Check for updates

Scand. J. of Economics 123(1), 33–59, 2021
DOI: 10.1111/sjoe.12406

The Anatomy of the Extensive Margin Labor Supply Response*

Spencer Bastani†

Linnaeus University, SE-351 06 Växjö, Sweden
spencer.bastani@lnu.se

Ylva Moberg‡

Swedish Institute for Social Research, Stockholm University, SE-106 91 Stockholm, Sweden
ylva.moberg@sofi.su.se

Håkan Selin§

Institute for Evaluation of Labour Market and Education Policy (IFAU), SE-751 20 Uppsala, Sweden
hakan.selin@ifau.uu.se

Abstract

We estimate how labor force participation among married women in Sweden responded to changing work incentives implied by a reform in the tax and transfer system in 1997. Using rich, population-wide, administrative data, we estimate an average participation elasticity of 0.13, thereby adding to the scarce literature estimating participation elasticities using quasi-experimental methods. We also highlight that estimated extensive margin responses necessarily are local to the observed equilibrium. Among low-income earners, elasticities are twice as large in the group with the lowest employment level, compared with the group with the highest employment level.

Keywords: Housing allowance; in-work tax credits; labor supply; social assistance; take-up of transfer programs

JEL classification: *H*20; *J*22

---

*We are grateful for helpful comments and suggestions from Andrea Weber, Björn Öckert, Lina Aldén, Mikael Elinder, Hilary Hoynes, Claus Kreiner, Che-Yuan Liang, Eva Mörk, Andreas Peichl, Jim Poterba, Mari Rege, Olof Åslund, two anonymous referees, and seminar participants at MIT, Mannheim/ZEW, Stockholm Institute for Transition Economics (SITE), Uppsala University, University of Nuremberg, DIW Berlin, the Nordic Tax Workshop in Helsinki, the IIPF Conference in Taormina, the CESifo Public Sector conference in Munich, the Linnaeus Conference on Discrimination and Labour Market Research in Kalmar, IEB Barcelona, Nordic Summer Institute in Labour Economics (Helsinki), and VATT in Helsinki. Financial Support from the Jan Wallander and Tom Hedelius Foundation is gratefully acknowledged. An earlier version of this paper circulated under the title “Estimating participation responses using transfer program reform”.
†Also affiliated with Uppsala Center for Fiscal Studies, Uppsala Center for Labor Studies, and CESifo.
‡Also affiliated with Uppsala Center for Fiscal Studies.
§Also affiliated with Uppsala Center for Fiscal Studies, Uppsala Center for Labor Studies, and CESifo.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.
This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

I. Introduction

How the labor force participation of secondary earners responds to work incentives is a question of great academic and policy interest. Labor supply responses along the extensive margin are often quantified by participation elasticities, measuring the percentage change in labor force participation in response to a percentage change in the financial reward from working. Participation elasticities determine the efficiency gains from tax breaks for secondary earners and represent a key concept in the literature on optimal tax and transfer systems (see Immervoll *et al*., 2011). Nonetheless, there are very few quasi-experimental estimates of participation elasticities, as evident from the meta-analysis by Chetty (2012).[1]

In this paper, we provide new quasi-experimental evidence on how labor force participation reacts when the work incentives of secondary earners change, and we also provide a first systematic analysis of how participation elasticities differ across groups with different initial employment levels. The latter is a key contribution of our paper, as the relationship between the labor supply response and the employment level has only previously been highlighted in the structural labor supply literature, which typically has relied on small, survey-based data sets.[2] This is an important gap to fill, because knowledge about heterogeneous elasticities is essential when designing tax reforms with the purpose of promoting labor force participation, as such reforms are often targeted to specific groups of the population.

Exploiting high-quality administrative data on the full population of Swedish taxpayers, we make two primary contributions. First, we present an estimate of 0.13 of the average participation elasticity in a population of women where the average labor force participation already is high. Second, exploiting our large sample size, we partition the sample and systematically investigate the participation responses for different subgroups of individuals with different baseline employment rates. We divide the sample into four equal groups based on the wife's potential income (predicted income)

[1]The enormous literature on in-work tax credit policies focuses on singles. Eissa and Hoynes 2004, Francesconi *et al*. (2009), Bosch and van der Klaauw (2012), and Ellwood (2000) are notable exceptions. To our knowledge, the only previous quasi-experimental studies that explicitly report the secondary earner's participation elasticity are Selin (2014) and Kosonen (2014). Related papers using quasi-experimental methods to estimate the effect of child-care prices on female labor supply are Lundin *et al*. (2008) for Sweden and Havnes and Mogstad (2011) for Norway. None of them has found an effect of child-care prices.

[2]When surveying a large number of elasticity estimates from the structural labor supply literature, Bargain and Peichl (2016) noted that married women's elasticities tend to be larger in countries with low female labor force participation. Bargain *et al*. (2014) find a similar pattern when using a coherent structural estimation approach on micro data from 17 European Union countries and the United States.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

and, interestingly, we find elasticities that are monotonically falling in the potential income of the wife, ranging from 0.24 to 0.09. The results suggest that cross-country and within-country comparisons of participation elasticities should always be made with reference to the relevant employment level. Our work complements, and is also broadly in line with, earlier structural labor supply studies on Swedish data. Flood *et al*. (2004), for example, also found fairly low elasticities for Swedish married women.

For identification, we use a reform in the Swedish system for housing allowances for couples with children in 1997. Before 1997, the housing allowance was means-tested based on family income – a family received maximal housing allowance if the joint income of the household did not exceed 117,000 Swedish kronor (SEK), which is approximately 15,000 US dollars (USD). For every SEK of household income in excess of SEK 117,000, the housing allowance was reduced by SEK 0.2. After the reform, the system was individualized so that every SEK of individual labor income earned in excess of SEK 58,500 reduced the housing allowance by SEK 0.2. The overall effect of the reform was to substantially lower participation tax rates (PTRs) for secondary earners married to low- and middle-income spouses, mainly by making it less attractive not to work.

Following earlier work on the labor supply of secondary earners in survey data (e.g., Eissa and Hoynes, 2004; Francesconi *et al*., 2009), we compare eligible households (with children) with ineligible households (without children) before and after the 1997 reform, and we provide graphical evidence of the reactions to the reform. As we have access to several pre-reform years of data, we can examine the parallel trends assumption. We focus on wives married to husbands with an income below the median, and we document that female employment increases in households with children relative to households without children in the post-reform period.[3] We carefully calculate PTRs in the treatment and control groups before and after the reform, and use the reform-based variation in these tax rates to estimate participation elasticities.

The paper is organized as follows. In the next section, we describe the 1997 reform in the Swedish housing allowance system. In Section III, we describe our data sources. In Section IV, we develop a model for interpreting the evidence, and in Section V, we present the empirical strategy. We provide

[3] From a different angle, the same reform has already been analyzed by Enström Öst (2012). Using data from the Swedish Social Insurance Agency, she compares earnings growth in households with different income compositions in 1996. She estimates significant earnings responses for women. In an experimental study on US data, Jacob and Ludwig (2012) estimated a negative effect of housing assistance on labor supply.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

a graphical analysis, present regression results, and report elasticities in Section VI. Finally, in Section VII, we offer concluding remarks.

II. The Reform

We begin by describing the reform in 1997 that we exploit to identify extensive margin labor supply responses.

