WORKING PAPER 2022:9

Filling in the blanks

How does information about the Swedish EITC affect labour supply?

Pär Nyman
Linuz Aggeborn
Rafael Ahlskog

IFAU
INSTITUTE FOR EVALUATION OF LABOUR MARKET AND EDUCATION POLICY

The Institute for Evaluation of Labour Market and Education Policy (IFAU) is a research institute under the Swedish Ministry of Employment, situated in Uppsala.

IFAU's objective is to promote, support and carry out scientific evaluations. The assignment includes: the effects of labour market and educational policies, studies of the functioning of the labour market and the labour market effects of social insurance policies. IFAU shall also disseminate its results so that they become accessible to different interested parties in Sweden and abroad.

Papers published in the Working Paper Series should, according to the IFAU policy, have been discussed at seminars held at IFAU and at least one other academic forum, and have been read by one external and one internal referee. They need not, however, have undergone the standard scrutiny for publication in a scientific journal. The purpose of the Working Paper Series is to provide a factual basis for public policy and the public policy discussion.

More information about IFAU and the institute's publications can be found on the website www.ifau.se

ISSN 1651-1166

Filling in the Blanks: How Does Information about the Swedish EITC Affect Labour Supply?[a]

Pär Nyman Linuz Aggeborn Rafael Ahlskog

April 19, 2022

Abstract

Information plays a key role in economics. According to the benchmark neoclassical model, agents require information in order to optimize their choices. Information, however, is sometimes incomplete or asymmetric in the real world. In this paper, we investigate the role of information for the labour–leisure choice. We conduct an information experiment wherein we distribute a leaflet about the Swedish Earned Income Tax Credit (EITC), and then study the effects with registry data. More specifically, we focus on the household decision to allocate between labour income and parental leave payments. The EITC, it bears noting, applies to the former but not to the latter. The construction of the Swedish EITC generates a strong economic incentive to have some labour income during a calendar year, because it is tailored to benefit low labour income earners more in relative terms. Yet, despite the substantial economic incentives involved, and despite the flexibility with which a person may earn labour income, we find that providing information about the features of the EITC has zero impact on labour supply.

Keywords: Labour Supply, Information, Experiment, Earned Income Tax Credit, Sweden
JEL-codes: J18, I21, I26

[a]The empirical analysis in this paper takes its point of departure from the registered pre-analysis plan: Nyman (2020) "An information experiment about the Swedish EITC." AEA RCT Registry. 10 September. https://doi.org/10.1257/rct.6430-1.1. The research project has been approved by the Regional Ethical Review Board in Uppsala (Dnr 2017/026 with supplement Dnr 2017/026/1). All authors have contributed to the planning of the research project, the analysis of the data, and the writing of the paper. Pär Nyman is the first author, because he came up with the original research idea, and he was listed as the contact person in the letter sent to the treatment group. We gratefully acknowledge financial support from IFAU. The authors are affiliated with the Department of Government at Uppsala University (PN, LA and RA), the Uppsala Center for Labour Studies (LA and RA), and the Uppsala Center for Fiscal Studies (LA). Pär Nyman, the corresponding author, has the following e-mail address: par.nyman@statsvet.uu.se. We thank Håkan Selin, Helge Bennmarker, and Arnaldur Stefánsson for their constructive comments. A special thanks to Arnaldur Stefánsson for feedback and discussions regarding the theoretical framework. We also thank seminar participants at IFAU and at SOFI at Stockholm University.

Contents

1 Introduction

Benchmark neoclassical economic models rely on the assumption of perfect information about the incentives faced by individuals, households, and firms. In order to maximize their utility (or profit), economic agents need to be aware of all existing budget constraints. This is the foundation for the standard economic view that agents are fully rational, meaning they act according to their objective function and take into account the economic limitations they face.

In many situations, this assumption is a reasonable one. However, when complex institutional systems are in question – in particular, where several systems interface – the assumption is rather more strained. The tax system is an example. Most agents probably have general information about tax levels, and act accordingly; but it is less likely they are aware in any detailed way of all of the aspects and nuances of the tax code. Acquiring information involves costs, and information is likely to be distributed unevenly between agents. Firms, for example, are likely to have a better understanding of the tax system than individuals do, because they need to engage with the tax code on a more regular basis. Indeed, large firms have accounting and legal departments for this very reason. Individuals, by contrast, do not generally have the resources to study the tax code in such detail. Recent empirical papers using American data have demonstrated that individuals do not have full knowledge about the tax system (Feldman et al. 2016; Gideon 2017; Ballard and Gupta 2018; Stantcheva 2021). In particular, they have problems understanding non-linearities in the system: they rely instead on some average tax rate, and mentally "linearize" their own tax rate from this (De Bartolome 1995; Rees-Jones and Taubinsky 2020). Such a lack of full information clearly has the potential to affect economic behaviour.

To date, the number of studies on the effects on labour supply of information about precise tax incentives has been quite small. This general lack of studies is likely a consequence of several factors. First, for information effects to be observable, the incentives need to be fairly large, and the adjustment time has to be fairly quick. Second, the causal mechanism must concern micro-level behaviour, as opposed to things like collective wage bargaining, for the identifying assumptions to be valid. Finally, the incentives need to be largely unknown prior to the informational intervention; they cannot be a major topic of public discourse. By their very nature, these things rarely coincide. Yet, we believe the institutional setting in the area

considered here provides a valuable case where these three features are present at the same time.

The particular institutional design of the Swedish Earned Income Tax Credit (EITC) as it interfaces with the parental insurance system provides a unique opportunity to evaluate information effects. In this paper, we therefore present results from an experiment in which we have randomized information about the EITC to parents in Sweden, in order to study its effects on labour supply. Simply put, the idea of the EITC is to incentivize people to prioritize work over other types of income, which do not yield any EITC. The non-linear elements of the EITC, in combination with the annualized nature of the tax system, means that these incentives are particularly large for people who are temporarily on leave from work. For example, due to the EITC, most of the mothers in our study have their marginal tax rate reduced by more than 10 percentage points, but only for incomes declared during 2018. Interestingly, the EITC is not generally well-known in Sweden. For example, while there were major discussions about it in connection with the 2006 election, a mere 40 per cent of respondents said in a survey three years later that they had heard of it (Riksrevisionen 2009, p. 10). In particular, we argue, the precise incentives generated by this complex system are opaque even to otherwise well-informed individuals when they decide on the distribution between labour income and parental leave income within their household.[1]

One central feature of the Swedish tax system is that all income taxes are personal; there is no joint taxation of mothers and fathers. Furthermore, the Swedish EITC differs from the American one in terms of its take-up. The EITC in the U.S. needs to be claimed, whereas it is automatically distributed in Sweden by the Tax Agency, on an annual basis. Hence, filing for the EITC in Sweden involves no costs, in line with the discussion in Benzarti (2020). As we discuss in greater detail later in this paper, household members in Sweden enjoy considerable flexibility in their labour-vs-leisure decision. They have several different ways of obtaining the EITC when one parent is on parental leave. Taken all together, these institutional features provide us with a most-likely case. Given the large but obscure economic incentives the EITC creates, information about it is likely to have an effect on labour supply.

The information letters we sent contained accurate but simplified information about the

[1] Albeit of little evidentiary value, it might be worth noting that, anecdotally speaking, even a seminar of Swedish labour economists were surprised to learn of the particularly obscure incentives we study during the presentation of this paper.

EITC, with concrete examples of how parents' taxes look depending on what labour–leisure choices they make. Then, using registry data, we have studied the effects of information on labour supply, on both the extensive and the intensive margin. The empirical results show that providing information about the Swedish EITC does not affect parents' labour supply, despite the clear economic incentives the EITC generates to earn some labour income during a calendar year, and despite the many different ways there are to earn labour income. The coefficients are precisely estimated, meaning that the upper bounds of the confidence intervals rule out any effects that could be considered economically relevant. In our discussion at the end of this paper, we discuss possible explanations for this zero effect.

Our study provides important evidence on the interaction between information, taxation, and labour supply; and more specifically on the household decision between labour income and non-labour income. We focus on the Swedish EITC schedule, mainly because of its helpful institutional features and the availability of registry data. It is also important, however, to evaluate one of the largest tax cuts in Swedish history, not least because it may affect the structure and functioning of the labour market. We believe our study contributes to such an effort as well.

2 Previous research

In an early review, Currie (2006) argues that more experiments are needed to investigate how information affects economic behaviour. But while there is a well-established theoretical understanding that information frictions *can* be very important for explaining such things as differences between macro- and micro-level elasticities (see, e.g., Chetty 2012), or even serve to decrease the distortionary effects of taxes (Goldin 2015), the empirical literature on the question is still fairly scant.[2]

There are four studies, however, which are of particular interest for this paper, since they also concern the effects of information about the EITC. First off, Chetty and Saez (2013) performed an information experiment in the United States in collaboration with tax preparers, in which

[2]It bears noting that, when we talk about information, we are referring to information about an actual fixed incentive structure. We are aware of the substantial game-theoretical literature on incomplete information, which typically investigates deviations from the assumption of complete information regarding the utility functions of other agents. Here, however, there is no game-theoretical element; rather, there are agents interacting with a fixed payoff schedule – a payoff schedule that may or may not be known.

half of a sample of 43,000 tax filers received a short presentation about their position on the tax schedule, and how much it would pay to work more. In this study, no effect on average earnings was found.[3]

Secondly, Bhargava and Manoli (2015) conducted an information experiment jointly with the IRS in the United States, in which they studied the take-up of EITC. The lack of information about the EITC, they contend, is one reason why the programme has a low take-up. The authors randomized information leaflets and worksheets by mail. The latter could be used to assess whether a person is eligible for EITC. The information material varied in its degree of complexity. The authors found that information on the EITC raised the claim rate.[4] One difference between our study on the one hand, and those of Chetty and Saez (2013) and Bhargava and Manoli (2015) on the other, is that the EITC in the United States needs to be claimed. In Sweden, by contrast, the EITC is automatically distributed to all eligible individuals. We argue, however, that public knowledge is far from perfect about the incentives the EITC generates.

Thirdly, Kostol and Myhre (2021) present (a) results on optimization frictions using notches in the EITC schedule in Norway, and (b) quasi-experimental evidence from an information letter about a new kink in that schedule. Finding substantial effects from the latter, they conclude that information frictions that can be ameliorated fairly easily account for at least 30 per cent of the optimization frictions.

Fourthly, Chetty et al. (2013) investigate in a related paper how the effect on wage earnings of changes in EITC eligibility due to childbirth differs between low- and high-knowledge areas. They find that individuals in high-knowledge areas change their labour supply in response, whereas their counterparts in low-knowledge areas only do so to a lesser extent. In this case the information was not randomized, but rather inferred by the degree of sharp bunching around existing EITC thresholds.

This study is also related to the research on the effects of the EITC on parents in general, and on mothers in particular. For example, Eissa and Liebman (1996) find that the Tax Reform Act of 1986 in the United States, which expanded the EITC, resulted in an increase in labour force participation among single mothers, but not on the intensive margin. Similarly, Meyer

[3]See also Jones (2010) for an information experiment regarding advanced participation in the EITC in the U.S.
[4]See also Manoli and Turner (2014) for an analysis on take-up of EITC due to reminders from the IRS.

and Rosenbaum (2001) argue that expansions in the American EITC between 1984 and 1994 increased labour supply among mothers, but on both the intensive and the extensive margin. On the other hand, Kleven (2019) finds in a re-analysis that these effects are generally not robust, and that information friction is the most natural explanation. Bastian (2020) concludes that the 1975 version of the EITC in the U.S. substantially increased employment among mothers.

In addition to relating to the literature on the EITC, this paper connects up with the literature on information experiments. Engström et al. (2019) conduct an experiment where they randomize information about the housing allowance to Swedish pensioners, the object being to ascertain if the effect of so doing is to increase the take-up of housing allowance. Their empirical design, like ours, involves randomizing leaflets with information about the housing allowance, and they find positive treatment effects. Chetty et al. (2009), on the other hand, investigate the effect on demand of different ways of posting information about tax-inclusive prices in a supermarket. Guthmuller et al. (2014) ask, in a study conducted in France, if the take-up of additional healthcare insurance for poor households is higher if an individual is randomly selected to receive a letter with a higher voucher amount than if he/she just receives a standard information letter. They find a small treatment effect. They also consider a second treatment, where individuals are also invited to an information meeting regarding the additional healthcare insurance. In the case of this subgroup, surprisingly, the treatment effect is negative. In a lab experiment, Abeler and Jäger (2015) demonstrate that individuals who are furnished with information about the tax system that varies in degree of complexity – but which otherwise includes the same economic incentives – underreact to information on economic incentives. Liebman and Luttmer (2015), in a study of the effect of information about certain social security provisions among older workers, find that the informational intervention led to an increase in labour supply, primarily among women. Karlan et al. (2016) focus on savings, and seek to evaluate the effect of a randomized information treatment by banks in Bolivia, Peru, and the Philippines. The treatment comes in a variety of forms, such as home visits, radio commercials, and text messaging. Overall, the authors conclude, information remainders about savings increase the probability that a person will begin saving; and some evidence suggests the amount of money saved increases. Matikka and Paukkeri (2016) conclude that, in Finland, mailing information letters together with an application form increases the take-up and results in earlier applications from pensioners with

low incomes. Stantcheva (2021) examines a randomized sample of survey participants who watch information videos that emphasize various aspects of the American tax system. These videos present arguments regarding efficiency, redistribution, and standard economic theory in connection with income and estate taxes. Stantcheva finds positive treatment effects from the redistribution and economic theory arguments, but no treatment effects from the efficiency argument.

3 Institutional background

After the electoral success of the centre-right "Alliance" in the 2006 election, one of the first orders of business for the new government was to carry out a major overhaul of the tax system. As a result, in 2007, the Swedish version of the Earned Income Tax Credit was first implemented. This reform was amended and extended in 2008, 2009, 2010, 2014, 2016, 2019, and 2022.

