Canadian Labour Economics Forum

WORKING PAPER SERIES

Strategic Self-employment and Family Formation

Neil Lloyd (VSE-UBC)

Conference Volume CLEF – 2020, WP # 20

Strategic Self-employment and Family Formation*

* Submission to Canadian Labour Economics Forum, 15 April 2020

Neil Lloyd*
University of British Columbia

April 15, 2020

Abstract: This paper highlights the strategic nature of self-employment decisions made by parents at the event of childbirth. Applying an event-study-design analysis of Canadian administrative data it shows that the increase in female self-employment attributable to childbirth is in fact an increase in joint self-employment. Contrary to the current position in the literature men do adjust their labour supply around childbirth by switching between wage- and self-employment. The paper argues that the observed increase in female self-employment is in fact a real labour supply decision, that adds value to the household business. However, the coordinated behaviour is partly incentivized by the tax savings that income-splitting self-employed households enjoy under an individual tax structure. Using a simulated instrument research design the paper shows that men and women respond symmetrically to a shock in the potential tax savings from income-splitting, after childbirth. In a parallel analysis of longitudinal survey data the paper shows that these mothers join their spouses in the same business, benefiting from increased employment at home, while maintaining above average hours of work.

Keywords: Self-employment, parent penalty, income splitting, childbirth, family labour supply.

*Vancouver School of Economics, 6000 Iona Drive, Vancouver, BC Canada V6T 1L4. E-mail: *neil_lloyd@outlook.com*

**** NOTE TO READER**

This paper represents a preliminary draft of my intended job market paper for later this year. For reasons related to COVID-19 and the closure of Statistics Canada's Research Data Centres I have been unable to release a number of new results as well as revisions of the early results presented here. The text largely tracks my current ongoing analysis leading to a few discrepancies with the included results. In particular, my ongoing analysis uses the proposed definition of self-employment, while the results presented here are from an earlier version. These changes do not affect the overall argument of the paper and I will do my best to be as transparent as possible. I hope to provide you with an updated version as soon as possible. **Finally, I ask that for these reasons you do not cite this version of the paper.**

If you have any questions pertaining to these differences, or any other matters, please do not hesitate to contact me at *neil_lloyd@outlook.com*. Thank you.

1 Introduction

The recent literature has highlighted the lasting impact of family formation for a woman's labour force participation and earnings (Angelov et al., 2016; Kleven et al., 2019a). Even twenty years after a women's first child she is significantly less likely to participate in the labour market, while men show no evidence of a short- or long-run parent penalty on either the earnings or participation dimension. Kleven et al. (2019a) argue that family formation explains what is left of the gender wage gap once you account for human capital differences, and differences in the shape and size of the parent penalty appear to be highly correlated with regional gender norms (Kleven et al., 2019b).

These recent studies have been made possible through increased access to large administrative longitudinal datasets. In this paper I provide similar evidence for Canada using the Longitudinal Administrative Dataset (1982-2016). Using information on the age of the seven youngest children in the family I identify the year that a woman has her first child and estimate changes in labour market outcomes in relation to this event. After corroborating the findings of Kleven et al. (2019a) I shift focus to the choice between wage- and self-employment (see Figure 1). This shift is motivated in the following way. The recent literature has emphasized "flexibility" as a key barrier to female participation in the labour market after childbirth (Goldin, 2014; Goldin and Katz, 2016). Kleven et al. (2019a) show that women select into more flexible forms of employment and even switch to 'family friendly' firms. Certainly, one would expect self-employment to demonstrate the same qualities given a self-employed worker's ability to more freely allocate their working hours and location. This is the central argument of Jeon and Ostrovsky (2019) recent discussion of female (maternal) self-employment in Canada. Indeed, I show that there is an increase female self-employment of 3 %-points following initial childbirth; an almost doubling from pre-childbirth levels (see Figure 2).

However, the flexibility advantage is inconsistent with the fact that female self-employment rates lag their male counterparts, at all stages in the lifecycle. Indeed gender discrepancies are larger in the self-employed labour market than the wage-paying sector Clain (2000); Hundley (2000). Many scholars have argued that this is because male and female self-employed workers (or enterpreneurship) are motivated by different factors. While men chase higher incomes, women desire flexibility

Georgellis and Wall (2005); Saridakis et al. (2014). This paper suggests that this rather simple comparison is incomplete, because self-employment among married adults often ends in joint self-employment (or co-employment). Indeed, I show it is only women married to self-employed men who demonstrate an increase in self-employment participation after childbirth; contrasting the position taken by Jeon and Ostrovsky (2019). Conditional on paternal self-employment, maternal self-employment increases by 10 %-points during the 10 years after initial childbirth in Canada (see Figure 7). For women married to wage-employed men there is no significant increase in self-employment associated with childbirth.

I find that paternal self-employment is itself related to childbirth (see Figure 2). In contrast to the existing literature which finds little to no evidence that men adjust their labour supply with family formation I find strong evidence of switching between wage and self-employment around child-birth for men in Canada. This switching begins before childbirth, while maternal selection into self-employment is delayed till after childbirth. The discontinuity observed for women can explained by the employment insurance structure in Canada. The increase in male self-employment is proportional to the increase in female self-employment across regions in Canada (see Figure 9 & 10). This pattern suggests that families coordinate their labour supply decisions and that there must be childcare advantages to co-employment. It also suggests that these advantages may not be available to individually self-employed women, or that women face significant barriers to entry in the self-employed labour market.

There are other policy incentives to consider. Primarily, the incentive under an individual tax structure to split income (Schuetze, 2006; Kleven and Waseem, 2013; Bauer et al., 2015). It is certainly true that households where the male spouse is self-employed demonstrate a more equal – even symmetric – relative income distribution with bunching at 50:50 income ratio; especially after childbirth (see Figure 12). I show in accordance with Zinovyeva and Tverdostup (2018) that this bunching explains the discontinuity in the relative income distribution at 0.5 (Bertrand et al., 2015; Doumbia and Goussé, 2019). This is because in Canada – as in any other individual tax jurisdictions - there are potential tax savings to doing so; bringing into question the real nature of

the self-employment labour supply decision I have thus far proposed. Given that my identification of an individual's self-employment status is based on their filing of self-employment income it is possible that the observed pattern simply resembles an accounting phenomenon.

I argue that this is a real labour response on the following grounds. First, these patterns are consistent with co-employment responses in the Canadian Labour Force Survey and the historical trend suggests that this was a labour supply decision before it become an income splitting one. Second, similar results are found in the Survey of Labour Income Dynamics (SLID) which includes both tax reported income and self-reported labour supply. The SLID shows that male self-employment results in a steep increase in industry matching within the household over the lifecycle, as well as an increase in the probability a woman working from home. In fact, self-employed households explain all of the rise in industry matching over the life-cycle: a 4% rise from 16%. Selection into co-employed is therefore a significant increase in the risk exposure of the household; especially as the self-employed sector is largely uninsured (Gruber and Cullen, 1996).

In the administrative data I find evidence to suggest that the self-employment induced by childbirth is value adding and does not simply reflect a splitting of the father's pre-childbirth income. However, there is also evidence of a delay in the timing of real and tax-reported maternal self-employment. The earnings benefit of maternal labour arise 3-4 years after initial childbirth, consistent with the self reported increase in at home work and self-employment in the SLID. Consistent with the immediate increase in tax-reported self-employment, I find that the average household tax rate of women married to self-employed men declines by an additional []% points immediately with childbirth. Thus, there are real, and immediate, tax savings associated with this labour supply path.

I contend that these tax savings - generated through income splitting - subsidize a very real labour supply adjustment. They do not just increase female self-employment - as has been argued (Schuetze, 2006)- but male self-employment too. It is also not the individual tax structure per se that does so, but rather the interaction of this tax structure with the asymmetric income shock of first childbirth and then childcare that creates an incentive for married couples to co-ordinate their

labour supply around co-employment.

Applying a simulated instrument research design I am able to show that parents respond symmetrically to exogenous changes in tax savings from income splitting around childbirth. While men respond to these savings before childbirth, women are only sensitive to them after the event, suggesting that (un)employment insurance (EI) benefits likely exceed the value of these savings (see Figure 15). The identifying variation in the instrument arises from provincial and federal tax changes to the income tax structure, as well as the Canadian Pension Plan (CPP) reform in 1996. The magnitude of the elasticity is sufficient to explain the rise in joint self-employment during the 1990s when the value to income splitting increased but is too small to explain the overall level of co-employment (see Figure 4 & 3). Since the 2000 federal income tax reform these savings have declined substantially while family induced co-employment persists.

I find substantial regional variation in co-employment and develop a model to evaluate how this decision interacts with regional family policies (see Figures 9 & 10). Alberta and Saskatchewan demonstrate the highest selection into co-employment. Throughout the 1990s Alberta had the lowest job-protected leave for new mothers (at 17 weeks), and some of the highest benefits to income splitting. Quebec is the only province which does not demonstrate this strategic behaviour. My model suggests that the expansion of childcare support in 1997, and the flatter tax structure, would explain the absence of co-employment and lower levels of self-employment in Quebec.[1] While I do not rule out culture as an explanatory factor, the regional policy variation seems adequate to explain these differences. Given the connection then between co-employment and income-splitting, it seems plausible that the relationship between cultural norms and relative income bunching in Canada can be explained by clear policy differences (Doumbia and Goussé, 2019).

Cultural gender norms appear in the parent penalty literature too, and have recently been shown to be highly correlated with the size and shape of the maternity earnings and labour supply shock (Kleven et al., 2019b). This paper therefore offers a novel lens through which to examine the role

[1]The higher share of common-law couples is another consideration. Quebec common-law couples have fewer property rights, even relative to common-law couples in the rest of Canada. This income insecurity may give rise to a lower willingness for co-employment (Goussé and Leturcq, 2018). While I include common-law status as a control variable in my empirical models, it does not feature at the forefront of this discussion.

of policy and childcare incentives in determining labour supply decisions at the point of family formation. It highlights the fact that men will, and do, adjust their labour supply around childbirth, by switching into self-employed roles with the intention of co-employment after childbirth. This decision highlights a few important points. The lack of a more substantial paternal adjustment to childbirth may be the outcome of constrained employee contracts. This creates a corner solution within the household; which is demonstrated by the fact that men appear willing to select out of such contracts to gain flexibility for their household through self-employment. A large share of male self-employed workers, may therefore be motivated by the childcare flexibility of self-employment; because self-employment is not always an individual decision. Second, the flexibility benefits of self-employment (as they pertain to childcare) may require household coordination and may therefore be limited to cases of co-employment. This latter point may change as the nature of self-employed ventures adjusts to changing, and more flexible, online marketplaces.

2 Data & Measurement

This paper is based on a parallel analysis of Canadian administrative and household survey data. The Longitudinal Administrative Dataset (LAD) forms the basis of the childbirth event-study-design. I benchmark my measurement of self-employment in the LAD using the monthly Labour Force Survey (LFS) and two household income surveys the Survey of Labour and Income Dynamics (SLID, 1996-2011) and the Canadian Income Survey (CIS, 2012-2016). I also replicate the event-study-design analysis of in the SLID, using its more qualitative information on labour supply to draw inference on the nature of the self-employment decisions taken at childbirth.

2.1 Identifying childbirth in the LAD

The LAD is a longitudinal administrative dataset based on a 20% sample of the T1 Family File (T1FF). It covers the years 1982 through 2016. Sampling is implemented at the individual level from the T1FF and once sampled an individual remains in the sample as long as they are observed in

the T1FF. Imputation is used in cases where a married individual does not file their own tax form, but is associated with their spouse's. Individual sampling means that an sample member's spouse may not be in the sample themselves, but for all married couples one observes the tax details of the spouse for each variable. The LAD includes basic demographics such as age, gender, and census family composition.

Crucial for this project, the T1FF collects information on the ages of the seven youngest children in the home. This can be used to identify the year of initial childbirth for each women. In Canada the federal child support grant creates an incentive for parents to complete tax forms; thereby improving the reliability of such data. One downside is that the T1FF attaches child information to one parent in the file; however, I find that in almost all cases the child is attached to the mother in the LAD dataset. This does mean that one cannot perform a separate event-study-design analysis of men and women, except in the case of married couples. I therefore limit my analysis to joint-filing married and common-law couples, but can also extend my analysis to all women.[2]

The working sample is built by identifying all women 20 years of age and older, regardless of marital status. These women may still reside with their parents at the time of sample entry. The year of first childbirth is then identified, and only women observed at least two years prior to childbirth remain in the sample. For this reason, the sample may have a slightly higher probability of working prior to initial childbirth than a random sample. On average I identify about 20,000 *first time* births a year, or 100,000 when weighted.[3] Canada's annual birth rate is around 300,000 for the respective period, but this *total* represents all births. Given the selection on age, I do believe that I am able to identify the majority of initial births.

NOTE: The results shown in this version of the paper correspond to women whose first child was born between 1992 and 2016. Future versions include all women who are observed at least two years prior to their first childbirth beginning in 1982.

[2] In general, I do not find that selecting on marriage changes the results of the event-study-design in a significant way. Moreover, this paper does primarily concern decisions made by co-habiting couples.

[3] Prior to selection on observation before childbirth.

2.2 Measuring self-employment

In survey data self-employment is typically identified by an individual's self-reported class of employment. Both the LFS and SLID include such a variable; for both an individual's main and second/additional jobs.[4] In the LFS an individual's self-employment status is disaggregated by their incorporation and employer status; where the latter identifies whether or not the self-employment enterprise has employees. In Canada, there has been a clear trend towards incorporated and own-account self-employment (see Figure 4). Own-account self-employment numbers increased sharply in 1996, while the number of employers has gradually decline since the mid-1990's.[5]

Incorporation has increased since the turn of the century as unincorporated self-employment numbers have declined, and (own-account) incorporated numbers have increased. This may be partly explained by the 'race to the bottom' of corporate tax rates (see Figure 5). However, occupation specific regulations matter too: for example, until a 2003 legislative change in Ontario self-employed physicians were not allowed to incorporate in the province, while British Columbian physicians have long been able to do so.[6] Indeed, provincial and federal regulations are an important determinant of both the self-employment and incorporation status of various occupations in Canada. For example, in the late 1980's real estate agencies in Canada were given permission to re-organize around self-employment.

There is no identifier of class of employment in Canada's administrative data and, as such, the identification of self-employed workers presents a measurement problem. In the LAD one observes separate filings of unincorporated self-employment income: grouped by business, professional,

[4]The CIS does not provide a single class of employment variable for an individual's main job. Instead, it identifies whether an employed individuals was an employee, self-employed, or unpaid worker during the past year. Naturally, an individual can report yes to more than one of these. The income structure in in the CIS is similar to the SLID and the LAD and as such can be used to identify receipt of unincorporated self-employment income. There is no reporting of CCPC ownership.

[5]These figures exclude unpaid family workers, who are typically women married to self-employed men. The number of unpaid family workers has been in a steady decline for the last four decades, and since the turn of century is largely obsolete.

