain their hours worked $h_{m,i}(w_{m,i}^*)$; and then (iii) computing the ratio of the percentage changes in the model averages of hours worked and the hourly wage, which can be rewritten as

$$\frac{\frac{\Delta \bar{h}}{\bar{h}}}{\frac{\Delta \bar{w}}{\bar{w}}} = \frac{\frac{\frac{1}{n} \sum_{i=1}^n \pi_i [h_{m,i}(w_{m,i} + 1) - h_{m,i}(w_{m,i})]}{\frac{1}{n} \sum_{i=1}^n \pi_i h_{m,i}(w_{m,i})}}{\frac{\frac{1}{n} \sum_{i=1}^n \pi_i w_{m,i}}{\frac{1}{n} \sum_{i=1}^n \pi_i \cdot 1}} \quad (5.1)$$


This yields an estimated elasticity of 0.67; the elasticity at the intensive margin is 0.27. (This is computed by restricting the averages on the r.h.s. of equation (5.1) to the subgroup of mothers who work in the baseline.) These estimates are similar to those obtained for sole mothers by (see Blundell, Costa Dias, Meghir, and Shaw, 2016).

5.2 Model fit: in-sample performance


Aggregate trends in household behaviour As discussed in Section 2.2 above, household behaviour changed systematically over 1993–2008 – possibly in response to the introduction of TANF – with in particular a very discernible shift from housework into market work (see also Meyer and Sullivan, 2008). Plots of observed and model-implied paths for

Figure 3: Trends in selected moments moments (means)


(a) Enrolled in AFDC/TANF
(full sample)*




(b) Enrolled in AFDC/TANF
(households with children aged 0-5)*




(c) Hours worked
(hr/wk)




(d) Proportion in
employment



(e) Time with mother
(hr/wk)



(f) Housework
(hr/wk)



* Solid line: sample moments; dashed line: simulated (fitted) moments.

these selected household choices are displayed in Figure 3. (Recall from Section 4.4 that the means of these choices, conditional on each of the year groups displayed, are among the moments used for SMM.)

The model fits the trend decline in the AFDC/TANF enrolment rate both overall, and for households with children aged 0–5 (panels (a)–(b)). Regarding maternal time use, the model is able to match most of the increase in hours worked over 1993–2002, although hours worked is underestimated in the earlier part of the sample by about 3 hr/wk. Labour force participation is systematically overestimated by around 5 percentage points. The model also matches around three-quarters of the 4 hr/wk decline in housework over 1993–2008. Thus the model is able to account for the main trends in household behaviour since the introduction of TANF, which lends greater credibility to the counterfactual exercises subsequently performed.

First and second moments First and second moments for mothers’ choices appear in Table 4: both simulated and sample moments are reported (standard errors of the sample moments in parentheses). On the whole, these moments are fit reasonably well. As noted in Section 4.4 above, I have included moments that condition on mothers’ wages being below the 20th percentile, because these households are of particular concern in this paper, being the most likely to receive government assistance. The model fits these moments well, with each simulated moment lying within (or very nearly within) the 95 per cent confidence interval for the corresponding sample moment in all cases (see rows 2 and 9).

Time use First and second moments for mothers’ time use appear in columns 1–6 of Table 4.

The model closely fits all moments describing mothers’ time with children (see cols 1–2). The unconditional simulated mean is within one standard deviation of the sample mean, and much of the variation across subgroups is reproduced by the model, with all the model-implied conditional means lying within two standard deviations of their sample counterparts (rows 2–7, cols 1–2). Housework moments (cols 3–4) are fit slightly less well. The unconditional mean of time spent in housework is over-estimated by one hour per week, and the model predicts there to be greater variation in housework time across subgroups than is actually present in the sample.

Although the unconditional mean of hours worked is fit well, the model does not fully capture the observed covariation of labour supply with observables (cols 5–6). For subgroups in which hours worked is less than the unconditional mean, the model somewhat exaggerates the difference between the conditional and unconditional means. For example, for households

Table 4: Time-use and consumption moments

	Time w/child. (hr/wk)		Housework (hr/wk)		Hours worked (hr/wk)		Public cons. (\$/wk)		Private cons. (\$/wk)			
	Siml.	Smpl.	Siml.	Smpl.	Siml.	Smpl.	Siml.	Smpl.	Siml.	Smpl.		
	1	2	3	4	5	6	7	8	9	10		
Mean												
1	<i>Full Sample</i>		6.16	6.00 (0.21)	12.48	11.66 (0.30)	28.66	29.46 (0.09)	252.96	250.60 (3.50)	302.65	301.67 (5.60)
2	<i>Wage in lowest quintile</i>		7.45	7.83 (0.46)	11.72	12.27 (0.61)	22.60	24.71 (1.37)	219.99	211.22 (5.76)	237.46	218.18 (8.79)
<i>Education</i>												
3	High school (at least)		6.09	6.14 (0.34)	13.68	12.06 (0.51)	29.27	30.07 (0.15)	260.42	241.83 (4.74)	297.08	297.92 (9.34)
4	Some tertiary		5.94	6.05 (0.31)	11.71	10.80 (0.39)	28.91	33.09 (0.13)	260.30	272.02 (5.92)	334.38	340.77 (8.67)
<i>Family structure</i>												
5	Children under 5		10.35	9.60 (0.41)	8.42	12.00 (0.48)	22.13	26.28 (0.18)	235.55	236.27 (7.15)	268.73	279.71 (13.6)
6	Children under 2		10.08	10.86 (0.58)	8.91	11.95 (0.62)			225.41	220.56 (9.91)	252.59	264.63 (20.4)
7	More than 1 child		7.43	6.97 (0.34)	12.87	11.88 (0.47)	25.70	28.34 (0.15)	244.84	256.40 (5.53)	316.85	302.90 (9.12)
Standard deviation												
8	<i>Full Sample</i>		4.53	4.49 (0.09)	7.75	7.66 (0.12)	18.52	17.55 (0.05)	112.81	103.90 (5.02)	167.49	171.66 (5.35)
9	<i>Wage in lowest quintile</i>		5.30	4.91 (0.18)	7.12	7.89 (0.27)	???	???	96.09	86.00 (6.35)	130.90	123.66 (8.99)
<i>Education</i>												
10	High school (at least)		3.86	4.41 (0.14)	7.48	7.63 (0.16)	18.79	17.06 (0.09)	111.26	90.71 (4.90)	166.25	169.27 (9.46)
11	Some tertiary		4.49	4.51 (0.14)	7.72	7.33 (0.17)	17.97	15.57 (0.11)	118.02	113.57 (8.86)	172.51	177.25 (7.79)
<i>Family structure</i>												
12	Children under 5		5.11	5.08 (0.13)	7.53	7.65 (0.18)	17.41	18.09 (0.08)	114.88	96.39 (5.80)	127.80	169.52 (11.8)
13	Children under 2		4.81	5.17 (0.18)	7.32	7.69 (0.23)			110.53	88.45 (7.89)	126.75	169.88 (24.4)
14	More than 1 child		5.30	4.99 (0.15)	8.67	7.65 (0.21)	17.72	17.77 (0.08)	110.11	107.89 (9.45)	167.11	178.94 (8.65)

• Simulated moments (*Siml.*) corresponding to the model parameter estimates, and sample moments (*Smpl.*) computed from the data. Standard errors of sample moments in parentheses. All dollar values are in year 2000 dollars.

with wages below the 20th percentile, average hours worked is under-predicted by two hours, as compared with 0.8 hours overall.

The standard deviations are also fit well, with the standard deviation of hours worked slightly over-predicted (rows 8–14, cols 5–6). There is little variation in the standard deviation of the time-use moments across the subgroups.

Consumption The model’s fit to the means and standard deviations of public and private consumption is reported in columns 7–10 of Table 4. Both unconditional means are fit closely (see row 1). The model is successfully matches the observed variation in consumption expenditures across subgroups. For example, the \$39/wk difference between the unconditional mean of public consumption, and that for the low-income subgroup, is exactly reproduced in the model (compare rows 1 and 2, cols 7–8), which also matches three-quarters of the corresponding \$83/wk difference in private consumption (compare rows 1 and 2, cols 9–10).

5.3 Model fit: external validation

This subsection demonstrates the model’s success in matching variation in household behaviour not targeted by the estimation procedure, which is likely driven by variation in tax and welfare policy. Firstly, recall from Section 2.1 that the introduction of TANF led to increased heterogeneity in the provision of cash welfare across states. Here I use this geographic variation, which is not matched in estimation, as a means of providing an external check on the validity of the estimates. Specifically, for a number of household choices – hours worked, housework, and public and private consumption – I compute the mean change between 1993–95 and 2004–08, for each of nine regions (as defined in Appendix A.6). The observed changes, and those implied by the estimated model, are particularly close for public consumption and housework (i.e. the two inputs used in the production of the public good): the correlation between the two being 0.73 and 0.79 respectively. This correlation is also positive, though not as pronounced, for private consumption (0.51) and hours worked (0.30).

These trends in household behaviour have been shaped by the large and widening regional disparities in welfare policies since 1996, and thus the model’s close match to these trends suggests that it is able to replicate households’ responses to variation in the tax and welfare system. From this I draw two conclusions. Firstly, this enhances the credibility of the counterfactual exercises performed using the model, which rely on it accurately predicting households’ responses to policy variation that is ‘out of sample’ by construction. Secondly, that the ‘in sample’ variation in the tax and welfare system – that is, that variation captured by the moments used to estimate the model – is highly influential in pinning down the estimated values of the model’s parameters.

Table 5: Percentage increase in asymptotic ‘standard error’
when indicated moments are excluded

		Public good	Mother’s utility	δ_k	Child. utility	TANF disutility		
	#					ψ_{st}	ψ_{it}	
	1	2	3	4	5	6	7	
<i>Means and standard deviations</i>								
1	t_m	16	1	5	14	28	11	-10
	h_m	21	-1	0	8	-4	11	3
	q_m	17	3	8	10	-9	-3	-21
4	c_m, c_k	2	1	15	30	2	-3	-4
	c_{pr}	14	0	-1	0	-1	0	-1
	c_q	14	0	1	1	4	1	0
<i>Correlations</i>								
7	$h_m \cdot (t_m, q_m)$	14	21	8	10	29	0	-7
	$h_m \cdot (c_{pr}, c_q)$	14	9	79	49	25	7	17
	$c_{pr} \cdot c_q$	7	4	109	10	-18	28	0
	$w_m \cdot (h_m, t_m, c_q)$	21	13	40	102	39	-1	5
<i>TANF enrolment rates</i>								
11	observed	9	-3	31	0	-4	98	732
	‘counterfactual’	12	6	5	46	-2	1420	1098
13	<i>A(H)TUS moments</i>	47	172	40	47	59	12	-18

- For the group of parameters indicated in each column, the table reports the percentage increase in the ‘standard error’ (square root of the the largest eigenvalue of the asymptotic variance submatrix pertaining to those parameters) that would result if the moments indicated in each row were ignored (rather than fitted) by the estimation procedure. ‘#’ denotes the number of moments.

5.4 Identification: a numerical analysis

It is often difficult to theoretically establish the identification of complex structural models. The usual approach to this problem is to discuss heuristically how the moments selected identify individual parameters, as per Section 4.4 above. Here I complement this approach, by determining numerically which moments convey the information that most substantially reduces the uncertainty with which parameters are estimated. The results of this exercise confirm that the correlations between household choices (and wages) are particularly important for pinning down the parameter estimates, as are the moments computed from time-use data (drawn from A(H)TUS).

The approach taken here involves computing the asymptotic variance according to the

formula given in (4.4) above, but with the weight matrix W replaced by another in which the entries corresponding to certain moments have been zeroed out. Comparing this with the asymptotic variance as originally computed (using W), indicates how the precision with which the parameters are estimated is sensitive to the matching of those (now-excluded) moments by SMM.⁴ For individual parameters, it is appropriate to compute how the associated ‘standard error’ – the square root of the appropriate diagonal entry of the asymptotic variance matrix – would increase if a subset of moments were excluded. For groups of related parameters – such as those pertaining to public good production or children’s utility – I define the ‘standard error’ to be the square root of the largest eigenvalue of the submatrix of V pertaining to those parameters. (This can be interpreted as the standard error associated with that linear combination of those parameters that would be least precisely estimated.)