When the EITC was first introduced, it was one of the largest reductions in the tax on labour income in modern Swedish history. It was widely debated in the 2006 election campaign. The central campaign message from the centre-right Alliance was that "it should pay to work". One of the major points of contention between the centre-right and centre-left blocs was the gap which the EITC introduced between the tax on labour income and that on other types of income.[5]

The Swedish EITC involves a simple idea: to reduce taxes on labour income, primarily for low- and middle-income wage-earners. Actually implementing it, however, is fairly complex. The deduction is calculated as the municipal tax rate multiplied by the difference between a specified "special amount" (*särskilt belopp*) and the "basic deduction" (*grundavdrag*). For the 2018 tax year, the special amount was a piecewise linear function of an individual's earned income, with kink points at 133,770, 367,640, and 616,070 SEK. The basic deduction, on the other hand, is a function of the *total* non-capital income (with some special provisions based on year of birth and on whether one has spent part of the year out of the country). Both functions are defined in constant prices, with adjustments being made on an annual basis. The 2018 schedule is illustrated graphically in Figure 1. Needless to say, it is hard for most people to get an accurate picture of the details around this tax policy – especially about the precise way in which their

[5]The centre-left Social Democrats, for instance, insisted on referring to the proposal as the "pensioner tax" due to this gap.

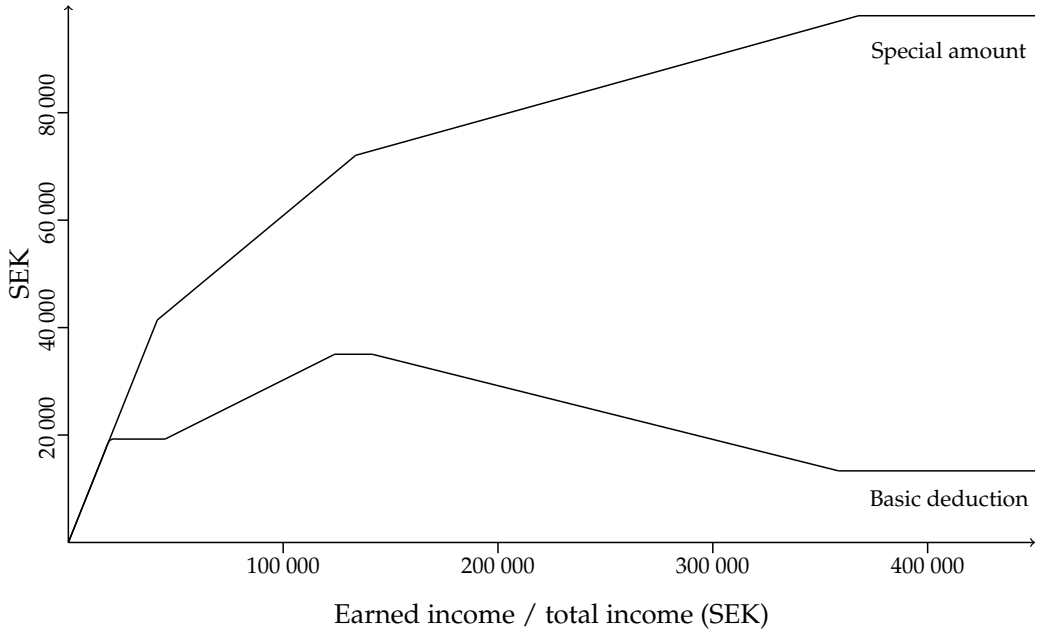

Note: The Earned Income Tax schedule for the 2018 tax year. The EITC is defined as the municipal tax rate multiplied by the difference between the "special amount" and the "basic deduction".

Figure 1: EITC schedule, 2018

after-tax income would be affected by marginal changes in labour supply.

It should be noted that the kink points in the upper curve (the special amount) are much more important than the ones in the lower curve (the basic deduction), whenever labour and non-labour income act as substitutes. This is likely to be the case when time out of work is compensated with reimbursements from the social security system, like unemployment benefits or paid parental leave. In Sweden, large parts of the social insurance system cover almost 80 percent of the income loss from not working. Hence, if a person decides to get back to work a month earlier, it may increase the person's earned income with 25,000 SEK, but the total income with only 5,000 SEK, effectively making any changes in the basic deduction much smaller than the ones in the special amount.

The most important kink point in the upper curve is the one found at approximately 135,000 SEK. After this point, the slope of the special amount is reduced by two thirds, which means that the marginal EITC is much smaller. Since most mothers in our study have a labour income of less than 135,000 SEK (see Figure 3), the incentives created by this kink point are relevant to the

majority of couples who receive our letter. More specifically, it means that the EITC strengthens the economic incentives for most mothers to increase their labour supply during the year after giving birth. The kink point at approximately 40,000 SEK also corresponds to a major change in the slope of the special amount. However, the EITC does not kick in until the labour income exceeds the basic deduction, which means that the income interval that corresponds to the maximum marginal EITC (identical with the municipal tax rate, thereby providing a marginal tax rate of zero) can be very narrow.

Both arguments made in the two preceding paragraphs are illustrated in Figure 2. These four graphs display the marginal EITC (a and b) and the marginal tax rate (c and d) for an incremental change in labour income. In the graphs to the left, other incomes are assumed to be unaffected, while in the graphs to the right, it is assumed that a marginal increase in earned income results in a fall in other sources of income that amounts to 80 per cent of the increase in labour income. As was described earlier, in situations when labour and non-labour income act as substitutes, tax rules based on total income are not as important as the kink points in the special amount. This can be seen by comparing the graphs on the left, with the ones on the right. By comparing the three lines in each graph, and focusing on earned incomes below 40,000 SEK, we can also see how the marginal incentives differ substantially depending on the level of non-labour income.

The very low tax rates in the bottom-right graph may require an explanation. The marginal tax rate is here defined as the change in taxes resulting from an increase in labour income with 100 SEK and a simultaneous fall in non-labour income with 80 SEK. For earned incomes just below 40,000 SEK, the marginal tax rate on labour income is equal to zero. This means that if an increase in labour income coincides with *any* decrease in other sources of income, the resulting tax change will be negative, which is why we find a negative marginal tax rate in this graph. And more generally, if an increase in earned income with 100 SEK only makes the total income go up with 20 SEK, then the resulting tax change will only be a fifth (ignoring the EITC) of what it would be without the fall in non-labour income. This is why the marginal tax rate shown in this graph is much lower than in the graph to the left.

The EITC is calculated based on annual labour income (cutoff per calendar year). In combination with its non-linear structure, this creates strong incentives for individuals to increase their labour supply during calendar years when they are partly on leave from work. For example,

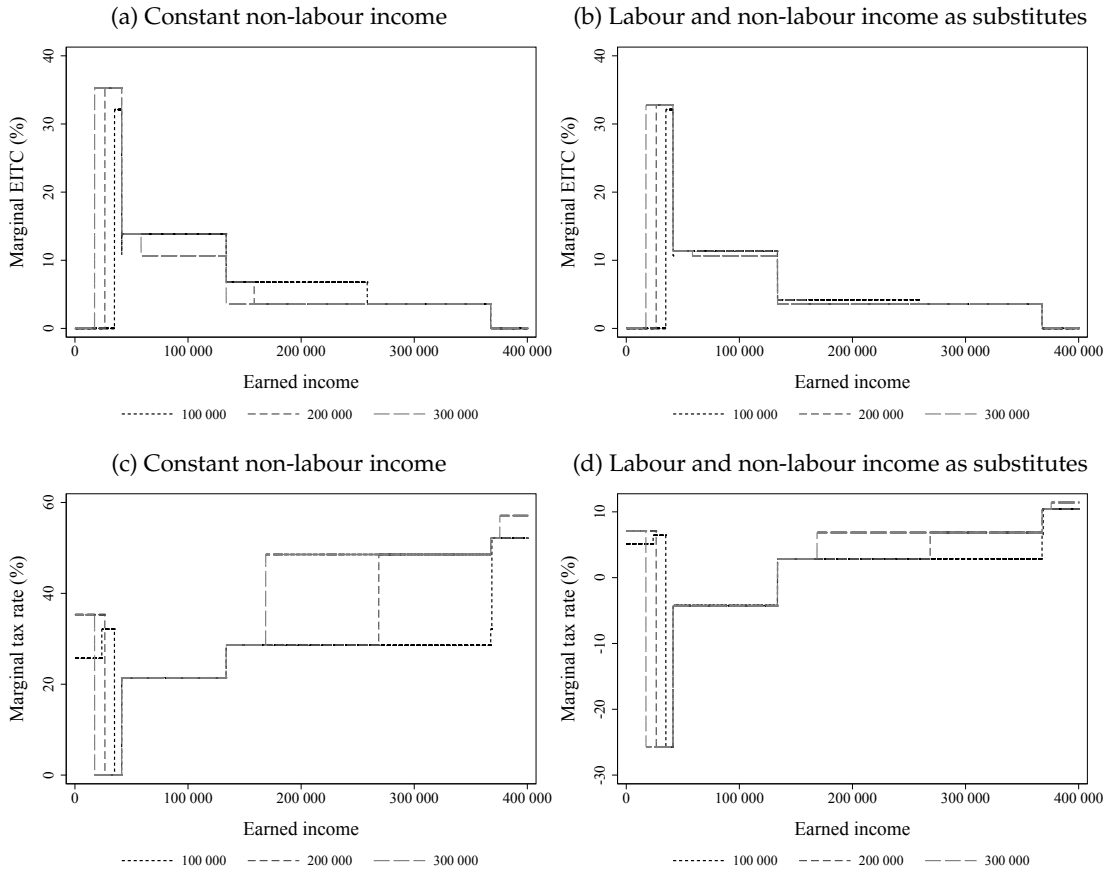

Figure 2: Marginal EITC (top) and marginal tax rate (bottom) in different settings

Note: The graphs display the marginal EITC (top) or the marginal tax rate (bottom), defined as the change in EITC/taxes that results from an incremental change in the earned income. In the left panels, other incomes are assumed to be unaffected by the increase in labour income; in the right panels, an increase in earned income by 100 SEK results in a fall in other sources of income by 80 SEK (since labour income and parental leave income are substitutes, for example). In all graphs, the different lines correspond to a baseline non-labour income of 100,000, 200,000 and 300,000 SEK, respectively.

in our study, most mothers have their marginal tax rate reduced by more than 10 percentage points due to the EITC. This also means that there are fairly strong incentives for people with comparable wage rates and shared economic arrangements to distribute labour income evenly between them in a given calendar year. If the higher earner works slightly less and the lower earner slightly more, the net effect is not zero but positive. This effect is amplified for the distribution of parental leave payments, which are not subject to the EITC. For new parents who both earn 35,000 Swedish crowns (SEK) per month, the joint tax deduction will be about 16,000 SEK larger (approximately 1,800 USD) if they split their time with labour income evenly (over a

calendar year) than it will be if one of them takes parental leave pay over the whole period.

Taxpayers in Sweden are informed about the EITC in their yearly tax declaration. The Tax Agency details the amount of EITC that a person has received. It also stipulates the income received from the employer (labour income) and from the Swedish Social Insurance Agency, which pays out parental leave benefits (non-labour income). A separate information box in the declaration explains that the total amount of labour income provides the base for the amount of EITC a person receives. A link to the website of the Tax Agency is also printed in the declaration, with additional information about the EITC, which in turn states that parental leave benefits do not result in any EITC.[6] This shows that Swedish policy-makers want taxpayers to be aware of the economic incentives the EITC generates.

Given the magnitude of the reform, the central place it had in political discourse at the time of its implementation, and the information contained in the yearly tax declarations, it *should* be relatively well-known. Nevertheless, a survey conducted in 2009 indicated that no more than 40 per cent of the Swedish population were aware of it (Riksrevisionen 2009, p. 10).[7] Even less well-known, arguably, is the way in which the tax reduction affects some of the economic incentives facing new parents. These incentives depend on the non-linear structure of the EITC, in combination with the annualized nature of the tax system. And we know from previous research that non-linearities in the tax system have a negative impact on optimization behaviour (Abeler and Jäger 2015).

It is important to point out that this does not mean the deduction is not claimed. The Swedish tax system is highly automatized, and the deduction is made transparently whether one claims it or not. This characteristic of the Swedish tax system – its "invisible" nature – reduces the incentive to learn the tax code, and it may be part of the reason why people are often unaware of how the tax system functions. Declaring one's taxes in Sweden is simple, compared to how it is for example in the United States. Most types of income are preprinted on the tax declaration;

[6] Website (in Swedish): https://www.skatteverket.se/jobbskatteavdrag. The Tax Agency has provided us with some aggregated data regarding their internet traffic. For 2020 and 2021, approximately 21,000 and 25,000 visits were made for this part of the website for each year respectively. In total, the website for the Tax Agency had around 22,000,000 and 25,000,000 visits for 2020 and 2021. The total Swedish population is around 10 million of which little more than 8 million declare their taxes annually. In conclusion, only a small fraction take part of the information regarding EITC.

[7] Andersson and Antelius (2010) also argues that individuals need to be aware of the EITC in Sweden if the incentives are going to play a role for economic behaviour.

and if you do not make any special deductions, you can simply confirm the tax declaration with a text message or in the Tax Agency's app on your phone. As explained above, the amount of EITC that a taxpayer receives is printed separately in the tax declaration. Since they do not have to engage with the actual numbers, however, taxpayers probably do not spend much time examining the amount.