[6]Similarly, Ontario gave law firms permission to incorporate in 2001. Incorporation has become a popular channel through which high income self-employed workers are able to defer taxes and increase their lifetime savings bay saving within the firm. The are little to know tax benefits if all corporate profits are passed through to the owner in a given year, as the mark-up of dividends adjusts for the corporate and personal income tax rate differential.

farming, fishing, and other self-employment income. Both gross and net income is observed, where the latter is taxable. Incorporated income is harder to identify, as corporate ownership of a Canadian Control Private Corporation (CCPC) is only observed in the LAD from 2000, and not observed for a spouse. Corporate income enters the household through one of two channels: dividend or wage income. In the case of wages, I rely on the fact that wage income paid to a shareholder of a CCPC owner (with equity $> 40\%$) is uninsured (i.e. ineligible for EI contributions). Likewise, wage payments paid to an individual within arm's length are also uninsured; this includes an individual's spouse.[7] I therefore use the observation of *large* uninsured wage receipts, as well as *large* dividend receipts to identify possible incorporation.[8]

Once I have identified potential sources of self-employment income, I define an individual as self-employed if their total self-employment income is more than half their total employment income. However, to account for the fact that many self-employed file negative taxable income I use the following definition.

$$Q = \mathbf{1}\left\{\frac{|income_q|}{|income_q| + income_w} > 0.5\right\}$$

In this way if an individual has insured wage receipts of \$40,000, but a self-employment loss of \$100,000 there are included as self-employed.[9] I also hope to avoid identifying cases where a wage-employee files a once-off amount of self-employment income from a short freelance job.[10]

The final challenge remains how to select a threshold with which to select uninsured wage

[7] The following link outlines Canada's policy concerning the hiring of family. There are circumstances in which the employment of a relative may be deemed at arm's length.

[8] From 2005 one can separately identify dividends received from a small and large, as after 2005 the two types receive a different mark-up to account for the small business tax exemption. As such, dividends from small and large firms are filed separately, helping to separate out dividends from smaller CCPC businesses and larger (potentially publicly traded) companies. However, this cannot be done for years prior to 2005 and therefore is not a viable means of consistently measuring self-employment.

[9] I am limited in my use of gross unincorporated income, because the LAD's measure of a couple's total grouse income is not the sum, but rather the max of the couple. This is because the LAD assumes that in all cases of joint unincorporated income (within the same category) they are derived from the same source.

[10] I retain a loser measure of self-employment income - a flag for self-employment - to use as a robustness check at later stages in the analysis.

$$Q_{flag} = \mathbf{1}\left\{income_q \neq 0\right\}$$

receipts and dividend receipts as self-employment income. Here I make use of the SLID's matched survey and tax data. I pick the threshold that gives me the best trade off between type one and two errors.[11] For wage receipts this is around $5,000 (in 2016 CAD), while for dividends it's closer to $8,000. In addition, I check the type one error for wage-employment and find that this is minimized by setting the wage threshold for uninsured wages to $0. According to this matched sample uninsured wages do not identify wage-employment. On this basis, uninsured wage receipts below the threshold are not regarded as employment income in this analysis. Regarding, dividend income I find that including dividends as a form of employment income vastly improves the type two error on employment status. That is, if you do not include dividends as a form of employment income you underestimate the number of employed individuals in the sample.

After selecting these thresholds I find that female type one errors are correlated with the reported self-employment status of a spouse. This is to be expected as my measure of self-employment is more broadly a measure of employment within a household business. It also suggests that self-reported self-employment (as per an individual's class of employment) potentially underestimates financial inclusion within a household business; be it through employment or income receipt alone. In this regard, the LAD provides a superior source of analysis when looking study self-employment and co-employment amongst married couples. This measurement difference between the LAD and LFS may also be exacerbated by the increased trend towards incorporation of small businesses. Figure 3 shows the LFS's measure of co-employment (joint self-employment) in decline since 1997 along with unincorporated self-employment (see Figure 4). In contrast, my measure of co-employment from the LAD does not decline, suggesting that incorporation may result in different reporting on class of employment within a married household.[12]

NOTE: The following results correspond to an earlier version, which made use of a slightly

[11] In this context, a type one error means that I identify an individual as self-employed when they are not, while a type two error means I fail to identify an individual as self-employed who is.

[12] Suppose an incorporated business owner employs their spouse paying them a wage. In the LFS this spouse may report their class of employment as private sector wage employee. If both individuals work within the same unincorporated business all business profits would automatically flow to the household without a need to pay out wages or dividends. In this scenario neither worker receives a fixed wage, and it is plausible that both will self identify as self-employed. That is, any hierarchy within a family business may result in different reporting dependent on incorporation status.

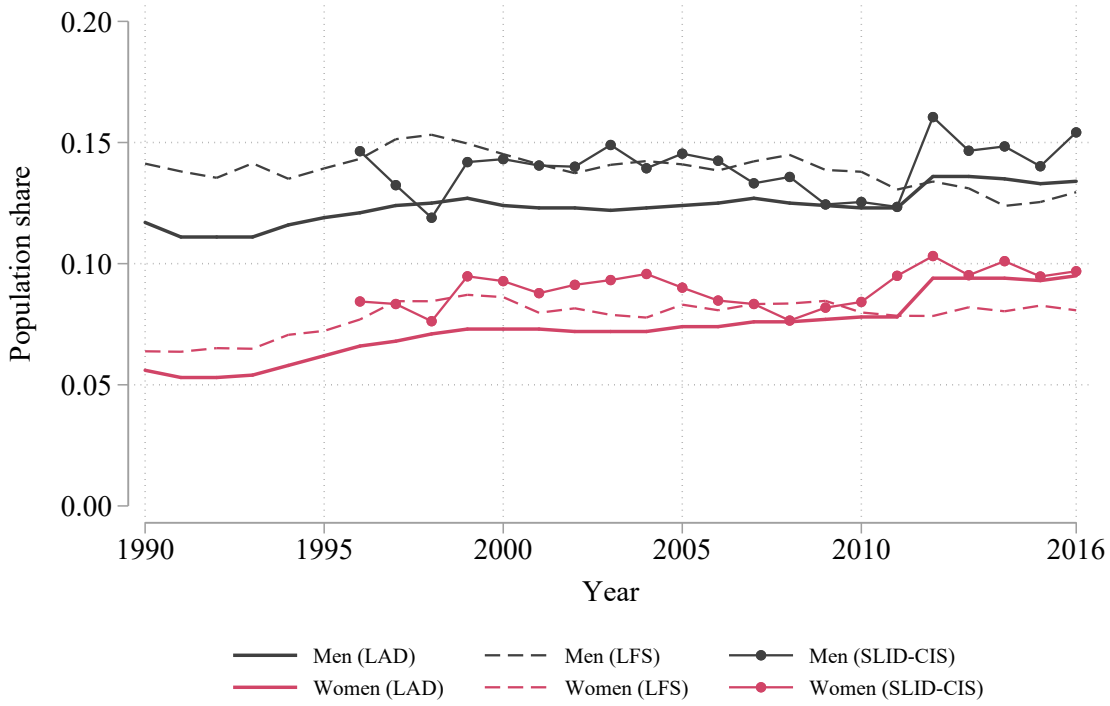

Figure 1: Comparison of self-employment definitions across LAD, LFS and SLID/CIS for men and women aged 25-44.

different definition of self-employment. In particular, a threshold was put on the absolute threshold of self-employment income. I suspect that this is why my measure underestimates the LFS's self-employment rate. Future releases correspond to the above discussion.

Figure 1 compares my measure of self-employment from the LAD with that of the LFS and SLID/CIS. Until 2010, the LAD's self-employment rate underestimates that of the LFS, but the indexes follow a similar trend. In particular, the rise in self-employment through the 1990s. In the LAD I measure a discontinuous rise in (incorporated) self-employment in 2012.[13] This increase does not appear in the LFS, but does appear in the SLID and CIS surveys. So long as this is just measurement error that is unrelated to childbirth then it will be absorbed by the year fixed-effects in the models I estimate.

This divergence between the LAD, SLID/CIS and LFS may also reflect important differences between the underlying measures. The LAD measure identifies income receipts from a arm's length business, be it unincorporated or incorporated. An additional important difference is that the LFS identifies monthly employment status and class of employment of the individual's main job. The SLID/CIS identifies reported participation in self-employment at any point in the past (tax) year; an individual could have been self-employed and a paid employee during the survey year. In contrast, the LAD definition identifies the largest annual income source.

3 Self-employment and family formation

This paper uses an event-study-design methodology to describe selection into self-employment as a function of family formation. This methodology has been popularized by Henrik Kleven and co-authors in a recent series of papers concerning the parent penalty (Kleven et al., 2019a,2019b). The estimating model is simple, but flexible, including a set of dummy variables for event-time

[13] This jump in the series is the result of an approximate 2% point increase in the incidence of uninsured wage income. When compared with CCPC ownership rates, these post 2012 figures best reflect ownership of an incorporation. I suspect, that the pre-2012 figures slightly underestimate incorporated self-employment. I am hoping to incorporate CCPC ownership into my measure in the future; however, for now it is only observable for the individual, and not their spouse.

(years since first birth), age, and year.

$$Y_{ist}^g = \sum_{j \neq -1} \alpha_j^g \cdot \mathbf{1}\{j = s\} + \sum_k \beta_k^g \cdot \mathbf{1}\{k = age_{it}\} + \delta_t^g + \varepsilon_{its}^g$$

Each event-time coefficient α_j is then rescaled by the expected value of the outcome variable under the counterfactual that the event did not occur (denoted here by $\tilde{Y}_{its}^g$). This is the standard methodology for estimating relative changes in such models, as opposed to the use of a log transformed outcome variable. It has the added benefit that the model can be estimated using level outcomes - such as annual taxable income - without the exclusion of 0 values. This is useful in the context of the LAD data as I do not observe weeks or hours worked, part-time or full-time status. I can only observe participation based on the filing of employment income.[14] As such it is better to model level employment income (including 0 outcomes) which will encapsulate both wage and labour supply responses to childbirth.

$$P_t^g = \frac{\alpha_s^g}{E\left[\tilde{Y}_{its}^g | t\right]}$$

In this paper I apply this methodology to earnings and discrete employment outcomes measured by the presence of non-zero employment earnings. Figure 6 shows the relative change in employment for new mothers and fathers in Canada, using event-time -1 as the base period.

Next, I consider the decision to be employed in the wage-paying sector versus self-employed. To do so, I repeat the same analysis replace the discrete outcome of employment with an indicator of wage and self-employment status.[15] Figure 2 shows an approximate 3% increase in maternal self-employment amongst women after childbirth. For their spouses there is a corresponding increase in self-employment which precedes childbirth, and evident switching from wage- to self-employment. When comparing across provinces we see that the magnitude of the increase in female self-employment after childbirth is met with a corresponding change in paternal self-employment

[14] I can however identify the number of T4 receipts an

[15] NOTE: in this version of the results the two outcomes are mutually exclusive: an individual is either wage or self-employed. Future results will relax this assumption.

(see Figure 9); suggesting that this is likely a joint decision of selection into co-employment.

Indeed, Figure 7 shows that when you separately compare the self-employment decisions of women married to wage- and self-employed men it is only those who are married to self-employed men that demonstrate an increase in self-employment associated with the event of childbirth. For women married to self-employed men there is no evidence that selection into self-employment is induced by the event of childbirth. This is not to say that there are no self-employed women married to wage-employed men, or that such self-employed women do not also benefit from the flexibility of self-employment when they have children; only that such cases of self-employment do not appear to be correlated with the event of childbirth. Given that selection into self-employment increases with age, it may be that women who individually select into self-employment after childbirth do so at an age that is common for all women to make such a switch.

The delayed response of women to co-employment is easily rationalized in the context of Canada's federal employment insurance (EI) structure under which self-employment earnings are not insured. In the Quebec the introduction of a provincial based Quebec Parental Insurance Plan (QPIP) in 2006 extended coverage to self-employed workers, and since 2011 self-employed workers in the rest of Canada have had the option of voluntarily contributing to the federal EI system. However, for the majority of Canadians, and the majority of the time period covered in these figures self-employment remained uninsured. For this reason, a mother must be wage-employed leading up to the event of childbirth in order to receive EI benefits. Moreover, this same EI structure suggests that women may delay co-employment until after a second child. Indeed, I find that in the SLID co-employment increases sharply with the second and third child. This may explain the gradual increase in co-employment after the event of initial child. I explore these policy interactions more in Section 5.

3.1 On the *real* nature of co-employment

Given that my identification of self-employment in the LAD is based on observed income tax filings, it begs the question as to whether the observed selection into co-employment is simply an

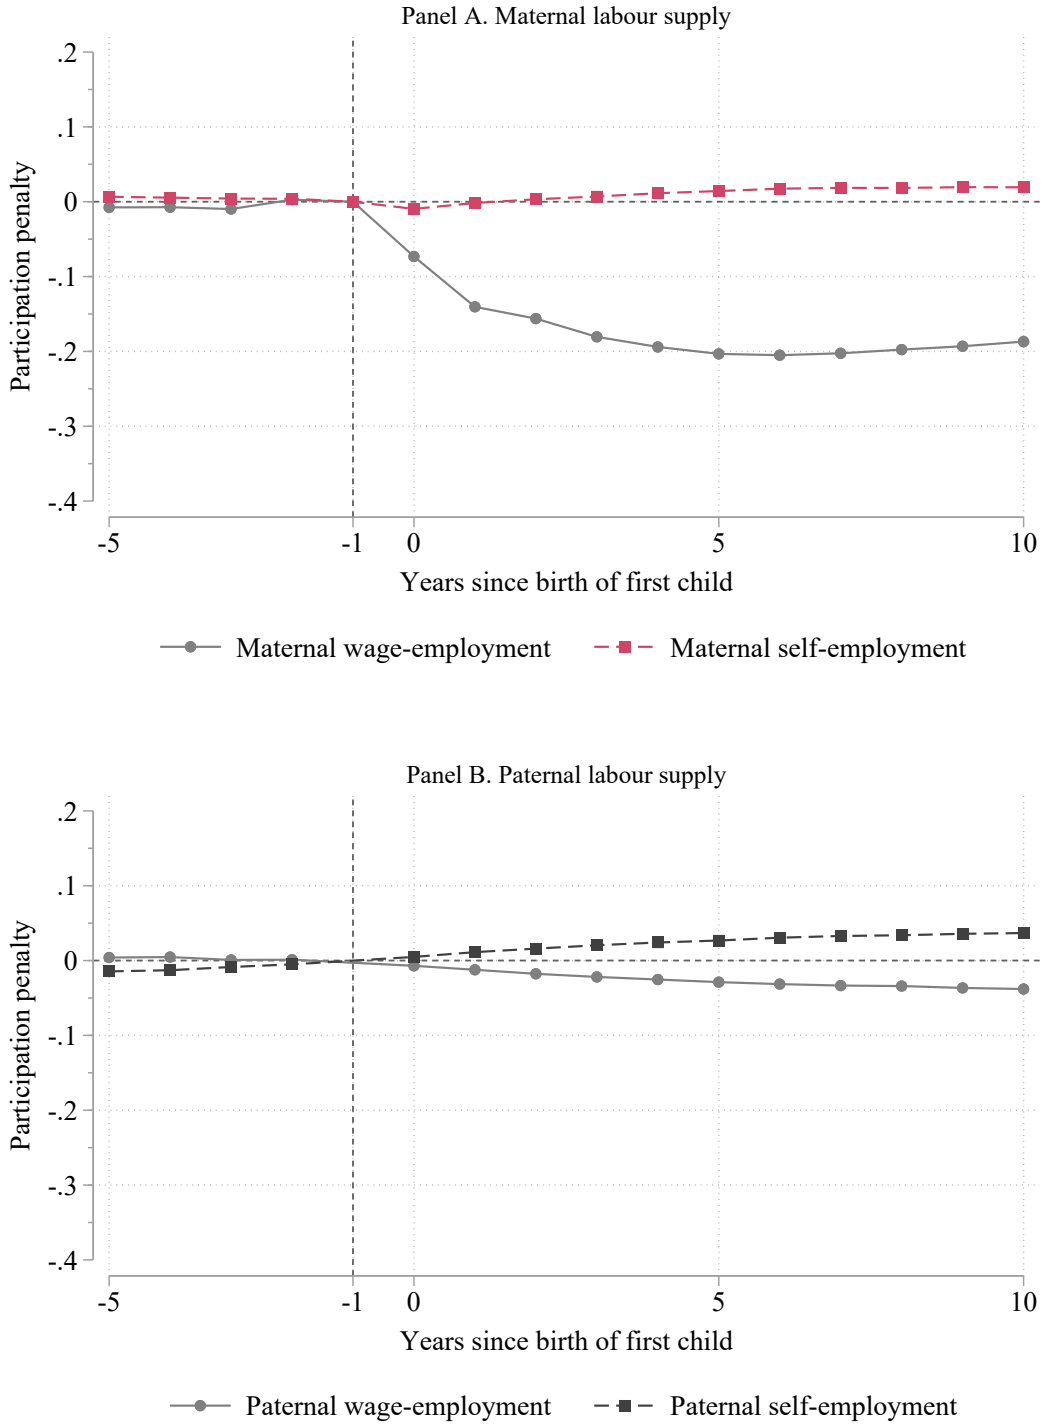

Figure 2: Change in the probability of self-employment and wage-employment relative to the counterfactual probability of employment. Estimated using the LAD (1990-2016)