Table 5 reports, for the group of parameters indicated in each column, the percentage increase in this ‘standard error’ that would result if the group of moments given in the row label were ignored by the estimation procedure. (Because the optimal weighting is not used by the SMM estimator, these figures need not be positive, though they overwhelmingly are.) The results reported for the TANF disutility parameters (cols 6–7) are included to verify that this approach yields sensible results. These parameters must surely be pinned down by the moments referring to TANF enrolment rates: if the results of this exercise suggested otherwise, then the whole approach taken in this section might be drawn into question. So it is reassuring that the results are overwhelmingly consistent with these priors: ignoring these moments is found to increase the ‘standard error’ associated with these parameters by an order of magnitude (see rows 11–12).

Consistent with the arguments given in Section 4.4, the results suggest that correlations between household choices (and wages; rows 7–10) are particularly important for identification. This is likely because these moments most clearly convey the households’ trade-offs between the alternative uses of their resources. More specifically, for the parameters of the public good production function, the correlations of hours worked with the two inputs into the public good (public consumption and housework), significantly influence the precision of the estimates. If the correlations of hours worked with housework and time with children were ignored, the ‘standard error’ of the estimates of these parameters would rise by 21 per

⁴An alternative measure of local identification is provided by the Jacobian matrix D : as is well known, a sufficient condition for θ_0 to be (locally) identified is that D should have full column rank. The further that D is from rank deficiency the better that θ_0 will be identified; and so the singular values of D – and of appropriately selected submatrices thereof – could also be used to measure identification. I have preferred to use V (re-computed for alternative groups of moments) rather than D because it is directly interpretable in terms of the precision with which the parameters are estimated. In any case, these two objects are closely related: e.g. it is clearly evident from equation (4.4) that the further D is from rank deficiency, the more precisely that θ_0 will be estimated.

cent (row 7, col 2); a deeper analysis reveals that most of this increase can be ascribed to the correlation between hours worked and housework. For the children's utility function, both these correlations and those between hours worked and consumption are important (see rows 7–8, col 5).

Regarding the mother's utility function, the correlations involving private consumption are notably influential (see rows 8–9, col 3); especially its correlations with hours worked and public consumption. The former may be indicative of the mother's trade-offs between leisure and consumption (since hours worked and leisure are inversely related); the latter of those between private consumption and the public good. Finally, the mother's weight on the children's utility is most sensitive to the correlations involving hours worked and (public and private) consumption, and wages (see rows 8 and 10, col 4). A further disaggregation of these results suggests that, of the latter group, it is the correlations of wages with hours worked and public consumption that are most important.

For the parameters discussed above – i.e. the parameters of the public good production function, children's and mother's utility functions, and the weight on children's utility – the correlations between wages and hours worked, time with children and the public consumption good also contribute importantly to the precision of the estimates (see row 10). Variation in wages alters the total resources available to the household, and thus these correlations help to capture how households allocate additional income between the mother, children and public consumption.

Recall that a major innovation in this paper, relative to much of the prior literature, is the use of *both* disaggregated consumption and time-use data to estimate a model of intra-household resource allocation. The final row of the table records how the precision of parameter estimates would be affected if the moments derived from time-use data were removed (47 from a total of 173). (Note that this measure does not decompose additively, so the figures in row 13 cannot be deduced from those in rows 1–12.) The results clearly indicate the importance of the information carried by the time-use moments for the parameter estimates. Excluding these moments would raise the standard errors of estimates of the parameters of the home production function by 172 per cent, and those of the children's utility function by 59 per cent.

6 Child welfare, TANF and alternative policies

In this section, I use the model to open the black box of household behaviour, so as to quantify the effects of TANF and alternative policies on the children of sole mothers, and to explain the mechanisms underlying these results. To evaluate the actual evolution of child welfare

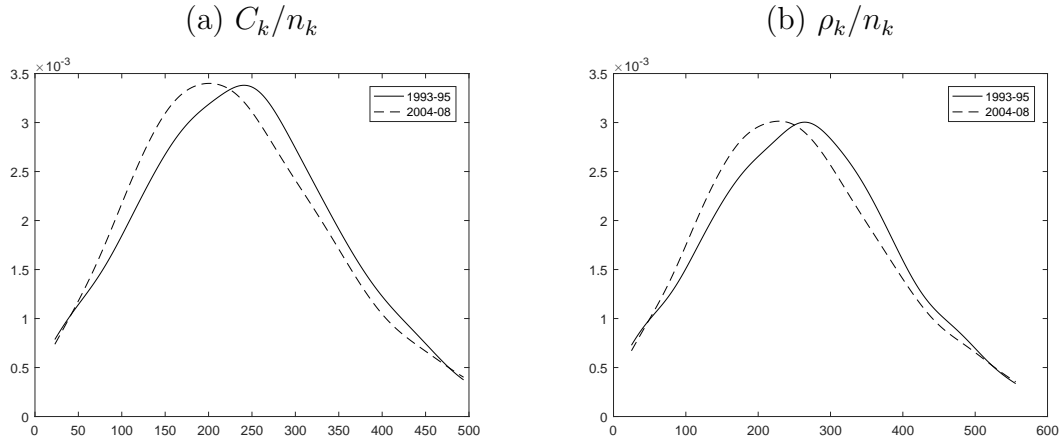
since the introduction of TANF in 1996, I first use the model to construct two measures of the household resources actually received by children (Section 6.1). The distributions of both measures exhibit a clear leftward shift over the sample period, and my calculations of child poverty rates based on each suggest that these increased by four percentage points between 1996 and 2008. These results stand in contrast to previous studies, which find evidence for something of an improvement in child poverty rates since 1996 (Blank, 2002; Meyer and Sullivan, 2004), but do so on the basis of household-level measures of income or consumption, which fail to take account of the allocation of those resources within the household.

Since it is unclear how much of this decline in children’s welfare can be ascribed to PRWORA, I subsequently use the model to perform a number of counterfactual experiments related to PROWRA and other welfare reforms – most notably, the expansion of the EITC – that have been implemented since 1996 (Section 6.2).⁵ The results indicate that replacing AFDC with TANF has indeed had a detrimental effect on child welfare, and goes some way toward explaining the decline in child welfare estimated for 1996–2008. More strikingly, while the PROWRA reforms taken as a whole had a positive effect on labour supply, this had much more to do with the withdrawal of AFDC than the introduction of TANF. If anything, TANF itself has tended to reduce maternal labour supply: its 30 hr/wk work requirements are set much too high for low wage mothers to elicit the desired response. In contrast, the EITC is rather more effective both at promoting maternal labour supply and in targeting resources to children.

The comparison of TANF with the EITC suggests that a policy that increases the net returns to work, rather than sanctioning enrollees who fail to meet mandated work requirements, is likely to better fulfil the objectives of welfare-to-work policies. Two policies of this kind are childcare subsidies and wage subsidies: for both of these I conduct counterfactual exercises in which TANF is alternately replaced by either policy (Section 6.3). The results clearly confirm the greater efficacy of these policies, relative to both TANF and the EITC. With regards to the latter, as I explain below, a shortcoming of the EITC is that it provides benefits on the basis of total household earnings, without taking into account the hours worked to obtain those earnings. It is thus equally generous to high-wage and low-wage

⁵As discussed in Section 3.3 above, in most of these counterfactuals TANF is eliminated and replaced with an alternative policy; the parameters of TANF itself are never adjusted. To this extent, the issue of how the (unmodelled) intertemporal trade-offs pertaining to TANF enrolment might be affected by these counterfactual interventions is irrelevant. However, I do always assume that savings behaviour is unaffected by these policy changes, which could be justified on the grounds that sole mothers tend to have very low asset holdings (Del Boca, Flinn, and Wiswall, 2014). Moreover, the hypothetical policy changes that I consider are intended to be permanent: and as such would induce less substitution in consumption across time periods than would a temporary policy change.

Figure 4: Densities for child resource allocation measures



- Simulated using the CPS.

households with the same earnings: creating an obvious inequity in the provision of benefits, because the greater time available to the high-wage households, which can be devoted to other productive activities (such as housework or time with children). In contrast, the wage and childcare subsidies proposed here would tend to be proportionately – and appropriately – more generous to low-wage households. Some further implications of my findings for welfare policies, and of the mechanisms underlying these findings, is given in Section 6.4.

6.1 Child welfare in the welfare-to-work era

In this section the structural model will be used to construct a money-metric measure of the household resources allocated to children, providing the basis for a new measure of child poverty (Section 6.1). This model-derived measure is intended to more accurately reflect material deprivation as experienced by children than the US Census Bureau’s official poverty measure (OPM), which is based purely on household before-tax income.⁶ In contrast, my measure accounts for: (i) taxes and transfer payments; (ii) home production; and most significantly, (iii) the household resources actually received by children. (This approach has some similarities to Aguiar and Hurst, 2007.) Contrary to much of the prior literature which has found no decline in the welfare of sole mothers and their children in the post-welfare-reform era, I find an increase in child poverty of 5 percentage points (Blank, 2002; Meyer and Sullivan, 2004).

⁶The official poverty measure was first constructed in 1963, based on the cost of a basket of goods, and has subsequently been updated every year using the Consumer Price Index. Institute for Poverty Research (2017)

Measures of children’s resources Evaluating the resources allocated to children requires an appropriate vector of relative prices: a natural choice here is the price vector that decentralises the household allocation, which is equal to each household member’s marginal rates of substitution. (This approach is an adaptation of that found in Chiappori and Meghir, 2014). These prices are denoted by $\tilde{p} = (\tilde{w}_m, \tilde{p}_m, \tilde{p}_k)$, where \tilde{w}_m denotes the price of the mother’s time – as distinct from the (pre-tax) market wage w_m – and \tilde{p}_m and \tilde{p}_k respectively denote the mother’s and children’s (Lindahl) prices for the public good. The total value of the resources allocated to children by the household is then defined as

$$\rho_k = c_k + \tilde{p}_k q + \tilde{w}_m t_m = C_k + \tilde{w}_m t_m \quad (6.1)$$

where C_k measures the value of the consumption resources (private and public) received by the children. For the calculations that follow, I always express ρ_k and C_k in per child terms, by dividing both by the number of children n_k in the household.

The estimated distributions of both ρ_k and C_k , for the years 1993–95 and 2005–08, appear in Figure 4. (Since all nominal magnitudes in the model are deflated to the year 2000, both measures are automatically adjusted for inflation.) Both distributions have shifted noticeably to the left, with the median of each declining by around \$15/wk over the sample period (from \$236/wk for C_k/n_k and \$246/wk for ρ_k/n_k).

Child poverty rates The preceding measures provides the basis for estimates of the child poverty rate, which – unlike existing estimates – take account of the totality of resources received by children. When calculating poverty rates, regardless of which measures is used, the choice of threshold below which children are considered to be ‘in poverty’ is unavoidably somewhat arbitrary. To ensure that my findings are reasonably robust to this choice, I consider two possible thresholds for each measure (ρ_k and C_k), corresponding to the 30th and 50th percentiles of that measure in 1993–95. The change in the poverty rate, using for example ρ_k and the 30th percentile threshold, is then computed as:

$$\Delta(\rho_k, 30) = \frac{\sum_{i \in \mathcal{I}_{08}} \pi_i n_{k,(i)} \mathbf{1}\{\rho_{k,(i)}/n_{k,(i)} \leq \tau_{93}(\rho_k, 30)\}}{\sum_{i \in \mathcal{I}_{08}} \pi_i n_{k,(i)}} - \frac{30}{100}.$$

where $\tau_{93}(\rho_k, 30)$ denotes the 30th percentile for ρ_k/n_k in 1993–95, and $i \in \mathcal{I}_{08}$ collects the households in the 2004–08 subsample; recall that $\{\pi_i\}$ denotes the inverse population weights from the CE.