Swedish parental insurance, which is often considered generous in an international comparison, provides a total of 480 days of paid parental leave for each birth (with an additional 180 days for twin births). For the first 390 of these days, approximately 80 per cent of previous labour income is provided (with a top-up to 90 per cent offered in most collective contracts). For the remaining days, 180 SEK per day is paid out. 90 days are reserved for each parent,[8] while the remaining days can be distributed freely between the two parents. Moreover, in addition to these 480 days of paid parental leave, parents have the right to unpaid leave, which many use to prolong their parental leave. For example, parents may not use paid parental leave days for all seven days of the week. They can also divide a day between parental leave income and labour income by working part time. Furthermore, parents earn vacation days while on parental leave, which they can use to extend their time off from work further. As for employment protection, all workers in Sweden (including those with temporary jobs) have the right to be away from work completely until their child is 18 months old, regardless of whether they are on paid or unpaid leave. Parents also have the right to combine part-time work with some parental leave income until the child is 8 years old. In sum, if one wants to stay home for a long time, one can. Historically, parental leave has been taken overwhelmingly by mothers, although a trend toward a larger take-up by fathers has been observed in recent decades. Parental leave payments from the Swedish Social Insurance Agency do not count as earned income, so persons receiving them do not qualify for the EITC.[9]

All children over 12 months old are entitled to a place at a preschool in the same municipality.

[8] Since the majority of parental leave days are still taken by the mother, these 90 reserved days are often referred to as "daddy days", although they are not exclusive to the father. For a study of what happened when daddy days were first introduced in 1995, see Avdic and Karimi (2018). These authors find that the divorce rate increased (or was rather re-timed) due to the daddy days, and that this effect was driven by an overall fall in the family income level due to changes in parental leave take-up between mothers and fathers. However, this reform was implemented long before the EITC was added to the Swedish income tax schedule.

[9] In a study of benefit levels in the Swedish parental leave insurance system, Ginja et al. (2020) focus on the speed premium. They conclude among other things that mothers' labour supply falls when benefit levels rise. See also Karimi et al. (2012) for an analysis on household labour supply from reforms in the Swedish parental leave system.

Most parents, however, stay home with their children for a longer period, the median age for preschool enrolment being approximately 18 months (Skolverket 2017, p. 192). Parents apply to the preschool where they want to place their child. The chances of getting their preferred choice are greater in August, and January seems a natural time to start, so many parents end their parental leave around those months (the child can usually begin one or two months later than the date applied for, so this can be seen as a soft constraint).

The most obvious way that parents can increase their labour income, and thus the size of their EITC, is of course to shorten the duration of their parental leave and to get back to work earlier – either by letting their partner extend his or her parental leave, or by enrolling their child in childcare at a younger age. However, the need to re-negotiate with their partner and to inform their employer about new plans may discourage individuals from increasing their time at work. For purposes of this experiment, it is important to note that there are ways of increasing earned income and the EITC that involve smaller frictions. And since the marginal EITC is large at low incomes, one's behavioural response need not be substantial for one to obtain a substantial tax reduction. First, the parent on parental leave may choose to work a couple of hours per day. Parents have great flexibility in this regard. In theory, parents may choose to work one hour of a day (but not necessarily more days), while being on parental leave for the rest of that day. Given that some professions make it possible to work from home, parents may for example work one hour while the baby is sleeping every second Friday during the parental leave period. Second, parents may start work a couple of weeks earlier and get help from grandparents before the child starts preschool. Third, parents can increase their hours of work by working overtime, or by working full-time instead of part-time once they return to work. Fourth, some may have the option of looking for work (if they were unemployed), of switching jobs, or of seeking promotion to a position with a higher salary. Fifth, those running their own business may withdraw a higher salary (instead of letting the money stay within the firm) during the calendar year.[10] Lastly, parents can use their vacation days strategically to increase their earned income, while saving days of paid parental leave. Vacation-day income, namely, is counted as labour income,

[10] Bastani and Selin (2014) find that the self-employed in Sweden bunch around the kink points for central government taxes. The foremost reason for this, they conclude, lies in reporting behaviour – not in actual changes in labour supply. Engström et al. (2018) argue that high income earners should deduct more in their tax declarations, since they face a higher marginal tax rate. They conclude, however that wage-earners do not in fact deduct around the central government tax kink.

and vacation days are not counted in the same way as days of parental leave. Making use of this option can in theory increase the number of days at home.[11] All in all, there are several ways parents may plan the distribution of labour income and parental leave income in order to increase their EITC. This flexibility is helpful for purposes of this paper, because it allows a multitude of responses to the information treatment depending on the parents' preferences regarding work and parental leave. In our empirical analysis below, we investigate labour supply on both the intensive and the extensive margin, in order to pinpoint the full response to our experiment.

Surprisingly little is known about the effects of the Swedish EITC – perhaps owing to the inherent difficulty of evaluating a policy of this character, which is implemented nationwide at the same time. Edmark et al. (2016), for example, note that exploiting differences in tax rates between municipalities in order to ascertain the EITC's impact on employment yields too little identifying variation; the policy may well, therefore, be impossible to evaluate. Bennmarker et al. (2014) look at the effect on wages of changes in the net replacement rate, and find elasticities in the range of 0.2–0.4; thus, the EITC may promote wage moderation.

4 Theoretical framework and hypotheses

In this section, we present and discuss a very simple theoretical framework for information, the EITC, and labour supply. Our overall purpose is to provide a theoretical foundation for the hypotheses we later present. We chose these hypotheses because of their theoretical underpinnings, but also because of the institutional features discussed in the last section.

The standard model of labour supply views the decision to allocate time between work and leisure as an individual optimization problem under a given budget constraint. In our case, the optimization problem takes place within the household between two individuals, whom we denote as the mother, m, and the father, f. During the first year after the birth of the child, at least one of these two persons stays home to take care of the infant. Our theoretical framework assumes a household utility which equals the sum of the mother's and father's utility functions

[11] A person on vacation need not use up vacation days in order to be on leave during public holidays or weekends; the opposite is true, however, for a person on parental leave. Placing vacation days around public holidays and weekends therefore yields extra days of paid leave which count as labour income for tax purposes.

(Browning et al. 2006). The framework presented below follows the theoretical part in Fortin and Lacroix (1997) closely, although our framework is much more simplistic.[12]

Let us assume both the mother and the father gain utility from consumption, c, and from leisure, l, where the latter equals spending time with the child. The utility function for each parent thus equals:

$$U_{m,f} = f(c_{m,f}, l_{m,f})$$

Each household, h, is assumed to have a joint utility where the mother and the father have equal weights, such that $U_h = U_m + U_f$. The household maximizes this utility function under the following budget constraint:

$$c_m + c_f = \sum_{j \in \{m,f\}} (h_j w)(1-\theta) + (bl_j)(1-\tau)$$

h equals hours worked, and l denotes leave hours. w is the hourly wage rate, and b is the hourly benefit level from parental leave. τ is a linear tax rate which is paid on all non-labour income. In our study this equals parental leave benefit, for which you pay the flat municipal tax rate in Sweden. θ is the marginal tax rate, including EITC, which is paid on labour income.

Up until now we have assumed perfect information, where both the mother and the father perceive the tax system perfectly with regard to the EITC. However, we argue that people do not in fact fully comprehend the tax system when it comes to the non-linear aspects of the EITC. Instead of observing θ perfectly, they observe this factor with a bias. Let us denote $\phi = \theta + \varepsilon$, where $E[\varepsilon] > 0$. The perceived budget constraint for the household may then be written as:

[12] There has been an intense discussion in the theoretical literature regarding labour supply on the household level – more precisely, whether households should be modelled with a unitary or a collective utility function (see Chiappori (1992) and Donni and Chiappori (2011) for reviews and discussions). We make no contribution to this theoretical discussion. Our purpose is instead to provide a simple point of departure for our empirical hypotheses.

$$c_m + c_f = \sum_{j \in \{m,f\}} (h_j w)(1-\phi) + (bl_j)(1-\tau)$$

On the household level, the mother and the father will have spent $L = l_m + l_f$ hours of parental leave in total during the calendar year. The time at work will equal $H = h_j + l_j$ on the household level.

A central assumption that underlies our theoretical expectations is that, without our letter, people in our study would on average *overestimate* their marginal tax rate on labour income (represented by the assumption that $E[\varepsilon] > 0$). Why do we believe that to be the case? When faced with complex tax schedules, individuals tend to form their beliefs about their relevant incentives using information that is easily available to them. For example, research carried out in the US has shown that people often use the average tax rate as if it were their marginal tax rate (De Bartolome 1995; Rees-Jones and Taubinsky 2020). In the Swedish setting, an alternative heuristic would be the municipal tax rate, which is a flat rate deducted from all kinds of income. For most people with a job, the municipal tax rate is also a reasonable approximation of their marginal tax rate. For our argument here, it does not matter much what heuristics people use, as long as they are based on previous experiences with the tax system or approximations applicable to a typical household. The reason for this is that most women in this study are located in the income brackets with the lowest possible marginal tax rates. Therefore, these tax rates will be much lower than the average and marginal tax rates that they usually face as well as the well-known municipal tax rate.

Pre-registered hypotheses

Let us now move on from the stylized labour economics framework to the hypotheses that we empirically test. These hypotheses are founded on the general theoretical framework of labour economics, but we also take into account the precise features of the Swedish EITC in our experiment, as well as the income distribution among parents. We presented all of our hypotheses in our pre-analysis plan, but we present them here in a different order (in order to connect more smoothly with the theoretical framework above).

As additional background to these hypotheses, Figure 3 shows the distribution of earned income in our sample of parents, together with our calculation of the marginal EITC. Since the income refers to the year 2018, and the tax declaration is based on the calendar year, this graph can be used to infer what incentives mothers and fathers face with regard to the EITC. As we would expect from the unequal distribution of parental leave, the distribution of earned income is heavily skewed toward zero for mothers. As a result, a substantial proportion of mothers (potentially a majority) will get more than 13 per cent of their additional wage back if they decide to increase their amount of labour income. Fathers, by contrast, face a marginal EITC of less than 5 per cent in most cases, and often of 0 per cent.

The Swedish EITC was implemented for the stated purpose of increasing employment, and it does increase incentives to work substantially, in particular for people who are able to work 3–4 months or more.[13] In our stylized framework above, we concluded that information about the EITC generates an incentive to enter the labour market, which leads to our first hypothesis.

Hypothesis 1: Receiving the letter increases labour supply on the extensive margin.

While previous empirical research suggests that information is crucial for people's ability to adjust their labour supply in response to the tax system, this primarily applies to responses on the intensive margin (Chetty 2012; Chetty et al. 2013). However, formulating our theoretical expectations regarding the intensive margin is a bit more complicated, because the information letter could in principle have different effects, depending on the relative size of the income and substitution effects. We utilized a general utility function in the theoretical reasoning above, and the effect on the intensive margin is ambiguous in this case. Yet we have two arguments for believing that the substitution effect dominates the income effect. We can therefore expect a positive effect on the intensive margin.

First, unlike EITCs in many other countries, the Swedish EITC increases with income up to income levels that in fact are slightly higher than the median income. It was designed this way in order to increase labour supply on the intensive margin as well (Bill 2006/07:1, p. 33), and several micro-simulations have concluded that this is what the tax reform actually does

[13] For smaller amounts of work, the EITC usually does not make much difference. The reason for this is that it does not kick in until the earned income exceeds the basic deduction, which is usually at somewhere between 20,000 and 35,000 SEK for people out of work.

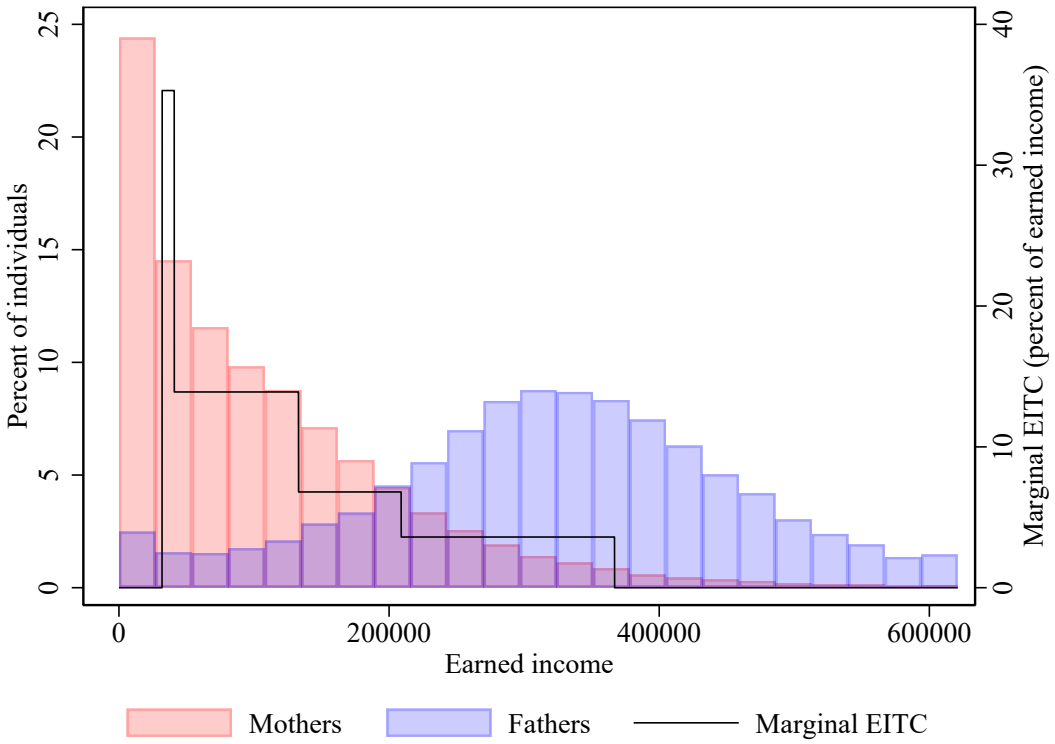

Note: The histogram shows the distribution of earned income in 2018 in our sample of employed mothers (red) and fathers (blue) of newborn children. Overlapping bars are displayed in purple. The black line shows how much the EITC increases as a percentage of additional earned income, calculated under the same assumptions we used in the information letter (i.e., an average municipality tax rate and a non-labour income of 150,000 Swedish SEK). Note that the exact location of the first (and fourth) kink – where the marginal EITC goes from 0 to 35 per cent – is dependent on the amount of non-labour income. We therefore chose to emphasize the kink at approximately 135,000 SEK in the letter.