General Description of the Transfer Program

The housing allowance system can be characterized as a guaranteed income program, as there is no work requirement for eligibility, and the associated cash transfer is reduced as a function of the income of the members of the household (means-testing). The program is administered by the Social Insurance Agency (*Försäkringskassan*) and payments are given on a monthly basis. To receive the transfer (which is a cash transfer), the household has to apply for it by the end of each year. In 1996, 180,000 Swedish couples received the housing allowance and the transfer made up an important budget share of many low-income households. The particular program that we analyze in this paper applies to low-income families with children under the age of 20. We show our motivations for our choice of control group in Section V.[4]

Incentive Effects

To ease the description of the incentive effects of the housing allowance, we introduce some notation. The housing allowance can be written as a function $B(\tilde{z}^p, \tilde{z})$ where $\tilde{z}^p$ and $\tilde{z}$ are, respectively, the qualifying income of the two spouses, or *bidragsgrundande inkomst*, which is the income concept used to assess eligibility for welfare programs in Sweden.[5] The function B is weakly decreasing in both its arguments, which reflects that the housing allowance is a means-tested program. The maximal level of the housing allowance is obtained when neither spouse has any qualifying income, and this is equal to $B(0, 0)$, which we denote B^{00}. The value of B^{00} depends on a number of non-income characteristics, such as the number of children in the household, housing costs, and the living space (m^2) of the household.[6]

Before the reform in 1997, the transfer was reduced as a function of the sum of the qualifying incomes of the two spouses. That is, the housing

[4]There is also a separate and different housing allowance system applying to young households without children that was not subject to reform, but we do not analyze this in this paper.
[5]Qualifying income includes not only earnings, but also capital income and a fraction of wealth.
[6]In Online Appendix B, we describe in more detail how the value of B^{00} is determined.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

allowance pre-reform could be written as $B(\tilde{z}^p, \tilde{z}) = B^{pre}(\tilde{z}^p + \tilde{z})$, and it took the following form,

$$B^{pre}(\tilde{z}^p + \tilde{z}) = \begin{cases} B^{00} & \text{if } \tilde{z}^p + \tilde{z} \leq 117,000 \\ \max\left\{B^{00} - h^{pre}(\tilde{z}^p + \tilde{z}), 0\right\} & \text{if } \tilde{z}^p + \tilde{z} > 117,000 \end{cases},$$

where $h^{pre}(x) = 0.2 \times (x - 117,000)$. Thus, a family received the maximum transfer if the joint income of the household did not exceed SEK 117,000. If the joint income exceeded this exemption level, then the transfer was reduced at a phase-out rate of 20 percent. Hence, if the family income was SEK 118,000, for example, then the transfer was reduced by SEK 200 $[= 0.2 \times (118,000 - 117,000)]$.

After the 1997 reform, the system was individualized so that the household received the maximum transfer only if the income of neither spouse exceeded SEK 58,500. The phase-out rate was kept at 20 percent.[7] Thus, the post-1997 housing allowance can be written as $B(\tilde{z}^p, \tilde{z}) = B^{post}(\tilde{z}^p, \tilde{z})$, defined as

$$B^{post}(\tilde{z}^p, \tilde{z}) = \begin{cases} B^{00} \\ \qquad \text{if } \tilde{z}^p \leq 58,500 \quad \text{and} \quad \tilde{z} \leq 58,500 \\ \\ \max\left\{B^{00} - h^{post}(\tilde{z}_p), 0\right\} \\ \qquad \text{if } \tilde{z}^p > 58,500 \quad \text{and} \quad \tilde{z} \leq 58,500 \\ \\ \max\left\{B^{00} - h^{post}(\tilde{z}_p) - h^{post}(\tilde{z}), 0\right\} \\ \qquad \text{if } \tilde{z}^p > 58,500 \quad \text{and} \quad \tilde{z} > 58,500 \end{cases},$$

where $h^{post}(x) = 0.2 \times (x - 58,500)$.

How did the 1997 reform affect work incentives? To answer this question, we need to make an assumption about how economic decisions within the family are organized. Even though there is individual taxation in Sweden, the transfer system depends on the income of both spouses. Hence, the total tax/transfer relevant for the labor force participation decision of one member of the family depends on the economic decision of his/her spouse. We analyze the incentive changes from the point of view of a sequential model, where the secondary earner decides whether to work or not conditional on the labor supply choice of the primary earner. For the

[7]The reform implied no change to the income thresholds, the level of the housing allowance or the phase-out rates for single parents. Therefore, single people with children could *a priori* be considered to serve as a control group to married people with children in the empirical analysis. However, because of differential employment trends and levels, we have not chosen this strategy.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

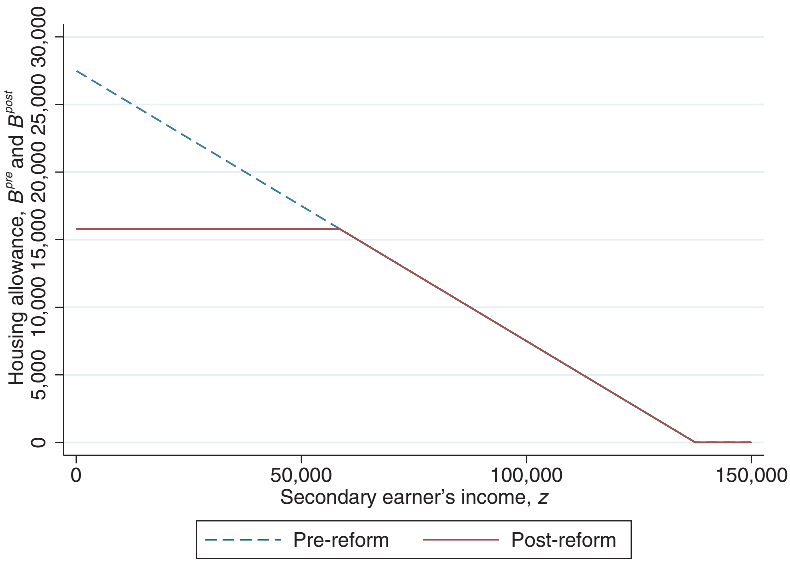

Fig. 1. Housing allowance before and after the reform
Notes: Housing allowance before and after the reform according to the functions $B^{pre}(\tilde{z}^p + \tilde{z})$ and $B^{post}(\tilde{z}^p, \tilde{z})$ as a function of secondary income $\tilde{z}$ for a family with two children. The primary earner's income is fixed at $\tilde{z}^p = 170,000$.

moment, we abstract from the take-up issue, and simply assume that the household always takes up the transfer when eligible. We discuss the model assumptions further in the second subsection of Section V.

In Figure 1, we illustrate the pre- and post-reform transfers $B^{pre}(\tilde{z}^p + \tilde{z})$ and $B^{post}(\tilde{z}^p, \tilde{z})$ for a family with two children as a function of the secondary earner's income $\tilde{z}$, while fixing $\tilde{z}^p$ to SEK 170,000 (a typical value of the primary earner's qualifying income in our estimation sample). We assume that if neither spouse were working, the household would be entitled to the maximum level of housing allowance for households with two children, $B^{00} = 38,100$. Given these assumptions, in the pre-reform scenario, the household is eligible for a transfer amounting to $38,100 - 0.2 \times (170,000 - 117,000) = 27,500$ when the secondary earner has zero earnings. According to the pre-reform rules, as soon as the secondary earner supplies any amount of positive earnings, the housing allowance is reduced. More specifically, it is reduced by SEK 0.2 for every SEK of secondary earnings up until the point where the total amount of SEK 27,500 is phased out (which happens at SEK 137,500). In the post-reform scenario, however, the transfer at zero earnings of the secondary earner is significantly smaller:

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

$38,100 - 0.2 \times (170,000 - 58,500) = 15,800$ but the phase-out does not kick in until the secondary earner exceeds the income level of SEK 58,500. At this point, the pre- and post-reform transfers are equal and the functions B^{pre} and B^{post} coincide for secondary earnings exceeding SEK 58,500.

The important lesson from Figure 1 is that if the potential earnings of the secondary earner are SEK 58,500 or more, the difference between the household's disposable income in the state of work and non-work, respectively, will entirely be driven by the difference in the transfer in the state of non-work. As most married women earn annual incomes above SEK 58,500 when working, we conclude that the variation used to recover participation elasticities in this paper is a variation in the housing allowance at zero earnings of the secondary earner. In summary, the reform makes not working much less attractive for the secondary earner. Accordingly, even though households might not be perfectly aware of the income splitting rules, households with a single earner will certainly recognize that the size of the transfer will be reduced after the reform.