Using the 30th percentile-based threshold, I find that the child poverty rate has increased by 3.7 and 4.5 percentage points, for estimates based on ρ_k and C_k respectively. If the 50th

percentile is used, these figures rise slightly to 4.5 and 5.2. I thus conclude that for the children of sole mothers (without a college degree), the child poverty rate rose by around four percentage points in the decade following the introduction of TANF. My results stand in contrast to calculations based on the OPM, which suggest that the child poverty rate has decreased by 6.5 percentage points to 43.9 per cent, among the children of sole mothers, over the same period (Child Trends Databank, 2017). They also differ from the findings of a number of authors who have attempted to address some of these shortcomings of the OPM, but have nevertheless found little evidence for an increase in child poverty rates since 1996 (Blank, 2002; Meyer and Sullivan, 2004).

The principal reason for these differences is that these authors still rely on household-level income or consumption to measure material well-being, whereas I take account of the how those resources are allocated within the household. Since 1996 the decline in the generosity of cash welfare has been effectively compensated for by an increase in time spent in paid work, and thus in before-tax incomes (recall Figure 2 above, and the accompanying discussion). This has come largely at the expense of time spent in housework, which has fallen at the median by just over 1 hr/wk between 1996 and 2008. Together with a decline in spending public consumption – the median of which fell by \$17/wk to \$259/wk in 2005–08 – this has resulted in a significant reduction in the output of the home-produced public good. It is this reduction that is largely responsible for the increase in child poverty, as I have measured it, something which measures based on household-level income or consumption will fail to capture.

The preceding calculations – similarly to those of Blank (2002) and Meyer and Sullivan (2004) – fail to account for any demographic shifts that might have occurred over the period, such as changes to the average educational attainment or fertility of sole mothers. To remedy this, I take the entire sample of sole-mother households from the CPS for 1993–2008, treating them as though they were drawn from 1993 (so applying the 1993 tax-and-welfare system, and imputing their wages using the 1993–95 estimates of the wage equation, etc.), and computing the poverty thresholds on this basis. I then counterfactually apply the 2008 tax and welfare system to these households; re-solving the households’ problems yields an estimate of $\Delta(\rho_k, 30)$ in which all demographic factors are held constant. The estimated change in child poverty rates is 1.4 and 1.6 for ρ_k and C_k respectively using the 30th percentile threshold, and 2.8 and 3.4 for the 50th percentile threshold. Demographic shifts would thus seem to account for only about 2 percentage points of the change in the poverty rate over this period, and so are insufficient to overturn the conclusion that TANF has had a deleterious effect on child poverty.

Table 6: Effect of AFDC/TANF on household choices and welfare
Mean responses for the subgroup affected by the policy change

		Δh_m	Δt_m	Δq_m	Δq	Net cost*	CV_k	Pass thru.*	Grp. size*
		1	2	3	4	5	6	7	8
<i>Policy counterfactuals</i>									
1	AFDC replaces TANF [†]	-13.3	0.7	1.4	1.4	106.9	13.6	12.7	23.3
	EITC expansion [‡]	0.2	0.0	1.0	1.0	19.6	6.7	34.2	31.7
<i>Introduction of policy, for years indicated[§]</i>									
3	TANF (1996–2008)	-11.7	0.6	1.6	1.9	102.3	17.3	16.9	8.3
	AFDC (1993–95)	-19.7	1.2	2.8	4.3	183.3	35.4	19.3	25.9
	EITC (1996–2008)	1.1	0.3	0.3	2.6	50.3	20.4	40.6	65.8

• Reports the results of counterfactual exercises conducted using the model, as described in Section 6.2. Units: *net cost* and CV_k are in \$/wk (2000 dollars); h_m , q_m and t_m in hr/wk; *pass thru.* and *grp. size* in percentage points.

* *Net cost* is the cost of the policy intervention to the government, calculated as the increase in net transfers (benefits paid less taxes received) paid to households as a result of the policy. *Pass thru.* is CV_k divided by the net cost (in percentage points; see (6.3)). *Grp. size* is the percentage of households (sole mothers without a college degree) affected by the policy change.

[†] Baseline is the tax and welfare system in place over 1996–2008; in the counterfactual, AFDC as it existed in 1995 is introduced in each of these years (suitably adjusting programme benefits for inflation).

[‡] Baseline is the tax and welfare system in place over 1996–2008; in the counterfactual, I introduce additional benefit payments for one-child families, as proposed by Hoynes (2014).

[§] In each case, the baseline is the tax and welfare system in place over the specified period, sans the indicated policy; this policy is (re-)introduced in the counterfactual.

6.2 PROWRA, child welfare and maternal labour supply

While it is clear that child welfare has declined since 1996 – even when changes in the composition of the sample of sole mothers are controlled for – the extent to which this is a direct result of the PROWRA reforms requires further analysis. I therefore use the model to perform counterfactual experiments that allow me to identify the consequences of the introduction of TANF, and also of the expansion of the EITC that occurred during the same epoch. As well as determining how the allocation of household resources has altered as a result of these policies, these exercises allow me to make a more reliable estimate of their welfare effects. (Whereas the comparisons of ρ_k and C_k made above involve the use of a different set of relative prices across the two comparison periods, below I am able to

construct a compensating variation-type measure that abstracts from such price changes.)

Results are displayed in Table 6. (This summarises the main results; a more complete record is given by Table 13 in Appendix E.) The first counterfactual experiment (row 1) reverses the PRWORA reforms: the baseline is the actual tax and welfare regime over 1996–2008, from which TANF is counterfactually replaced by AFDC (as it was in 1995, adjusting for inflation). Columns 1–4 report the change in mother’s time use and in the production of the public good, by subtracting the outcome in the baseline (TANF) from the outcome under the counterfactual policy (AFDC), and averaging this over the subgroup of households affected by the policy change (those which either receive TANF in the baseline, or AFDC in the counterfactual). Columns 5–7 summarise the overall cost (to the government) of the policy change, and quantify its effect on children’s welfare; column 8 reports the proportion of households affected by the intervention. I shall now give a brief discussion of the methodology used to compute the figures in columns 4–7, before discussing the results.

Measuring welfare effects and programme costs To quantify the welfare consequences of PROWRA (and other counterfactual interventions) in money-metric terms, I use the following compensating variation-type measure. Let \tilde{p}^0 denote the price vector that decentralises the household’s allocation in the baseline; ρ_k^0 the associated value of the children’s allocation, computed as per (6.1) above; and \bar{K}^1 the children’s utility under the counterfactual policy. I then compute

$$CV_k = \min_{\Delta\rho \in \mathbb{R}} \left\{ \max_{BC(\tilde{p}^0, \rho_k^0 + \Delta\rho)} K(c_k, t_m, q) \geq \bar{K}^1 \right\}, \quad (6.2)$$

where $BC(\tilde{p}, \rho) = \{(c_k, t_m, q) \in \mathbb{R}_+ \mid c_k + \tilde{w}_m t_m + \tilde{p}_k q \leq \rho\}$.⁷ CV_k can thus be interpreted as the amount the children would need to be paid in the baseline, to be as well-off as under the counterfactual policy, if they were able to purchase consumption, time with their mother, and the public good at the prices \tilde{p}^0 with an ‘income’ of $\rho_k^0 + \Delta\rho$.⁸

⁷I also construct an analogous measure for the mother: but since my focus here is on children this is only reported in the tables in Appendix E, where it is denoted as CV_m . To compute this, I first determine the utility \bar{U}^1 that the mother enjoys under the counterfactual policy (here, AFDC), and then calculate the additional income (provided as a lump-sum transfer) that the household would need in the baseline (i.e. under TANF) to achieve a utility of exactly \bar{U}^1 . (Note that CV_m will be negative for households that are worse off under AFDC.)

⁸Because the household’s objective is separable in children’s utility, a second interpretation of CV_k is also available (see Browning, Chiappori, and Weiss, 2011, Ch. 4). Suppose I were to provide the household with an additional lump-sum transfer Δy , which it could allocate however it liked, but which was chosen so as to be exactly sufficient to raise the children’s utility to \bar{K}^1 . Then using \tilde{p}^0 to evaluate the resultant change in the child’s resources would yield $\Delta c_k + \tilde{w}_m^0 \Delta t_m + \tilde{p}_k^0 \Delta q = CV_k$.

The policy change considered in my first counterfactual exercise is far from being revenue neutral, as AFDC is considerably more generous than TANF. The overall cost of the intervention can be computed by subtracting net government tax revenues (counting transfers as a diminution of those revenues) in the baseline from those in the counterfactual; I term this the *net cost*, and report its average in column 5. This measure takes account of households' behavioural responses to the policy change, and so in this case may differ markedly from a simple comparison of the benefits paid out under AFDC with those paid out under TANF. (Insofar as the introduction of AFDC tends to reduce labour supply, the 'net cost' as computed here will tend to be higher than might be suggested by such a naive calculation.) To facilitate the comparison of AFDC and TANF with other policies, in terms of their relative efficacy at improving child welfare, I define the *pass-through rate* of a policy as

$$\text{pass-through rate} = \frac{CV_k}{\text{net cost}}, \quad (6.3)$$

(expressed as a percentage). This can be interpreted as the proportion of the welfare spending that gets through to children.

Replacing TANF with AFDC; and some other counterfactuals The results suggest that the 23 per cent of sole mothers (without a college degree) who are affected by the policy change would work (on average) 13 fewer hours per week. Most of this newly freed time would be enjoyed as additional leisure, but they also spend just over half an hour more with children, and an hour and a half extra per week in housework (cols 1–3). Qualitatively, this mirrors the increase in market work and reduced time spent in housework over 1996–2008, but only explains a part of those changes.

Reintroducing AFDC is estimated to improve the welfare of the affected children by almost \$8/wk, in terms of CV_k (col 6). Recall from the previous section that, in terms of ρ_k/n_k , the median child was \$17/wk worse off in 2004–08 relative to 1993–95. Given that the second figure controls for neither demographic shifts in the sample of sole mothers, nor changes in the relative (decentralising) prices of household resources, whereas the former does, the two figures seem broadly comparable. (Note that there are on average only 1.3 children per family in my sample, so that the fact that CV_k is not expressed in per-child terms matters relatively little here.)

In view of this, it might be argued that reintroducing AFDC would go a considerable way to remedying the decline in children's welfare that has occurred since 1996. However, such a conclusion must be weighed against the cost of such a policy change, which is estimated to be on the order of almost \$100/wk (col 5). This implies a pass-through rate of 8 per cent (col 7).

To give some context for this figure, I conduct two further counterfactual exercises. In the first, the the baseline is the existing tax and welfare policies without TANF (for 1996–2008), which is then (re-)introduced in the counterfactual. This is reported in row 3 of the table; row 4 reports the outcome of an analogous exercise, performed with respect to AFDC (and for the years 1993–1995). The pass-through rates associated with both programmes are between 15 and 20 per cent.

In fact, relative to the EITC, neither AFDC nor TANF is particularly effective at targeting resources to children. Repeating the same exercise for the EITC yields an estimated pass-through rate of almost 41 per cent, more than twice that of AFDC or TANF (row 5, col 7). The results also indicate the EITC is in fact far more effective at increasing maternal labour supply than TANF, despite the latter being regarded as the archetypal ‘welfare to work’ programme. Because only a third of TANF enrollees meet the work requirements, for most it has more the character of a lump-sum transfer – albeit one that may involve some disutility to the recipient – and so has the expected income effect of *reducing* labour supply. The EITC, on the other hand, by reducing the returns to work for households with a broad range of incomes, has a positive (albeit small) effect on maternal labour supply.⁹

6.3 Alternatives to welfare-to-work programmes

In light of the preceding results, the EITC appears to be considerably more effective than either AFDC to TANF in terms of promoting maternal labour supply and targeting household resources to children (per dollar spent). This finding motivates the search for alternatives to TANF, which might better fulfil these two objectives. Below I examine several policies that encourage labour supply by increasing the net returns to working, rather than by sanctioning enrollees who fail to meet mandated work requirements: these take the form of childcare subsidies and wage subsidies.