accounting trick for self-employed households to lower their tax burden. Certainly, if a spouse is otherwise out of the labour market, shifting income under an individual tax structure will do so. This is the view taken in Schuetze (2006) of Canada's self-employed: one of 'fake' job creation. This is also consistent with other findings in the public policy literature which suggests that the self-employed have far more elastic behaviour (Saez, 2010) and do engage in tax avoidance behaviour (Kleven and Waseem, 2013). I argue that in the context of family formation this is not the case.[16]

In the LFS the *difference* in co-employment rate of married women with and without children increased sharply during the 1990's, and has remained remarkably stable since the early 2000's (see Figure 3). When you include women unpaid family workers (who are almost all women married to self-employed men) the family formation gap appears to be remarkably constant throughout the 1976-2016 period. This suggests, that co-employment was a labour supply decision before it became an accounting one, and the LAD may historically underestimate the true extent of this decision.[17]

When measured by reported labour force status in the SLID, female self-employment is increasing in family formation for women married to self-employed men as in the LAD (see Table 1).[18] Switching to a measure of self-employment using the observation of self-employed tax receipts does not change this result. One would expect the coefficient on family size to increase if the tax reported self-employment over estimates the true labour supply decision. In addition, matching on industry, and to a lesser extent occupation, increases in a monotonic fashion with childbirth suggesting that this is co-employment in the same firm, and not a spillover into a second venture as has been suggested might be the case (see Table 2). The scale of this increase is proportional to

[16]Even if the observed maternal self-employment after childbirth is 'fake', this still represents a very real labour supply shift along two dimensions. First, both parents' decisions cannot be an accounting phenomenon. If it is believed that the observed maternal self-employment income is 'fake' income-splitting and the mothers remain inactive it remains significant that fathers are selecting into self-employment for this tax windfall. Second, the increased selection of women married to self-employed men out of their wage paying job after child-birth is identifiable and economically significant.

[17]This decline in unpaid family workers appears in the US too; however, in Canada it may also be underpinned by policy changes. The 1988 tax reform removed a family tax benefit that automatically generated tax savings for *all* married couples through income splitting (Crossley and Jeon, 2007). For this reason, there was less of a tax incentive for two co-working spouses to both report taxable income. The 1988 tax reform therefore increases the incentive for tax-reported co-employment; as in LaLumia (2008).

[18]NOTE: This current specification does not replicate the ESD model. Future versions will include a ESD using the longitudinal component of the SLID.

the ESD results.

Indeed, Figure 11 shows that the age profile of industry matching amongst two-earner married couples is essentially flat for wage-employed households. Self-employed households (male spouse self-employed) explain all of the increase in industry matching among married couples over the life-cycle. A consequence of this will be increased exposure to industry specific shocks, and the absence of an "added worker effect" (Gruber and Cullen, 1996). Figure 11 also show that women married to self-employed mean are far more likely shift to working from home, suggesting that many co-employed couples take advantage of the flexibility benefits associated with co-employment with regards to childcare. On the more intensive margins of weeks worked, weekly hours worked, and part-time status there is no evidence in the SLID data to suggest that the labour supply of these women is less intensive than those married to wage-employed men. The main differentiating feature of their labour supply is that they work in the same business and often from home.

There is some evidence of a missed timing between the immediate increase in maternal self-employment in the LAD, and the reported change in labour supply in the SLID. Table 1 suggests that conditional on a woman's spouse being self-employed, the increase in maternal self-employment associated with the presence of a single child is not significantly different from that of a woman married to a wage-employed man. This may be because of the difference between the two estimating equations: in Table 1 I replace the event-time of childbirth with number of children[19]. However, in results not shown here I am able to replicate the ESD model using the SLID's reported labour variables and I find that maternal self-employment increases only 3 years after a women's first child is born; consistent with the significant second child result in the cross-sectional data. Thus, tax reported self-employment may precede the real maternal labour supply shift.

3.2 The long-run payoff to co-employment

Examining the financial benefit to co-employment in the short and medium term is not only a way of assessing whether this joint strategy of labour supply and family formation pays off, but

[19] These equations are estimated using the cross sectional public release files, and as such event-time is not observed.

also an additional way of assessing the real nature of the decision. If the observed pattern of co-employment is simply an accounting trick it should result in a redistribution of the father's income - a simple 'splitting of the same pie' - not an increases the total employment income of the household. This next section provides evidence that maternal self-employment is indeed value adding in co-employed households. The same event-study-design methodology can be used calculate the present-value cost of child-birth. Davis and von Wachter (2011) estimate the present value cost of job loss by aggregating the event-time coefficients after the event of a lay off. I apply this methodology here to calculate the difference in the income loss associated with childbirth between self-employed and wage-employed households.

NOTE: These results are still waiting to be released from the RDC.

The results suggest that 10 years after childbirth households where the spouse is self-employed earn an estimated $[] more per annum relative to the average wage-employed household. This increase corresponds with the increased selection of mothers into co-employment, suggesting that the labour force contribution is value adding.[20] However, the increase in employment earnings is not immediate. Self-employed households appear to be better around 3-4 years after childbirth, consistent with the above findings from the SLID that the real labour supply effect is not immediate.

What is immediate are the tax savings. I estimate that the average household tax rate of self-employed households falls by additional []% points relative to wage-employed households immediately with the birth of a first child. Such savings are consistent with the my simulated savings from income splitting (see Figure 14). Taken together, these results suggest that there may be a timing difference between selection into real co-employment and tax reported co-employment. This timing difference is only observed for the mother.

It remains significant that fathers choose to switch into self-employment with the birth of their first child. After all, it cannot be the case that both the maternal and paternal self-employment

[20] However, this does not account for the dynamic selection of the spouse into self-employment. Spouses selecting into self-employment after childbirth may have a higher pre-childbirth earnings, to the extent that this is correlated with labour supply decisions of the mother it will bias the simple ESD estimates. To deal with this I use a matching estimator to compare co-employed spouses 6 years after childbirth with similar wage-employed households, matching on pre-childbirth (event-time = -2) income of the spouse, and both parents. **NOTE:** Results to follow.

decision are mere tax reporting changes. In Section 5 I develop a simple model that gives rise to this exact result. In the model households select into co-employment because of the increased flexibility (i.e. lower cost of childcare), and fathers choose to do so leading up to childbirth in order to take advantage of income splitting tax savings that arise when a money gives birth and there is a short term loss of income.

4 Tax incentivise for co-employment

This paper contends that in the medium term the observed pattern of co-employment is a very real labour supply phenomenon. That is, women who take up self-employment together with their spouse after childbirth supply positive hours of work to the business and by doing so add value to the enterprise and household. However, such a position does not rule out the possibility that savings from the splitting of self-employment income acts as an incentivize to do so; thereby increasing both the level of co-employment and (male) self-employment in the economy. In this section I develope a simulated instrument research design that exploits policy variation in the savings from income splitting. The results find that the savings from income splitting do indeed increase paternal self-employment, even before childbirth. Maternal self-employment does not respond to these savings before childbirth, but the elasticity increases monotonically after childbirth. Three years from the event of initial childbirth the maternal and paternal elasticities are symmetric which is consistent with strategic behaviour.

The magnitude of the elasticity is sufficient to explain approximately 20-30% of the increase in co-employment after childbirth. This increase is proportional to the average difference in male self-employment in Canada and the US during this period, suggesting that individual tax structures may subsidize male self-employment. In what follows I first discuss the relevant regulatory framework and provide an approximation of potential income splitting tax savings. I then develop a measure for these savings, and design an identification strategy based on changes in to Canada's provincial income tax structures.

4.1 Regulatory setting

Within the context of a married household, income splitting refers to the shifting reported income between spouses (or dependents) for the sake of reducing the household's overall tax burden. For this reason it is typically associated with tax jurisdictions that have an individual tax structure, as under a joint tax system the holder of the income is largely irrelevant (for joint filers). Income splitting is therefore the default in jurisdictions such as the US and Germany. Within the Canadian context, shifting income to a family member purely for the sake of lowering one's tax burden is deemed tax avoidance.[21] Moreover, the third party reporting of wage income means that income splitting is largely confined to the domain of self-employed income (Duff, 1999).[22]

Acting to dissuade the shifting of income through child dependents, the Canadian government adjusted the tax treatment of dividends paid to the children of incorporated business owners in 2000 (Bauer et al., 2015; Macnaughton and Matthews, 1999; Donnelly et al., 2000). However, limiting income splitting between adults in a household, and in particular married or common-law couples, is more difficult. Joint ownership of a business, even employing a spouse, is certainly legal. In contexts where both spouses participate in a household business it may not be obvious how to allocate the taxable profits in accordance with individual effort. The simplest solution may be the only practical solution: to split the business income. This just so happens to be the most tax-efficient solution too. Finally, payments to a spouse for services rendered need not be at the market rate but must be "reasonable"; leaving room for tax optimization.

The question remains then, how much can a couple save through income splitting? The answer depends on a number of factors, including the households relative and absolute incomes, as well as

[21] Income splitting is governed by the following act:

> *"A person who, after the first day of August, 1917, has reduced his income by the transfer or assignment of any real or personal, movable or immovable property, to such person's wife or husband, as the case may be, or to any member of the family of such person, shall, nevertheless, be liable to be taxed as if such transfer or assignment had not been made, unless the Minister is satisfied that such transfer or assignment was not made for the purpose of evading the taxes imposed under this Act or any part thereof. "* 4(4) of the Income War Tax Act, S.C. 1917, c. 28

[22] There are many ways in which households may act to lower their average tax burden by shifting non-employment income - such as investment or rental income - even in wage-employed households. This paper concerns the shifting of employment income.

the progressivity of the tax structure. It also fundamentally depends on the incorporation status of the business. For my sample, I estimate that the average household tax rate of self-employed households falls by an additional []% after childbirth relative to wage-employed households, consistent with many of the simulations I've performed. Only in cases where one spouse has no income, can the savings be as high as 10%.[23]

Consider the example of an unincorporated business. In 1990 an Ontario household with a combined income of $61,060 ($100,000 in 2016 CAD's) and a 30:70 income split could lower their income tax burden by $1,200 by splitting their income 50:50; a 2.7% increase in after tax income. In 2000 the same household with a combined earnings of $74,300 ($100,000 in 2016 CAD's) would save only $152; an increase of only 0.3% in disposable income. Similarly, in 2010 the potential savings are only 0.4% of disposable income. In contrast, a wealthier Ontario household with a combined income of $150,000 (in 2016 CAD's) and the same 30:70 split could increase their disposable income 0.8% in 1990, 2.3% in 2000, and 1.7% in 2010.

It is clear that the value of income splitting has changed dramatically during the past four decades and that these changes are highly dependent on the households absolute and relative income. This variation also arises from a number of key policy changes. First, as a component of the extensive tax reform of 1988 the federal government removed a number of family tax benefits that effectively generated income splitting tax savings for married Canadian households (Crossley and Jeon, 2007). Then, beginning in 1996 the government embarked on an extensive reform of the Canadian Pension Plan. The reform introduced a staggered increase to the contribution rate from 2.8% in 1996 to 4.95% in 2001.[24] CPP contributions are paid on all income above a basic exempt amount and below a maximum pensionable threshold (adjusted annually): $3,500 and $54,900 in 2016, respectively. Because the self-employed pay both the employer and employee component of the CPP the reform increased their contribution rate from 5.6% to 9.9%; a 4.3% increase on the marginal tax rate of all pensionable earnings.[25] This dramatically changed the progressivity of tax

[23]Such savings were only possible prior to the 2000 federal income tax reform.

[24]In 1980 the contribution rate was 1.8% an increased by 0.1% on annual basis between 1986 and 1996. Until the recent reform in 2019, the contribution rate has remained constant at 4.95%.

[25]The 1996 reform of the CPP was required to ensure the sustainability of the fund, but did not increase the pension

structure for the self-employed.[26]

A third important source of variation is the federal income tax reform of 2001, which coincided with the introduction of an independent provincial tax structure. In 2001 the federal government introduced a third income tax bracket above $100,000 (at 29%), lowering the marginal tax rate on income between $60,000 and $100,000 from 29% to 26%. The combined effect of this reform and the CPP reform is demonstrated by Figure 13. Between 2000 and 2001 Canada's provinces adopted reforms to switch from a tax-on-tax structure to a fully independent income tax structure. In particular, these reforms allowed provinces to set tax brackets and rates independent of the federal structure. For example, Alberta adopted a flat 10% provincial tax rate. For self-employed Albertans this flat rate exacerbated the lack of progressivity in the underlying federal (and CPP) structure, resulting in a very flat income tax rate with regressive sections. The result was a very large drop in the savings from income splitting.