Time Line and Anticipation Issues

The main objective of the 1997 reform was to cut government expenditures related to the housing allowance program. The size of the program more than doubled between 1990 and 1995 (Boverket, 2006). In April 1995, when the annual expenditures were projected to amount to more than SEK 9 billion, the Social Democratic government appointed a government committee (*Kommittédirektiv* 1995:65). The mandate of the committee was straightforward – the committee was supposed to propose expenditure reductions, for example, by changing the rules for means-testing. The committee issued their report in December 1995. The committee's proposal was similar to the reform that was to be implemented on 1 January 1997. The Social Democratic government presented a government bill in March 1996 and the bill was passed in parliament on 8 May 1996.[8]

Did households anticipate the 1997 reform? This is a key issue when interpreting the estimated elasticities (Blundell *et al.*, 2011). In principle, well-informed households could have adjusted their behavior already in December 1995 when the committee's report became publicly known.[9]

[8]The Social Democratic party was in minority in the parliament, but was supported by the Centre (agrarian) party (*Centerpartiet*).

[9]As discussed by Blundell *et al.*, it is not *a priori* clear in which direction such anticipatory responses would go. If inter-temporal substitution is the dominating mechanism, we would observe people working less in anticipation of the reform. However, if labor market friction is the key mechanism, we would expect people to start searching for new jobs already in the pre-reform period.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

However, we think that large-scale pre-reform anticipatory responses are unlikely. As far as we can tell, there was no public discussion about the income limits when the committee's report was presented.[10] According to Enström Öst (2012), the Social Insurance Agency (*Försäkringskassan*) informed beneficiaries about the reform by sending out letters in June and October 1996. Accordingly, it is likely that the vast majority became aware of the new earnings limits close to the implementation of the reform on 1 January 1997.

III. Data

This study primarily exploits large population-wide administrative data sets provided by Statistics Sweden. We have access to all key variables from 1991 and onwards. These include earned income (which we define as the sum of wage income and self-employment income), education level, geographical indicators, the number of children in the household, and region of origin. Our graphical analysis covers the years 1991–2010, whereas, as we explain in Section IV, we focus on the years 1994–2001 in the regression analysis.

Because the variables that we use are collected from administrative registers, the overall quality is very good. There are, however, two caveats with these data. One is that the 1990s data on education level for many non-natives (who obtained their education degrees from other countries) are missing. We have been able to correct the missing values by using leads of the education variable. Later, the Swedish authorities actively sent questionnaires to immigrants in which they were asked to report their education level.[11] Another caveat is that non-married, cohabiting couples without common children are observed as singles in the administrative data. Therefore, even though the housing allowance system applies to both married and cohabiting couples, we limit the sample to formally married couples. We simply do not observe cohabiting couples without children.

In line with previous literature (e.g., Eissa and Hoynes, 2004), we assume that the wife is the secondary earner and that the husband is the primary earner.[12] Accordingly, we restrict the sample to households where the husband has positive earnings. Our main identification strategy relies on

[10]A search on *bostadsbidrag* in the media archive Newsline suggests that the main media focus was on actions against fraud in the system for housing allowances, rather than work incentives when the committee presented their report. The media coverage was larger when the reform was legislated on 8 May 1996, but the focus was not on the earnings limits.
[11]Unless the individual died or migrated between year t and year 2000, we use education information as of 2000 when constructing the variable for education level.
[12]In our data, the vast majority of secondary earners are women.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

comparing households with and without children. Therefore, we restrict the sample to households where the husband's actual qualifying income falls below the median level of qualifying income, which is an income concept used by the government to compute eligibility for transfers.[13] The cut-off at the median income was chosen because it corresponds to an income level of around SEK 230,000 in 1996, and households with levels of qualifying income exceeding this threshold were not eligible for any sizable housing allowances prior to the reform.[14] As described below in Section IV, we also run placebo regressions on a separate sample of high-income couples who were unaffected by the reform. This sample is identical to the main sample in all other respects.

We drop households where any of the two spouses are aged below 30 or above 55. As described in Section II, households with two spouses aged below 30 were subject to different housing allowance rules both before and after the reform. The upper age limit is imposed as we are interested in the labor supply behavior of prime-aged individuals, and not in retirement behavior.

Our regression analysis focuses on the time period 1994–2001, while the graphical analysis covers the years 1991–2010. The reason for focusing on the time period 1994–2001 in the regression analysis is that reliable estimates from the micro-simulation FASIT, which we use to calculate PTRs, are available from 1994 and onwards. There was also a severe macro-economic crisis at the beginning of the 1990s in Sweden. The reason for not using years after 2001 is that a large child-care fee reform was implemented in 2002 (see Lundin *et al.*, 2008).

In line with the theoretical framework presented in Online Appendix A, we construct so-called household types, h, based on the ages of the two spouses (five groups) and their education (four groups), giving rise to a total of $4^2 \times 5^2 = 400$ household types. In the empirical analysis, we use the household types as fully saturated controls for age and education. In the final subsection of Section V, we also estimate heterogeneous elasticities based on dividing the potential (predicted) income of the secondary earner into four groups. As described in Online Appendix A.3, these groups correspond to a partitioning of the household types according to the potential income of the secondary earner.

[13] In the register data, we compute qualifying income based on information on earnings and capital income, and imputing financial assets from information on capital income.

[14] The upper limits of qualifying income (i.e., the income level where the entire housing allowance was phased out) differed depending on the number of children below 20 in the household. In 1997, the upper limit was SEK 267,000 for one child, SEK 307,500 for two children, and SEK 351,000 for three or more children. As we pool all households in the main analysis, we cannot use separate income cut-offs.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

IV. Difference-in-Difference Analysis

Model

Our empirical analysis consists of two parts. In this section, we begin with the simplest and most transparent specification, a difference-in-difference analysis, by comparing the evolution of the average labor force participation for secondary earners in families with and without children before and after the reform. In Section V, we proceed to estimate participation elasticities by relating the change in labor force participation of secondary earners to the change in PTRs induced by the reform.

Here, we focus on the following specification:

$$e_{ihkt} = \mu_{kt} + \mu_t + \mu_k + \mu_h + \mu_{hk} + \mu_{ht} + \delta X_{ihkt} + \upsilon_{ihkt}, \tag{1}$$

where e_{ihkt} is a dummy equal to one if an individual i in household type h with parental status k at time t is working, and zero otherwise, μ_k is a dummy variable for having at least one child aged below 20 in the household, μ_t is a time fixed effect, μ_h is a household type fixed effect, μ_{kt} is a shorthand for the interactions between the child dummy and the time dummies, and μ_{ht} represents the interaction between the household type dummies and the time dummies.

In equation (1), we are interested in the interactions μ_{kt} between the indicator variables for having children and the year dummies. For the post-reform years, these interactions capture the reform effect and its dynamics over time. The dynamic dimension is crucial: in the presence of adjustment costs, we expect the long-run response to be larger than the short-run response.

The identifying assumption in the difference-in-difference specification is that the labor supply behavior of secondary earners with children (the treatment group) would have evolved in a similar way to the labor supply behavior of secondary earners without children (the control group), in the absence of the reform. This cannot, for obvious reasons, be verified directly (we do not know how treated individuals would have behaved if they had not been treated). Instead, as is customary in the literature, we proceed in the following way.

First, we use the fact that we have access to several years of pre-reform data to test if the labor force participation trends in the treatment and control group were similar in the years before the reform. The pre-reform trends are reflected in the coefficients of the μ_{kt} interactions in equation (1) for the years before the reform. If the trends are parallel in the years before the reform, this increases the likelihood that the post-reform trends also would have been parallel in the absence of the reform.