Since eligibility for the EITC depends only on total earnings, the EITC does not discriminate between households in which the mother: (a) earns a low wage and works full time; and (b) earns twice that wage, and only works part time. From the perspective of the model developed in this paper, there are other valuable activities to which a mother’s time might be devoted – in particular, housework and time with her children. Once these

⁹Hoynes (2014) has recently suggested that the EITC benefits be increased for one-child families: this would be particularly beneficial for sole mothers, 66 per cent of whom only have one child. I evaluate this proposal by counterfactually expanding the EITC from the baseline of the actual tax and welfare system in place in 1996–2008, in the manner proposed by Hoynes (2014; see that paper for the exact details of her proposal). Her proposed changes would indeed be effective at targeting household resources to the children of sole mothers: the estimated pass-through rate is 32 per cent, and the policy reaches roughly 50 per cent of one-child households (see row 2, col 7–8).

are accounted for, household (b) would seem to be much better off than household (a): thus it must be counted a shortcoming of the EITC that it would pay exactly the same subsidy to both. In effect, the EITC ignores the value of time spent in activities outside of market work. One way of resolving this is to provide a subsidy that is proportionately more generous to lower-*wage* – as distinct from lower-earnings – households. Hence I am led to consider policies that increase the net return to an hour's work: for example, the \$1/hr wage subsidy proposed below would pay twice as much to household (a) as to household (b).

Childcare subsidies For sole mothers with no access to sources of free informal childcare (such as might be provided by relatives), the pre-tax net return to an hour's work is the wage, less the cost of an additional hour's worth of childcare. On average, sole mothers who work more than 20 hr/wk spend \$32/wk on childcare for each child ages under 5, and \$15 per week for each child aged 6 to 13 (in the SIPP). Childcare subsidies thus provide an effective means of raising the net returns to work for this group. At present, means-tested subsidies are available in the form of a tax credit, the CDCTC, which delivers a maximum 30 per cent rebate to eligible households.

I conduct three counterfactual experiments related to childcare subsidies. In the first, the baseline is the actual tax and welfare system in 1996–2008, without the CDCTC, which is then (re-)introduced in the counterfactual. In the second and third, the baseline includes the CDCTC, which is then alternately replaced by universal free childcare, or expanded along the lines proposed by Ziliak (2014), who has suggested that CDCTC should provide more generous rebates to low-income households (see that paper for the specific details of his proposal). A summary of the main results appear in the first two rows of Table 7 (a complete listing appears in Table 14 in Appendix E).

Most striking are the high pass-through rates for children: almost 100 per cent for CDCTC, 50 per cent for free childcare, and 40 per cent for the Ziliak (2014) expansion (see rows 1-2, col 7). The last two are twice the pass-through rate for TANF, and comparable to that of the EITC. Children benefit because the household's income increases – not only directly, because childcare is cheaper, but also indirectly, via an increase in maternal labour supply (rows 1–2, col 1 in Table 7). This additional income is directed both to their private consumption, and to public consumption. Although children do receive less maternal time as a result of CDCTC, the magnitude of this decline is only 25 per cent of the increase in maternal labour supply. All three policies are also considerably more effective than either TANF or the EITC at promoting maternal labour supply.

Previous work on the effect of childcare on child development has focussed on the quality of care provided at home relative to childcare centres (Dustmann and Schönberg, 2012;

Table 7: Alternatives to welfare-to-work programmes
Mean responses for the subgroup affected by the policy change

		Δh_m	Δq_m	Δt_m	Δq	Net cost*	CV_k	Pass thru.*	Grp. size*
		1	2	3	4	5	6	7	8
<i>Childcare policies</i> [†]									
1	CDCTC	3.1	-0.5	-0.8	1.2	14.8	14.8	99.7	11.3
	Free childcare	15.2	-1.0	0.2	5.7	101.4	50.2	49.5	38.0
	Expansion	6.2	0.2	0.0	5.7	51.7	19.8	38.3	31.0
<i>TANF eliminated; replaced by indicated policy</i> [†]									
4	Free childcare	10.3	-0.9	-0.9	3.7	86.5	35.3	40.8	46.2
	Flat wage subsidy	1.2	0.0	0.0	0.9	10.6	8.1	76.5	93.5
	Targeted wage subsidy	5.0	0.5	0.3	4.8	67.9	33.6	49.5	64.1

- Reports results of counterfactual exercises described in Section 6.3. Only households in years 1996–2008 of the sample are exposed to the policies. Units: *net cost* and CV_k are in \$/wk (2000 dollars); h_m , q_m and t_m in hr/wk; *pass thru.* and *grp. size* in percentage points.

* For definitions of *net cost*, *pass thru.*, and *grp. size*, see the corresponding note to Table 6.

[†] In *CDCTC*, the baseline is the tax and welfare system in place over 1996–2008, sans CDCTC, which is (re-)introduced in the counterfactual. *Expansion* additionally expands the CDCTC in the manner proposed by Ziliak (2014). In *free childcare*, the baseline additionally includes CDCTC, which in the counterfactual is replaced by universal free childcare.

[‡] In each exercise, the baseline is the tax and welfare system in place over 1996–2008; in the counterfactual, TANF is eliminated and replaced by the policy indicated.

Fiorini and Keane, 2014). Quality of care is certainly one channel through which childcare subsidies may affect child outcomes, but my results highlight another channel, which works through the effect that these subsidies have on the allocation of household resources. These two channels may help account for the apparently contradictory findings of two recent studies examining the introduction of universal childcare: Baker, Gruber, and Milligan (2008) find substantially negative effects on child outcomes among middle- and high-income two-parent families (in Canada), while Havnes and Mogstad (2014) find positive effects for lower-income households (in Norway). For less educated (and likely low-income) sole mothers, the additional income that the household receives may be highly beneficial for children – as my results forcefully contend – whereas for highly educated (and likely high-income) households, childcare quality might be more important.

Wage subsidies Wage subsidies directly raise the net returns to work. Rows 4–6 in Table 7 summarise the results of three counterfactual exercises that replace TANF with a wage subsidy. In each case, the baseline is the actual tax and welfare regime over 1996–2008, and the counterfactual consists of eliminating TANF and introducing one of the following policies: universal free childcare (row 4); a flat-rate wage subsidy of \$1/hr (row 5); or a wage subsidy targeted at low-wage earners, which brings all wages up to a minimum of \$11.60/hr (row 6; in 2000 dollars). (\$11.60 in 2000 corresponds to \$15 in 2018 dollars; this is motivated by the \$15/hr minimum wage that is planned for New York in 2018.)

Each of these policies raise household income directly by increasing the returns to work, and indirectly by also promoting maternal labour supply. Free childcare increases labour supply by 10 hr/wk on average, the flat-rate wage subsidy by a little over 1 hr/wk and the targeted wage subsidy by 5 hr/wk (rows 4–6, col 1; note that these averages are computed for the sub-populations affected by the respective policy changes, and express changes relative to a baseline in which TANF is in place). Since these labour supply responses raise taxable earnings, the net cost of all these programmes (measured by the net change in government taxes and transfers) is much lower than the gross benefits paid to households (compare cols 8 and 9, rows 4–15, in Table 14 in Appendix E).

Wage subsidies also achieve high pass-through rates. For example, the \$1/hr wage subsidy raises child welfare (in terms of CV_k) by \$8/wk with a pass-through rate of 77 per cent (row 5, cols 7–8 in Table 7). The targeted wage subsidy costs \$34/wk per recipient, and has a pass-through rate of 50 per cent (row 6, cols 7–8).

6.4 Implications for welfare policies

Overall, my results suggest that policies that increase the net returns to working, are more effective than TANF both at encouraging maternal labour supply, and in targeting household resources to children. TANF provides sole mothers with the option of either: (a) working 30 hr/wk and claiming benefits, if their incomes are still sufficiently low; or (b) claiming benefits without meeting that work requirement, at the cost of having to engage in some other time-consuming activity, such as a job training programme. Most TANF enrollees (66 per cent in 2008) do not meet the work requirements, and so for them TANF has more the character of a lump-sum transfer, and thus tends to *reduce* their labour supply.

As to why TANF fails to meet one its principal objectives, my results suggest that the 30 hr/wk work requirement is set too high to induce enough sole mothers to increase their labour supply to this level: although 9 per cent of my sample (both in the CE and CPS) have wages low enough that they would be eligible for TANF if they worked 30 hr/wk, only 22 per cent of these households do so. (They would be entitled an average of \$58/wk from TANF, relative to earnings of \$180/wk at the 25th percentile.) For the remainder of these women, their wages are so low that for the most part their time is better spent in alternative activities, even with the additional work subsidy offered by TANF (at 30 hr/wk); average hours worked by the entire group being only 23hr/wk. For many women, having to work 30 hr/wk week would be too much of a distortion of the optimal allocation of their time.

In view of this, it is perhaps not surprising that policies which increase the net returns to work, by a small amount for everyone, are more effective at promoting maternal labour supply than is TANF. Such policies are also much more beneficial for children, in the sense of achieving higher pass-through rates. A key reason for this lies in their effect on home production. Recall from Section 6.1 that we observed an increase in child poverty rates after 1996, and that this was associated with a simultaneous decline in the production of the public good. The policies considered in Section 6.3 would be able to help reverse this decline, by stimulating home production at a much lower cost than can either TANF (or AFDC, for that matter). For example, when TANF is replaced with free childcare or the targeted wage subsidy, it increases the output of the public good at an effective cost of under \$15/wk, whereas when TANF is replaced with AFDC, although public good production also rises, the additional units of the public good cost almost \$200/wk (these figures are arrived at by dividing the increment in the output of the public good by the net cost of the policy change).

7 Conclusion

The introduction of TANF was the most significant change to the US welfare system in the postwar era. TANF was intended to reduce welfare dependency and promote labour supply, through work requirements and lifetime limits. But while caseloads have indeed fallen, and hours worked have increased, my results indicate that these apparent successes have come at the expense of children's welfare.

Using a structural model of household decision-making and resource allocation, I find that the welfare of the children of sole mothers has been materially worsened by the introduction of TANF. Because of the prohibitively high work requirements (30 hr/wk), low-wage women typically do not find it worthwhile to work sufficiently many hours to be eligible for TANF, even though their weekly earnings put them well below the poverty line. For this group, the main consequence of PRWORA was simply a withdrawal of the benefits they had previously been entitled to (under AFDC). This loss of income has been partly offset by an increase in their labour supply, but the net result of this has been detrimental to children, because it has left mothers with less time to devote to home production. Using the model to value the household resources allocated to children, and thus account for the effects of the tax and welfare system, the intra-household allocation of resources, and home production, I find that child poverty rates have risen by four percentage points since the introduction of TANF, among the children of sole mothers (without a college degree).

To the extent that work requirements are not rigidly enforced – currently, two-thirds of TANF recipients fail to meet these requirements – TANF continues to provide something like a traditional safety net for low-income households. However, states are now coming under increasing pressure to ensure that a higher proportion of enrollees meet the work requirements: if they were to do so, then even this already very diminished safety net would disappear, with further negative consequences for low-income households.

I consider a number of alternatives to welfare-to-work programmes, which encourage labour supply by instead increasing the return from working, rather than by sanctioning non-compliers: these are free childcare, and two types of wage subsidy. Each of these alternatives is more effective at promoting labour supply than is TANF, and raises household income by increasing not only the returns from working, but also hours worked. These policies are also far more successful at targeting household resources to children, with all having estimated pass-through rates in excess of 40 per cent (as compared with 17 per cent for TANF).

New York City, California, Chicago, Seattle, and Washington D.C. all plan to raise minimum wages to \$15/hr during the next decade. (The targeted wage subsidy of \$11.60 in 2000 dollars, considered above, is equivalent to a \$15 wage subsidy in 2016.) If these higher

minimum wages do not significantly attenuate the demand for low-skilled labour, nor raise prices so far as to absorb the real gains from the wage increase (for this group), my results suggest that children in sole-parent households might be the unexpected beneficiaries of such a policy, as much of the additional income will translate into tangible resources for these children.