Figure 3: Sample income distribution by sex

(Riksrevisionen 2009). If this is true for the Swedish population, it is arguably so for our sample as well, because the lower incomes therein imply that the marginal EITC is larger on average than for the population as a whole (i.e., the substitution effect should be larger), and that the average EITC is smaller (i.e., the income effect should be smaller).

Second, we argue that an information treatment like ours is much more likely to result in substitution effects than in income effects. One reason for this is that people are relatively well-informed about how much they earn. So, while we do believe we provide new information about the marginal EITC, we find it unlikely that the letter will make recipients re-evaluate their own income levels substantially. Another reason is that income effects are not based solely on

income levels during a specific period of time, but rather on life-cycle income, at least to the degree that people are able to allocate their consumption across time. The marginal incentives we inform about, on the other hand, are tailored to the specific situation in which the people in question find themselves, as parents of an infant. The effect can be compared to that of a temporary tax reduction on labour income, which would give rise to the same substitution effect as a permanent tax reduction, but to much smaller income effects.

Even if the information letter increases labour supply on average, this is not necessarily true for fathers. As seen in Figure 2, the marginal EITC tends to be larger for mothers than for fathers in our sample, since most mothers in the sample spend most of the year with their child; and increasing the duration of the father's parental leave is one of the ways that a mother can increase her labour income. Since a couple can increase their EITC by moving parental leave from the mother to the father, we restrict our second hypothesis regarding intensive labour supply to mothers.

Hypothesis 2: Receiving the letter increases mothers' labour supply on the intensive margin.

We have concluded already that the marginal EITC tends to be larger for mothers than for fathers, because the former have lower incomes during this specific year. We expect our treatment to result in an increase in mothers' labour supply, but it is an open question whether both parents will work more, or whether fathers will instead reduce their time at work in order to enable mothers' increased labour.

While mothers face a larger marginal EITC than fathers, and women in general have a higher labour supply elasticity than men, it is not obvious that this translates into a stronger response among mothers to our information treatment. For example, it may be the case that mothers value the first year with their child so much that they have a lower elasticity in this specific setting than their partners do. We acknowledge this possibility, but we still find it probable that mothers' behavioural change dominates. This brings us to our third hypothesis:

Hypothesis 3: Receiving the letter makes parental leave more equally distributed within the couple.

We have stated repeatedly that the EITC strengthens the incentive to work. We would therefore expect couples' joint labour supply to increase. It is rare that children are already enrolled in

pre-school when they turn 12 months, so it is *possible* for most couples to prepone the preschool start for their child. However, many parents may find it difficult to shorten their joint parental leave, either because they do not want their child to be enrolled in childcare too early or because they worry that doing so will make it harder to gain admission to the specific preschool for which they applied. This would then attenuate their behavioural response. As we discussed in relation to the institutional background, however, there are also many other ways whereby parents can increase their labour income without decreasing the time spent with their child. This leads to our last hypothesis:

Hypothesis 4: Receiving the letter increases the couple's joint labour supply.

5 Conducting the experiment

To investigate the hypotheses derived above, we conducted an information experiment wherein we sent out a leaflet about the EITC to a randomized sample of new parents. A copy of the leaflet (in Swedish with an English translation) is attached in the appendix. The overall timing of our research project is summarized in Figure 4. We proceed now to describe how our study was designed, and what types of data we used to evaluate the effects.[14]

The sampling frame for the experiment was all biological mothers to children born between 1 June 2017 and 31 January 2018. This resulted in a total sample of 74,667 individuals. A total of 37,334 mothers were randomized to receive the leaflet (the treatment group), and the remainder were left as controls.[15] 151 individuals were removed from the mailing list due to emigration, death, protected identity, or being underage, leaving a total of 37,183 mothers (and their partners) who ended up getting the leaflet.

The treatment consisted of a simple information leaflet sent out by regular mail. This leaflet

[14]The study was approved by the Regional Ethical Review Board in Uppsala, and subsequently approved by the internal review board at Statistics Sweden, which both delivered our treatment and compiled the data. The researchers were not involved in that process. It should be noted that the pre-analysis plan was published prior to analysis of the data, but after the treatment. The reason is that we needed to discuss available variables and detailed variable names with Statistics Sweden before finalizing the pre-analysis plan. Statistics Sweden also needed to approve the delivery of registry data ("sekretessprövning") for research purposes. However, the study design was already set prior to the treatment, in the Ethical Review Board application.

[15]The sampling frame was created using the latest version of The Total Population Register (RTB). Because it was the mothers who were randomized, and not the children, parents to twins or triplets could not be assigned to both the treatment and the control group.

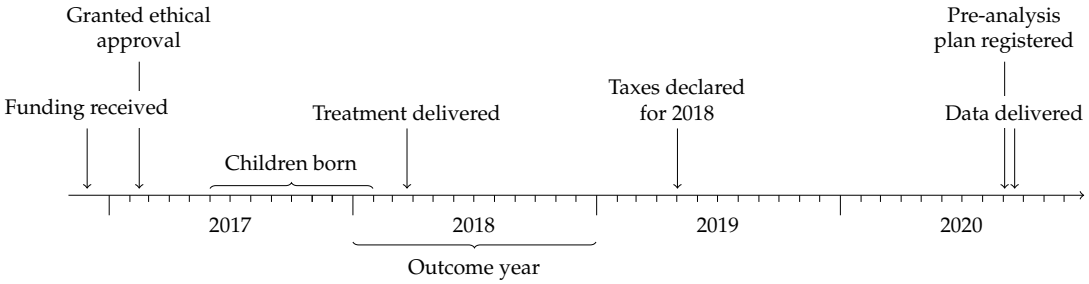

Note: The figure displays the timeline for the planning and execution of the research project.

Figure 4: Timeline of the study

informed parents about the Swedish EITC, and about how the size of this tax credit increases with the size of an individual's labour income. Specifically, the letter highlighted the effect of parental leave on the household's total tax credit, and stressed that the EITC increases rapidly with labour income up to 135,000 SEK (16,000 USD). This was illustrated both with a graph and with a written example showing how many months a fictitious person would need to work in order to reach that level. The letters were individually tailored, based on the parents' previous tax records. The income of the fictitious person was based on the mother's income; moreover, if either of the parents had registered a high enough income during previous years, the couple were also informed of the income level at which the EITC no longer increases (incomes larger than 365,000 SEK), and in some cases where it starts decreasing (incomes larger than 615,000 SEK).[16] Recipients were also informed that the leaflet was sent out as part of a project at Uppsala University in collaboration with Statistics Sweden. The letter was signed by the researcher in charge (name and title) and contact information was provided.

The letter was sent out by Statistics Sweden during March 2018. The timing of the letter was chosen so as simultaneously to maximize adaptation time and to minimize information decay – admittedly a difficult balance to strike. Since the incentives we highlighted made it particularly profitable to work 4–5 months during 2018, this left a post-treatment period of 4–5 additional months for possible behavioural adjustments to take effect, with enough time to change plans for the second half of the year (following summer vacation). So much time should not have

[16] We did include one sentence that, in general, a family that splits the parental leave evenly receive more EITC than a family where one parents is off work for the whole year. The reason for including this sentence is that we wanted to include one concrete example of EITC payout and how it changes with labour income on the household level to pinpoint that EITC is tilted toward low labour income.

passed, on the other hand, for the information to be forgotten. The letter was directed to both parents unless information about the biological father was missing, and sent to the home address registered for the mother.

The treatment consisted of accurate information about the EITC. Because we want to pinpoint the effect of providing information about this particular tax reform, we did not include any other information about the tax system or the household economy.[17] Nor did we make any normative appeal to act in any particular way. We are solely interested, namely, in the effect of information; we do not seek to induce parents to behave in any specific manner. This goes back to the main motivation behind our paper, which is to address the role of information for economic behaviour. As discussed previously, parents have several options if they wish to earn tax credit during a year. However, we could not explicitly state all of these options, given that we needed to strike a balance between the length of the leaflet (which was one page long) and the number of details in the tax code. As shown by Abeler and Jäger (2015) and Bhargava and Manoli (2015), people's ability to react to tax changes decreases with the complexity of information about the tax system; accordingly, we needed to write the leaflet in such a way that the majority of Swedes would be able to understand it. This meant excluding some technical definitions and tax code jargon. Furthermore, we did not want to provide a specific road map on how to think about the matters in question, given that we wanted the parents themselves to consider the economic incentives in play. The exact action to be taken had to be chosen by the parents themselves.[18]

We cannot know to what extent the parents engaged with the actual information. In other words, the persons in question were not all necessarily compliers in terms of engaging with the information. We had no interaction with recipients of the letter, and the whole sample frame was in the hands of Statistics Sweden.[19] We chose this strategy because we wanted to focus solely

[17] There are two reasons for our focus on the EITC. The first is that the tax reform is relatively unknown, which is a prerequisite for any information experiment with accurate information. The second is that the incentives we inform about are quantifiable, which means that we can relate them to the size of the incentives.

[18] Persons employed at a workplace with a collective agreement often have a top-up from their employer, so that the parental leave provides 90 per cent of their earlier wage, rather than 80 per cent. This is not true for all workers in Sweden, however. Furthermore, the rules for the top-up vary as between agreements, and may depend on how long one has worked at a particular workplace. The information letter was not individually tailored to take this into account, since we did not have data on whether a person had received a top-up. It should be noted, however, that a person does get an EITC on this small top-up, since it is paid out by the employer, and that such top-ups are included in our measure of earned income.

[19] Some individuals e-mailed or telephoned the researcher in charge, given that contact information was provided in the information leaflet.

on information, and to avoid giving any sense of direct engagement with us or with Statistics Sweden. Such engagement could result in social desirability bias, which we wanted to avoid. Naturally, furthermore, we could not monitor social interactions between the treatment group and the control group.

Our design rests on the assumption that our information letter has not affected labour supply in the control group. However, it is theoretically possible that some individuals in the treatment group discussed the content of the letter with persons in the control group. For example, parents on parental leave could have met each other in connection with various activities (like the so-called open preschools which parents and their child may have attended). It is important to note, however, that both the treatment group and the control group corresponded to minor fractions of the population (approximately 0.7 per cent each). Interaction between individuals in the two groups was therefore quite unlikely, and the risk was small that the treatment would result in any general public discussion.

6 Data and empirical specification

The data we have used come from two main registers maintained by Statistics Sweden: the Total Population Register (RTB), and the Integrated Database for Labour Market Research (LISA). These include information on parental leave, income, occupation, and education for both parents, as well as on the birth month of the child. The full set of variables can be found in Table 1.

To reduce residual variation in the dependent variables, thereby increasing the precision of our estimates, we include a number of control variables in the models. The birth month and gender of the child is included in all specifications below. Parental characteristics – more specifically years of education, potential income, year of birth, employment status in 2017, and previous parental leave – are entered in slightly different ways for the different models.

6.1 Model specifications

The preregistered models for evaluating the four main hypotheses are outlined as follows. In all cases, robust standard errors are used (cluster-robust for the first hypothesis), and one-tailed hypothesis tests are performed. We chose a one-tailed test, because this is how our hypotheses

Description	Years
Anonymous personal ID	–
Marker for reused ID	–
Anonymous ID of biological children	–
Treatment indicator	–
Year and month of birth (children)	–
Year of birth (everyone)	–
Days with parental allowance	2000–2013
Days with parental allowance	2014–2018
Earned income	2000–2018
Biological sex (binary)	–
Municipality of residence	2017–2018
Occupation	2017–2018
Year for info about occupation	2017–2018
Occ. congruence with Nov. work	2017–2018
Employment status	2017–2018
Labour market status	2017–2018
Industry of workplace	2017–2018
Employees at workplace	2017–2018
Highest level of education	2017
Source for education	2018
Year when edu. was finished	2018
Wage income	2018
Per cent of full-time	2018
Monthly wage	2018
Monthly wage excl. "rörliga tillägg"	2018
Enumeration factor	2018

Note: The table lists the variables in our data set and indicates the year for which each variable is available.

Table 1: Variables used

were formulated. We had planned in our pre-analysis plan to conduct multiple comparison corrections using the free step-down resampling method described in Westfall and Young (1993), but we found this to be superfluous when we obtained our unadjusted p-values. Controlling the rate of false positives is only necessary when there are positives.

The first hypothesis states that the treatment will increase labour supply on the extensive margin for both parents. To define extensive labour supply E_i, we use data on employment status to see whether the person was employed at all during 2018. This variable can take on the values Employed, Not employed but with taxable income, and Not employed and no taxable income. We have defined being employed as the categories Employed and Not employed but

with taxable income (the latter means one has had labour income previously during the year), and not being employed as the category Not employed and no taxable income. Regressing E_i on the treatment indicator as well as a control gives us the following model, where H_2 implies that $\beta_1 > 0$:

$$E_t = \alpha + \beta_1 \times T_j + \boldsymbol{\beta\chi_i} + \epsilon_i. \tag{1}$$

For this model, standard errors are clustered at the household level.

The second hypothesis states that receiving the letter will increase the intensive labour supply of the mother. To test this, we define intensive labour supply as annual labour income I_i.[20] We then regress I_i for 2018 on the treatment indicator as well as a vector of controls, limiting the sample to mothers:

$$I_t = \alpha + \beta_1 \times T_j + \boldsymbol{\beta\chi_i} + \epsilon_i \tag{2}$$

where i denotes the individual and j denotes the couple. In this model, H_1 implies that $\beta_1 > 0$.