Estimating the potential savings from income splitting for an incorporated business owner is more complex. The channel through which income is parsed from the business to the household and the extent to which income is saved within the firm will matter. Parsing income through wages will trigger CPP payments on both the firm and employee side, while dividends are not considered employment income and are therefore exempt from CPP contributions.[27] The income tax implications of wages versus dividends are relatively minor. Wages are deducted from corporate profits and taxed at individual income tax rate, while dividends are first adjusted to take into account the corporate taxes that have been paid before being taxed at the individual level. The net result should be that regardless of the channel the income is taxed at the same rate.[28] For

benefits of participants. Hence, there was no corresponding increase to the lifetime savings of contributing workers, only an increase in contributions, amounting to an effective tax increase.

[26] In certain provinces the regressivity that the CPP introduces can exceed the difference in the marginal tax rate brackets, resulting in a locally regressive tax structure and higher taxes under income splitting (see Figure 14). However, these losses are generally very small, and a 50:50 income split generally yields the lowest tax burden.

[27] In Section 2 I make the case that excluding dividend income as a signal of tax payer's employment status results in the underestimation of total employment. Avoiding CPP payments may be a primary reason why incorporated business owners choose to allocate profits through dividends.

[28] Slippage can occur when at the provincial level the dividend adjustment rates are not corrected for changes in the provincial or federal corporate tax rates. There therefore may be periods where arbitrage is possible; however, these savings are typically not large.

small business owners, the primary tax benefit to incorporation lies in long term savings. If a household can afford to leave profits within the incorporated firm it can accumulate interest at the lower (corporate) tax rate, which in the long run generates significantly higher returns. For this reason, incorporation is popular among high income self-employed professionals such as doctors and lawyers. However, I am unable to identify these savings in my data, and as such, focus my attention on the savings from income splitting within an unincorporated firm.[29]

4.2 Identification

This path of enquiry must overcome a number of measurement (observation) and identification challenges. First and foremost, the extent to which a household has split their income (for tax purposes) is unobserved. One only observes the filed version of a couple's income statement, not the business account's prior to distribution. To account for this, I develop an identification strategy based on a household's predicted path of income (through the event of childbirth) and restrict my sample to couples that were both engaged in the wage-paying sector prior to childbirth. I then simulate the value of income splitting as the proportional change in after-tax income were a household to split their income as unincorporated self-employed filers.

Given that this measure of savings is itself endogenous to the household's income structure - which includes the labour market decisions they make at childbirth - it cannot be used directly as a source of identifying variation. Adopting an approach similar to Gruber and Saez (2002) I then simulate the same savings under a particular historical counterfactual: the income tax structure of a reference period prior to childbirth. Including this historical counterfactual as an additional control in the model, the identifying variation from the contemporaneous measure arises only from tax changes experienced during the event of childbirth.

Let y^1 and y^2 be the true taxable income of individuals 1 and 2 in the household respectively, prior to any redistribution. Then the savings from income splitting can be measured as the change in after income taxes under the given earnings structure and under income splitting. This is essentially

[29] One could also think of savings from income splitting under unincorporation as the lower bar for incorporation.

a modified income effect (Gruber and Saez, 2002). I denote this measure as Δ. The Δ function varies with time, while the arguments are denoted with both a time (t) and event-time (s) subscript.

$$\Delta_t \left(y^1_{ts}, y^2_{ts} \right) = \mathrm{T}_t(y^1_{ts}) + \mathrm{T}_t(y^2_{ts}) - 2\mathrm{T}_t(0.5(y^1_{ts} + y^2_{ts}))$$

From the outset it is important to emphasize that the value of Δ will depend on both individuals' incomes. That is, we could write $\Delta_t(y^1, y^2) = \Delta_t(\mu, y)$, where μ is the relative income of the household, and y the absolute income of the household. The value of Δ therefore depends on both the absolute and relative incomes within the household. This will result in a number of identification challenges, as it is insufficient to know just one of these dimensions.

Consider then the change in Δ around childbirth. For notational simplicity I will use event-time 0 as the pre-childbirth reference period. Relative to this base period the change in delta at time t and event-time $s > 0$ has both a tax and earnings component, distinguishable by the counterfactual value $\Delta_{t-s} \left(y^1_{ts}, y^2_{ts} \right)$. From the perspective of the household in the base period, the value of Δ_{t-s} may be predictable as the policy structure is fixed.[30] A household considering its labour supply decisions around family formation may be able to predict this counterfactual value under variety of labour supply scenarios, thereby informing the observed labour supply path. Given that childbirth is typically associated with an asymmetric shock to income one expects that in the majority of households the 'labour supply' component will be positive. While this variation is certainly key to understanding the observed pattern of co-employment, it is not a useful source of exogenous identifying variation.

[30] The alternative counterfactual $\Delta_t \left(y^1_{t-s,0}, y^2_{t-s,0} \right)$ is unintuitive from the perspective of the household as it assumes an ability to foresee policy changes.

$$\Delta_t\left(y_{ts}^1, y_{ts}^2\right) - \Delta_{t-s}\left(y_{t-s,0}^1, y_{t-s,0}^2\right) = \underbrace{\Delta_t\left(y_{ts}^1, y_{ts}^2\right) - \Delta_{t-s}\left(y_{ts}^1, y_{ts}^2\right)}_{\text{tax structure}} + \underbrace{\Delta_{t-s}\left(y_{ts}^1, y_{ts}^2\right) - \Delta_{t-s}\left(y_{t-s,0}^1, y_{t-s,0}^2\right)}_{\text{labour supply}}$$

From the perspective of the household in event-time s, any deviation in the realized value of Δ_t relative to the counterfactual constitutes a policy shock (denoted above as the 'tax structure' component). Note, this is only a shock relative to the base period. By simulating both realized and counterfactual value of Δ for each household one can then examine the correlation between the decision to enter self-employment and the realized value of Δ_t, conditional on the counterfactual value of Δ_{t-s}. See Figure 14 for a simulation of this variation. The use of such simulated counterfactual is fairly common in the taxable income elasticity literature when attempting to instrument for the income effect of a tax change (Gruber and Saez, 2002).

At what incomes should one simulate both the realized and counterfactual value of Δ? Again, it cannot be the observed values as these are endogenous to the labour supply decisions taken in a given year as well as any unobserved income redistribution. Suppose that the income path of each individual relative to an observed reference period follows the path,

$$y_{its}^g - y_{it-s,0}^g = v_{ts}^g + \tilde{v}_{its}^g$$

The change in income between event-time s and the base period includes both a common and idiosyncratic component. Then we can replace y_{its}^g with

$$\tilde{y}_{its}^g = y_{it-s,0}^g + v_{ts}^g$$

This combination of base period characteristics - including the household's base period income

share - and a common (event-)time shock has elements of a Bartick style instrument, which is generally constructed using the inner product of base-period shares and a common growth rate (Goldsmith-Pinkham et al., 2019). However, in this context the substitution takes place within a highly non-linear function, Δ. However, given the stepped nature of the tax structure, the derivative of Δ with respect to a single argument will be locally linear (Gruber and Saez, 2002).

$$\begin{aligned}\Delta_t\left(y^1_{its}, y^2_{its}\right) &= \Delta_t\left(\tilde{y}^1_{its}, y^2_{its}\right) + \frac{\partial \Delta_t\left(\tilde{y}^1_{its}, y^2_{its}\right)}{\partial y^1}\tilde{v}^1_{its} + o(1) \\ &\approxeq \Delta_t\left(\tilde{y}^1_{its}, y^2_{ts}\right) + \delta^1_t \tilde{v}^1_{its}\end{aligned}$$

By this approximation I contend that,

$$\begin{aligned}\Delta_t\left(y^1_{its}, y^2_{its}\right) &= \Delta_t\left(\tilde{y}^1_{its}, \tilde{y}^2_{its}\right) + \delta^1_t \tilde{v}^1_{its} + \delta^2_t \tilde{v}^2_{its} + o_p(1) \\ \text{and} \qquad \Delta_{t-s}\left(y^1_{its}, y^2_{its}\right) &= \Delta_{t-s}\left(\tilde{y}^1_{its}, \tilde{y}^2_{its}\right) + \delta^1_{t-s} \tilde{v}^1_{its} + \delta^2_{t-s} \tilde{v}^2_{its} + o_p(1)\end{aligned}$$

Ideally we would observe both $\Delta_t\left(y^1_{its}, y^2_{its}\right)$ and $\Delta_{t-s}\left(y^1_{its}, y^2_{its}\right)$ leading to the following linear probability model specification,

$$Q^g_{its} = \gamma^g_{0s} + \gamma^g_{1s}\Delta_t(y^1_{its}, y^2_{its}) + \gamma^g_{2s}\Delta_{t-s}(y^1_{its}, y^2_{its}) + \varepsilon^g_{its} \qquad \forall\, s \neq 0$$

where γ^g_1 is the elasticity of selection into self-employment with respect to unexpected changes in income splitting tax savings, given the inclusion of the counterfactual. Identification of γ^g_1 therefore relies on the assumption that,

$$Cov(\varepsilon^g_{its}, \Delta_t(y^1_{its}, y^2_{its}) | \Delta_{t-s}(y^1_{its}, y^2_{its})) = 0 \qquad \forall\, s \neq 0$$

Replacing y^g_{its} with $\tilde{y}^g_{its}$ we get,

$$Q^g_{its} = \gamma^g_0 + \gamma^g_1 \Delta_t(\tilde{y}^1_{its}, \tilde{y}^2_{its}) + \gamma^g_2 \Delta_{t-s}(\tilde{y}^1_{its}, \tilde{y}^2_{its}) + \tilde{\varepsilon}^g_{its} \qquad \forall\ s \neq 0$$

$$\text{where} \qquad \tilde{\varepsilon}^g_{its} = \varepsilon^g_{its} + \delta^1_{t-s}\tilde{v}^1_{its} + \delta^2_{t-s}\tilde{v}^2_{its}$$

It is evident then that the linearity assumption is now needed for identification. If δ^g_t is correlated with $\Delta_t(\tilde{y}^1_{its}, \tilde{y}^2_{its})$ then the identifying assumption of exogenous deviations in Δ fails.

Whether these tax savings incentivize joint selection into self-employment around child-birth can be thought of a test for the symmetry between γ_{1s} across partners:

$$H_0 : \gamma^1_{1s} = \gamma^2_{1s}$$

Alternatively, if we think of the above LPM for the father and mother as a first-stage and reduced form equations respectively. The above test can then be reframed in terms of the IV estimator. Indeed, in this context the local average treatment effect is the object of interest, and can be used as a test for the mechanism. A couple can only be benefiting from income splitting if both report taxable self-employment income. For every father who is observed selecting in or out of self-employment as a result of the shock to the value income-splitting we should observe an symmetric change in the mother's behaviour. However, given the tendecy for women to delay such decisions till after childbirth I reframe the null hypothesis as,

$$H_0 : \gamma^{LATE}_s = 1 \qquad \text{for some } s > 0$$

To summarize, there are a number of measurement and identification challenges that must be dealt with. First, pre-income splitting income values are not necessarily observed and the observed path of income is endogenous to a couple's labour supply response to childbirth. Second, the

value of Δ depends on both the absolute and relative incomes of the couple, with the result that it will typically increase in the event of an asymmetric, negative income shock such as childbirth. For this reason, the value of Δ must be simulated at an unobserved level of income that will vary with childbirth. From a measurement perspective, I contend that if one assumes both a common and stochastic component to the change in income relative to this base period one can replace the unobserved income path with base period income structure adjusted by a common shock. Moreover, if we assume that there is no income splitting in the base period - by selecting on households that are jointly wage-employed in the base period households- then this expected income path can be used to simulate Δ.

Regarding identification I show that changes in Δ relative to a base period will include both a labour supply and tax change component. From the perspective of the forward looking family the policy shock is arguably unpredictable, and can therefore be used to identify tax-induced selection into self-employment around childbirth. However, given the absence of the true underlying income, one does require an additional linearity assumption: deviations between the predicted and realised income path can be approximated by a linear change in the value of Δ.

4.3 Estimation

In practice I make two important adjustments. First, in a similar vain to Gruber and Saez (2002) I replace the level measure of Δ with the proportional change in after tax income.

$$\Delta_t\left(y^1_{ts}, y^2_{ts}\right) = \ln\left(y^1_{ts} + y^2_{ts} - 2\mathrm{T}_t(0.5(y^1_{ts} + y^2_{ts}))\right) - \ln\left(y^1_{ts} + y^2_{ts} - \mathrm{T}_t(y^1_{ts}) - \mathrm{T}_t(y^2_{ts})\right)$$

Second, I simulate both values of Δ using the income series $\{\tilde{y}^1_{its}, \tilde{y}^2_{its}\}$, based on the reference period s=-2,

$$\{\tilde{y}^1_{its}, \tilde{y}^2_{its}\} = \{\tilde{a}^1_{ts} y^1_{i,t-s-2,-2}, \tilde{a}^2_{ts} y^2_{i,t-s-2,-2}\}$$

where

$$\tilde{a}^g_{ts} = \frac{E[y^g_{its}|X_{it}, t, s]}{E[y^g_{it,-2}|X_{it}, t, s = -2]}$$

This can be done with and without inflation adjustments. Canada's income tax brackets have not always remained pegged to inflation. Hence, it is not obvious which choice is better; although for simulations of Δ_s long after childbirth inflation adjustment seems the conservative choice. For each parent $\tilde{a}^g_{ts}$ is estimated using the full sample of observed parents as the relative change in taxable income associated the birth of a first child. That is, $\tilde{a}^g_{ts}$ is estimated from the equation

$$y^g_{ist} = \sum_{j \neq -2} \sum_{k \neq 2000} \alpha^g_{jk} \cdot \mathbf{1}\{s = j\} \cdot \mathbf{1}\{t = k\} + \sum_l \beta^g_l \cdot \mathbf{1}\{age_{its} = l\} + \delta^g_t + \xi^g_{its}$$

The IV model can be estimated separately by event-time, or in a pooled model where paternal self-employment is interacted with the event-time dummies. The second stage equation is given by,

$$Q^2_{its} = \sum_{j \neq -2} \alpha_j \cdot \mathbf{1}\{j = s\} + \beta Q^1_{its} + \sum_{j \neq -2} \beta_j Q^1_{its} \times \mathbf{1}\{j = s\} + X'_{its}\gamma + \varepsilon_{its}$$

In addition to age and time fixed-effects I include a quadratic in the spouse's age, province fixed-effects, and an indicator for formal marriage as the sample includes common-law couples. I also include the household's simulated after tax household income, under both counterfactuals, as a control for the income effect these policy changes may induce. The coefficients of interest (β_j) are simple difference-in-difference coefficient's comparing the the additional channel maternal self-employment of women married to wage- and self-employed relative to a base period before childbirth.

$$\beta_j = \Big[E[Q^2_{itj}|Q^1_{itj} = 1, X_{itj}] - E[Q^2_{i,t-j-2,-2}|Q^1_{i,t-j-2,-2} = 1, X_{i,t-j-2,-2}]\Big]$$
$$- \Big[E[Q^2_{itj}|Q^1_{itj} = 0, X_{itj}] - E[Q^2_{i,t-j-2,-2}|Q^1_{i,t-j-2,-2} = 0, X_{i,t-j-2,-2}]\Big]$$

Because the identification strategy relies on deviations in the tax structure relative to the base-period policies the average difference in maternal self-employment of women married to wage- and self-employment men in the base period (β) is not identified. Only deviations from this level are identified using this strategy. In practice, I employ a block-bootstrap two-step estimation procedure, clustering at the individual level.[31]

4.4 Results

Figure 15 shows the first-stage, reduced form, and IV results from the above specification when estimated individually by event-time. These results are include inflation adjustments.[32] The first-stage and reduced form coefficients correspond to a simpler model including only the realized and counterfactual values of Δ. Prior to the base period the first-stage is insignificant, which acts as a valuable placebo test of the instrument. The first-stage coefficient, which represents paternal responses to income-splitting tax shocks is positive, even before childbirth, and remains stable through the event of child-birth. The scale coefficient (~ 0.5) is reasonable; given the simulated tax savings of ~ 0.02 this would account for an increase of 1% in paternal self-employment. This is approximately the difference between US and Canadian male self-employment rates.