8 References

- AGUIAR, M., AND E. HURST (2007): “Life-cycle prices and production,” *American Economic Review*, 97(5), 1533–1559.
- APPS, P., AND R. REES (1996): “Labour Supply, Household Production and Intra-family Welfare Distribution,” *Journal of Public Economics*.
- BAKER, M., J. GRUBER, AND K. MILLIGAN (2008): “Universal Child Care, Maternal Labor Supply, and Family Well-Being,” *Journal of Political Economy*, 116(4), 709–745.
- BLANK, R. M. (2002): “Evaluating welfare reform in the United States,” *Journal of Economic Literature*, 40(4), 1105–1166.
- BLUNDELL, R., M. COSTA DIAS, C. MEGHIR, AND J. SHAW (2016): “Female labor supply, human capital, and welfare reform,” *Econometrica*, 84(5), 1705–1753.
- BLUNDELL, R., AND A. SHEPHARD (2011): “Employment, Hours of Work and the Optimal Taxation of Low-Income Families,” *Review of Economic Studies*.
- BLUNDELL, R., AND I. WALKER (1986): “A Life-cycle Consistent Empirical Model of Family Labour Supply Using Cross-Section Data,” *Review of Economic Studies*.
- BROOKS-GUNN, J., AND L. MARKMAN (2005): “The contribution of parenting to ethnic and racial gaps in school readiness,” *The Future of Children*, 15(1), 139–168.
- BROWNING, M., P.-A. CHIAPPORI, AND Y. WEISS (2011): *Family Economics*. Online Book.
- CHANG, Y., AND F. SCHORFHEIDE (2003): “Labor-supply shifts and economic fluctuations,” *Journal of Monetary Economics*, 50(8), 1751–1768.
- CHERCHYE, L., B. DE ROCK, AND F. VERMEULEN (2012): “Married with children: A collective labor supply model with detailed time use and intrahousehold expenditure information,” *American Economic Review*, 102(7), 3377–3405.
- CHIAPPORI, P.-A., AND C. MEGHIR (2014): “Intra-household welfare,” Discussion paper, National Bureau of Economic Research.
- CHILD TRENDS DATABANK (2017): “Children in poverty,” <https://www.childtrends.org/indicators=children-in-poverty>, Accessed: 2017-08-10.

- CROUSE, G., AND S. MACARTNEY (2016): “Welfare Indicators and Risk Factors,” Discussion paper, U.S. Department of Health and Human Services.
- DAHL, G., AND L. LOCHNER (2011): “The impact of family income on child achievement: evidence from changes in the Earned Income Tax Credit,” *American Economic Review*.
- DEL BOCA, D., C. FLINN, AND M. WISWALL (2014): “Household Choices and Child Development,” *Review of Economic Studies*.
- DUNBAR, G., A. LEWBEL, AND K. PENDAKUR (2013): “Children’s Resources in Collective Households: Identification, Estimation and an Application to Child Poverty in Malawi,” *American Economic Review*.
- DUSTMANN, C., AND U. SCHÖNBERG (2012): “Expansions in maternity leave coverage and children’s long-term outcomes,” *American Economic Journal: Applied Economics*, 4(3), 190–224.
- FIORINI, M., AND M. KEANE (2014): “How the Allocation of Children’s Time Affects Cognitive and Non-cognitive Development,” *Journal of Labor Economics*.
- GOURIEROUX, C., A. MONFORT, AND E. RENAULT (1993): “Indirect inference,” *Journal of Applied Econometrics*, 8(S1), S85–S118.
- GROGGER, J., AND C. MICHALOPOULOS (2003): “Welfare Dynamics under Time Limits,” *Journal of Political Economy*, 111(3).
- HAVNES, T., AND M. MOGSTAD (2014): “Is universal child care leveling the playing field?,” *Journal of Public Economics*.
- HOYNES, H. (1996): “Welfare Transfers in Two Parent Families: Labor Supply and Welfare Participation under AFDC-UP,” *Econometrica*.
- (2014): “Proposal 11: Building on the Success of the Earned Income Tax Credit,” in *Improving Safety Net and Work Support*. The Hamilton Project.
- INSTITUTE FOR POVERTY RESEARCH (2017): “How is poverty measured in the United States,” <http://www.irp.wisc.edu/faqs/faq2.htm>, Accessed: 2017-08-10.
- ISAACS, J., K. TORAN, H. HAHN, K. FORTUNY, AND C. STEUERLE (2012): “Kids’ share 2012: report on federal expenditures on children through 2011,” Discussion paper, The Urban Institute.
- KIMMEL, J., AND R. CONNELLY (2007): “Mothers’ Time Choices: Caregiving, Leisure, Home Production, and Paid Work,” *The Journal of Human Resources*.
- LISE, J., AND S. SEITZ (2011): “Consumption inequality and intra-household allocations,” *The Review of Economic Studies*, 78(1), 328–355.
- MCGRATTAN, E. R., R. ROGERSON, AND R. WRIGHT (1997): “An equilibrium model of the business cycle with household production and fiscal policy,” *International Economic Review*, pp. 267–290.
- MEYER, B. D., AND J. X. SULLIVAN (2004): “The effects of welfare and tax reform: the material well-being of single mothers in the 1980s and 1990s,” *Journal of Public Economics*, 88(7), 1387–1420.

- (2008): “Changes in the consumption, income, and well-being of single mother headed families,” *American Economic Review*, 98(5), 2221–2241.
- PHILLIPS, M. (2011): *Whither opportunity? Rising Inequality, Schools, and Children’s Life Chances*. Parenting, time use, and disparities in academic outcomes, pp. 207–228. Russell Sage Foundation.
- SCARBOROUGH, H. S., AND W. DOBRICH (1994): “On the efficacy of reading to preschoolers,” *Developmental Review*, 14(3), 245–302.
- SCHOTT, L., L. PAVETTI, AND I. FINCH (2012): “How states have spent federal and state funds under the TANF block grant,” Discussion paper, Center of Budget and Policy Priorities.
- SNOW, C. E., AND D. E. BEALS (2006): “Mealtime talk that supports literacy development,” *New Directions for Child and Adolescent Development*, 2006(111), 51–66.
- TAMIS-LEMONDA, C. S., M. H. BORNSTEIN, AND L. BAUMWELL (2001): “Maternal responsiveness and children’s achievement of language milestones,” *Child development*, 72(3), 748–767.
- ZILIAK, J. (2014): “Proposal 10: Supporting Low-Income Workers through Refundable Child-Care Credits,” in *Improving Safety Net and Work Support*. The Hamilton Project.

A Data sources and construction

A.1 ATUS

The American Time Use Survey (ATUS) randomly selects one individual from each family, and asks participants to complete a time diary for one full day, recording their main activity, secondary activity and who the activity was with. All observations with incomplete time diaries are dropped, as per Guryan, Hurst, and Kearney (2008).

Sample The sample (2003–08) of single mothers is restricted to households in which the mother

- (i) is not self-employed or in the armed forces;
- (ii) is aged between 18 and 58;
- (iii) has children below the age of 18 who she lives with;
- (iv) does not share a household with any other family;
- (v) has all demographic variables required by my estimation procedures; and
- (vi) has hourly wages between \$2.50–\$250 (in 2000 dollars), if reported.

Mothers' time with children Defined as an aggregate of: caring for and helping children, including general care of children; activities related to children's health, including providing and obtaining medical care for children; activities related to children's education; time spent on animals and pets with children; time spent using childcare services; time spent eating and drinking with children; time spent socialising and communicating and attending or hosting social events with children; time spent listening to music, playing games and doing hobbies with children; time spent attending performing arts, museums or any arts and entertainment with children; sports, exercise and recreation with children; religious and spiritual activities with children; volunteer activities with children; and travelling associated with any of the above activities.

Mother's time in housework Defined as an aggregate of: food/meal preparation; cleaning; laundry and clothes repair; home repairs and vehicle maintenance; domestic work; purchasing goods and services; gardening; and pet care.

Variance estimates for time-use measures For time with children and housework, daily time diaries provide a consistent estimate of the average weekly time spent in these activities, but yield an upwardly biased estimate of the weekly variance (Bruins and Duffy, 2015). Using a unique Dutch dataset (Tijdsbestedingsonderzoek, 2005), which provides a complete set of daily time diaries for each respondent (i.e. one for every day of the week), I estimate these variance measures to be upwardly biased by a factor of 1.5. The sample variance moments for time with children and housework are therefore computed by multiplying the (daily) sample variance computed from the ATUS time diaries by a factor of 0.67 (equivalently, the sample standard deviation is scaled down by a factor of 0.82).

A.2 AHTUS

Data on time spent in housework in 1995 is based on time diaries provided by the American Heritage Time Use Study (AHTUS). Bad or incomplete diaries, as indicated in the AHTUS documentation, were dropped; the sample is otherwise constructed exactly as for the ATUS. Time spent in housework is aggregated from the same set of activities as for the ATUS.

A.3 CE

In constructing the Consumer Expenditure Survey (CE) sample I follow Blundell, Pistaferri, and Preston (2008). The sample (1993–2008) is constructed identically to that of the ATUS except that, owing to the problems discussed in Blundell, Pistaferri, and Preston (2008), the age range in (ii) is modified to 25–58.

Public consumption expenditure This is defined to include food expenditures, property taxes, mortgage interest, rent, household operations, utilities, fuel, and public services.

Non-labour income less savings Recall from the household budget constraint (3.7) that ‘non-labour income less savings’, $y - s$, is the difference between household non-durables consumption and after-tax earnings. Non-durables expenditure is defined as total consumption expenditure less expenditures on: house furnishings and equipment; outlays for other lodging; mortgage principal outlays; maintenance, repairs and insurance; educational expenditures; and personal insurance and pensions.

A.4 CPS

The CPS data is taken from the ‘CPS Merged Outgoing Rotation Groups’ provided by the NBER website. The sample is restricted analogously to the ATUS.

A.5 SIPP

The Survey of Income and Programme Participation (SIPP) sample (1996, 2001, 2004, and 2008) is constructed identically to the CE, except that I retain only those households in which the mother works (otherwise there is no need for childcare) and (vi) above is not imposed.

For each household, I calculate the total amount spent on childcare, including expenditures on family daycare, sports and club events. The price of childcare per child, per hour, is computed by dividing this expenditure by the total number of hours spent in childcare. These calculations are performed separately for children 5 and under, and those aged 6–13. The amount of free childcare available to the household, per child, is calculated as $[h_m - h_{c,u5}]_+$ for children 5 and under, where h_m denotes mother’s hours worked, and $h_{c,u5}$ the (average) time spent in childcare by children 5 and under ($[a]_+ = \max\{a, 0\}$). For children aged 6–13, the preceding is modified to $[h_m - h_{c,u5} - 30]_+$ since these children attend school for 30 hours each week.

A.6 Regional groupings

Throughout the paper, I use the following nine regional groupings, taken from the US Census Bureau: New England (CT, ME, MA, NH, RI, VT); Middle Atlantic (NJ, NY, PA); East North Central (IL, IN, MI, OH, WI); West North Central (IA, KS, MN, MO, NE, ND, SD); South Atlantic (DE, DC, FL, GA, MD, NC, SC, VA, WV); East South Central (AL, KY, MS, TN); West South Central (AR, LA, OK, TX); Mountain (AZ, CO, ID, MT, NV, NM, UT, WY); and Pacific (AK, CA, HI, OR, WA).

B Estimation procedure: further details

B.1 Censored regression models for childcare variables

Price of childcare The price of childcare is affected by a sample selection problem, being observed only for those households that purchase childcare. Let

$$c = \max\{x'_c \beta_c + u_c, 0\}$$

denote the hours of childcare purchased for children aged 0–5. Assuming that (u_c, u_p) are jointly Gaussian, I may write

$$p - p_{\min} = \max\{x'_p\beta_p + u_p, 0\} = \max\{x'_p\beta_p + \rho_c u_c + \sigma_p \epsilon_p, 0\}$$

for the price of childcare for children aged 0–5, where $\epsilon_p \perp u_c$ and $x_c = (x'_p, z'_c)'$. x_p consists of a constant, the age of the mother, two race dummies (white and black), a dummy for tertiary education, and (nine) region dummies. z_c consists of a Hispanic dummy and the age of the youngest child.