Second, we run placebo regressions on a sample of high-income households (which were essentially all unaffected by the reform,

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

Table 1. *Reduced-form effects (in percentage points)*

	Low income				High income
	(1)	(2)	(3)	(4)	(5)
Year 1994 × children	−0.060	−0.152	0.000	−0.097	−0.264
	(0.130)	(0.129)	(0.163)	(0.159)	(0.171)
Year 1995 × children	−0.097	−0.121	−0.095	−0.140	0.016
	(0.111)	(0.110)	(0.140)	(0.137)	(0.149)
Year 1997 × children	0.120	0.154	0.348**	0.404***	−0.117
	(0.114)	(0.113)	(0.144)	(0.141)	(0.153)
Year 1998 × children	0.129	0.245*	0.331**	0.392**	0.000
	(0.134)	(0.132)	(0.169)	(0.164)	(0.178)
Year 1999 × children	0.652***	0.833***	0.681***	0.813***	0.189
	(0.145)	(0.144)	(0.181)	(0.177)	(0.192)
Year 2000 × children	0.976***	1.24***	0.790***	0.992***	0.245
	(0.154)	(0.152)	(0.189)	(0.185)	(0.202)
Year 2001 × children	1.214***	1.485***	0.863***	1.120***	0.385*
	(0.160)	(0.159)	(0.196)	(0.193)	(0.211)
Household type dummies	No	Yes	Yes	Yes	Yes
Household type × children	No	No	Yes	Yes	Yes
Household type × year dummies	No	No	Yes	Yes	Yes
Additional controls	No	No	No	Yes	Yes
No. of observations	2,770,100	2,770,100	2,770,100	2,770,100	1,385,071

Notes: The dependent variable is the probability of having positive earnings. The low-income sample consists of wives married to husbands with a positive qualifying income, which falls below the 50th percentile. The high-income sample consists of wives married to husbands with a positive qualifying income that falls above the 75th percentile. All specifications contain a dummy for having children and a full set of year dummies. 400 household types are defined based on five age dummies for each spouse and four education-level dummies for each spouse. The additional control variables are specified in Section V. Standard errors reported below the estimates are robust to heteroskedasticity and clustered at the household level. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

independently of whether they had children or not). If the labor force participation of secondary earners with and without children in high-income households evolved similarly after the reform, this increases the likelihood that the post-reform trends for low-income households with and without children would have evolved similarly in the absence of the reform. Our placebo test amounts to estimating equation (1) on the sample of women married to husbands with qualifying incomes above the 75th percentile. In 1996, this corresponded to an income level of around SEK 310,000.[15]

The complete set of results for the reduced-form effects analysis are presented in Table 1, and Figure 2 presents graphical evidence. The housing allowance reform occurred in 1997, which means that the estimation sample

[15] Here, it should be noted that some households with three or more children could be eligible for housing allowance up to SEK 351,100.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

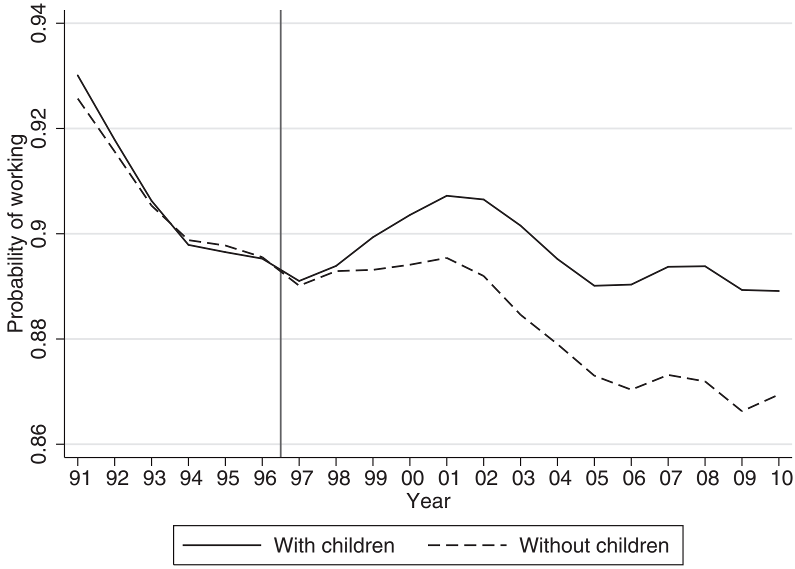

Fig. 2. Graphical reduced-form evidence and long-term trends
Notes: Female participation (share with positive earnings) in low-income households where the husband participates in the labor force.

contains three pre-reform years and five post-reform years. We choose 1996 as the reference year.

Graphical Evidence and Regression Results

Beginning with the graphical analysis in Figure 2, we can see how the employment of married women (defined as having positive earnings) evolved in couples with and without children between 1991 and 2010. A nice feature of Figure 2 is that it illustrates the evolution of employment outside the more narrow time period of our regression analysis.[16] We make the following observations. At the beginning of the 1990s, there was a sharp decline in employment as a result of a deep economic recession. Figure 2 suggests that female employment decreased slightly more among households with children during the period 1991–1993. However, between 1993 and 1996, the two lines moved in parallel. Note

[16]In the graphical analysis and in the regressions, we employ the same sample restrictions; that is, we focus on households where the husband's qualifying income falls below the 50th percentile and where the husband reports positive earnings.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

also that the employment levels are strikingly similar. After the 1997 reform, employment continued to evolve similarly until 1998. Then, there is a clear employment increase for women with children relative to women without children, an increase that continued until 2001. After 2001, the two graphs appear again to follow more or less similar trends. The fact that the response appears to be increasing in the post-reform years suggests that adjustment costs (e.g., costs associated with finding a new job) could be important. As discussed already at the end of Section II, information about the reform became publicly available just before its implementation, and it probably took some time for households to adjust.

Columns 1–4 in Table 1 show the coefficients for the main low-income sample where most households with children were eligible for housing allowances (at zero earnings of the wife). First of all, it should be noted that the results appear to be quite robust across the different specifications in these four columns. The first column reports the results of a difference-in-difference specification without any control variables. In this column, the first thing to notice is that the coefficients for the pre-reform years, 1994 and 1995, are statistically insignificant, confirming the visual evidence of Figure 2 that the pre-reform trends were very similar for the treatment and control groups. In fact, the coefficients for the pre-reform years remain insignificant for all the specifications that we have considered, as evident from Columns 1–4. Moreover, also consistent with Figure 2, we see that there is a statistically significant response to the reform in 1999 and that the response grows monotonically across the post-reform years. For 2001, the estimated effect amounts to 1.2 percentage points.

In Column 2, we have added household type controls and the estimated effects become slightly larger. In Column 3, we control for trends in a flexible way including the full set of interactions between the time dummies and the household type dummies as well as the interactions between the household type dummies and the dummy for having children. Interestingly, in this specification, the reduced-form effect estimates are also significant for the two post-treatment years 1996 and 1997 (at the 5 percent level). Finally, when the full set of controls is included in Column 4, the overall pattern of coefficients is similar to Column 3, but the reform effect estimate for 2001 is more in line with that obtained in the specification without controls in Column 1. Our preferred estimate of the reform effect is the coefficient for 2001 in our most ambitious specification of Column 4, and amounts to a 1.12 percentage point increase in the probability that married women will participate in the labor force.[17]

[17]These results are robust to excluding cells (defined based on year × children × household type) that contain fewer than 100 observations.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

In Column 5, we report the results from a placebo regression with the full set of controls, where we have estimated equation (1) on a sample consisting of women married to husbands with qualifying income above the 75th percentile (which were essentially all untreated by the reform). In all other respects, the selection criteria are identical to the main low-income sample. It is reassuring that all estimated coefficients are insignificant at the 5 percent level. One interaction, the interaction for 2001, is significant at the 10 percent level, but the coefficient estimate is considerably smaller than the corresponding point estimate in the low-income sample. The results of this placebo regression, considered in conjunction with the results in Columns 1–4 (showing that the trends before the reform were parallel), and the visual evidence in Figure 2, allow us to be reasonably confident that the identifying assumption in our difference-in-difference set-up is satisfied.

Discussion

Before moving on to estimating participation elasticities, we briefly discuss a few aspects of the preceding analysis. First, we discuss some potential concerns with the use of child status as an indicator for treatment status, and the possibility of using within-individual variation to estimate the effect of the reform. Second, we discuss the standard errors in the difference-in-difference set-up. Finally, we discuss an analysis of male responses that we have done to test the validity of the primary–secondary earner assumption.

A growing body of research has stressed the importance of children and the onset of parenthood for labor market outcomes at the individual level (see Angelov *et al*., 2016; Kleven *et al*., 2019). The strategy in this paper is essentially to follow groups rather than individuals, and we have demonstrated that the pre-reform trends are parallel at the group level. When groups are defined in a coherent way over time, employment dynamics associated with child birth is not necessarily a problem. Moreover, in our analysis, child status is associated with having a child under the age of 20, so families with very small children represent only a small subset of our treatment group.