The third hypothesis is that the treatment will make parental leave take-up more equally distributed within couples. To get as accurate a measure of staying at home as possible, we define parental leave days as all days not financed with labour income, rather than as registered reported days of parental leave.[21] This, however, requires us to get a measure of days financed with labour income, D_i. We define D_i as labour income I_i as a proportion of potential labour income I_i^*, multiplied by 365 to give the number of days at work, with the restriction that it must not exceed 365 minus the number of days of registered parental leave P_i. To calculate potential labour income, we use the highest relative income during the preceding three years multiplied by the average income for 2018, in order to compensate for general wage growth. If this yields a figure lower than working 75 per cent time on minimum wage (i.e.,180,000), we use the latter figure instead. This yields the following definition of extensive labour supply:

[20] In connection with footnote 12 regarding top-ups from employers due to a collective agreement, it bears noting that a top-up of this kind (from 80 per cent to 90 per cent of the earlier wage) also counts as earned income, since it is paid out by the employer.

[21] Our reason for not simply choosing registered reported days is that it is fairly common to stay at home for longer than the allotted parental leave days, without compensation.

$$D_i = \min\left(\frac{I_i}{I_i^*} \times 365, 365 - P_i\right) \tag{3}$$

where the potential income is:

$$I_i^* = \max\left(180000, \frac{I_{i,15} \times \hat{I}_{18}}{\hat{I}_{15}}, \frac{I_{i,16} \times \hat{I}_{18}}{\hat{I}_{16}}, \frac{I_{i,17} \times \hat{I}_{18}}{\hat{I}_{17}}\right). \tag{4}$$

Finally, we regress the absolute within-pair difference in days financed by labour income – i.e., $|D_j^M - D_j^F|$ – on the treatment indicator and a vector of controls, where H_4 implies that $\beta_1 < 0$:

$$|D_j^M - D_j^F| = \alpha + \beta_1 \times T_j + \boldsymbol{\beta}\boldsymbol{\chi_i} + \epsilon_i. \tag{5}$$

The fourth hypothesis states that the joint labour supply will be increased by the treatment. Here we set joint labour supply to the sum of the two, such that $I_j = I_j^M + I_j^F$. We then regress I_j on the treatment indicator and controls, where H_3 implies that $\beta_1 > 0$:

$$I_j = \alpha + \beta_1 \times T_j + \boldsymbol{\beta}\boldsymbol{\chi_i} + \epsilon_i. \tag{6}$$

It should be noted that our four outcome variables can all be manipulated without actual changes in the labour supply. For example, people can increase their earned income and thus their EITC by substituting days of paid parental leave with vacation days; and those who are self-employed can decide to increase their own salary without putting in more work. If our experiment causes such behaviour, it will be picked up in our empirical analysis, given that we observe the total effect. While such behaviour would be interesting to study, it is not possible for us to separate it from actual changes in labour supply (i.e. being at work).

7 Main results

We now look at how parents' labour supply was affected by our information letter. We begin with a balance check, before presenting the outcome in relation to our four hypotheses. We then conduct a more open-ended analysis with a focus on possible bunching behaviour.

If the treatment assignment is truly random, the treatment group and the control group are,

Variable	Control	Treat	Dif	p
Variables related to the child				
Female	0.49	0.48	−0.00	0.47
Variables related to the mother				
Years of education	13.08	13.08	−0.00	0.97
Potential income (thousands)	280.93	280.65	−0.28	0.78
Year of birth	1986.23	1986.24	0.01	0.78
Employed 2017	0.81	0.82	0.00	0.79
Previous parental leave	128.97	128.73	−0.24	0.84
Missing prev. par. leave	0.49	0.49	0.00	0.65
Variables related to the father				
Years of education	12.68	12.65	−0.03	0.11
Potential income (thousands)	408.68	410.13	1.45	0.43
Year of birth	1983.22	1983.28	0.06	0.19
Employed 2017	0.92	0.92	0.00	0.92
Previous parental leave	29.08	29.76	0.69	0.12
Missing prev. par. leave	0.50	0.50	0.00	0.93

Note: The table displays a balance test of the covariates included. The p-values indicate that there are no statistically significant differences between the treatment and the control group.

Table 2: Balance tests

in expectation, identical with regard to everything not affected by the treatment. To evaluate the treatment assignment, we have compared the mean values for all binary and continuous variables in the control vector. These results are presented in Table 2. As it turns out, all differences between the treatment group and the control group are small, and none are statistically significant at the 95 per cent level. In other words, this analysis does not raise any concerns about the treatment assignment, and we can move on to analysing how the treatment affected parents' labour supply.[22]

The results of our preregistered analyses for all of the four main hypotheses can be found in Table 3. From left to right, we see the estimated effect of the letter on (i) the parents' employment rate, (ii) the mother's earned income, (iii) the parents' joint earned income, and (iv) the absolute difference in days at home between the two parents.

[22]Our main preregistered analysis adjusts for these background factors; in the appendix, however, we also report models without these adjustments. The results here are in line with our main results; as can be expected, however, they are estimated less precisely. It should also be mentioned that we have conducted balance tests not mentioned in the pre-analysis plan. These tests revealed two statistically significant differences: parents in the treatment group were more likely to have a child born in June and less likely to have a child born in January.

	(1)	(2)	(3)	(4)
Treatment	−0.19	−870.87	−1124.60	0.27
	(0.14)	(585.28)	(873.61)	(0.78)
Units of analysis	Parents	Mothers	Couples	Couples
Outcome	Employed	Income	Joint income	Dif. in days
Mean	84	85,055	393,844	193
Observations	143,925	74,439	74,689	68,926
Adjusted R2	0.491	0.575	0.845	0.194

Note: Main results for our four pre-registered hypotheses. From left to right, the outcomes are (i) an indicator for being employed, (ii) the mother's earned income, (iii) the parents' joint earned income, and (iv) the absolute difference in days at home between the two parents. Robust standard errors in parentheses (cluster-robust for the first hypothesis). The estimated treatment effects are close to zero for all tested hypotheses.

Table 3: Pre-registered hypotheses

A brief but conclusive summary is that we see a rather precisely estimated zero effect in each case. All of the hypotheses, in other words, should be considered falsified. None of the estimates are statistically significant; point estimates are practically trivial; and all of the coefficients have the wrong sign relative to those in our hypotheses. This means labour supply is actually lower among those receiving the information letter (although, again, the estimated coefficients are not statistically significant). In terms of precision, even the upper bounds of the confidence intervals correspond to trivial effect sizes. For example, the letter is estimated to have reduced the earned income among mothers, as reported in model two, by approximately 900 SEK annually. If we calculate a 95 per cent confidence interval around this estimate, the letter could at most have increased earned income by 300 SEK. Compare this to the average income of 85,000 SEK in the sample. This means our point estimate corresponds to a fall of 1 per cent, and according to the confidence interval we can rule out increases larger than 0.4 per cent. This point estimate is also close to the figure for joint labour supply, at a fall of approximately 1100 SEK annually. The point estimate for model four, meanwhile, corresponds to an increase of less than half a day in within-pair differences in parental leave. In all respects, this is a zero effect.

We have also investigated heterogeneous treatment effects with regard to our four hypotheses and to subgroups in the data. The results are presented in Figure A2 in the appendix. These are also null findings, all in all, although the confidence intervals are wider due to the smaller size of the sample for each subgroup.

To rule out the possibility that major behavioural adjustments in this group have already taken place, we also report an analysis of changes in our outcome measures potentially resulting from the introduction of the EITC. While the construction of the EITC is such that a robust causal analysis is impossible, large behavioural effects should be visible when we compare those who became parents at the time of its introduction with those who had become parents the year before. These tests, reported in the appendix, reveal no clear evidence for such behavioural adjustments. It is important to point out that this evaluation only concerns the group facing the particular – and unusually opaque – configuration of incentives investigated here; it does not necessarily generalize to behavioural responses to the EITC in general.

8 Supplementary analysis: Bunching

As mentioned in our pre-analysis plan, we have also conducted a more exploratory analysis of the possible treatment effects, in which we compare the density of the treatment group and the control group over the range of possible incomes. This analysis allows us to detect other effects than those set out in our pre-specified hypotheses, and it can also yield more detailed information about such things as where in the income distribution changes in labour supply take place. However, the exploratory nature of this analysis also carries an increased risk of getting false positives: with an indefinite number of comparisons to be made, the probability increases dramatically that sampling variability will give rise to large differences between the two groups. It is therefore important to treat this analysis as exploratory, as well as to interpret all results against what we should expect based on economic theory and the information in our letter.[23]

We have done this separately for mothers and fathers, and the results are presented in Figure 5. We have marked the thresholds at 135,000 (solid) and 365,000 (dashed) using vertical lines, and the largest area above or below zero using a grey background. Figure 5a a shows the estimated density difference for mothers. Here we see a peak just before the 135,000 threshold, indicating that mothers in the treatment group were more likely to end up with an earned income in this interval, compared to mothers in the control group. This is interesting, seeing this is the threshold

[23] In the pre-analysis plan we mention a number of potential exploratory analyses in addition to the carefully specified hypotheses. However, as was also foreseen in the pre-analysis plan, we have adjusted the set of supplementary analyses in response to our main null results. For example, in this section we have integrated the vizualization that was described in the first part of the pre-analysis plan with a discussion about bunching, which was mentioned at the end.

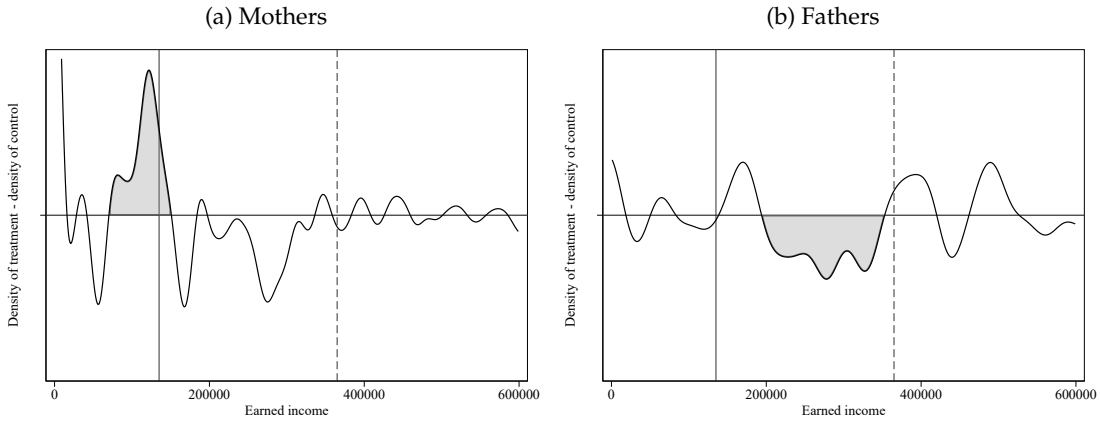

Figure 5: Differences in density

Note: The graphs display the difference in density between the treatment group and the control group. The left graph shows the results for mothers; the right graph shows them for fathers. The density is estimated with a Gaussian kernel function.

we emphasized in the letter, in which we pointed out that the EITC increases rapidly up to this point. The corresponding difference for fathers is shown in Figure 5b. Here two small peaks close to the thresholds can perhaps be discerned, but the largest area is negative and is found between those peaks.

The main question is how large these density differences are. The shaded areas in Figure 5 correspond to 0.5 and 0.6 per cent of the entire treatment group, respectively, or to 2.0 and 1.5 per cent of the individuals in the control group located at the same interval. One way to present these differences is in an overlapping histogram, as we do in Figure 6. Expressing effect sizes in words is difficult, but we believe most people would agree these effects are very small.[24] All in all, we conclude, there is no clear evidence of any significant bunching.[25]

[24]How likely is it, then, that the small differences described above correspond to causal effects of the treatment? On the one hand, bunching close to the 135,000 threshold is exactly what we can expect to see if female recipients try to optimize. Moreover, such optimization is dependent on detailed knowledge about the tax system, which most subjects did not have before receiving our letter. On the other hand, the effect is barely statistically significant, and only in a naive sense: the probability of getting a difference this large at *this point in the income distribution* will be small if the treatment has no effect (simulations assuming a zero effect provide us with a p-value of 0.058 for the difference at the fifth bar being of at least this size). But the probability of obtaining a difference this size at *some* place in the income distribution is rather large, and focusing on differences where they appear prominently increases the risk for Type I errors. Related, a Kolmogorov–Smirnov test, which uses the largest observed difference in the *cumulative* density functions as its test statistic, cannot reject the null-hypothesis that the samples are drawn from the same distribution. Moreover, the increased density in the treatment group comes from a smaller density at higher incomes, which is not what we would have expected.

[25]In the appendix, we present a similar histogram for those who are self-employed. Because they tend to have more control over their declared labour income, we would expect the bunching to be more pronounced in this group, similar to what Bastani and Selin (2014) find. The results in the appendix indicate that this may be the case, although our

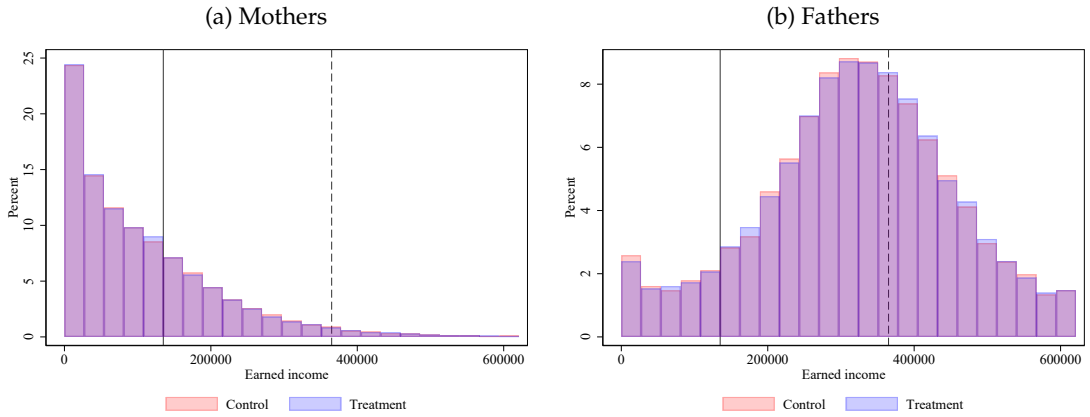

Figure 6: Distribution of labour income

Note: The graphs show the distribution for labour income among those employed. The graph to the left corresponds to mothers, the graph to the right to fathers. The control group is coloured in red, the treatment group in blue. The bars are purple for overlapping bars. Note that the scale of the y-axis is different in the two subfigures.