I adopt the following control function approach to estimating β_p :

- (i) Tobit regression of c on x_c : yields residuals \hat{u}_c .
- (ii) Tobit regression of $p - p_{\min}$ on (x_p, \hat{u}_c) , using the subsample $c > 0$.

To produce separate estimates of β_p for children aged 6–13, I run the above procedure a second time, with (c, p) now recording the hours purchased and price of childcare for children aged 6–13.

Hours of free childcare Estimation of the parameters in (4.7) is complicated by sample selection and two-sided censoring. Suppose mother’s hours worked, h_m , is generated according to:

$$h_m = \max\{x'_h\beta_h + u_h, 0\}.$$

Assuming that (u_h, u_a) are jointly Gaussian, I may write

$$a = \max\{x'_a\beta_a + u_a, 0\} = \max\{x'_a\beta_a + \rho_h u_h + \sigma_a \epsilon_a, 0\},$$

for the hours of free childcare available for children aged 0–5, where $\epsilon_a \perp u_h$ and $x_h = (x'_a, z'_h)'$. x_a consists of the same variables as x_p above, and additionally records whether the family owns their residence. z_h consists of a Hispanic dummy and the age of the youngest child.

Now note that:

- (i) a is observed only if $h_m > 0$; and
- (ii) even if $h_m > 0$, the observed values of a are censored to lie in the interval $[0, h_m]$;

i.e. I observe $\tilde{a} := \min\{a, h_m\}$ for these households. I therefore adopt the following control function approach to estimating β_a :

- (i) Tobit regression of h_m on x_h : yields residuals \hat{u}_h .
- (ii) ‘Double tobit’ regression of \tilde{a} on (x, \hat{u}_h) , with h_m as the right censor point (and zero as the left censor point), using the subsample $h_m > 0$.

To estimate the parameters of (4.7) for children aged 6–13, exactly the same approach is used, except that now a refers to the hours of free childcare available for children aged 6–13, and is observed only if $h_m - 30 > 0$ (since these children spend 30 hr/wk in school).

B.2 Household members’ private consumption

Although the CE records the total private consumption of the household, it does not provide a breakdown of that consumption across family members. However, it does provide data on some private assignable goods, in particular clothing, which has been used in prior literature to estimate household members’ individual private consumption. Here I adapt the procedure developed by Dunbar, Lewbel, and Pendakur (2013) to estimate the breakdown of private consumption. This is in turn used to construct (unconditional) sample means for mother’s and children’s consumption, which are matched in estimation (as discussed in Section 4.4).

Suppose that there are N consumption goods: let \tilde{c}_{ij} denote expenditure on the j th consumption good for person $i \in \{m, k\}$, and $c_i = \sum_{j \in J} \tilde{c}_{ij}$ his (or her) total private consumption expenditure. Under certain conditions on individual preferences, given by Dunbar, Lewbel, and Pendakur (2013, Sec. II), the optimal choice of \tilde{c}_{ij} will be related to c_i via

$$\frac{\tilde{c}_{ij}}{c_i} = (\beta_{0i} + \beta_{1i} \log c_i). \quad (\text{B.1})$$

Dunbar, Lewbel, and Pendakur (2013) further show that if preferences are ‘similar across types’, in the sense that $\beta_{1m} = \beta_{1k}$, then

$$\frac{\tilde{c}_{ij}}{C} = \eta_i \frac{\tilde{c}_{ij}}{c_i} \quad (\text{B.2})$$

for $i \in \{m, k\}$, where $C = c_m + c_k$ total private consumption expenditure.

With the aid of (B.1) and (B.2), it is possible to recover c_i from observations on C and $(\tilde{c}_{m1}, \tilde{c}_{k1})$ alone, i.e. from data on total household private consumption expenditure and on a single assignable good – in my case, clothing. This involves first estimating the slope coefficients in the regression

$$\frac{\tilde{c}_{i1}}{C} = \alpha_i + (\beta \eta_i) \log C = \alpha_i + \gamma_i \log C,$$

which follows from substituting (B.1) into (B.2), and recognising that $c_i = \eta_i C$. I then compute $\hat{\eta}_i = \hat{\gamma}_i / (\hat{\gamma}_m + \hat{\gamma}_k)$, whence $\hat{\eta}_i C$ provides an estimate of c_i . Separate estimates are calculated for families with different numbers of children.

B.3 Inference: estimation of Σ

Let n index the sample size in the CE, which therefore also indexes the size of the model, and n_A the sample size of the ATUS. (Although some sample moments are constructed from the CPS, this is ignored here to simplify the exposition.) I suppose that $n_A = \lambda n$, where λ is held constant as $n \rightarrow \infty$; the number of simulations draws, S , is also held fixed.

Σ is the limiting variance of $n^{1/2}[\hat{\mu}_n(\hat{\theta}) - \bar{\mu}_n]$. To construct an estimate of Σ , recall from equation (4.2) and the surrounding text that the simulated moments (at θ) are given by $\hat{\mu}_n(\theta) = \varphi[\hat{m}_n(\theta)]$, where

$$\hat{m}_n(\theta) = \frac{1}{n} \sum_{i=1}^n \frac{1}{S} \sum_{s=1}^S \pi_i \mathbf{m}[g(\mathcal{Y}_{(i)}, x_{(i)}, \xi_{(is)}; \theta); x_{(i)}] =: \frac{1}{n} \sum_{i=1}^n \frac{1}{S} \sum_{s=1}^S \psi_{is}(\theta) =: \frac{1}{n} \sum_{i=1}^n \Psi_{iS}(\theta).$$

The sample counterparts \bar{m}_n of these moments are computed using the CE and the ATUS: to reflect this, I partition these as

$$\bar{m}_n = \begin{bmatrix} \bar{m}_{n1} \\ \bar{m}_{n2} \end{bmatrix} = \begin{bmatrix} \frac{1}{n} \sum_{i=1}^n \psi_{i1}^* \\ \frac{1}{n_A} \sum_{i=1}^{n_A} \psi_{i2}^* \end{bmatrix} \quad \begin{array}{l} \text{(CE)} \\ \text{(ATUS)} \end{array}$$

Under the assumption of a correctly specified model, both \bar{m}_n and $\hat{m}_n(\theta_0)$ will converge in probability to $m_0 = \mathbb{E}\bar{m}_n$. Since the CE and ATUS samples are independent,

$$n^{1/2} \begin{bmatrix} \hat{m}_n(\theta_0) - m_0 \\ \bar{m}_n - m_0 \end{bmatrix} = n^{1/2} \begin{bmatrix} \frac{1}{n} \sum_{i=1}^n [\Psi_{iS}(\theta_0) - m_{01}] \\ \frac{1}{n} \sum_{i=1}^n [\psi_{i1}^* - m_{01}] \\ \frac{1}{n_A} \sum_{i=1}^{n_A} [\psi_{i2}^* - m_{02}] \end{bmatrix} \xrightarrow{d} N \left[0, \begin{bmatrix} U & R_{12} & 0 \\ R_{21} & V_{11} & 0 \\ 0 & 0 & \lambda^{-1} V_{22} \end{bmatrix} \right]$$

where $R_{12} = R'_{21}$, and $m_0 = (m'_{01}, m'_{02})'$ is partitioned conformably with $(\psi_{i1}^*, \psi_{i2}^*)'$. Let Ω denote the limiting variance matrix on the r.h.s. of the preceding. To estimate the nonzero blocks of Ω , I use

$$\begin{bmatrix} \hat{U} & \hat{R}_{12} \\ \hat{R}_{21} & \hat{V}_{11} \end{bmatrix} = \frac{1}{n} \sum_{i=1}^n \begin{bmatrix} \Psi_{iS}(\hat{\theta}) - \hat{m}_n(\hat{\theta}) \\ \psi_{i1}^* - \bar{m}_{n1} \end{bmatrix} \begin{bmatrix} \Psi_{iS}(\hat{\theta}) - \hat{m}_n(\hat{\theta}) \\ \psi_{i1}^* - \bar{m}_{n1} \end{bmatrix}'$$

$$\hat{V}_{22} = \frac{1}{n_A} \sum_{i=1}^{n_A} [\psi_{i2}^* - \bar{m}_{n2}] [\psi_{i2}^* - \bar{m}_{n2}]'$$

which are assembled to construct $\hat{\Omega}$ (note that λ is computed as n_A/n). Finally, letting $J_\varphi(m) = \partial_m \varphi(m)$ denote the Jacobian of φ with respect to m , I can estimate Σ using

$$\hat{\Sigma} = \begin{bmatrix} J_\varphi(\hat{m}_n) & -J_\varphi(\overline{m}_n) \end{bmatrix} \hat{\Omega} \begin{bmatrix} J_\varphi(\hat{m}_n)' \\ -J_\varphi(\overline{m}_n)' \end{bmatrix}.$$

C Parameter estimates

C.1 Models used for the imputation of missing variables

Table 8: Log wages, and non-labour income less saving

	log w_m			$(y - s)/100$		
	1993–98	1999–03	2004–08	1993–98	1999–03	2004–08
<i>Education</i>						
High school diploma	0.28 (0.09)	0.35 (0.09)	0.35 (0.11)	82.24 (27.70)	18.37 (31.95)	-17.64 (35.30)
Tertiary	0.46 (0.10)	0.58 (0.10)	0.58 (0.12)	75.33 (28.81)	-7.86 (31.08)	-0.43 (32.58)
<i>Race</i>						
White	1.08 (0.15)	1.29 (0.15)	1.34 (0.18)	52.80 (60.67)	63.58 (68.50)	-21.29 (54.93)
Black	-0.04 (0.05)	-0.11 (0.05)	0.05 (0.05)	22.30 (61.92)	41.16 (70.15)	-83.01 (56.82)
<i>Mother's age</i>						
Age	0.15 (0.04)	0.04 (0.02)	0.13 (0.03)	-5.08 (13.00)	5.85 (11.16)	-12.06 (10.69)
Age ² /100	-0.18 (0.04)	-0.04 (0.03)	-0.15 (0.04)	2.83 (16.67)	-7.65 (14.43)	13.81 (13.94)
Case-Shiller house price index				-20.21 (6.26)	-1265.90 (413.62)	287.96 (137.32)

- Coefficient estimates for (4.5), obtained from the CE. Standard errors in parentheses. Units: w_m in \$/hr, $y - s$ in \$/wk (year 2000 dollars). All explanatory variables are interacted with dummy variables for the years 1993–98, 1999–2003, 2004–08. Coefficients on region (nine) and year dummies are also estimated (suppressed from the table).

Table 9: Price of childcare and availability of free childcare

	Price		Free care	
	p_{u5}	p_{o5}	a_{u5}	a_{o5}
Tertiary education	-0.80 (0.75)	2.29 (1.03)	-13.26 (1.60)	-11.66 (1.37)
<i>Mother's age</i>				
Age	0.02 (0.23)	-0.4 (0.53)	6.75 (0.45)	5.76 (0.44)
Age ² /100			-0.09 (0.01)	-0.07 (0.01)
<i>Race</i>				
White	0.25 (0.56)	-0.40 (1.09)	49.29 (2.51)	35.64 (1.45)
Black	1.16 (1.09)	1.10 (1.67)	44.13 (2.74)	32.14 (1.75)
<i>Regions*</i>				
New England	-0.20 (0.87)	0.16 (1.62)	-8.02 (1.69)	-5.70 (1.85)
Mid Atlantic	-0.64 (1.00)	-0.67 (1.62)	0.78 (1.46)	1.00 (1.66)
East North Central	-0.34 (0.78)	0.42 (1.75)	1.02 (2.16)	-6.65 (1.98)
West North Central	-3.01 (1.29)	-0.71 (2.03)	7.37 (1.47)	5.16 (1.76)
South Atlantic	-2.69 (0.94)	2.03 (1.92)	-7.17 (3.74)	-0.20 (2.17)
East South Central	-1.44 (1.09)	1.79 (1.83)	2.08 (2.03)	4.52 (2.06)
West South Central	-1.04 (0.87)	0.99 (1.80)	3.24 (4.40)	4.55 (1.67)
Mountain	0.11 (0.03)	0.07 (0.03)	-2.94 (1.54)	-7.14 (1.69)

• Coefficient estimates for (4.6) and (4.7), obtained from the SIPP. Standard errors in parentheses. Units: p_{u5} , p_{o5} in \$/hr; a_{u5} , a_{o5} in hr/wk.