An alternative empirical strategy would be to add individual fixed effects to equation (1). In this case, the identifying assumptions would be different as one would be comparing individuals with and without children before and after the reform. However, a major concern in this case is that individuals would be observed at different stages in their life cycle before and after the reform. We have estimated equation (1), adding individual-level fixed effects, and we have confirmed these concerns. In contrast to the baseline analysis, the results are in this case very sensitive to the inclusion of control variables (e.g., controls for age). When adding the full set of control variables, the post-reform results are reasonably similar to our

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

baseline estimates in Column 4 of Table 1, but we also estimate significant pre-reform trends. The details are contained in Table A2 of Online Appendix D.[18]

Throughout the paper, we report standard errors that are clustered at the individual level. If employment shocks primarily operate on the individual level, our standard errors will be robust to serial correlation and arbitrary heteroskedasticity. Recall that we are using individual-level data on the entire population, and individuals will move in and out from the treatment group over time (as new children are born and children grow up). If one instead worries about aggregate shocks to women with and without children, then inference becomes more challenging as we essentially only have 16 clusters (two groups in eight time periods). In that case, it is well known that the cluster-robust covariance estimator is likely to be incorrectly estimated. Furthermore, if shocks are autocorrelated at the group level, the estimated standard errors would also be wrong in the presence of many clusters. Therefore, as a robustness check, we performed a randomization inference exercise influenced by Bertrand *et al.* (2002, section 4.6), by comparing the observed treatment effect estimate with a distribution of placebo treatment effect estimates obtained in Monte Carlo simulations. This exercise strongly indicates that inference is robust to both clustering and serial correlation at the group level. The details of the procedure are given in Online Appendix F.

In order to examine the validity of the primary–secondary earner assumption, we have estimated equation (1) on a sample of males. Our idea has been to construct the male sample as a mirror image of the female low-income sample by conditioning the male sample on the wife's qualifying income falling below the 50th percentile. The results are presented in Column 1 of Table A1 in Online Appendix D, where it can be inferred that the estimated coefficients for this male sample are very different from the female sample. For 1994–2000, none of the interaction terms is statistically significantly different from zero. For 2001, we estimate a negative effect on male employment equal to −0.36 percentage points, which is significant at the 5 percent level.

To dig deeper into the potential mechanisms at play, we have also examined male earnings responses at the intensive margin. We transformed earnings into log earnings in the standard way, thereby excluding

[18]As a robustness check, we have also estimated triple-difference models with individual fixed effects, incorporating also the high-income sample (which we previously exploited in the placebo regression reported in Column 5 of Table 1). In these regressions, we did not obtain significant pre-reform trends, and with the full set of control variables we estimated treatment effects that are in the same ballpark as in the baseline difference-in-difference analysis. The details are contained in Table A3 of Online Appendix D. Still, the estimates are more sensitive to control variables than in the baseline analysis.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

observations with zero earnings. We found no clear evidence of a response in log earnings after including the full set of controls, see Column 2 of Table A1. Finally, we have also estimated equation (1) on the main female sample with log earnings instead of employment on the left-hand side. The estimation of this pure intensive margin response resulted in significantly positive coefficients, especially for the years 2000 and 2001 (see Column 3 of Table A1). It is worth noting that even a small intensive margin response can have important fiscal consequences when the baseline employment rate is high.[19]

V. Participation Elasticities

In this section, we estimate participation elasticities by relating the change in labor force participation to the size of the change in the financial gain from working, as measured by the PTR.

Participation Tax Rates

To calculate PTRs, detailed information on individuals' budget sets is needed. As the housing allowance interacts with other parts of the transfer system, most notably social assistance, it is important to take into account the entire tax and transfer system when constructing households' budget sets. To achieve this, we use the microsimulation model, FASIT, developed by the Swedish Ministry of Finance and Statistics Sweden. FASIT relies on a larger set of variables than is available in the administrative registers, and therefore employs a smaller supplementary data set called HEK (*Hushållens ekonomi*) that also contains survey data.[20] Therefore, as described in more

[19]We also experimented with log(earnings + 1) as the dependent variable, thereby including females with zero earnings in the regression. We found that the estimated coefficients were significant in all post-reform years and around three times as large in Column 3 of Table A1. Even though this log transformation is controversial, the results indicate that women primarily reacted to the reform along the extensive margin (i.e., they went from zero earnings to a positive amount of earnings).

[20]After having imposed the same sample restrictions on HEK as on the administrative data, the size of the HEK sample varies between 1,000 and 2,000 observations each year. The sample is too small to be used in the labor supply analysis described in Section IV, but it is still very useful for the purpose of estimating PTRs. Remember that the households' budget sets are given deterministically by the microsimulation model and the variables in the HEK data. Of course, this does not mean that the sample size of HEK is unimportant, because the precision of the estimated group means becomes more precise the larger the number of households represented in the HEK sample. A detailed comparison between the HEK and population-wide data is provided in the Online Appendix for earnings (Tables A5 and A6) and for labor force participation (Tables A7 and A8).

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

detail below, we compute PTRs for individuals in our population-wide data using an imputation procedure based on the variables that are available in both data sets.[21]

We let $T^{total}(z^p, z)$ refer to all taxes paid and benefits received by a household with primary earnings z^p and earnings of the secondary earner equal to z, assuming the household takes up all transfers.[22] The PTR for the secondary earner is defined as

$$\tau(z^p, z) = \frac{T^{total}(z^p, z) - T^{total}(z^p, 0)}{z}. \tag{2}$$

This is the key independent variable that will appear in our estimation equations (4) and (5). Importantly, we compute PTRs for all households, assuming full take-up of housing allowance and social assistance (for eligible households).

The PTR concept implies that the household chooses between two hypothetical disposable income states: the household disposable income when the secondary earner is working, and the household disposable income when the secondary earner is not working. Note that to be able to estimate the effect of PTRs on labor force participation, we need to compute PTRs for all individuals, both labor force participants (with positive earnings) and labor force non-participants (with zero earnings) in our population-wide register data. This implies that we need to compute the potential (hypothetical) income that secondary earners with zero income would have if they started working.

We proceed in the following way. We start by calculating PTRs for all secondary earners with positive earnings in the HEK data. This is done by first computing the disposable income of each household at zero earnings of the secondary earner. We then subtract this number from the household's actual disposable income to obtain the household's financial gain from the employment of the secondary earner. Finally, we divide this financial gain by the secondary earner's earnings to obtain the PTR according to equation (2).[23]

[21] HEK includes the full set of variables that determine eligibility for the housing allowance program (such as housing costs and dwelling space) as well as the size of the benefit actually received (from registers). We also use HEK to compute the take-up of the housing allowance.

[22] The function T^{total} corresponds to $T + B$ in Online Appendix A.

[23] We acknowledge that earnings in the state of work might differ between employed and unemployed women, even conditional on observable characteristics, which might induce a selection bias. However, we have not been able to find any valid instruments that enable us to use a selection correction term. In this respect, our approach has some similarities to Gelber and Mitchell (2012) and Meyer and Rosenbaum (2001). When researchers can use variation from several tax reforms, it is, in principle, feasible to estimate the extensive and intensive margins simultaneously (see Alpert and Powell, 2014).

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

Next, pooling the HEK data for the years 1994–2001, we regress PTRs on four dummies based on the actual qualifying income of the husband (year-specific quartiles), four dummies based on the number of children in the household, and eight year dummies, as well as the full set of interactions between the income, children, and year dummies. Additionally, we include three dummies for the educational level of the wife, which we also interact with the year dummies, of course. These variables explain a very large share of the variation in PTRs, and the R^2 value exceeds 90 percent. The estimated coefficients from these regressions are then used to impute PTRs for all secondary earners (with either zero or positive earnings) in the population-wide register data.[24]

Econometric Framework

Our aim is to estimate the following relationship on secondary earners in (formally) married couples where both spouses are aged 30–55, using data for the years 1994 to 2001:

$$e_{ihkt} = \alpha + \beta\tau_{ihkt} + \eta_{ihkt}. \tag{3}$$

Here, the dependent variable e_{ihkt} is a dummy that takes the value of one if individual i with k children in household type h in year t is employed, and zero otherwise. In our baseline specification, we define employment as having positive earnings. Moreover, k will be binary in the analysis and equal to one if there is at least one child aged below 20 in the household, and zero otherwise. The independent variable τ_{ihkt} is individual i's PTR, which is calculated assuming that eligible households take up the housing allowance. Finally, η_{ihkt} is an error term. The parameter of interest is β, the participation elasticity.