However, even if the bunching-peaks are small and economically insignificant, they contribute to our assessment of the main results. If we find no differences at all over the entire distribution, concerns can perhaps be raised whether something went wrong with the experiment: e.g., the recipients did not read the letter, there was an error in the delivered data, etc. However, if the only effect we find is bunching around the 135,000 threshold, we may conclude that it dominates the hypothesized effects, at least for the subset of people who changed their behaviour in response to the information letter; and we may be more inclined to reconsider our hypotheses than the way in which the experiment was implemented.

9 Discussion and conclusion

Our empirical analysis shows that the effect of receiving a letter with personalized information about the EITC is zero across the board. The high power of this experiment, together with point estimates in the opposite direction of our hypotheses, allows us to rule out anything but negligible effects: receiving the letter may, at most, have increased the employment rate by 0.08 percentage points, or the couple's joint income by 600 SEK (0.15 per cent) per year. This holds despite the fact that the economic incentives are large and parents have multiple options for

sample of self-employed is too small to say anything for sure.

acquiring tax credit from labour income.

So why did this information experiment not affect any of the measures we set out to explore? Granted, most parents would answer that they have other priorities than maximizing their net income, and that they do what is best for their child regardless of how it affects their economic situation. It may be they believe it would be bad for the human capital development of their child to prepone the start of preschool. It may also be they value their own time with their child so much that they are unwilling to reduce it. Such arguments, however, are unconvincing for explaining our findings, because there are many different ways to increase one's labour income, and not all of them require parents to sacrifice time with their child. We estimate the total effect with regard to labour income in our empirical analysis, and we would pick up such behavioural responses if they were present.[26] Furthermore, putting a heavy stress on one objective (spending time with one's child) does not require an individual to ignore other aims. After all, most parents in our study do end their parental leave during the year under study – presumably to earn income, which we interpret as meaning that the point in time where the expected marginal utility from labour exceeds the marginal utility of parental leave falls within this year. Under the assumption of a continuous utility function, this point would be reached earlier during the year if parents were aware of how the EITC affects their incentives for acquiring labour income. Having dismissed this account of parents' behaviour, we think there are three possible explanations for our zero findings. We discuss each one in some detail.

The first explanation is that there were problems with how we conducted the experiment. Was the timing wrong, so that the letter was sent out too late, allowing too little time to adapt? Or, on the contrary, was the letter sent out too early, leaving enough time for the recipients simply to forget about the information given? While both of these explanations doubtless have *some* leverage, we think it unlikely that either is strong enough to explain the precisely estimated zero effects. As discussed previously, we chose the timing (March) strategically – roughly two to three months before the summer holidays (generally taken between June and August), allowing that amount of time for parents to discuss the post-holiday parental leave with each other and with employers. The response time permitted between the letter and an actual adjustment is

[26] In theory, the information on the EITC could even be used to decrease the time spent at work, if that is what the parents desire. We have referred here, for example, to the planning of vacation days. Strategic behaviour of this kind requires some effort, and the information treatment could perhaps been seen as a nudge to do the necessary planning.

therefore roughly half a year (March–September); or slightly less in the maximum effect scenario, which involves going from zero work during the year to the optimal 4–5 months. At the same time, the time frame needed for recipients to react to the information is substantially shorter, since in most cases they will need to discuss these adjustments before the summer. Whether this time frame (March–May) is still long enough to render the letter "inert" can be debated, and we have no strong priors on what to expect in that regard. Given the incentives involved, however, we would argue that a two- to three-month period for (re-)negotiation of the parental leave distribution is not too long.

Is it wrong to assume that recipients even read (or, if they did, understood) the contents of the letter? We think it is unlikely that many people would discard a letter sent from a national authority (Statistics Sweden) without even opening it, especially in a country like Sweden, where trust in public institutions is very high by international standards. In a similar study, Engström et al. (2019) estimated positive treatment effects in an information experiment where the leaflet was mailed out by the Swedish Pension Agency. Whether or not the respondents understood the information we sent them is hard to know; however, we made sure to include concrete examples with ballpark figures derived from the recipients' own previous income, in order to make the matter as transparent as possible. It is worth noting that similar information interventions have shown fairly large effects earlier (e.g., Liebman and Luttmer 2015; Kostol and Myhre 2021), which indicates that this type of information is probably not too hard to understand. One possibility would have been for us to include a number of questions at the end of the information letter, and to ask the individuals to mail it back. That would have furnished us with a lower bound regarding how many parents actually read and understood the letter. On the other hand, answering a survey would also introduce a sense of engagement with the researcher, and our goal was to stick with neutral information about the tax system, in order to avoid introducing effects that might be attributable to other factors than the reduction of information frictions (such as social desirability bias induced by a desire to be seen as conforming to social norms regarding the equal sharing of parenting duties).[27]

The second explanation is that parents were already sufficiently well-informed – even in the

[27] Such a survey would also cost a lot of money, given that Statistics Sweden would need to compile the answers. This was not possible, for monetary and practical reasons.

absence of the leaflet – about the incentives arising from the interface between the EITC and parental insurance. This would imply the additional information is superfluous, since the parents are already at their utility-maximizing levels (minus adjustment frictions *other* than information frictions) of joint labour supply. In terms of the theoretical framework presented earlier in this paper, households would then be observing the EITC and the tax rate without any bias (i.e., $E[\varepsilon] = 0$). We think this explanation strongly lacks face validity: as outlined in the background section, it is unlikely to hold due to the complexity of the EITC's functioning and the "invisibility" of the tax system, and also because of existing evidence on the poor state of knowledge among the general population about the EITC. Moreover, if a large behavioural response had already occurred due to the introduction of the EITC, in line with our hypotheses, then we would expect to see it in the data. Our analyses, however, turned up no such indications (see the appendix).

This leads us to the explanation we think makes the most sense, given the context: the elasticity is simply very low for parents with infants who face the decision of how to distribute a parental leave benefit. One way of putting this is to say that parents, at least in this specific case, are so-called satisficers: they have arrived at a good enough solution, and they will expend very little energy to optimize further. Such satisficing behaviour can be understood in slightly different but not incompatible ways. It can be seen as a strategy for making decisions when optimizing is costly, and such costs may be unusually high when all of one's energy is focused on taking care of a child. Or perhaps other objectives, or norms, crowd out monetary considerations as soon as the couple has secured an income which is good enough. Put in economic terms, one might say there is a cost associated with bringing up the topic, especially when the parents have a shared understanding that their child comes before everything else.

In light of previous evidence on labour supply elasticity, our results may seem somewhat paradoxical, inasmuch as several existing studies argue that such elasticity is in fact fairly high, especially among mothers (e.g., Eissa and Liebman 1996; Meyer and Rosenbaum 2001). Our results are however in line with later studies claiming to show that this phenomenon is an artefact of such things as other changes in welfare benefits that coincide in time (Kleven 2019). Pace Kleven (2019), however, the null results in our case probably cannot be considered a consequence of information friction. Comparing our results with those found by Kostol and Myhre (2021), on the other hand, may reveal the difference between labour supply responses in different types

of compensation system. Whereas disability insurance is often a long-term benefit, parental leave by its nature is limited to a fixed and shorter period, the expectation being that the person in question will return to his/her main source of income. This may increase the likelihood of satisficing behaviour: the immediate incentives are large, yet parents may be less likely to optimize in response to what they perceive as a "temporary" situation.

References

Abeler, Johannes, and Simon Jäger. 2015. "Complex tax incentives." *American Economic Journal: Economic Policy* 7 (3): 1–28.

Andersson, Christian, and Jesper Antelius. 2010. "Jobbskatteavdraget – bra tänkt men illa känt." *Ekonomisk Debatt* 38 (2): 6–17.

Avdic, Daniel, and Arizo Karimi. 2018. "Modern family? Paternity leave and marital stability." *American Economic Journal: Applied Economics* 10 (4): 283–307.

Ballard, Charles L, and Sanjay Gupta. 2018. "Perceptions and realities of average tax rates in the federal income tax: Evidence from Michigan." *National Tax Journal* 71 (2): 263–294.

Bastani, Spencer, and Håkan Selin. 2014. "Bunching and non-bunching at kink points of the Swedish tax schedule." *Journal of Public Economics* 109:36–49.

Bastian, Jacob. 2020. "The rise of working mothers and the 1975 earned income tax credit." *American Economic Journal: Economic Policy* 12 (3): 44–75.

Bennmarker, Helge, Lars Calmfors, and Anna Seim. 2014. "Earned income tax credits, unemployment benefits and wages: Empirical evidence from Sweden." *IZA Journal of Labor Policy* 3 (1): 1–20.

Benzarti, Youssef. 2020. "How taxing is tax filing? Using revealed preferences to estimate compliance costs." *American Economic Journal: Economic Policy* 12 (4): 38–57.

Bhargava, Saurabh, and Dayanand Manoli. 2015. "Psychological frictions and the incomplete take-up of social benefits: Evidence from an IRS field experiment." *American Economic Review* 105 (11): 3489–3529.

Bill 2006/07:1. *Budget Bill for 2007.* Finansdepartementet, Regeringskansliet.

Browning, Martin, Pierre-Andre Chiappori, and Valerie Lechene. 2006. "Collective and unitary models: A clarification." *Review of Economics of the Household* 4 (1): 5–14.

Chetty, Raj. 2012. "Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply." *Econometrica* 80 (3): 969–1018.

Chetty, Raj, John N Friedman, and Emmanuel Saez. 2013. "Using Differences in Knowledge across Neighborhoods to Uncover the Impacts of the EITC on Earnings." *American Economic Review* 103 (7): 2683–2721.

Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and taxation: Theory and evidence." *American Economic Review* 99 (4): 1145–77.

Chetty, Raj, and Emmanuel Saez. 2013. "Teaching the tax code: Earnings responses to an experiment with EITC recipients." *American Economic Journal: Applied Economics* 5 (1): 1–31.

Chiappori, Pierre-André. 1992. "Collective labor supply and welfare." *Journal of Political Economy* 100 (3): 437–467.

Currie, Janet. 2006. "The take-up of social benefits." In *Poverty, the distribution of income and public policy,* edited by A. Auerbach, D. Card, and J.M. Quigley, 80–148. New York (NY): Russell Sage.

De Bartolome, Charles AM. 1995. "Which tax rate do people use: Average or marginal?" *Journal of Public Economics* 56 (1): 79–96.

Donni, Olivier, and Pierre-André Chiappori. 2011. "Nonunitary Models of Household Behavior: A Survey of the Literature." In *Household Economic Behaviors,* edited by J. A. Molina, 1–40. New York (NY): Springer.

Edmark, Karin, Che-Yuan Liang, Eva Mörk, and Håkan Selin. 2016. "The Swedish Earned Income Tax Credit: Did It Increase Employment?" *FinanzArchiv* 72 (4): 475–503.

Eissa, Nada, and Jeffrey B Liebman. 1996. "Labor supply response to the earned income tax credit." *The Quarterly Journal of Economics* 111 (2): 605–637.

Engström, Per, Eskil Forsell, Johannes Hagen, and Arnaldur Stefánsson. 2019. "Increasing the take-up of the housing allowance among Swedish pensioners: A field experiment." *International Tax and Public Finance* 26 (6): 1353–1382.

Engström, Per, Katarina Nordblom, and Arnaldur Stefánsson. 2018. *Multiple misbehaving: Loss averse and inattentive to monetary incentives.* Working Paper 2018:8. Department of Economics, Uppsala University.

Feldman, Naomi E, Peter Katuščák, and Laura Kawano. 2016. "Taxpayer confusion: Evidence from the child tax credit." *American Economic Review* 106 (3): 807–35.

Fortin, Bernard, and Guy Lacroix. 1997. "A test of the unitary and collective models of household labour supply." *The Economic Journal* 107 (443): 933–955.

Gideon, Michael. 2017. "Do Individuals Perceive Income Tax Rates Correctly?" *Public Finance Review* 45 (1): 97–117.

Ginja, Rita, Jenny Jans, and Arizo Karimi. 2020. "Parental leave benefits, household labor supply, and children's long-run outcomes." *Journal of Labor Economics* 38 (1): 261–320.

Goldin, Jacob. 2015. "Optimal tax salience." *Journal of Public Economics* 131:115–123.

Guthmuller, Sophie, Florence Jusot, and Jérôme Wittwer. 2014. "Improving takeup of health insurance program: A social experiment in France." *Journal of Human Resources* 49 (1): 167–194.

Jones, Damon. 2010. "Information, preferences, and public benefit participation: Experimental evidence from the advance EITC and 401 (k) savings." *American Economic Journal: Applied Economics* 2 (2): 147–63.

Karimi, Arizo, Erica Lindahl, and Peter Skogman Thoursie. 2012. *Labour supply responses to paid parental leave.* Working paper 2012:22. The Institute for Evaluation of Labour Market and Education Policy (IFAU).

Karlan, Dean, Margaret McConnell, Sendhil Mullainathan, and Jonathan Zinman. 2016. "Getting to the top of mind: How reminders increase saving." *Management Science* 62 (12): 3393–3411.

Kleven, Henrik. 2019. *The EITC and the Extensive Margin: A Reappraisal.* NBER Working Paper No. 26405. National Bureau of Economic Research.