* Regions taken from the US Census Bureau; omitted region is Pacific.

C.2 TANF enrolment and lifetime limits

Table 10: Probit estimates for TANF receipt (sole mothers)

<i>Demographic variables</i>			
Mother's age	0.14 (0.00)	Has a child aged 0–2	0.23 (0.02)
Mother's age ² /100	-0.19 (0.00)	Has a child aged 0–5	0.32 (0.02)
White	-0.18 (0.03)	Has a child aged 0–13	0.29 (0.02)
Black	0.28 (0.03)	Has more than one child	-0.05 (0.01)
High school diploma	-0.52 (0.02)		
<i>Exposure to lifetime limits*</i>			
z_1 ('never')	0.28 (0.04)	z_2 ('partial')	0.13 (0.02)
<i>Years</i>			
1996–99	0.15 (0.03)	2003–05	0.21 (0.02)
2000–02	0.24 (0.03)		
<i>Regions[†]</i>			
New England	-0.38 (0.03)	South Atlantic	-0.70 (0.04)
Mid Atlantic	-0.55 (0.03)	East South Central	-0.77 (0.03)
East North Central	-0.35 (0.03)	West South Central	-0.63 (0.03)
West North Central	-0.72 (0.03)	Mountain	-0.27 (0.03)

• Coefficient estimates for (4.9), obtained from the CPS (1996–2008, $n = 25,000$). Standard errors in parentheses.

* $z_1 = 1$ if a mother's youngest child was aged over 12 when time-limits were introduced in her state, so that she was never never subject to time limits. $z_2 = 1$ if the mother's youngest child was born before time limits were introduced, in which case she was only partially exposed to time limits.

† Regions taken from the US Census Bureau; omitted region is Pacific.

C.3 Structural model

Table 11: Structural model parameter estimates

Mother's utility			
$\tilde{\gamma}_{m,c} = \exp(\mathbf{x}'_m \beta_{m,c} + \sigma_{m,c} \epsilon_{m,c})$			
$\beta_{m,c}$: constant	-1.12 (0.95)	$\beta_{m,c}$: some tertiary education	0.73 (0.50)
$\sigma_{m,c}$	0.10 (0.95)		
$\tilde{\gamma}_{m,l} = \exp(\mathbf{x}'_m \beta_{m,l} + \sigma_{m,l} \epsilon_{m,l})$			
$\beta_{m,l}$: constant	-1.27 (0.89)	$\beta_{m,l}$: some tertiary education	1.78 (0.61)
$\beta_{m,l}$: wage equation residual	1.92 (0.95)	$\sigma_{m,l}$	0.27 (0.46)
$\eta_m = 1 - \exp(\mathbf{x}'_{m,\eta} \beta_{m,\eta})$			
$\beta_{m,\eta}$: constant	-1.74 (0.38)	$\beta_{m,\eta}$: some tertiary education	0.69 (0.24)
$\beta_{m,\eta}$: youngest child aged 0–5	0.07 (0.38)		
$\delta_k = \exp(\mathbf{x}'_{\delta_k} \beta_{\delta_k} + \sigma_{\delta_k} \epsilon_{\delta_k})$			
β_{δ_k} : constant	-0.84 (0.21)	β_{δ_k} : youngest child aged 0–5	0.29 (0.20)
β_{δ_k} : more than one child	0.05 (0.32)	σ_{δ_k}	0.24 (0.16)
τ (utility cost of working)	0.14 (0.02)		

- Standard errors in parentheses, calculated as per Section 4.2. ϵ always denotes an i.i.d. standard Gaussian disturbance.

Table 12: Structural model parameter estimates

Children's utility			
$\tilde{\gamma}_{k,c} = \exp(\mathbf{x}'_{k,c}\beta_{k,c} + \sigma_{k,c}\epsilon_{k,c})$			
$\beta_{k,c}$: youngest child aged 0–5	0.24 (0.45)	$\beta_{k,c}$: youngest child aged 6–18	-0.37 (0.29)
$\sigma_{k,c}$	0.44 (0.17)		
$\tilde{\gamma}_{k,t} = \exp(\mathbf{x}'_{k,t}\beta_{k,t} + \sigma_{k,t}\epsilon_{k,t})$			
$\beta_{k,t}$: youngest child aged 0–5	-1.62 (0.36)	$\beta_{k,t}$: youngest child aged 6–18	-1.01 (0.15)
$\sigma_{k,t}$	0.45 (0.10)		
$\eta_k = 1 - \exp(\mathbf{x}'_{k,\eta}\beta_{k,\eta})$			
η_k : constant	-0.18 (0.41)	η_k : youngest child aged 5 or under	1.43 (1.54)
η_k : two or more children	0.47 (0.26)		
Home production			
$\tilde{\gamma}_{q,c} = \exp(\mathbf{x}'_{q,c}\beta_{q,c} + \sigma_{q,c}\epsilon_{q,c})$			
$\beta_{q,c}$: youngest child aged 0–5	0.94 (0.16)	$\beta_{q,c}$: youngest child aged 6–18	2.93 (1.00)
$\sigma_{q,c}$	0.55 (0.14)		
$\eta_q = 1 - \exp(\mathbf{x}'_{q,\eta}\beta_{q,\eta})$			
$\beta_{q,\eta}$: youngest child aged 0–5	-0.50 (0.66)	$\beta_{q,\eta}$: youngest child aged 6–18	0.62 (0.15)

- Standard errors in parentheses, calculated as per Section 4.2. ϵ always denotes an i.i.d. standard Gaussian disturbance.

D Calculation of taxes and welfare payments

The calculation of household’s after-tax earnings is an important component of the model. (Recall the discussion of the household’s budget constraint given in Section 3.2.) Details of the relevant tax and benefit calculations, and the sources that are drawn upon to make these calculations, are given below. In general, after-tax earnings will vary according to: before-tax earnings; the number of children (aged 0–2, 0–5, 0–13, 0–16 and 0–18) in the household; the age of youngest child; childcare expenditures; and the hours worked by the mother.

D.1 Taxes

Federal taxes Federal tax parameters are taken from the Urban Institute’s *Tax Policy Centre* (TPC) database. I first determine the tax-free threshold as the sum of the standard deduction and the dependent exemption (multiplied by the number of dependants). Then the usual bracketing formula is applied to income above this threshold.

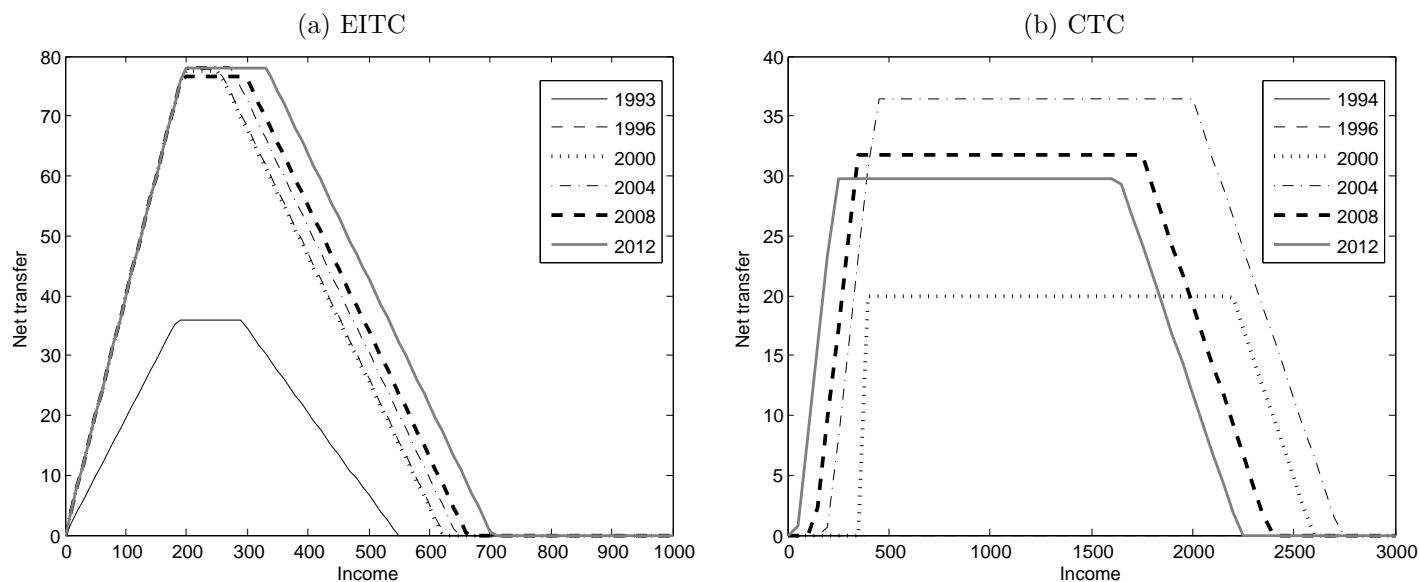
State taxes State tax parameters, for those states that do have income taxes, are drawn principally from the *State Tax Handbook*. These either take the form of a flat rate on before-tax incomes (CT, IL, IN, NH, PA), or involve a similar bracketing calculation to that used to compute federal taxes.

D.2 Tax credits

Parameters for the tax credits are drawn from the TPC. Each of these are means tested, and are calculated on the following basis. As before-tax earnings y increases from zero, the value of the credit increases linearly from zero, at rate r_{in} . This is the phase-in region: once a certain income level is reached, denoted \bar{y}_{in} , the credit stops increasing and remains at $b = r_{\text{in}} \cdot \bar{y}_{\text{in}}$. Finally, beyond a certain threshold \bar{y}_{out} , the value of the credit declines linearly with income to zero, at rate r_{out} (the phase-out region). Mathematically, the value of the credit for which a household is eligible can be expressed as:

$$\text{credit}(y) = \begin{cases} r_{\text{in}} \cdot y & \text{if } y \leq \bar{y}_{\text{in}} \text{ (phase-in)} \\ b & \text{if } y \in [\bar{y}_{\text{in}}, \bar{y}_{\text{out}}] \\ \max\{r_{\text{out}} \cdot [b - (y - \bar{y}_{\text{out}})], 0\} & \text{if } y \geq \bar{y}_{\text{out}} \text{ (phase-out)} \end{cases} \quad (\text{D.1})$$

Figure 5: Federal tax credits
Weekly transfers and income (year 2000 dollars)



- Calculated for a two-parent family with two children aged 0–16.

The household's tax liabilities are then reduced by the magnitude of the credit. If the credit is refundable, then the household receives any excess of the credit over its tax liabilities as a payment.

Earned Income Tax Credit This credit is refundable. The formula for the federal EITC is as in (D.1); the parameters (b , r_{in} , r_{out} , \bar{y}_{in} and \bar{y}_{out}) vary over time and with family size (the number of children aged 0–16). The value of the credit for a [**two-parent**], two-child family, is displayed in Figure 5. The state EITC is an additional credit that is calculated as a proportion of the federal EITC benefit; this proportion varies across time and states (and is zero in most states).

Child Tax Credit This credit is non-refundable. The formula for the credit is (D.1); again its parameters vary over time and family size. The value of the credit is depicted in Figure 5, which illustrates that the credit is not as strictly means tested as is the EITC.