The reduced form coefficient - representing the maternal response - is insignificantly different from 0 prior to childbirth. However, it increases monotonically after childbirth reaching a level that is statistically indistinguishable from the paternal elasticity. The result is an IV estimate the converges to 1 by event-time 3. This delayed response is consistent with the evidence from the

[31] NOTE: The SE's reported here are not bootstrapped. They are clustered at the individual level, while I am waiting to release the bootstrapped SE's.

[32] See 16 for the same results without inflation adjustment.

SLID (see Table 1) which suggests that co-employment increases sharply with the a second child. Women are also less likely to respond to these shock prior to childbirth because it would result in a loss of EI support in the period of childbirth.

Similar results are obtained when pooling all event-times and instrumenting for the above interaction of event-time and paternal self-employment. Table 3 provide the coefficients for event-time (α_j's) and event-time interacted with paternal self-employment (β_j's) for the models with inflation adjustment respectively. Table 4 provides the same coefficients without inflation adjustment. In columns (1) and (4) I include a model without paternal self-employment as a base line for the sample. Relative to the overall sample, the increase in maternal self-employment associated with childbirth is greater. This is likely due to the selection on couples who are wage-employees in the base period. In columns (2) and (5) paternal self-employment (and the interaction with event-time) is included, and in columns (3) and (6) the interactions terms are instrumented for using the simulated Δ savings.

From the OLS results we see that even in the sample of previous wage-employees the paternal self-employment increases the probability that a mother enters self-employment by around 25% four years after childbirth. Thus, approximately a quarter of the increase in paternal self-employment in the sample is associated with selection into co-employment. The local average treatment effect is much larger. It increases monotonically from 0 before childbirth to a value that is indistinguishable from 1 by event-time 3. This is robust to controls for province fixed-effects, formal marital status, age of spouse, and individual (predicted) earnings (model (6)). Moreover, inflation adjustment appears to improve the estimation (see Table 4).

These results then are consistent with the symmetric responses to the savings from income splitting suggesting that such savings increase paternal self-employment and maternal self-employment after childbirth. These elasticities are large enough to explain a 1-1.5% increase in paternal self-employment, but the associated increase in paternal self-employment from childbirth is closer to $\sim 4\%$ (see Figure 6). Therefore, while income splitting may be a factor in determining selection into co-employment, it cannot be the primary driver.

For this reason, I develope a simple model that gives rise to these patterns of selection into self-employment and co-employment in Canada. The model integrates key policy mechanisms, including income splitting. Consistent with these empirical findings the model will predict that income splitting increases paternal self-employment rates even before childbirth. However, by simulating some of the regional policy reforms in Canada, the model predict that the childcare benefits of co-employment are arguably a more important factor in explaining the regional variation in parental self-employment (see Figures 8 and Figures 9).

5 Finite horizon, unitary household model

Over the past 40 years Canada's federal and provincial government's have taken a number of steps to increase their support for families and children, while encouraging further female labour force participation. Here, I develop a unitary family model with discrete labour supply decisions to demonstrate how some of these policies may have interacted with families' decisions to enter co-employment. The model assumes that one of the primary benefits to co-employment is a positive childcare externality. With a fixed start-up cost, this childcare benefit leads fathers to increase their self-employment prior to childbirth and mothers to join them in co-employment immediately after.

Self-employed parents also benefit from income splitting and as a result I simulate both the real and tax-reported self-employment status of parents. The model also accounts for the fact that upon childbirth employment insurance benefits are dependent past labour supply decisions. The predictions of this model under various counterfactuals are then related to the regional and inter-temporal variation in Canadian co-employment. Consistent with the lack of empirical evidence that co-employment increases with childbirth in Quebec, I find that the model predicts such an outcome under the Quebec childcare reforms of 1997. Quebec's introduction QPIP in 2006 did less to alter the co-employment decision.

5.1 Set-up

Consider a model where households maximize their combined, finite-lifetime utility through a sequence of discrete labour supply decisions. To begin, each household comprises two potential earners who face symmetric labour supply choices. In each period they must make a discrete labour supply decision (l_t),

$$l_t = \begin{cases} \text{not economically active} \\ \text{wage-employee} \\ \text{self-employed} \end{cases}$$

$$\begin{aligned} \text{where} \quad & h_t = \mathbf{1}\{l_t = \text{wage-employee}\} \\ \text{and} \quad & s_t = \mathbf{1}\{l_t = \text{self-employed}\} \end{aligned}$$

Asymmetry enters the model through the event of childbirth (ι_t) which takes place in a predetermined period ($t = 0$), and requires the birthing parent (i.e. 'mother': $g = 2$) to leave the labour market for a single period.[33]

$$\iota_t = \mathbf{1}\{t = 0\}$$

After the event of childbirth the, now parents, once again face symmetric labour supply decisions. Earnings are fixed and known to the worker. In the wage-paying sector the parent will receive wages w_t^g. Alternatively, if they select into self-employment their earnings will be v_t^g respectively. For the non-birthing parent (i.e. 'father': $g = 1$) this yields the earnings function,

$$y^1(l_t^1, W) = w_t^1 \cdot h_t^1 + v_t^1 \cdot s_t^1$$

[33] To simplify matters I denote time in event-time of childbirth.

The mother's earnings function is state dependent. During the period of childbirth the mother is unable to work, and receives an income supplement (i.e. employment insurance) based on her previous period's earnings. However, only wage earnings are insurable. This assumption is in align with the employment insurance policy that governed Canada for the majority of the past four decades.[34] The mother's earnings function is given by,

$$y^2(l^2_{t-1}, l^2_t, W) = (1 - \iota_t) \cdot (w^2_t \cdot h^2_t + v^2_t \cdot s^2_t) + \iota_t \cdot (\psi \cdot w^2_{t-1} \cdot h^2_{t-1})$$

where ψ is the income replacement rate. For example, Canada's federal employment insurance policy has a standard replacement rate of 55%.

The household's budget constraint depends on the after tax income of the household under an individual tax structure. There are no savings, but the budget includes three additional elements: tax savings from income splitting (Δ), the cost of child-care (N), and the start-up cost to self-employment (K).[35]

$$c_t = y^1_t - T(y^1_t) + y^2_t - T(y^2_t) + \Delta(l_t, W) - N_t(l_t) - K(l_{t-1}, l_t)$$

If either parent is self-employed, then the household is able to lower their tax burden through income splitting. Only self-employment income can be split, with the result that there are no gains to income splitting in households where the wage-employee earns more than the self-employed parent. The principal formula for these savings is given by,

$$T(y^1_t) + T(y^2_t) - 2 \cdot T(\bar{y}_t)$$

where $\bar{y}_t$ denotes the average income of the two earners in the household in period t.

[34] Quebec included the self-employed in their QPIP reform of 2006 and since 2011 federal EI contributions have been optional for the self-employed. These can be claimed for maternity leave; although business revenue must decline by at least 40% to do so.

[35] To simplify the notation the start-up costs are not tax deductible. This also implies that they are equal across income brackets.

After the event of childbirth the household must solve the problem of childcare. I assume that there is a fixed cost to childcare that must be paid in the event that neither parent leaves the labour force. This cost declines with time to reflect the increased access to education services. In addition, I assume that the cost of childcare is lower for co-employed households. This is one of the primary benefits to co-employment: increased flexibility that allows both parents to supply labour to an enterprise while solving the problem of childcare. These benefits do not extend to individually self-employed parents.

Finally, there is a fixed start-up cost to self-employment entry. These costs can be paid by either parent and are shared when both parents enter simultaneously. Thus, joint selection into self-employment is assumed to be co-employment: employment within the same enterprise. In addition, these start-up costs can be transferred between parents. A parent who enters self-employment when their spouse was self-employed in the previous period does not pay theses costs. For this reason, parents may switch between self-employment states and because mothers leave the labour force during the event of childbirth it will be up to fathers to 'maintain' any enterprise in period $t = 0$.

The value function of the household is given by,

$$V_1(l_0, W, \Theta) = \max_{l_1} \left[U(c_1(l_0, l_1, W), \Theta(l_1)) + \beta \cdot V_2(l_1, W, \Theta) \right]$$

and the household's utility function is assumed to be linearly separable between consumption and the dis-utility of labour.

$$\text{where} \qquad U(c_t, \Theta) = u(c_t) - \theta^1(l_t^1) - \theta^2(l_t^2)$$

If not for the event of childbirth this model would yield an allocation of workers between wage and self-employment based entirely on relative earnings in both sectors. In addition, if the earnings of 'mothers' and 'fathers' were equal across sectors there would be no 'gender' difference in self-

employment and wage-employment rates in equilibrium. It is the event of childbirth and the subsequent cost of childcare that interacts with other aspects of the model to generate staggered selection into self-employment along 'gendered' lines.

Co-employment after childbirth has obvious childcare benefits. However, selection into self-employment after childbirth is more costly, as there is less time to overcome the fixed start-up costs. For this reason, selection into self-employment will take place early on in life cycle. If mothers enter self-employment before childbirth they lose out on employment insurance benefits when they give birth. Thus, fathers will will select into self-employment first; even if the mother has higher self-employment earnings. They can always switch over after the child is born.

The decline in maternal earnings during the childbirth creates an additional incentive for fathers to enter self-employment before childbirth through income splitting. A self-employment father is able to lower the household's tax burden by shifting self-employment income to the economically inactive mother. This will create a wedge between real and tax-reported self-employment. However, these savings are likely too small to induce a mother to abandon a wage paying position before childbirth, as it is not just one period of earnings she would forego. Thus, future savings from income splitting act as a subsidy to paternal self-employment leading up to childbirth. This together with the employment insurance structure insure that fathers will lead the household's selection into self-employment. In addition, the childcare savings from co-employment will drive a wedge between the relative earnings in each sector and increase paternal self-employment in all periods and maternal self-employment after childbirth.

5.2 Simulation

I solve the model under the parameters described in the Appendix A. Households live for five periods ($t = -2, -1, ..., 2$), with the event of childbirth occurring in period 0. Prior to period -2 all parents are assumed to be employed in the wage-paying sector. In period 0 the mother is unable to work and automatically claims EI based on her prior employment status. The EI replacement rate is assumed to be 55%. This is a simplification of Canada's real EI structure which includes a

Working While on Claim (WWC) component. Under WWC the benefits are gradually clawed back if someone has other employment income. In periods 1 and 2 employment is again optional, but the household must pay for childcare if both parents are employed. If the parents are co-employed the cost of childcare is halved. In period 2 the cost of childcare falls to simulate the decline in childcare costs as children reach school age.

All self-employed households automatically practice income splitting, in addition to knowing each individual's labour market status I keep track of reported self-employment income. This gives me two measures of self-employment status: one real and one tax reported where the former is included in the latter. I am then able to replicate the measure of self-employment I use in my empirical analysis: taxable self-employment income is more than 50% of the individual's total taxable employment income.

Wage earnings are drawn from a normal distribution ($\mu_w^g = 10, \sigma_w^g = 1$) and held constant over time. In addition, they are uncorrelated within a couple. For each individual self-employment earnings are equal to wage earnings plus a 'entrepreneurial' shock. This shock is normally distributed ($\mu_v^g = -0.5, \sigma_v^g = 1$) to ensure that self-employment rates are lower than wage-employment rates in equilibrium. In addition, the 'entrepreneurial' shock is assumed to be positively correlated within a couple ($\sigma_v^{12} = 0.5$) thereby ensuring a positive level of co-employment even in the absence of children.

When solving the model I assume a log utility function for consumption and linear dis-utility of labour. The dis-utility parameters are drawn from a uniform distribution and are uncorrelated within a couple. These ensure a level of inactivity in the labour market; in particular, after childbirth when there are additional childcare costs to working.

Figure 17 plots the equilibrium levels of maternal and paternal self-employment, and co-employment by period (i.e. event-time). It also shows the tax reported value for each series (>50% of taxable income is from self-employment); denoted by a dashed line of the same colour. As there is no under reporting of earned income, the real level of self-employment cannot exceed the tax-reported level. However, because the model excludes the WWC components of EI, it overstates

the return to income splitting during period 0 and therefore over estimates the level of tax-reported self-employment in period 0. This is less of a concern if one considers that the periods in the model need not represent a tax year, in which case the model still demonstrates how households might exploit income splitting opportunities during the early years of intensive childcare.

Figure 17 demonstrates a distinct increase in paternal self-employment prior to childbirth. In contrast, maternal self-employment falls to zero before childbirth. There is income splitting during the period 0, but no more than the employment insurance replacement income. In period 2, both real and tax reported maternal self-employment rises sharply. Income splitting clearly benefits certain households where a parent leaves the labour market for childcare reasons. The real level of paternal self-employment declines slightly after childbirth, while the tax reported level remains level. This reflects the assumed symmetry of the childcare decision which pushes (relatively) less productive self-employed men - who paid the fixed start-up cost - to switch places with their spouse after childbirth and select into full time childcare instead; while the family benefits from income splitting. Between period 1 and period 2 the measurement error in tax-reported co-employment is largely erased as the benefits to labour supply outweigh the declining cost to childcare.

5.3 Discussion: Counterfactuals

Figure 18 shows the equilibrium level of self-employment by period under various counterfactuals. In Panel A I remove the option to split income with a spouse. This lowers the level of paternal self-employment before childbirth, as the savings from income during the period of childbirth is removed. However, the level of real co-employment actually increases slightly after childbirth as income splitting acts as a subsidy to labour market inactivity. For the same reason male self-employment does not decline in period 1, as it does in the standard model. Notably, beyond the absence of proxy maternal self-employment in period 0 removing income splitting barely changes the level of maternal self-employment.

Under the counterfactual of free childcare (Figure 18 Panel B) the childcare benefit to co-employment is redundant. Paternal self-employment falls dramatically; even two periods prior to

childbirth; reiterating the fact that a large share of paternal self-employment is childcare motivated. Still, paternal self-employment falls after childbirth once the income splitting benefits of self-employment during childbirth are passed; suggesting that pre-childbirth levels of self-employment remain elevated because of income splitting - even without the childcare savings. For mothers, self-employment does not rise above its pre-childbirth (period -2) level. In fact, it drops below suggesting that the fixed start-up costs are a hindrance to maternal self-employment after childbirth as there is less time to recuperate them.