Kostol, Andreas Ravndal, and Andreas S. Myhre. 2021. *Labor Supply Responses to Learning the Tax and Benefit Schedule.* IZA Discussion Papers 13900. Institute of Labor Economics (IZA).

Liebman, Jeffrey B, and Erzo FP Luttmer. 2015. "Would people behave differently if they better understood social security? Evidence from a field experiment." *American Economic Journal: Economic Policy* 7 (1): 275–99.

Manoli, Dayanand S, and Nicholas Turner. 2014. *Nudges and learning: Evidence from informational interventions for low-income taxpayers.* Working paper No. w20718. National Bureau of Economic Research (NBER).

Matikka, Tuomas, and Tuuli Paukkeri. 2016. *Does information increase the take-up of social benefits? Evidence from a new benefit program.* Working paper 83. VATT Institute for Economic Research.

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

Rees-Jones, Alex, and Dmitry Taubinsky. 2020. "Measuring 'schmeduling'." *The Review of Economic Studies* 87 (5): 2399–2438.

Riksrevisionen. 2009. *Jobbskatteavdraget.* Report 2009:20. Stockholm: Riksrevisionen.

Skolverket. 2017. *Beskrivande data 2016: Förskola, skola och vuxenutbildning.* Rapport 452. Stockholm: Skolverket.

Stantcheva, Stefanie. 2021. "Understanding tax policy: How do people reason?" *The Quarterly Journal of Economics* 136 (4): 2309–2369.

Westfall, Peter H, and S Stanley Young. 1993. *Resampling-based multiple testing: Examples and methods for p-value adjustment.* New York (NY): John Wiley & Sons.

Appendix to "Filling in the Blanks: How Does Information about the Swedish EITC Affect Labour Supply?"

Pär Nyman, Linuz Aggeborn and Rafael Ahlskog

To test the robustness of our results and facilitate our interpretation of the same, we have conducted a number of analyses in addition to those presented in the main paper. In this appendix, we present (i) results from regressions without the vector of covariates, (ii) an analysis with the sample restricted to self-employed, (iii) an analysis of heterogeneous treatment effects, and (iii) an analysis of possible behavioural changes prior to the experiment. At the end, we also include an example of what the information letter looked like, together with an English translation.

Where our empirical analysis is concerned, we have run the same regressions as in the main paper, but without the vector of control variables. These results are shown in Table A1. As expected, the estimated treatment effects are in the same ballpark as when they were estimated with controls. All of the estimates are small, and none is even close to being statistically significant.

Yet there is one difference which stands out a little bit: the coefficient of interest for the third hypothesis has changed sign, although it remains small in substantive terms and is far from statistically significant. This difference should not be exaggerated; however, in order to understand the reasons behind it, we have added the control variables one at a time (not shown here). This exercise showed that including no other covariate than the father's potential income is enough to move the coefficient from 812 (bivariate) to -918 (fairly close to the main

	(1)	(2)	(3)	(4)
Treatment	−0.14	−441.28	811.97	−0.35
	(0.21)	(896.92)	(2201.61)	(0.87)
Units of analysis	Parents	Mothers	Couples	Couples
Outcome	Employed	Income	Joint income	Dif. in days
Mean	84	85,055	393,844	193
Observations	143,925	74,439	74,689	68,926
Adjusted R2	-0.000	-0.000	-0.000	-0.000

Note: This table displays the results for our four pre-registered hypotheses, but without the vector of covariates. From left to right, the outcomes are (i) an indicator for being employed, (ii) the mother's earned income, (iii) the parents' joint earned income, and (iv) the absolute difference in days at home between the two parents. Robust standard errors in parentheses (cluster-robust for the first hypothesis).

Table A1: Main analysis without covariates

specification). This in turn is explained by a small number of fathers in the treatment group who had a very large income before the experiment was conducted. In our view, controlling for this imbalance is the reasonable thing to do. More importantly, however, the effect we estimate should be interpreted as close to zero, regardless of model specification and the resulting sign of the coefficient.

Second, in accordance with the pre-analysis plan, we have also run our models on a sample restricted to those who are self-employed. These individuals tend to have more control over their income. Partly because they can decide for themselves how much they want to work, and partly because they, to a varying degree, can adjust their labour income without adjusting their amount of work. Since our experiment could have affected whether a person is counted as self-employed or not, we use a measure from 2017.

The regression results for the self-employed are presented in Table A2. The first thing that stands out is the small number of observations. Only a few per cent of our sample are self-employed, and as we can see from the standard errors, restricting the analysis to this sub-sample severely affects our precision. We should therefore be careful when interpreting the estimated treatment effects of the letter. For most of the outcomes, the coefficients are also relatively similar to the main results, suggesting that labour supply among the self-employed was unaffected by the experiment. With that said, there is one effect which deserves extra attention. According to the estimate in the second column, the information letter caused self-employed mothers to on

	(1)	(2)	(3)	(4)
Treatment	0.02	−7448.39	−2191.09	3.58
	(0.49)	(4597.51)	(7039.71)	(4.55)
Units of analysis	Parents	Mothers	Couples	Couples
Outcome	Employed	Income	Joint income	Dif. in days
Mean	95	119,869	434,889	166
Observations	7,411	2,205	2,210	2,040
Adjusted R2	0.022	0.463	0.886	0.100

Note: This table displays the results for our four pre-registered hypotheses, but estimated on a sample restricted to the self-employed. From left to right, the outcomes are (i) an indicator for being employed, (ii) the mother's earned income, (iii) the parents' joint earned income, and (iv) the absolute difference in days at home between the two parents. Robust standard errors in parentheses (cluster-robust for the first hypothesis).

Table A2: Main analysis with sample restricted to the self-employed

average reduce their earned income with approximately 7500 SEK. However, the standard error is large and the effect is not statistically significant ($p = 0.105$).

To better understand this possible effect, we have created a similar overlapping histogram as in Figure 6a, but this time with the sample restricted to the self-employed. Here, too, we find indications of strategic bunching. The density in the treatment group is higher at two locations in the income distribution: just below the 135,000 threshold, after which the marginal tax rate is increased, and at the very small incomes represented by the first bar.

It is impossible to say whether these differences are caused by a behavioural response to the information letter, or by chance alone. On the one hand, the probability of obtaining differences this large by sampling errors alone is small ($p < 0.01$), and at least the bunching below 135,000 would be a rational response to the letter. On the other hand, the probability of finding at least one false positive increases with the number of tests, and therefore also small p-values are inconclusive – especially for sub-group analyses that were only vaguely described in the pre-analysis plan. Even if it was the letter that resulted in these differences, the effect would still be of minor economic importance. However, it would speak in favour of our preferred interpretation of the null results: the letter was read and understood, but people chose not to adjust their labour supply in the ways we had hypothesized.

Third, we have also estimated our main models on a number of other sub-samples of our data, in order to analyse whether the treatment effects differ in a systematic way across groups.

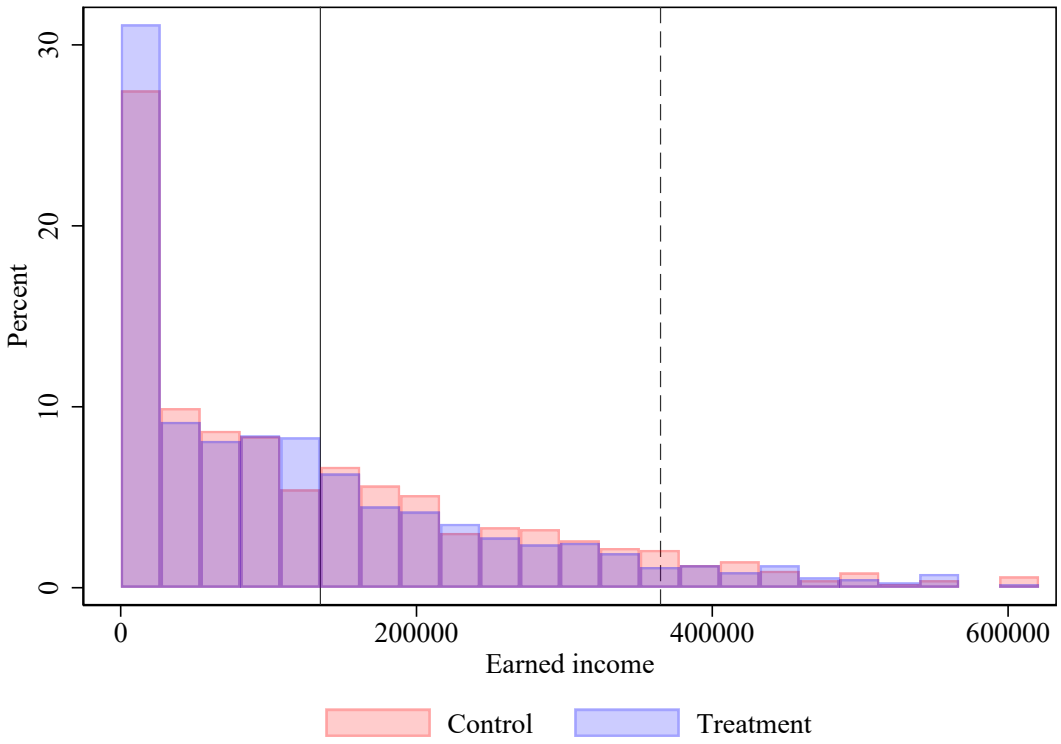

Figure A1: Distribution of labour income among self-employed mothers

Note: The figure shows the distribution for labour income among self-employed mothers. The control group is coloured in red, the treatment group in blue. The bars are purple for overlapping bars.

These results are summarized in Figure A2. The sub-figures a–d correspond to hypotheses 1–4; and each dot shows an estimated treatment effect within a specific sub-sample, together with the surrounding confidence interval.

The first sub-samples are defined by when the child was born. These groups differ with respect to how old their child is when they receive the letter, as well how much the parents earn during 2018 (the earlier the child was born, the sooner will the parents get back to work). The next set of sub-samples are based on the parents' level of education (higher or lower than average), because we expect people with a long education to be better equipped at understanding and acting on the information in our letter. We then create sub-samples defined by the parents' potential income (higher or lower than average) as well as the mother's share of the couple's total potential income (less than 40 per cent, between 40 and 60 per cent or more than 60 per

(a) H1: Parents' employment

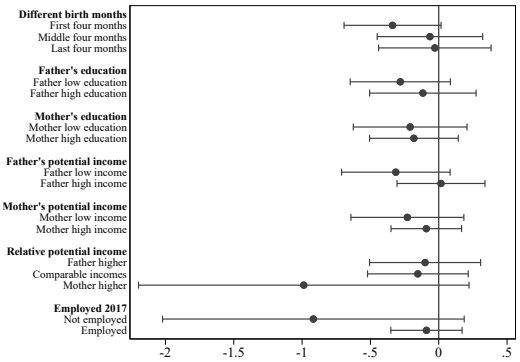

(b) H2: Mother's income

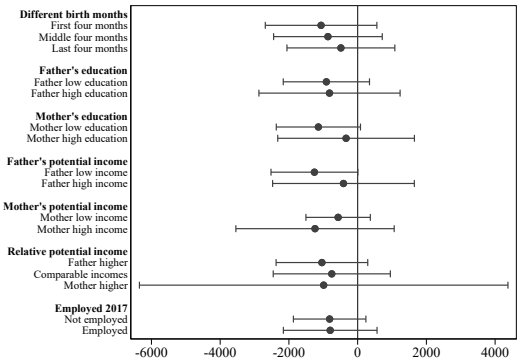

(c) H3: Joint income

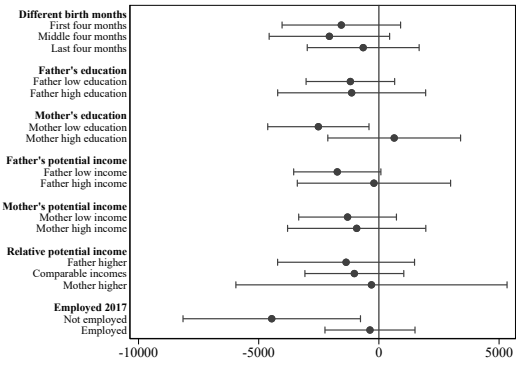

(d) H4. Dif. in parental leave

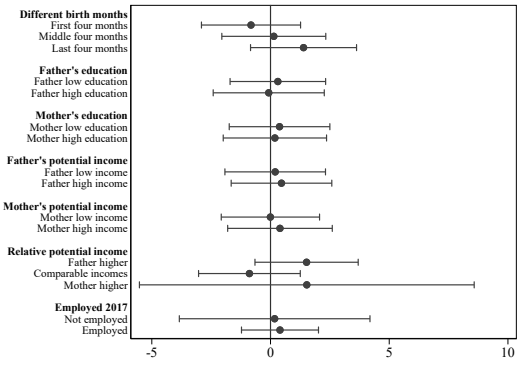

Figure A2: Differences in density

Note: The figures display the estimated treatment effect when the analysis is limited to a specific sub-sample, using the same model specification as in the main paper. These sub-samples are based on when the child was born (June–Sep, Aug–Nov, Oct–Jan), the parents' level of education (above or below average), the parents' potential income (above or below average), and their labour market status in 2017 (for the last two hypotheses, we use the mother's status).

cent), because their respective wages affect their economic incentives in general as well as the incentives we informed about in the letter. Lastly, we also compare those who were employed 2017 with those who were not, because we believe that it is easier to quickly adjust your labour income if you have a job.

In short, we find no major differences between groups, and there is no sub-sample for which any of our hypotheses is supported. Of the 52 estimated coefficients, only two are statistically significant at the 95 per cent level, and they have the wrong sign relative to that hypothesized: the experiment is estimated to have reduced the joint income of couples where the mother lacks

employment and/or has a low level of education. However, the differences between these effects and the ones in the reference group (employed and high education) are not statistically significant. Additionally, and most importantly, the expected number of false positives with 52 tests is 2–3, given a significance threshold of 95 per cent.