Child and Dependent Care Tax Credit This credit is non-refundable. Its formula is given by

$$\text{cdctc}(y) = \begin{cases} \min\{r_{\text{in}} \cdot c, \bar{c}_{\text{max}}\} & \text{if } y \leq \bar{y}_{\text{in}} \\ \min\left\{\left(r_{\text{in}} - 0.01 \cdot \frac{y - \bar{y}_{\text{in}}}{2000/52}\right) \cdot c, \bar{c}_{\text{max}}\right\} & \text{if } y \in [\bar{y}_{\text{in}}, \bar{y}_{\text{out}}] \\ \min\{r_{\text{out}} \cdot c, \bar{c}_{\text{max}}\} & \text{if } y \geq \bar{y}_{\text{out}} \end{cases}$$

where c denotes the household's childcare costs, and \bar{c}_{max} the maximum value of the benefit.

D.3 Welfare payments

Supplemental Nutritional Assistance Program Parameters for SNAP are taken from Eslami, Leftin, and Strayer (2012, Appendix G). Letting $y_{\text{net}} = 0.8y - d_{\text{std}}$, where d_{std} denotes the standard deduction, the benefits (in the form of food stamps) for which a household is eligible is given by:

$$\text{snap} = \begin{cases} \bar{b} - 0.3 \cdot y_{\text{net}} & \text{if } y \leq 0.3 \cdot y_{\text{pov}} \text{ and } y_{\text{net}} \leq y_{\text{pov}} \\ 0 & \text{otherwise} \end{cases}$$

where \bar{b} denotes the maximum benefit payable, and y_{pov} the official poverty threshold.

Temporary Assistance for Needy Families Parameters for TANF are taken from the Urban Institute's *Welfare Rules Database*. The benefits for which a household is eligible under TANF is calculated, for most states (excepting CT, HI, MO, NV, NJ, ND and WI), from the positive part of

$$\text{tanf} = \min\{\bar{b}, r \cdot [b - (y - d)]\} \quad (\text{D.2})$$

where \bar{b} denotes the maximum benefit claimable, r is the rateable percentage, b the benefit standard, y before-tax earnings, and d the disregard. Both the benefit standard and the maximum benefit vary with the number of individuals in the family. In most states, TANF can only be claimed for a maximum of five years, and the value of the disregard depends on the number of years for which TANF has already been claimed. To simplify the calculations – and because I do not observe the number of times that an individual has claimed TANF – I set d to the average of these values. The disregard also depends on the amount that the household spends on childcare, which I take account of in my calculations.

Aid to Families with Dependent Children AFDC is calculated using an identical formula to (D.2); its parameters are drawn from the Urban Institute's *Trim3* database.

E Policy counterfactuals: full listing of results

Table 13: Effect of AFDC/TANF on household choices and welfare

		Δc_m	Δc_k	Δc_q	Δh_m	Δt_m	Δq_m	Δq	Net cost*	CV_m	CV_k	Pass thru.*	Grp. size*
		1	2	3	4	5	6	7	8	9	10	11	12
AFDC (from 1995) replaces TANF[†]													
<i>Mean responses for subgroup which received benefits under:</i>													
1	AFDC	-8.2	6.1	-2.0	-13.3	0.7	1.4	1.4	106.9	49.3	13.6	12.7	23.3
	AFDC only	-12.5	6.2	-6.3	-16.7	0.9	1.7	1.4	128.5	57.2	14.8	11.5	21.2
	TANF	18.8	13.2	32.0	4.5	0.2	4.5	3.0	16.0	36.4	21.3	132.9	8.4
AFDC (from 1995) replaces TANF; EITC set to 1995 parameters[†]													
<i>Mean responses for subgroup which received benefits under:</i>													
4	AFDC	-6.7	7.2	0.5	-12.8	0.7	1.4	1.6	110.0	53.3	15.1	13.7	24.5
	AFDC only	-10.7	7.5	-3.2	-16.3	0.9	1.7	1.7	132.6	61.8	16.4	12.4	22.3
	TANF	18.1	12.6	30.6	4.1	0.2	0.2	2.9	16.1	38.6	21.5	133.3	8.4
Hypothetical policy interventions[‡]													
<i>Mean responses for subgroup affected by the hypothetical policy</i>													
7	EITC expansion	4.0	3.4	7.4	0.2	0.0	1.0	19.6	0.0	6.7	34.2	31.7	31.7
Introduction of policy, for years indicated[§]													
<i>Mean responses for subgroup affected by the policy introduced</i>													
8	TANF (1996–2008)	-2.4	9.4	7.0	-11.7	0.6	1.6	1.9	102.3	63.8	17.3	16.9	8.3
	AFDC (1993–95)	-1.0	17.9	16.9	-19.7	1.2	2.8	4.3	183.3	100.3	35.4	19.3	25.9
	EITC (1996–2008)	10.0	11.2	21.3	1.1	0.3	0.3	2.6	50.3	16.1	20.4	40.6	65.8

• Reports the results of counterfactual exercises conducted using the model, as described in Section 6.2. Units: c_m , c_k and c_q , *net cost*, CV_m , CV_k are in \$/wk (2000 dollars); h_m and t_m in hr/wk; *pass thru.* and *grp. size* in percentage points.

* *Net cost* is the cost of the policy intervention to the government, calculated as the increase in net transfers (benefits paid less taxes received) paid to households as a result of the policy. *Pass thru.* is CV_k divided by the net cost (in percentage points). *Grp. size* is the percentage of households (sole mothers without a college degree) in the specified subgroup: for all rows except row 10, the relevant universe is all households in years 1996–2008 of the sample; for row 10, the universe is households in years 1993–95.

† In both these ‘AFDC’ counterfactuals, the baseline is the tax and welfare system in place over 1996–2008; in the counterfactual, AFDC as it existed in 1995 is introduced in each of these years (suitably adjusting programme benefits for inflation). In the second of these counterfactuals, the parameters of the EITC are simultaneously set to those in place in 1995 (again, adjusting for inflation).

‡ In both cases, the baseline is the tax and welfare system in place over 1996–2008. *EITC expansion* introduces additional benefit payments for 1-child families, as proposed by Hoynes (2014). *TANF: no work requirements* keeps the existing benefit schedule for TANF in place, but relaxes work requirements (following Bitler and Hoynes, 2016) by eliminating the (additional) disutility incurred by TANF enrollees who fail to meet these requirements.

§ In each case, the baseline is the tax and welfare system in place over the specified period, sans the indicated policy; this policy is (re-)introduced in the counterfactual.

Table 14: Alternatives to welfare-to-work programmes

	Δc_m	Δc_k	Δc_q	Δh_m	Δq_m	Δt_m	Δq	Bene- fits*	Net cost*	CV_m	CV_k	Pass thru.*	Grp. size*
	1	2	3	4	5	6	7	8	9	10	11	12	13
Childcare policies: mean responses among affected households[†]													
1	CDCTC	5.1	12.5	17.6	3.1	-0.5	-0.8	1.2	18.0	14.8	12.9	14.8	11.3
	<i>Expansion</i>	28.7	12.8	41.5	6.2	0.2	-0.2	2.6	—	51.7	—	19.8	31.0
	Free childcare	64.2	29.5	93.7	15.2	-1.0	0.2	5.7	125.6	101.4	21.0	49.5	38.0
TANF eliminated; replaced by free childcare:[‡] mean responses for subgroups[§]													
4	Either	44.7	0.2	3.7	10.3	-0.9	-0.9	3.7	103.7	86.5	15.0	35.3	44.1
	TANF	53.9	-0.2	3.2	24.4	-3.0	-3.0	3.2	98.5	40.9	11.0	26.9	13.1
6	Wage subsidy	65.3	0.3	5.9	15.3	-1.1	-1.1	5.9	127.1	109.9	21.0	52.7	37.5
	AFDC	28.4	0.2	2.1	8.6	-1.0	-1.0	2.1	74.4	66.1	11.9	20.7	33.7
TANF eliminated; replaced by flat wage subsidy (\$1/hr):[‡] mean responses for subgroups[§]													
8	Either	7.5	4.4	11.9	1.2	0.0	0.0	0.9	31.3	10.6	-12.0	8.1	66.7
	TANF	19.5	1.7	21.3	16.7	-1.7	-0.4	0.5	22.9	-71.6	-24.2	-0.1	16.9
10	Wage subsidy	9.7	5.7	15.4	6.3	0.0	0.1	1.4	30.8	12.6	-9.8	10.7	59.7
	AFDC	1.9	0.0	1.9	2.4	-0.3	0.0	0.0	21.9	-8.0	-13.8	1.3	34.3
TANF eliminated; replaced by targeted wage subsidy (all wages raised to \$11.60):[‡] mean responses for subgroups[§]													
12	Either	32.3	18.9	51.2	5.0	0.5	0.3	4.8	134.9	67.9	44.4	33.6	64.1
	TANF	29.2	16.1	45.4	6.3	0.0	0.3	4.1	138.0	7.1	42.7	37.3	8.4
14	Wage subsidy	39.2	22.6	61.8	6.3	0.6	0.3	5.7	145.0	74.5	48.9	40.4	81.0
	AFDC	50.8	23.4	74.3	20.4	-1.2	0.1	5.5	128.2	57.2	43.8	31.5	19.6

• Reports results of counterfactual exercises described in Section 6.3. Only households in years 1996-2008 of the sample are exposed to the policies. Units: c_m , c_k and c_q , *benefits*, *net cost*, CV_m , CV_k and *earnings* are \$/wk (2000 dollars); h_m , q_m and t_m in hr/wk; *pass thru.* and *grp. size* are in percentage points.

* For definitions of *net cost*, *pass thru.*, and *grp. size*, see the corresponding note to Table 6; for the latter two, the relevant universe is all households in years 1996-2008 of the sample. *Benefits* is the average benefit paid out, to the indicated subgroup, by the policy introduced under the counterfactual.

† In *CDCTC*, the baseline is the tax and welfare system in place over 1996-2008, sans CDCTC, which is (re-)introduced in the counterfactual. *Expansion* additionally expands the CDCTC in the manner proposed by Ziliak (2014). In *free childcare*, the baseline additionally includes CDCTC, which in the counterfactual is replaced by universal free childcare. Table reports means responses for those households affected by the counterfactual policy change (in row 1, this would be all households receiving the CDCTC; in row 2, all households receiving free childcare).

‡ In each exercise, the baseline is the tax and welfare system in place over 1996-2008; in the counterfactual, TANF is eliminated and replaced by the policy indicated.

§ *Either* refers to those households who either received benefits from TANF in the baseline, or from the policy introduced under the counterfactual, i.e. this subgroup collects all households that were affected by the counterfactual exercise. *Free childcare* refers to those households that would have received subsidised childcare under the counterfactual; *wage subsidy* those that would have received the specified wage subsidy under the counterfactual. *AFDC* here refers to those households that claimed AFDC benefits in the first counterfactual exercise displayed in Table 6, under which TANF was replaced by AFDC (as in 1995) throughout 1996-2008.

References for appendices

- BITLER, M., AND H. HOYNES (2016): “Strengthening Temporary Assistance for Needy Families,” in *Improving Safety Net and Work Support*. The Hamilton Project.
- BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): “Consumption inequality and partial insurance,” *American Economic Review*, 98(5), 1887–1921.
- BRUINS, M., AND J. DUFFY (2015): “The econometrics of time-use data,” *Working paper*.
- DUNBAR, G., A. LEWBEL, AND K. PENDAKUR (2013): “Children’s Resources in Collective Households: Identification, Estimation and an Application to Child Poverty in Malawi,” *American Economic Review*.
- ESLAMI, E., J. LEFTIN, AND M. STRAYER (2012): “Supplemental Nutrition Assistance Program Participation Rates: Fiscal Year 2010,” Discussion paper, United States Department of Agriculture, Appendix G.
- GURYAN, J., E. HURST, AND M. KEARNEY (2008): “Parental Education and Parental Time with Children,” *Journal of Economic Perspectives*.
- HOYNES, H. (2014): “Proposal 11: Building on the Success of the Earned Income Tax Credit,” in *Improving Safety Net and Work Support*. The Hamilton Project.
- ZILIAK, J. (2014): “Proposal 10: Supporting Low-Income Workers through Refundable Child-Care Credits,” in *Improving Safety Net and Work Support*. The Hamilton Project.