Panel C shuts down the employment insurance channel: households receive no EI benefits during the year of childbirth, and mothers remain unable to work during this period. However, the childcare and income splitting benefits to self-employment remain. Because mothers have no incentive to leave self-employment before childbirth maternal self-employment remains at its period -2 level in period -1. As a result, paternal self-employment does not rise between periods -2 and -1. The income splitting and childcare benefits to self-employment with childbirth remain, but there are no longer EI incentives for the start-up cost to be born by the father. Paternal self-employment now rises at, and not before, childbirth (to avoid re-starting the business) and still declines afterwards. Maternal self-employment and co-employment still rise sharply after childbirth.

Figure 19 includes three additional counterfactuals. Panel A removes the income splitting, childcare, and employment insurance incentives, thus combing the counterfactuals in Figure 18. Panels B and C remove the event of childbirth; with Panel B maintaining the income splitting benefit to self-employment. Removing all policy distortions (Panel A) reveals that childbirth still acts to reduce maternal self-employment through the fixed start-up cost channel. The interruption of childbirth lowers the chance that mothers are able to recuperate these costs, and maternal self-employment falls with childbirth. Notably, co-employment does not fall because the spouse is able to maintain the enterprise through childbirth relinquishing the need to 'restart' the business. The slight increase in paternal self-employment during childbirth demonstrates an attempt to shift the maintenance of the business to avoid these 'restart' costs. This is optimal because the model

assumes no frictions in the wage-paying sector.

Panels B and C demonstrate and important element. In the absence of the event of childbirth income splitting has very little impact on either real and proxy self-employment. It is the shock to female labour supply that generates the savings from income splitting and the increase in paternal self-employment leading up to childbirth. Note, the observed differences in the equilibrium levels of paternal and maternal self-employment are a result of approximation error (i.e. sample size), and not a parametric or modelling difference.

5.4 Discussion: Policy Reforms

Next, I use this simple model to shed light on how a number of policy reforms in Canada may have affected selection into self-employment and co-employment. First, I consider the major federal reforms at the turn of the century. These include the federal income tax reform of 2000 and the EI reform of 2001. Second, I examine the Quebec family reforms including the subsidization of childcare from 1997 and the introduction of an independent Quebec Parental Insurance Plan (QPIP, 2006).

Figure 20 depicts the federal reforms. Panel A simulates the outcome levels of equilibrium self-employment and co-employment under the federal policy structures of the 1990's;[36] while Panel C depicts the levels under the cumulative changes of the income tax reform of 2000 and EI reform of 2001. Note, Panel C is equivalent to the standard model depicted in Figure 17. As previously discussed, the federal income tax reform of 2000 introduced an additional upper income tax bracket, thereby lowering the marginal tax rate for much of the middle class in Canada (REF). I simulate the pre reform equilibrium by equating the 3rd marginal tax rate in the standard model to the 4th marginal tax rate. The federal EI reform of 2001 extended parental leave benefits from 10 to 35 weeks. For a mother this increased the maximum paid coverage from 25 weeks to 50 weeks.

[36]These differ substantially from the pre-1988 income tax structure and the pre-1991 EI structure. The 1988 income tax reform removed a family tax benefit that generated automatic income splitting tax savings for all couples (see REF). The 1991 EI reform introduced a shared 10 weeks of parental leave. Prior to 1991 mothers had access to 15 weeks of maternity leave, but no shared parental leave.

Thus, prior to 2001 the annualized income replacement rate was not $50/52 \cdot 55\% \simeq 53\%$, but rather $25/52 \cdot 55\% \simeq 26\%$. Panel B shows the equilibrium after the tax reform, but before the EI reform.

Under the 1990's federal policy parameters the lower income replacement rate during childbirth and the steeper income tax gradient increases the value to income splitting during the child-raring years. Wanting to take advantage of these savings, while facing the fixed start-up costs, fathers increase their selection into self-employment leading up to childbirth. For this reason, both reforms have effect of reducing paternal self-employment prior to childbirth which is consistent with the empirical decline in male self-employment since the late 1990's. The higher level of paternal self-employment in the 1990's lowers the cost to maternal self-employment after childbirth. Thus, both the real and tax reported levels of maternal self-employment and co-employment are higher in the 1990's. The EI reform lowers the savings from income splitting during childbirth. This has the effect of reducing paternal self-employment before childbirth and tax reported maternal self-employment after childbirth. However, real maternal self-employment only declines in period 2. Taken together these two reforms had the effect of reducing tax-reported co-employment, as well as real paternal self-employment.

The province of Quebec has introduced a number of policies to increase female labour supply while supporting young families (REF). Here I examine the introduction of a childcare subsidy in 1997 and the shift towards an independent Quebec Parental Insurance Plan (QPIP) to replace the federal EI structure in 2006. Figure 21 Panel A depicts the equilibrium levels of parental self-employment and co-employment in Quebec after the 1997 reform. I assume that prior to the reform Quebec had the same policy structure to the rest of Canada during this time (i.e. Figure 20 Panel A). I assume that the childcare reform lowered the cost of childcare for all working parents thereby removing the childcare savings from co-employment. Childcare remains costly, but is reduced by 50% and is equal across employment status. The EI structure remains the same as the rest of Canada. Panel B then shows the cumulative effect of the federal income tax and EI reforms with these lower childcare costs (as depicted in Figure 20). In Panel C I demonstrate the effect of the QPIP reform of 2006. In addition to maintaining the lower cost of childcare, the introduction of

QPIP changed two important elements. First, it forced self-employed workers to contribute to the policy and made them eligible for EI benefits based on past self-employment income. Second, it increased the income replace rate from a flat 55% to 70% for this first 25 weeks and 55% for the second 25 weeks (or an annualized $\simeq$ 60%).

It is evident that the simulation of the 1997 Childcare reform had the largest effect on co-employment rates (using Figure 20 Panel A as the baseline). After the reform paternal self-employment increases only slightly with childbirth; as in Figure 18 Panel B. Tax reported maternal self-employment rises with childbirth, but falls immediately as higher-paying wage-employment increases. Equilibrium co-employment is actually lower in period 2 relative to period -2. In this setting the federal reforms acted to the further flatten the paternal self-employment trend as there are neither childcare nor income splitting savings from co-employment. As the introduction of QPIP made self-employment insurable had a similar effect to removing the EI benefit entirely (Figure 18 Panel C). Maternal self-employment is sustained in period -1, and as a result paternal self-employment falls even further.

6 Conclusion

This paper makes an important contribution to the literature on the long run parent penalty. It finds that while men remain employed through the event of child-birth they do make infra-marginal labour adjustments; in particular, selecting into self-employment leading up to the birth of their first child. This adjustment is then associated with a comparable increase in maternal self-employment after the event of childbirth. Thus, childbirth induces co-employment which is confirmed by evidence on industry matching. This selection into co-employment has important consequences for the risk exposure of the household; in particular, given the uninsured nature of much of the self-employed labour market (Gruber and Cullen, 1996).

I argue that this shift into co-employment is a real labour supply phenomenon and not simply an accounting trick, in the medium term. Certainly, both self-employment decisions cannot reflect *pure*

income splitting. I do also find evidence that there is a delay in the real labour supply adjustment of mothers. For this reason, there may be a period of 2-3 years during which couples take advantage of tax savings from income splitting, while new mothers focus on childcare. In the medium term (3-4 years of first child) there is a real maternal self-employment adjustment that results in an increase in the employment income of the household. This timing difference between the tax reported and real labour supply decision appears in the model too, and depends partly on the generosity of the EI structure.

Regardless, self-employed households do split their income. There is definite bunching in the relative income distribution at 0.5%; only in self-employed households (Zinovyeva and Tverdostup, 2018; Bertrand et al., 2015). As argued by Zinovyeva and Tverdostup (2018), co-employment explains the 'break-winner' hypothesis. Moreover, estimate that the average household tax rate of a self-employed couple as falls by an additional []% immediately after childbirth. These savings are in line with my simulations of the return to income splitting.

Such tax optimization is also consistent with other findings on the self-employed (Kleven and Waseem, 2013; Saez, 2010). While the existing literature shows that the self-employed are more likely to bunch and are typically more elastic to policy changes, in these settings an individual's self-employment status is typically taken as given. This paper shows then that selection into self-employment depends on policy incentives too. Moreover, it highlights the event of family formation as crucial moment in this decision. As demonstrated by my model, these policy incentives remain relatively dormant at other stages in the life cycle. It is the asymmetric shock of family formation that interacts with EI and income tax structures to incentivize this timed and coordinated selection into co-employment.

I show that a significant share of paternal selection into self-employment in the event of childbirth can be attributed to the tax savings from income splitting. The estimated a paternal and maternal elasticities are in the range of 0.5, and are symmetric after childbirth. The estimated local average treatment effect of 1 provides evidence of the mechanism at play in the simulated instrument.

This results of this paper also suggest that the flexibility benefits of self-employment - as they

pertain to childcare - may not be easily accessed (Jeon and Ostrovsky, 2019). I find no evidence of an increase in female self-employment associated with childbirth in households where the spouse is a wage-employee. The co-employed appear to primarily benefit through increased work from home which allows a mother to balance childcare and work obligations. One may then conclude that a significant share of male self-employment is in fact motivated by childcare and a demand for flexibility as the event of child birth is associated with a $\sim 4\%$ increase in paternal self-employment among married men.

References

Nikolay Angelov, Per Johansson, and Erica Lindahl. Parenthood and the Gender Gap in Pay. *Journal of Labor Economics*, 34(3):545–579, 2016. URL http://www.journals.uchicago.edu/t-and-c.

Henrik Kleven, Camille Landais, and Jakob Egholt Søgaard. Children and gender inequality: Evidence from Denmark. *American Economic Journal: Applied Economics*, 11(4):181–209, 2019a. ISSN 19457790. doi: 10.1257/app.20180010.

Henrik Kleven, Camille Landais, Johanna Posch, Andreas Steinhauer, and Josef Zweimüller. Child Penalties across Countries: Evidence and Explanations. *AEA Papers and Proceedings*, 109:122–126, 2019b. ISSN 2574-0768. doi: 10.1257/pandp.20191078.

Claudia Goldin. A grand gender convergence: Its last chapter. *American Economic Review*, 2014. ISSN 00028282. doi: 10.1257/aer.104.4.1091.

Claudia Goldin and Lawrence F. Katz. A most egalitarian profession: Pharmacy and the evolution of a family-friendly occupation. *Journal of Labor Economics*, 34(3):705–746, 2016. ISSN 0734306X. doi: 10.1086/685505.

Sung Hee Jeon and Yuri Ostrovsky. Balancing family and work: Transition to self-employment among new mothers. *Oxford Economic Papers*, 71(1):47–72, 2019. ISSN 14643812. doi: 10.1093/oep/gpy061.

Suzanne Heller Clain. Gender differences in full-time self-employment. *Journal of Economics and Business*, 52:499–513, 2000. URL https://ac.els-cdn.com/S0148619500000321/1-s2.0-S0148619500000321-main.pdf?{_}tid=f84b044d-d183-4b1e-99ca-81d124206982{&}acdnat=1535731783{_}53d17280d76c515463ccc5abdcebd594.

Greg Hundley. Male/female earnings differences in self-employment: The effects of marriage, children, and the household division of labor. *Industrial and Labor Relations Review*, 2000. ISSN 00197939. doi: 10.1177/001979390005400106.

Yannis Georgellis and Howard J. Wall. Gender differences in self-employment. *International Review of Applied Economics*, 2005. ISSN 02692171. doi: 10.1080/02692170500119854.

George Saridakis, Susan Marlow, and David J Storey. Do different factors explain male and female self-employment rates? *Journal of Business Venturing*, 29:345–362, 2014. doi: 10.1016/j.jbusvent.2013.04.004. URL http://dx.doi.org/10.1016/j.jbusvent.2013.04.004.

Herbert J Schuetze. Income splitting among the self-employed. *The Canadian Journal of Economics*, 39(4):1195–1220, 2006. URL http://www.jstor.org/stable/4121800http://about.jstor.org/terms.

Henrik J Kleven and Mazhar Waseem. Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from Pakistan. *The Quarterly Journal of Economics*, 128 (December):669–723, 2013. ISSN 00335533. doi: 10.1093/qje/qjt004.Advance.

Andrew M Bauer, Alan Macnaughton, and Anindya Sen. Income splitting and anti-avoidance legislation: evidence from the Canadian "kiddie tax". *International Tax and Public Finance*, 22: 909–931, 2015. doi: 10.1007/s10797-014-9342-z. URL https://link.springer.com/content/pdf/10.1007{%}2Fs10797-014-9342-z.pdf.

Natalia Zinovyeva and Maryna Tverdostup. Gender Identity , Co-Working Spouses and Relative Income within Households. *IZA Discussion Paper*, (11757), 2018.

Marianne Bertrand, Emir Kamenica, and Jessica Pan. Gender Identity and Relative Income within Households. *Quarterly Journal of Economics*, pages 571–614, 2015. ISSN 15314650. doi: 10.1093/qje/qjv001.Advance.

Maéva Doumbia and Marion Goussé. Gender Identity and Relative Income within Households: Evidence from Canada. 2019.

Marion Goussé and Marion Leturcq. More or Less Unmarried. The Impact of Legal Settings of Cohabitation on Labor Market Outcomes. *SSRN Electronic Journal*, (April), 2018. ISSN 1556-5068. doi: 10.2139/ssrn.3211428.

Emmanuel Saez. Do Taxpayers Bunch at Kink Points? *American Economic Journal: Economic Policy*, 2 (3):180–212, 2010. doi: 10.1257/pol.2.3.180. URL http://www.aeaweb.org/articles.php?doi=10.1257/pol.2.3.180.

Thomas F. Crossley and Sung Hee Jeon. Joint taxation and the labour supply of married women: Evidence from the Canadian tax reform of 1988. *Fiscal Studies*, 28(3):343–365, 2007. ISSN 01435671. doi: 10.1111/j.1475-5890.2007.00059.x.

Sara LaLumia. The effects of joint taxation of married couples on labor supply and non-wage income. *Journal of Public Economics*, 92(7):1698–1719, 2008. ISSN 00472727. doi: 10.1016/j.jpubeco.2008.01.009.

Jonathan Gruber and Julie Berry Cullen. Spousal Labor Supply as Insurance: Does Unemployment Insurance Crowd Outthe Added Worker Effect? *National Bureau of Economic Research Working Paper Series*, No. 5608, 1996. doi: 10.3386/w5608. URL http://www.nber.org/papers/w5608{%}5Cnhttp://www.nber.org/papers/w5608.pdf.

Steven J. Davis and Till von Wachter. Recessions and the costs of job loss. *Brookings Papers on Economic Activity*, (2):1–72, 2011. ISSN 00072303. doi: 10.1353/eca.2011.0016.

David G. Duff. Neuman and Beyond: Income Splitting, Tax Avoidance, and Statutory Interpretation

in the Supreme Court of Canada. *The Canadian Business Law Journal*, 32(3):345–383, 1999. ISSN 02729490. doi: 10.3366/ajicl.2011.0005.