Note that the participation elasticity of secondary earners is a structural parameter, and is defined in relation to a particular model of household labor supply behavior. We outline such a model in Online Appendix A, where we formally define β (see equations (13)–(15)). A key assumption in our framework is that households behave sequentially; that is, secondary earners decide about their labor supply taking the behavior of the primary

[24]Because the main purpose of FASIT has been to assess revenue effects of changes in the tax and transfer system, we had to rewrite some parts of the code to serve our purposes. Most importantly, there were no modules for computing social assistance benefits for the years 1994–1995. Hence, for these years, we wrote the code ourselves based on national guidelines for social assistance. Rules for social assistance differ across municipalities. For some, but not all, years, we can compute social assistance as a function of both municipality-specific parameters and national guidelines. For coherency, we have chosen to use national guidelines for all years. We have verified that the two methods produce similar results for the years for which both methods are available to us.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

earner as given. In line with earlier literature (see Eissa, 1995; Eissa and Hoynes, 2004), we treat the behavior of the primary earner as exogenous.[25]

The model also features an endogenous take-up decision and we specify sufficient conditions under which reforms in transfers (that are subject to take-up decisions) can be used directly to assess the sensitivity of employment to taxes (see Proposition 1 in Online Appendix A). In this paper, we transform transfer elasticities to participation elasticities by scaling the transfer elasticities with the inverse of the take-up rate. Our model clarifies the conditions under which this approach is valid.[26]

As already described in the previous subsection, we estimate τ_{ihkt} on a smaller survey data set that contains all variables necessary to compute the household's taxes and transfers accurately. Let W denote a vector of variables that are contained both in the main (population wide) data set and in the smaller survey data set (W is a subset of the variables needed to compute the PTR). We refer to the coefficient vector in the regression of τ_{ihkt} on W_{it} in the smaller data set as ρ, and we focus on the following regression model for the population wide data set,

$$e_{ihkt} = \alpha + \beta\hat{\tau}_{ihkt} + \eta_{ihkt}, \tag{4}$$

where $\hat{\tau}_{ihkt} = \hat{\rho}W_{it}$. To account for the fact that $\hat{\rho}$ is estimated with uncertainty, we have checked that the standard errors are robust to the corrections suggested by Murphy and Topel (1985); see the following subection. The fundamental condition for the validity of this imputation procedure is that ρ, and its covariance matrix, is consistently estimated on the auxiliary data source (HEK).

If we were to estimate equation (4) in a cross-section without any control variables, there would be a concern that $\widehat{\beta}$ would be biased. The reason is, of course, that $\widehat{\beta}$ would also capture the direct effects of W on

[25]The model of household behavior builds on Immervoll *et al*. (2011) and assumes Pareto efficiency, no income effects, and a sharing rule (dictating how resources are divided in the family) that is unaffected by taxes. To simplify the interpretation of our empirical results, we assume that the extensive margin of the primary earner is inelastic. This seems *ex ante* reasonable given the high participation rate of primary earners in Sweden. The non-responsiveness of primary earners along the extensive margin is also supported by our empirical results in Table A1 in Online Appendix D. The omission of income effects is not without loss of generality, but simplifies the analysis considerably and has also become standard practice in the literature (see Brewer *et al*., 2010). The secondary earner assumption is also common in the literature, but has been criticized by Gelber (2014). He questions the implication of the unitary model that a married individual's pre-tax earnings should react equally to an increase in that individual's unearned income as it reacts to an increase in the unearned income of the spouse. In our model, both of these effects are assumed to be zero as utility is linear in consumption.

[26]A similar model of labor force participation and take-up has recently been developed by Gelber *et al*. (2018).

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

e. However, if we were to include controls for W in a very flexible way, then identification would be lost. The leading idea of our paper is to exploit the 1997 housing allowance reform to address the potential endogeneity of $\hat{\tau}_{ihkt}$ in equation (4). The housing allowance reform substantially reduced PTRs for households with children in certain income intervals, but left households without children unaffected. Hence, if there are no direct effects on the outcome variable of the interactions between the children dummy λ_k and the time dummy λ_t (conditional on λ_k and λ_t), then the housing allowance reform can be used as an instrument for τ.

The richness of the data enables us to control for covariates and time trends in a very flexible way. We let λ_{kt}, the full set of interactions between the child and time dummies, be the vector of excluded instruments. Ultimately, we wish to estimate the equation

$$e_{ihkt} = \alpha + \beta\hat{\tau}_{ihkt} + \lambda_t + \lambda_k + \lambda_h + \lambda_{hk} + \lambda_{ht} + \gamma X_{ihkt} + \eta_{ihkt}, \qquad (5)$$

where X_{ihkt} is a rich set of pre-determined control variables not used to construct the household types. In the X vector, we include seven dummies for region of origin as it is well known that foreign-born individuals, on average, exhibit lower employment rates than natives.[27] In addition, we include 21 dummies for county of residence to account for regional employment differences. Moreover, we interact the dummies for region of origin and the county dummies with the children and the time dummies. Finally, we also include detailed age dummies (one dummy per age), which we interact with the children dummy.

Notice that, at the individual level, the imputed PTR $\hat{\tau}$ in equation (4) will be measured with error. The reason is that the imputations are made at the group level (see the previous subsection). However, as we instrument $\hat{\tau}$ with λ_{tk}, the requirement for consistent estimation of β in equation (5) is that the year-specific group averages are correct.

Before moving on to the results, we mention that an alternative identification strategy would have been to focus only on households with children and to define treatment status according to the income of the husband. That is, the wives with low-income husbands would be assigned to the treatment group and wives married to high-income husbands would be assigned to the control group. One main reason why we have not taken this route is that, in order for the structural interpretation of β to hold, we need to impose the assumption that the marginal effect of τ on e is the same in the treatment and control groups. In practice, this means that

[27]These regions are (i) Sweden, (ii) Western Europe, North America and Oceania, (iii) Eastern Europe and former Soviet Union, (iv) South America, (v) Sub-Saharan Africa, (vi) Northern Africa and Middle East, and (vii) Asia.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

we not only have to consider common trends for households with and without children, but we also need to check that the employment levels are reasonably similar between the groups.[28] As is apparent from Figure 2, this is indeed the case for couples with and without children. In contrast, female employment is systematically higher in high-income households than in low-income households.[29]

Results

In Figure 3, we plot the evolution of the average PTR for the treatment and control groups (households with and without children) over the time period 1994–2001.[30] As can be seen, the reform in 1997 implied a sharp drop in the average PTR for the treatment group. This drop was caused by the housing allowance reform and demonstrates the strength in the first stage of our instrumental variable (IV) strategy. Before the housing allowance reform of 1997, the gap in the average PTRs for couples with and without children, respectively, exceeded 10 percentage points. After the reform, the gap was substantially smaller.

We now turn to our IV estimates of participation elasticities. In Online Appendix A.4, we describe how we construct the participation elasticities based on the regression coefficients (marginal effects). The results are presented in Table 2. Columns 1–4 show estimates using different sets of control variables. The instruments are strongly correlated with the PTR, as a lot of the variation in PTRs comes from the interaction between time and child status. In the two-stage least-squares (2SLS) regressions presented in Table 2, the first-stage F-statistic of the excluded instruments is always very large. In each case, we obtain precise estimates of the participation elasticity. Our preferred estimate is obtained for our most ambitious set of controls (Column 4), in which case the elasticity estimate is 0.13. The exact magnitude of the elasticity estimate varies somewhat depending on the set of control variables used in the regressions. This is perhaps not too surprising in light of the results for the reduced-form effects in Table 1.