Fourth, as we discussed in the concluding section of the main paper, one possible explanation for the small treatment effects may be that people had sufficient knowledge about the EITC and had already adapted their behaviour in response to the reform. That is not the explanation we believe in, mainly because we strongly doubt people are so well-informed. It is difficult, however, to rule out this explanation altogether, absent a robust estimate of the adjustments following the introduction of the EITC.

Unfortunately, as others have pointed out, the way that the Swedish EITC was implemented makes it impossible to estimate its causal effects rigorously. Nevertheless, looking for such effects in the data is helpful for interpreting the results: if there actually were strong effects, that would constitute a strong argument against our interpretation of the null results in our main analyses. The following cursory analysis proceeds on the assumption that *if* the reform gave rise to substantial behavioural responses, we should be able to detect them simply by comparing the situation before and after the EITC was implemented.

In Figure A3, we show how our four outcomes developed between 2002 and 2011: i.e., within a five-year window around when the EITC was implemented (the decision was taken in late 2006, and incomes started being affected in 2007, as illustrated by the vertical dashed line). For each year, the sample consists of parents with the same amount of time between giving birth and the outcome year used in our experiment. If we look carefully at these graphs, we find it is only for the first hypothesis (employment) that we see a change in behaviour in line with the incentives generated by the EITC. Mothers' income increased more slowly than before (A3b), while fathers' incomes increased faster than they did before the reform (not shown directly, but this can be inferred by comparing A3c with A3b, and it is also behind the plateau in A3d). As for the employment rate, a closer examination shows the rise in it cannot be attributed to the implementation of the EITC. According to the Swedish labour force statistics, the rapid rise in the employment rate began already during the summer of 2006, well before the right-wing alliance won the election.

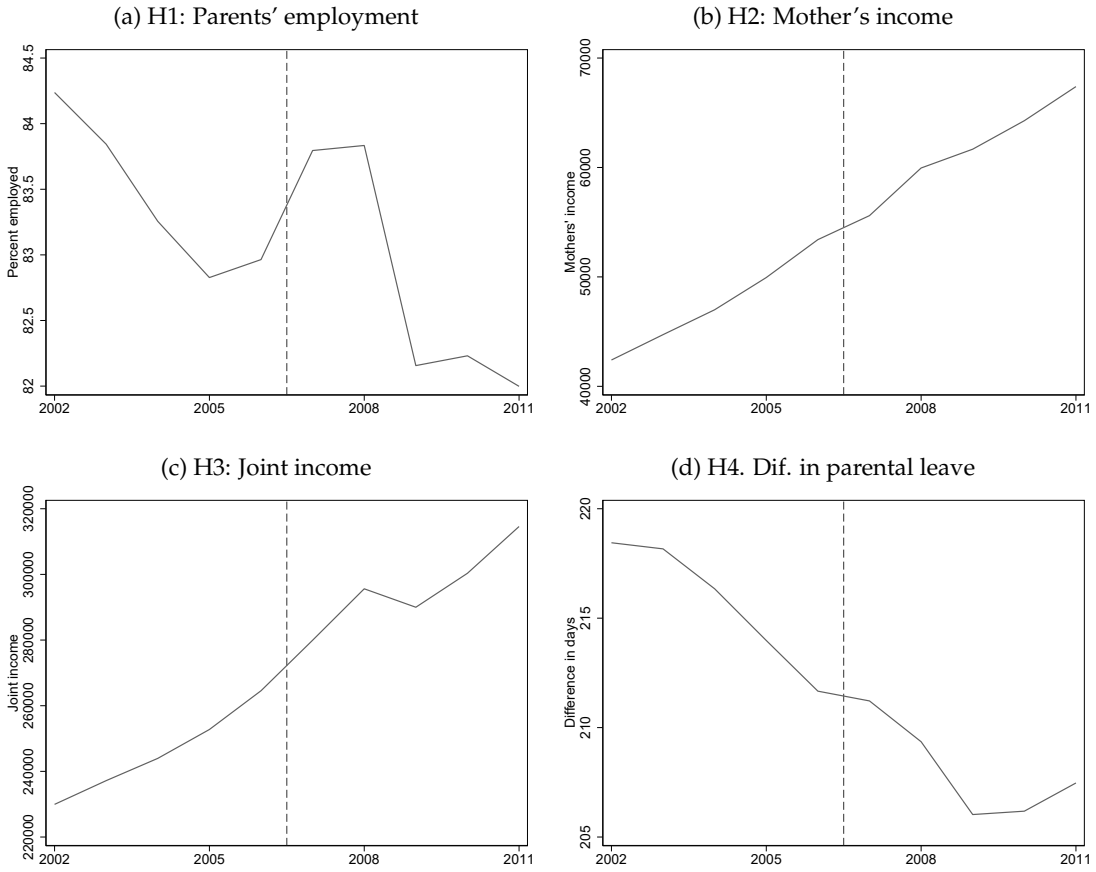

Figure A3: Outcomes before and after the EITC introduction

Note: The figures display our four outcomes – measured during the five years before and after the EITC was introduced – for parents with a child born between June of the preceding year and January in the outcome year.

The main problem with making inferences from observational data like these, of course, is that there are so many other things happening that also affect the outcomes of interest. This is not, unfortunately, something that we can solve. One way to alleviate the problem, however, may be to find a comparison group that has many features in common with the sample we presented above, but which is less affected by the EITC. We have therefore chosen to complement the picture presented above with a comparison between the mothers in that sample and those who had given birth one year earlier. The latter tended to have much larger labour incomes, simply because they had plenty of time to return to work. Their marginal EITC, and thus their incentive to increase their labour income, therefore tended to be much smaller.

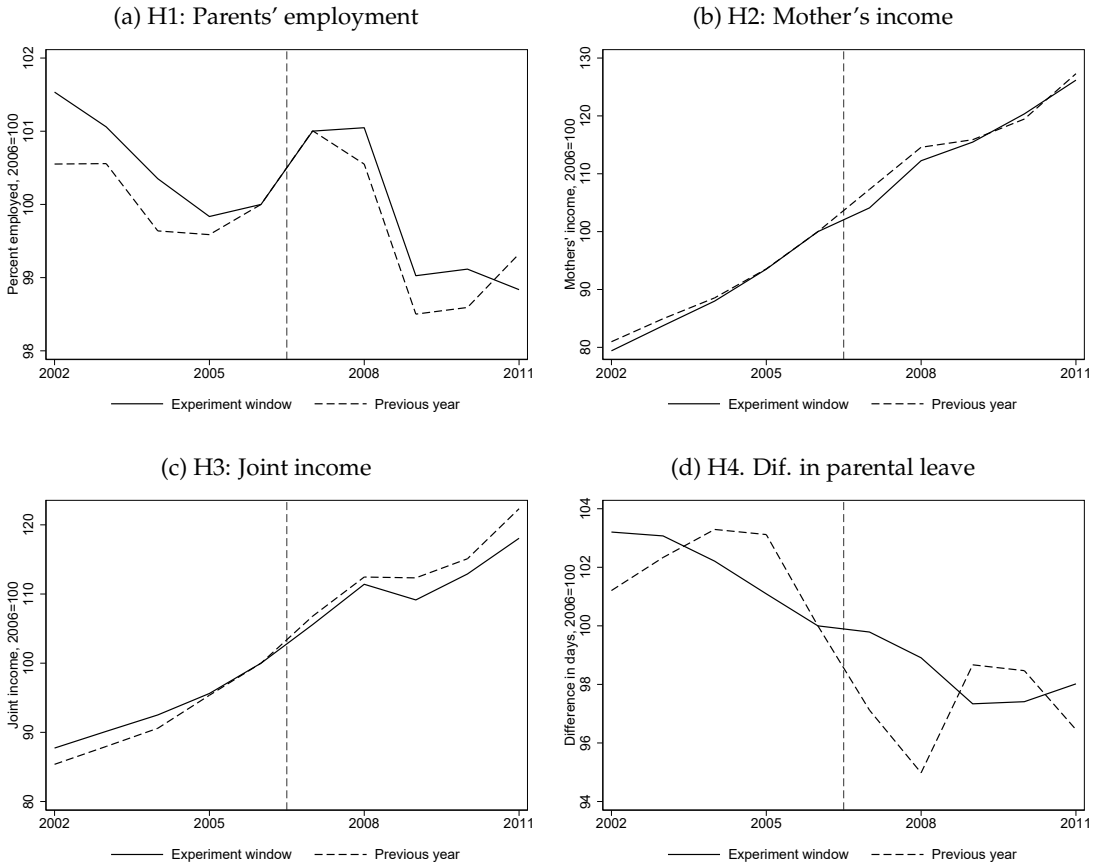

Figure A4: Outcomes before and after the EITC introduction

Note: The figures display our four outcomes, measured during the five years before and after the EITC was introduced. The solid line represents parents with a child born between June of the preceding year ($t-1$) and January in the outcome year (t). The dashed line represents parents with a child born one year earlier (between June, in $t-2$, to January, in $t-1$).

The results are presented in Figure A4. In A4a–c, we see that both employment and income for both parents move rather precisely in parallel, regardless of cohort. Differences in parental leave trend downward for both groups, as shown in A4d, consistent with the general time trend. The trend is stronger for the earlier cohort, however – contrary to what we would observe if the introduction of the EITC had already had this effect prior to treatment. Furthermore, this stronger downward trend starts years before the EITC's introduction. The overall picture appears to be that no marked behavioural adjustments takes place at the time of the EITC's introduction in the demographic group (the parents of infants) under study.

UPPSALA UNIVERSITET

23 mars 2018

Pär Nyman
Forskare

Statsvetenskapliga institutionen
Box 514
SE-751 20 Uppsala

Besöksadress
Gamla torget 6

Telefon:
018-471 33 76
+46 18 471 33 76

www.parnyman.com
par.nyman@statsvet.uu.se

Det här utskicket genomförs i samarbete med Statistiska centralbyrån (SCB).

Till nyblivna föräldrar

Har ni funderat på hur fördelningen av föräldraledigheten påverkar hur stort jobbskatteavdrag ni får? Det handlar faktiskt om ganska stora summor. I ett projekt vid Uppsala universitet informerar vi nyblivna föräldrar om hur hushållskassan påverkas av att en person är ledig från arbetet under ett helt kalenderår.

I Sverige får alla som arbetar ett jobbskatteavdrag. Storleken på avdraget ökar snabbt med arbetsinkomster som under ett kalenderår uppgår till 135 000 kr. Med en sådan årslön får man ungefär 15 000 kr mer i jobbskatteavdrag jämfört med den som är ledig hela året. **För en person som tjänar 35 000 per månad är det därför mycket lönsamt att arbeta minst 4 månader under 2018.**

Storleken på jobbskatteavdraget fortsätter att öka med inkomsten även efter 135 000 kr, men då betydligt långsammare. I regel innebär det att ett par som delar lika på föräldraledigheten under ett kalenderår får ett större jobbskatteavdrag än ett par som låter den ena föräldern vara hemma hela året.

Figuren nedan visar hur storleken på jobbskatteavdraget varierar med din arbetsinkomst under 2018.[1] Det är kanske någonting ni vill ta hänsyn till när ni planerar hur mycket ni ska arbeta under resten av året?

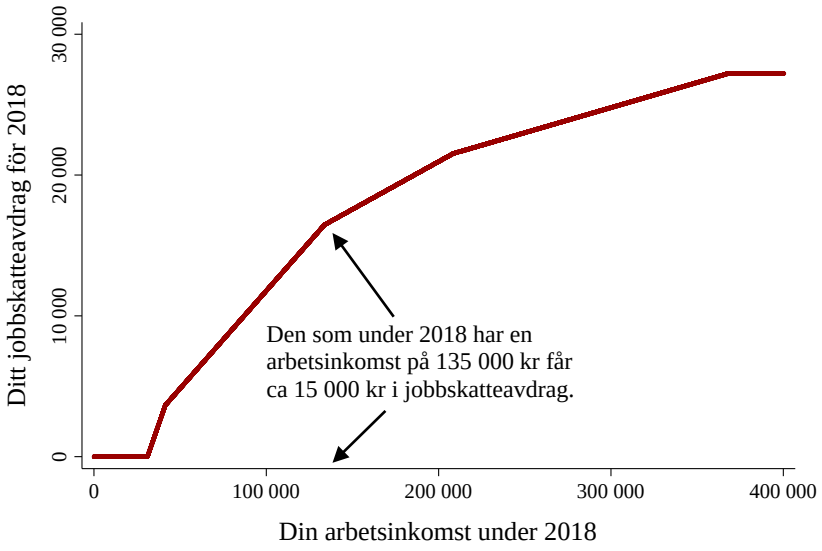

[1] Exakt hur stort jobbskatteavdrag du får beror på dina individuella förutsättningar, men uppskattningarna vi presenterar stämmer relativt väl för alla under 65 år.

Organisationsnr:
202100-2932

Translation of letter

To new parents

Have you thought about how your distribution of parental leave affects how large your tax credit is? It actually involves rather large sums. In a project at Uppsala University, we are informing new parents how their pocketbook is affected when one of them is off work over a whole year.

Everyone who works gets a tax credit in Sweden. The size of the credit increases quickly for labour incomes lower than 135,000 SEK per year. At that annual salary, you will get about 15,000 SEK more in tax credit than will someone who is off work all year. For a person who earns at least 35,000 per month, therefore, it will be highly profitable to work for at least 4 months during 2018.

The size of the tax credit keeps increasing with incomes above 135,000 SEK, but much more slowly. Generally this means that a couple who split their parental leave equally during a calendar year will receive a tax credit larger than that received by a couple where one of the parents stays at home the whole year.

The figure below shows how the size of the tax credit varies with your labour income during 2018.[1] Perhaps this is something you will want to take into account when planning how much to work the rest of the year?

1. The exact size of your tax credit will be influenced by your individual circumstances, but the estimate we present here is relatively accurate for everyone below the age of 65.