Alan Macnaughton and Thomas Matthews. Is the Income-Splitting Tax Needed? Some Empirical Evidence. *Canadian Tax Journal*, 1164(5), 1999. URL http://staging.fcf-ctf.ca/ctfweb/Documents/PDF/1999ctj/1999CTJ5{_}Macnaughton.pdf.

Maureen Donnelly, Joanne Magee, and Allister Young. Income Splitting and the New Kiddie Tax: Major Changes for Minor Children. *Canadian Tax Journal*, 48(4):979–1018, 2000. URL http://staging.fcf-ctf.ca/ctfweb/Documents/PDF/2000ctj/2000ctj4{_}donnelly.pdf.

Jon Gruber and Emmanuel Saez. The elasticity of taxable income : evidence and implications. 84: 1–32, 2002.

Paul Goldsmith-Pinkham, Isaac Sorkin, and Henry Swift. Bartik Instruments: What, When, Why, and How. 2019.

Tables and Figures

List of Figures

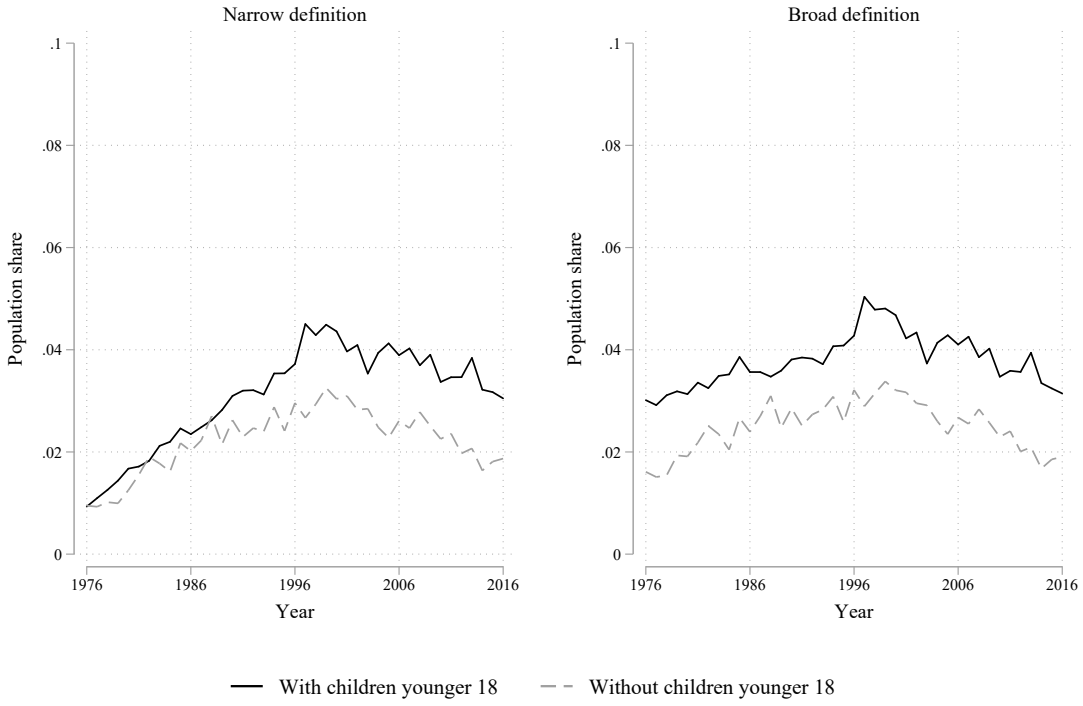

Figure 3: Incidence of joint self-employment amongst married women, aged 25-44, with and without children. Broad definition includes women who identify as "Private, unpaid family work".

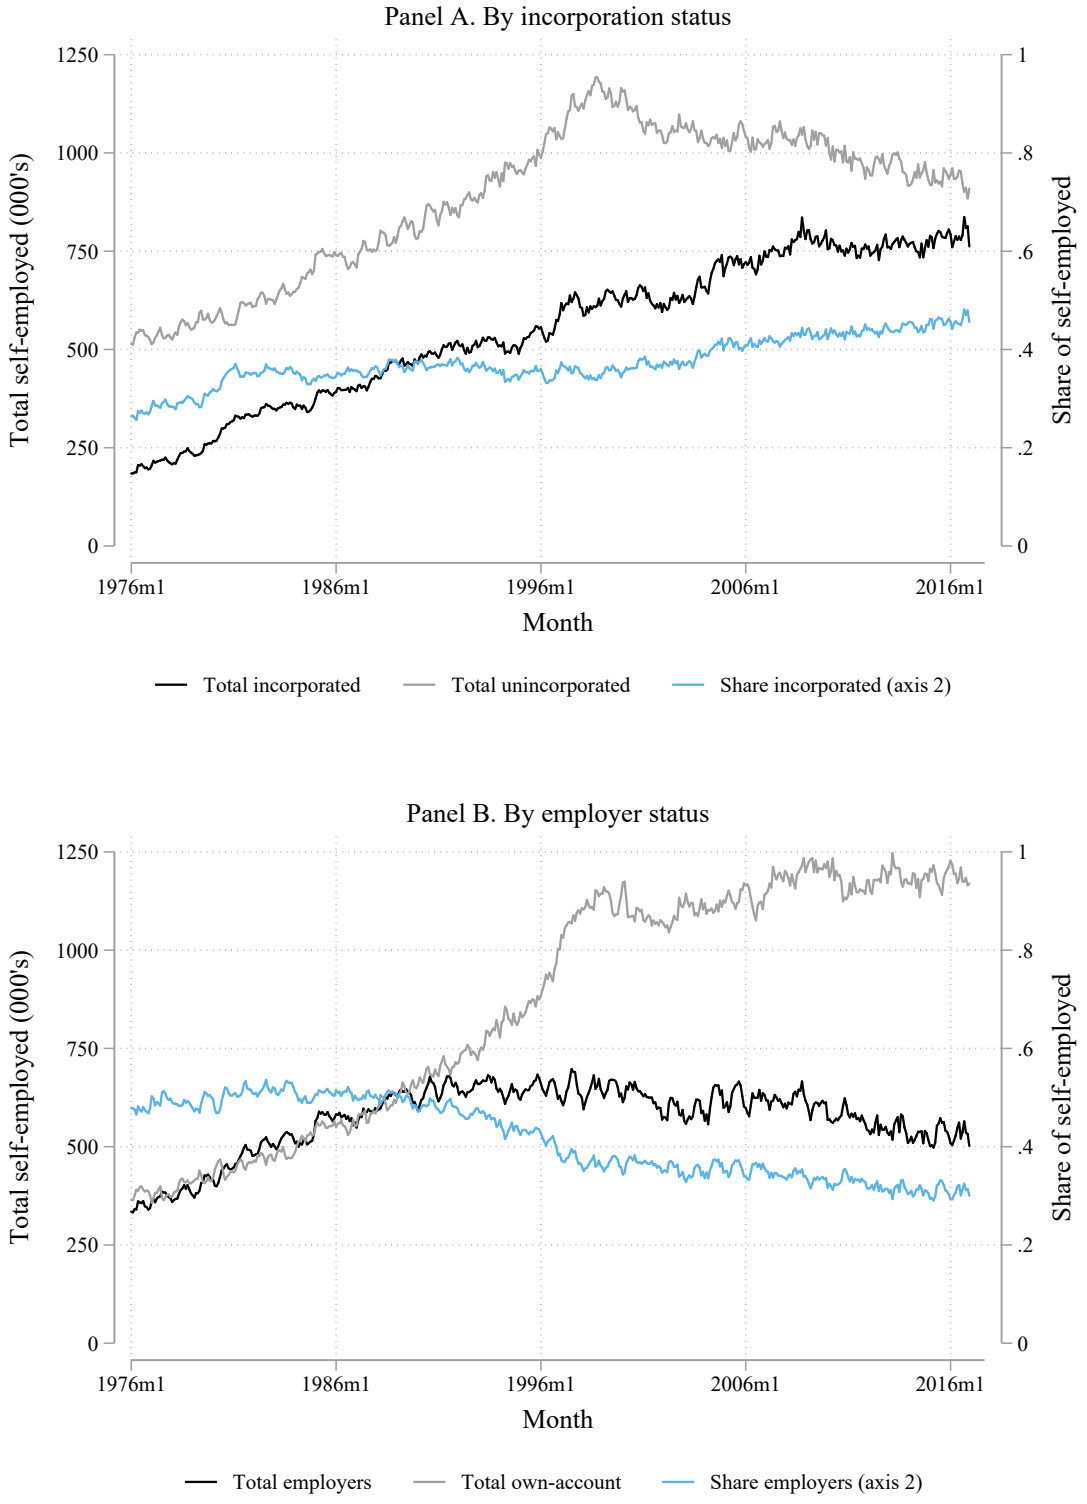

Figure 4: Trend in incorporation and employer status of Canadian self-employed, ages 25-54. Labour Force Survey 1976-2016

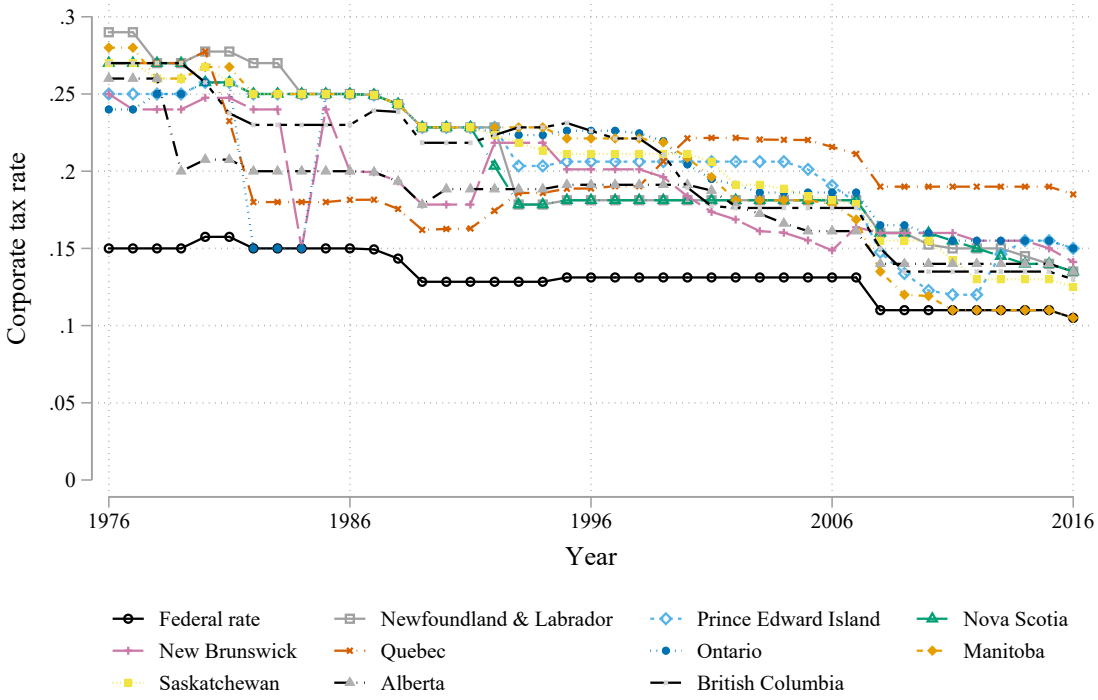

Figure 5: Canadian federal and provincial (+ federal) small business corporate tax rate, 1976-2016.

Figure 6: Change in the probability of employment relative to the probability of under the counterfactual of no childbirth. Estimated using the LAD (1990-2016)

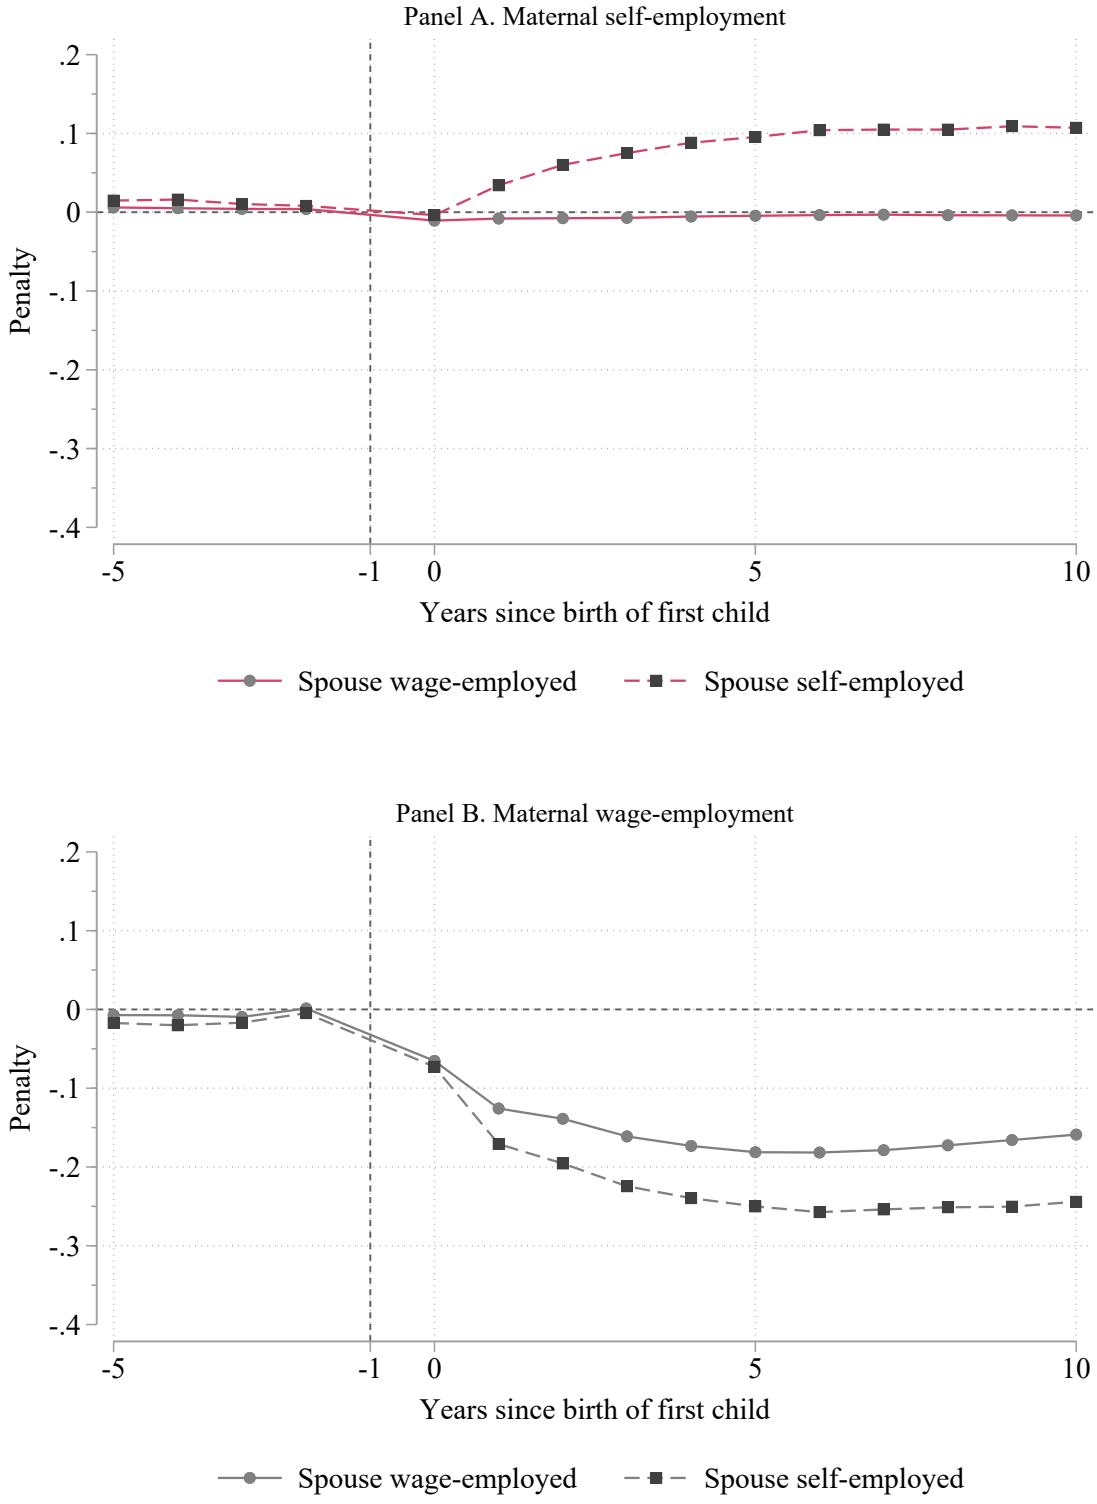

Figure 7: Change in the probability of maternal self-employment and wage-employment relative to the counterfactual probability of employment, by spouse's class of employment. Estimated using the LAD (1990-2016)