Before closing this section, we would like to remark on the generated regressor problem. The PTRs have been estimated in a separate step with

---

[28]As emphasized in Online Appendix A.3, we expect the employment response to depend on the employment level.

[29]For this reason, as explained in Section IV, we instead exploit untreated high-income households for making placebo tests. Reduced-form results are, however, quite similar if we keep low-income households with children as the treatment group, but instead use high-income households with children as the control group.

[30]We maintain the same sample restrictions as in Section IV.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

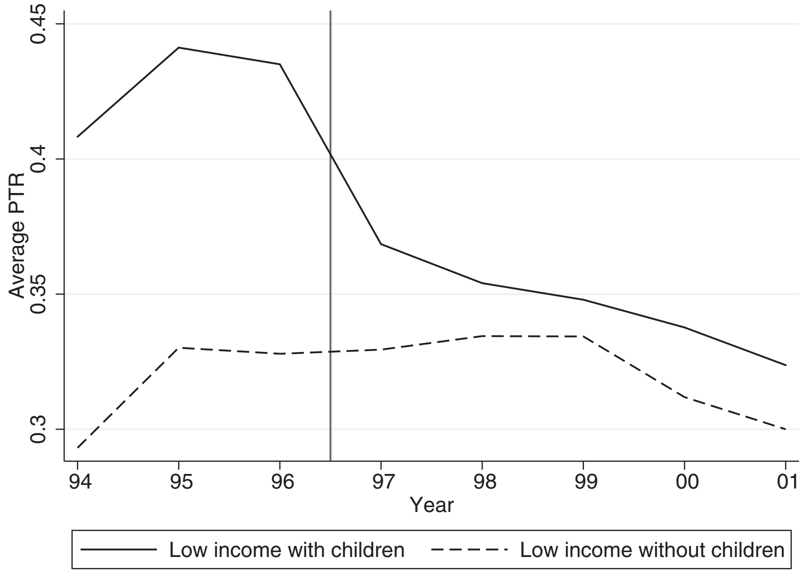

Fig. 3. Graphical first-stage
Notes: Average PTRs by child status on HEK data, calculated using FASIT. The sample is restricted to households where the husband's qualifying income falls below the 50th percentile and where the husband reports positive earnings.

some uncertainty, and our standard errors in the main analysis might be biased due to the presence of the generated regressor in equation (5). Therefore, we made a correction of the standard errors following Murphy and Topel (1985) for the specification without control variables reported in Column 1.[31] The correction did not have any substantial impact on the standard errors. The standard error increased only slightly from 0.013 to 0.014.[32] This does not come as a surprise given the large R^2 value from the imputation regression (the included variables explain more than 90 percent of the variation, so the uncertainty in the PTR predictions is small). Thus, we conclude that the imputation procedure is not problematic from the perspective of statistical inference.

[31] Performing a proper correction of the covariance matrix for the full specification, which contains a huge amount of dummy variables, would be computationally very burdensome.
[32] We make this correction while estimating equation (5) with the control function method. By construction, 2SLS and the control function method give identical coefficients under linearity, but in general the standard errors differ. In our case, the standard errors are very similar with 2SLS and the control function method.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

Table 2. *Participation elasticity estimates*

	(1)	(2)	(3)	(4)
Participation elasticity	0.088***	0.116***	0.098***	0.127***
	(0.013)	(0.013)	(0.020)	(0.019)
Household type dummies	No	Yes	Yes	Yes
Household type × children	No	No	Yes	Yes
Household type × year dummies	No	No	Yes	Yes
Additional controls	No	No	No	Yes
No. of observations	2,770,100	2,770,100	2,770,100	2,770,100

Notes: Elasticities are evaluated at the mean values of employment (0.897) and (1-PTR) (0.659) over the years 1994-2001 in the total 'low income sample'. 2SLS regressions are run on 'low income sample', which consists of wives married to husbands with a qualifying income below the 50th percentile. The average take-up rate is set to 0.6. The interactions between the year dummies and the dummy for having children are the excluded instruments. All specifications contain a dummy for having children and a full set of year dummies. 400 household types are defined based on 5 age dummies for each spouse and 4 education level dummies for each spouse. The additional control variables are specified in section V. Standard errors reported below the estimates are robust to heteroscedasticity and clustered at the household level. Standard errors for elasticities are obtained by the delta method. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

Heterogeneous Responses

In the past, extensive margin responses to taxes have been estimated on relatively small data sets. Since we have access to population wide registers we are able to examine how the elasticity differs across subpopulations with different baseline employment rates in a systematic way. More specifically, we divide the low income sample into four quartile groups based on the imputed log earnings of the secondary earner. As already mentioned in section III, and formally explained in Online Appendix A.3, this corresponds to a particular grouping of the household types. Of course, predicted earnings may reflect other differences than differences in skills, and our analysis should not be interpreted as an attempt to assess the causal effect of the skill level on the elasticity.

After partitioning the sample into four quartile groups, we rerun equation (5) on each group and evaluate the elasticity at the subsample-specific mean values of employment and $(1-\tau)$.[33] The results are shown in Table 3 where we use the full set of control variables. As we move across the four quartile groups, we see that the elasticities fall monotonically in the wife's potential income, mirrored by a corresponding monotonic increase in the employment level. In line with our expectations, the elasticity is the largest in the first quartile group, where the employment level is substantially smaller than in the other three quartile groups. The elasticity

[33] For details, see equation (16) in Online Appendix A.3.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

Table 3. *Heterogenous response*

	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Participation elasticity	0.226***	0.119*	0.110**	0.088***
	(0.056)	(0.046)	(0.037)	(0.026)
Mean employment level	0.808	0.903	0.923	0.955
PTR coefficient	−0.178	−0.102	−0.094	−0.078
Household type dummies	Yes	Yes	Yes	Yes
Household type × children	Yes	Yes	Yes	Yes
Household type × year dummies	Yes	Yes	Yes	Yes
Additional controls	Yes	Yes	Yes	Yes
No. of observations	692,559	692,542	692,476	692,523

Notes: Elasticities are evaluated at the mean values of each subsample. 2SLS regressions are run on 'low income sample', which consists of wives married to husbands with a qualifying income below the 50th percentile. Quartiles are created based on the wife's predicted income. The average take-up rate is set to 0.6. The interactions between the year dummies and the dummy for having children are the excluded instruments. All specifications contain a dummy for having children and a full set of year dummies. 400 household types are defined based on 5 age dummies for each spouse and 4 education level dummies for each spouse. The additional control variables are specified in section V. Standard errors reported below the estimates are robust to heteroscedasticity and clustered at the household level. Standard errors for elasticities are obtained by the delta method. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

estimate for the first quartile group (0.226) and the fourth quartile group (0.088) are statistically different at a level of 95 percent.[34]

Before closing this section, it is useful to compare our results with Selin (2014) who exploited the switch from joint to individual taxation in Sweden in 1971 to estimate participation elasticities for married Swedish women. Selin found estimates in the range 0.5–1.0, which are well above our estimates. However, the results are actually completely consistent when adopting the perspective of our paper. Selin (2014) reports that the pre-reform employment level for married women was 67 percent (Table 8) whereas the corresponding share in our study is 90 percent. This further highlights the important relationship between the participation elasticity and employment level that we have emphasized in this paper.

[34] Following, for example, Clogg *et al.* (1995, p. 1276), we test this using the fact that differences between the coefficients from a regression run on two independent large samples x and y can be assessed by the statistic $Z = (\hat{\beta}_x - \hat{\beta}_y)/\sqrt{se_x^2 + se_y^2}$, which follows a standard unit normal distribution. Here, $\hat{\beta}_j$ and se_j are the coefficient and the standard error of sample $j = x, y$. Because we are interested in testing for differences in elasticities, we have made the proper adjustments by multiplying the coefficients and standard errors by different constants. Using the values for the elasticities and standard errors in Columns 1 and 4 of Table 2, we obtain a Z-ratio of 2.235, which is larger than the critical value 1.96.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

VI. Concluding Remarks

In this paper, we have investigated how the labor force participation of secondary earners responded to a large reform in the tax and transfer system in Sweden. Using detailed information about individuals' budget sets, and a specific economic framework, we have estimated participation elasticities, exploiting the reform for identification. Our central estimate of the participation elasticity is 0.13, arguably a lower value than many earlier estimates obtained in the literature. Crucially, we have also presented quasi-experimental estimates of participation elasticities for subgroups of the population with different employment levels. This exercise was made possible by virtue of our large sample size. Dividing up the population into four quartile groups based on the wife's potential income, we find participation elasticities ranging from 0.23 at the bottom to 0.09 at the top. The point estimates of the elasticities fall monotonically across these groups, and the elasticity differences between the bottom and the top are statistically significant. These results are quite intuitive: The higher the employment level, the smaller is the pool of unemployed that can be incentivized to enter the labor force. In line with the public finance literature, we have assumed that employment is voluntary and focused on the decision to enter the labor force. If involuntary unemployment is more common among those with low potential incomes, we potentially underestimate the participation elasticity in this group.

The key insight from this paper is that the participation elasticity is fundamentally different in nature from the intensive margin labor elasticity. When designing tax reforms targeted to specific groups, it is important to consider the employment level in the subpopulation of interest.[35] This point has been made before, see e.g. Chetty *et al*. (2013); our contribution is to examine this feature of the participation response using administrative data and a quasi-experimental identification strategy.

[35]Our quasi-experimental estimates provide a useful contrast against estimates obtained using microsimulation models. Immervoll *et al*. (2007) analyze welfare reforms in 15 European countries including Sweden, and calibrate the average participation elasticity for the whole economy to 0.2, but decreasing across deciles. In a related exercise, which is more focused on participation responses, Immervoll *et al*. (2011) assume participation elasticities for secondary earners in the range 0.3-0.7. In light of this paper these elasticities appear to be too large, at least for a country like Sweden.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Online Appendix
Replication Files

References

Alpert, A. and Powell, D. (2014), Estimating Intensive and Extensive Tax Responsiveness Do Older Workers Respond to Income Taxes?, Working paper, Rand Corporation.

Angelov, N., Johansson, P., and Lindahl, E. (2016), Parenthood and the Gender Gap in Pay, *Journal of Labor Economics 34*, 545–579.

Bargain, O. and Peichl, A. (2016), Own-Wage Labor Supply Elasticities: Variation Across Time and Estimation Methods, *IZA Journal of Labor Economics 5*, 10.

Bargain, O., Orsini, K., and Peichl, A. (2014), Comparing Labor Supply Elasticities in Europe and the United States: New Results, *Journal of Human Resources 49*, 723–838.

Bertrand, M., Duflo, E., and Mullainathan, S. (2002), How Much Should We Trust Differences-in-Differences Estimates?, NBER Working Paper No. 5023.

Blundell, R., Francesconi, M., and van der Klaauw, W. (2011), Anatomy of Welfare Reform Evaluation: Announcement and Implementation Effects, Institute for the Study of Labor (IZA) Discussion Paper No. 6050.

Bosch, N. and van der Klaauw, B. (2012), Analyzing Female Labor Supply: Evidence from a Dutch Tax Reform, *Labour Economics 19*, 271–280.

Boverket (2006), Bostadsbidrag, ett rättvist bostadsstöd för barnen? - långsiktiga effekter av 1990-talets besparingar, *Diarienummer: 212-3795/2006*.

Brewer, M., Saez, E., and Shephard, A. (2010), Means-testing and Tax Rates on Earnings, in S. Adam *et al*. (eds), *Dimensions of Tax Design: The Mirrlees Review*, Oxford University Press, Oxford.

Chetty, R. (2012), Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply, *Econometrica 80*, 969–1018.

Chetty, R., Guren, A., Manoli, D., and Weber, A. (2013), Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities, in D. Acemoglu, J. Parker, and M. Woodford (eds), *NBER Macroeconomics Annual 2012*, Volume 27, University of Chicago Press, Chicago, IL, 1–56.

Clogg, C. C., Petkova, E., and Haritou, A. (1995), Statistical Methods for Comparing Regression Coefficients between Models, *American Journal of Sociology 100*, 1261–1293.

Eissa, N. (1995), Taxation and Labor Supply of Married Women: The Tax Reform Act of 1986 as a Natural Experiment, NBER Working Paper No. 5023.

Eissa, N. and Hoynes, H. W. (2004), Taxes and the Labor Market Participation of Married Couples: The Earned Income Tax Credit, *Journal of Public Economics 88*, 1931–1958.

Ellwood, D. T. (2000), The Impact of the Earned Income Tax Credit and Social Policy Reforms on Work, Marriage, and Living Arrangements, *National Tax Journal 53*, 1063–1105.

Enström Öst, C. (2012), Ekonomiska drivkrafter i bostadsbidragssystemet: En utvärdering av individuella inkomstgränser för makar med barn, Technical report, Rapport 2012:6. Inspektionen för socialförsäkringen.

Flood, L., Hansen, J., and Wahlberg, R. (2004), Household Labor Supply and Welfare Participation in Sweden, *Journal of Human Resources 39*, 1008–1032.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.

Francesconi, M., Rainer, H., and van der Klaauw, W. (2009), The Effects of In-Work Benefit Reform in Britain on Couples: Theory and Evidence, *Economic Journal 119*(535), 66–100.

Gelber, A. M. (2014), Taxation and the Earnings of Husbands and Wives: Evidence from Sweden, *Review of Economics and Statistics 96*, 287–305.

Gelber, A. M. and Mitchell, J. W. (2012), Taxes and Time Allocation: Evidence from Single Women and Men, *Review of Economic Studies 79*, 863–897.

Gelber, A. M., Jones, D., Sacks, D. W., and Song, J. (2018), Using Non-Linear Budget Sets to Estimate Extensive Margin Responses: Method and Evidence from the Social Security Earnings Test, NBER Working Paper No. 23362.

Havnes, T. and Mogstad, M. (2011), Money for Nothing? Universal Child Care and Maternal Employment, *Journal of Public Economics 95*, 1455–1465.

Immervoll, H., Kleven, H. J., Kreiner, C. T., and Saez, E. (2007), Welfare Reform in European Countries: A Microsimulation Analysis, *Economic Journal 117*(516), 1–44.

Immervoll, H., Kleven, H. J., Kreiner, C. T., and Verdelin, N. (2011), Optimal Tax and Transfer Programs for Couples with Extensive Labor Supply Responses, *Journal of Public Economics 95*, 1485–1500.

Jacob, B. A. and Ludwig, J. (2012), The Effects of Housing Assistance on Labor Supply: Evidence from a Voucher Lottery, *American Economic Review 102* (1), 272–304.

Kleven, H., Landais, C., and Søgaard, J. E. (2019), Children and Gender Inequality: Evidence from Denmark, *American Economic Journal: Applied Economics 11*, 181–209.

Kosonen, T. (2014), To Work or not to Work? The Effect of Childcare Subsidies on the Labour Supply of Parents, *B.E. Journal of Economic Analysis and Policy 14*, 32.

Lundin, D., Mörk, E., and Öckert, B. (2008), How Far Can Reduced Childcare Prices Push Female Labour Supply?, *Labour Economics 15*, 647–659.

Meyer, B. D. and Rosenbaum, D. T. (2001), Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers, *Quarterly Journal of Economics 116*, 1063–1114.

Murphy, K. M. and Topel, R. H. (1985), Estimation and Inference in Two-Step Econometric Models, *Journal of Business and Economic Statistics 3*, 370–379.

Selin, H. (2014), The Rise in Female Employment and the Role of Tax Incentives: An Empirical Analysis of the Swedish Individual Tax Reform of 1971, *International Tax and Public Finance 21*, 894–922.

First version submitted November 2017;
final version received February 2020.

© 2020 The Authors. *The Scandinavian Journal of Economics* published by John Wiley & Sons Ltd on behalf of Föreningen för utgivande av the *SJE*/The editors of *The Scandinavian Journal of Economics*.