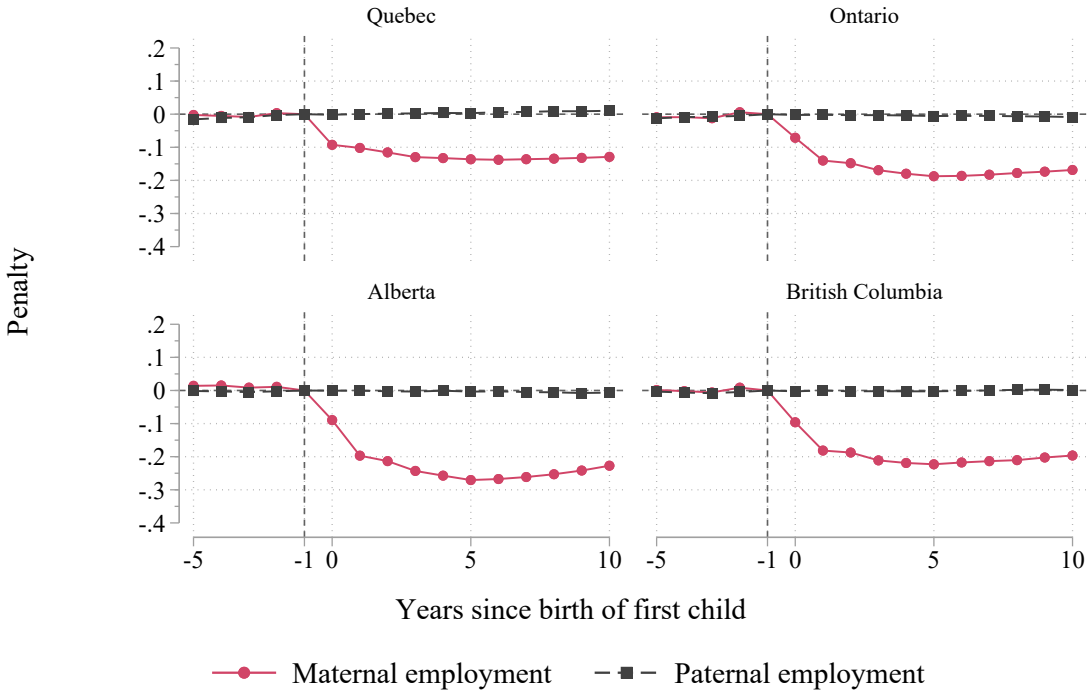

Figure 8: Change in the probability of employment relative to the probability of under the counterfactual of no childbirth. Estimated using the LAD (1990-2016)

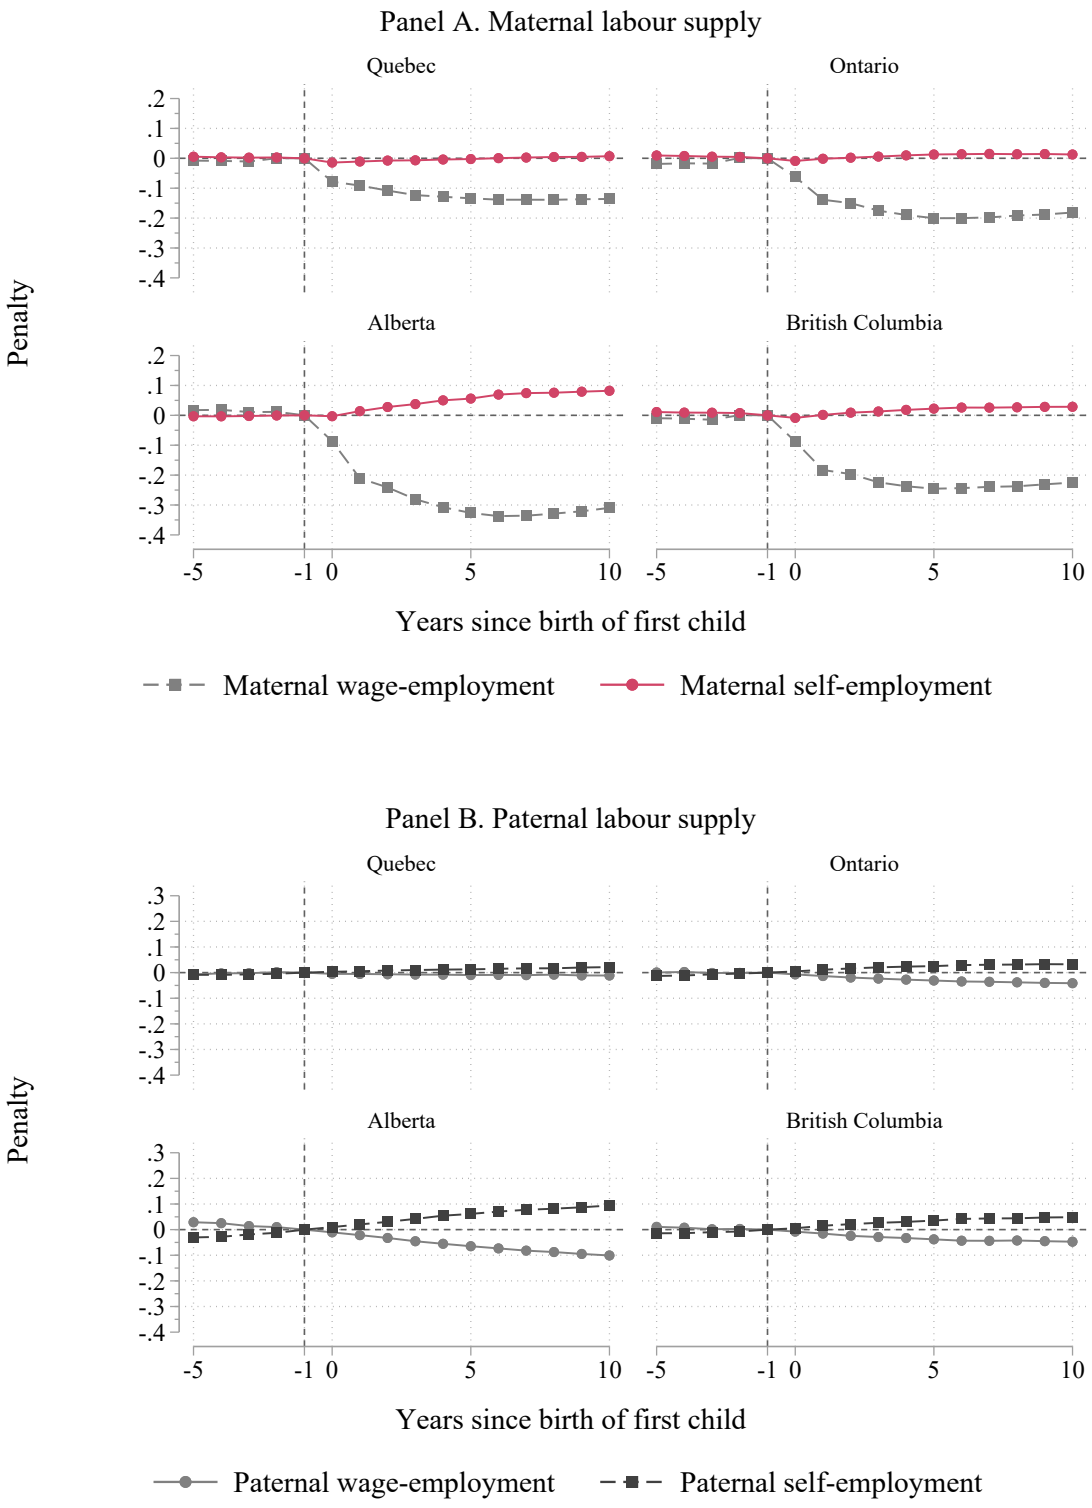

Figure 9: Change in the probability of self-employment and wage-employment relative to the counterfactual probability of employment. Estimated using the LAD (1990-2016)

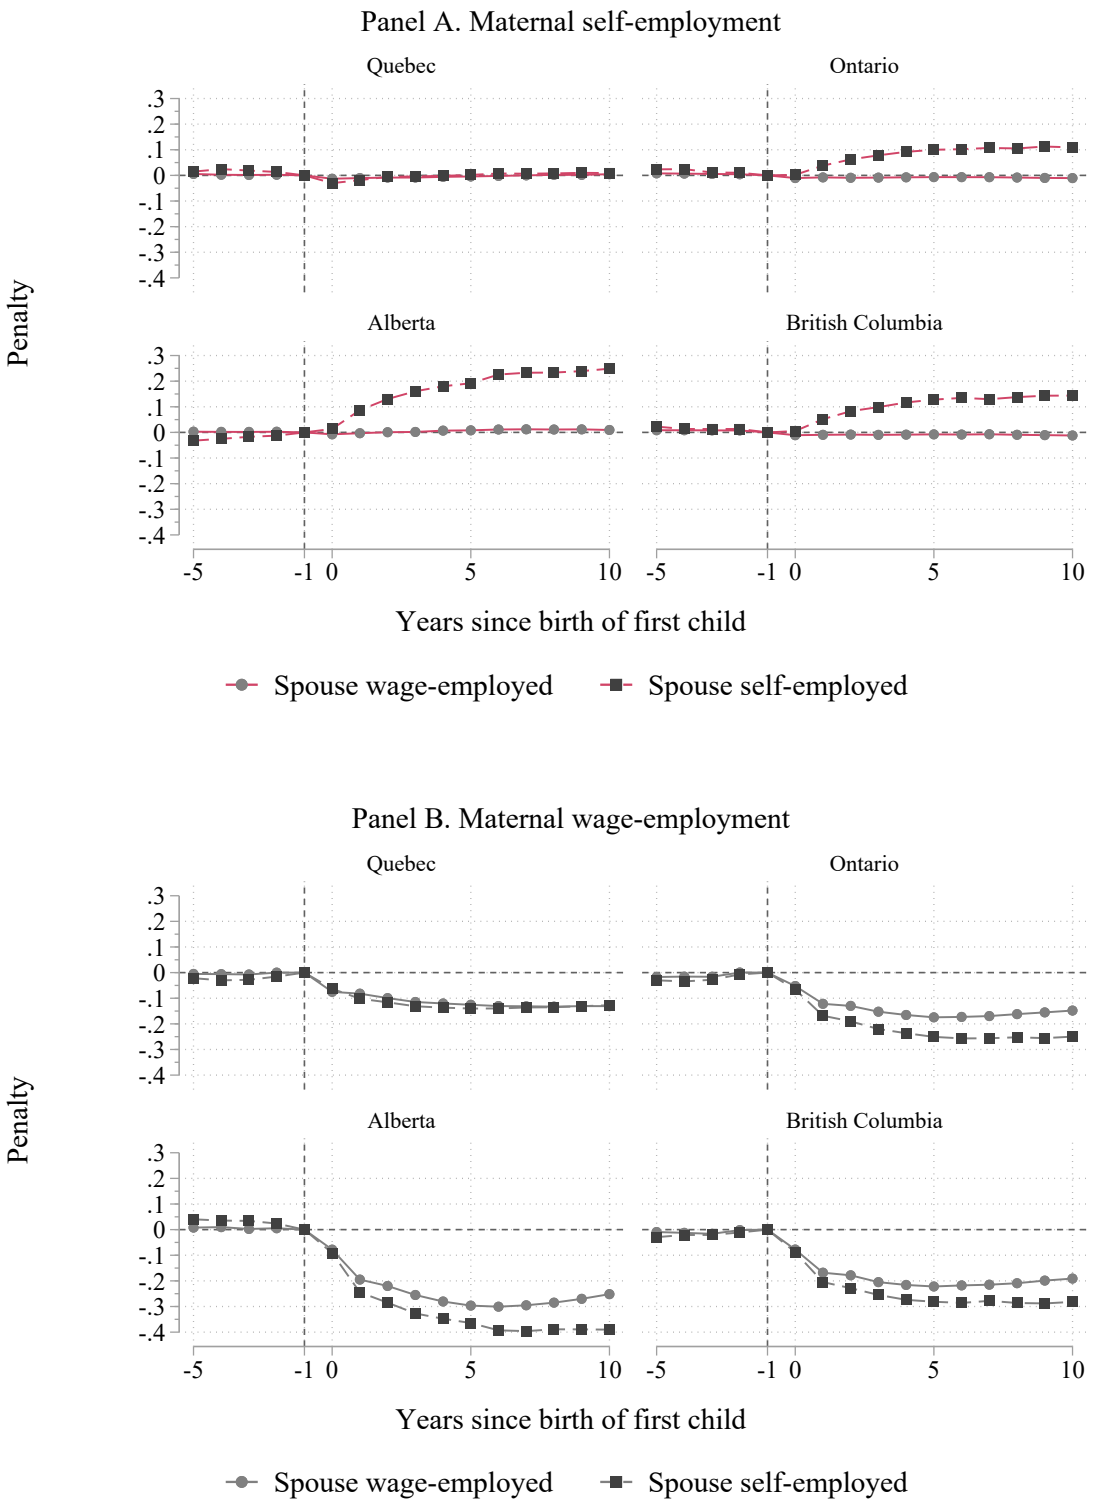

Figure 10: Change in the probability of maternal self-employment and wage-employment relative to the counterfactual probability of employment, by spouse's class of employment. Estimated using the LAD (1990-2016)

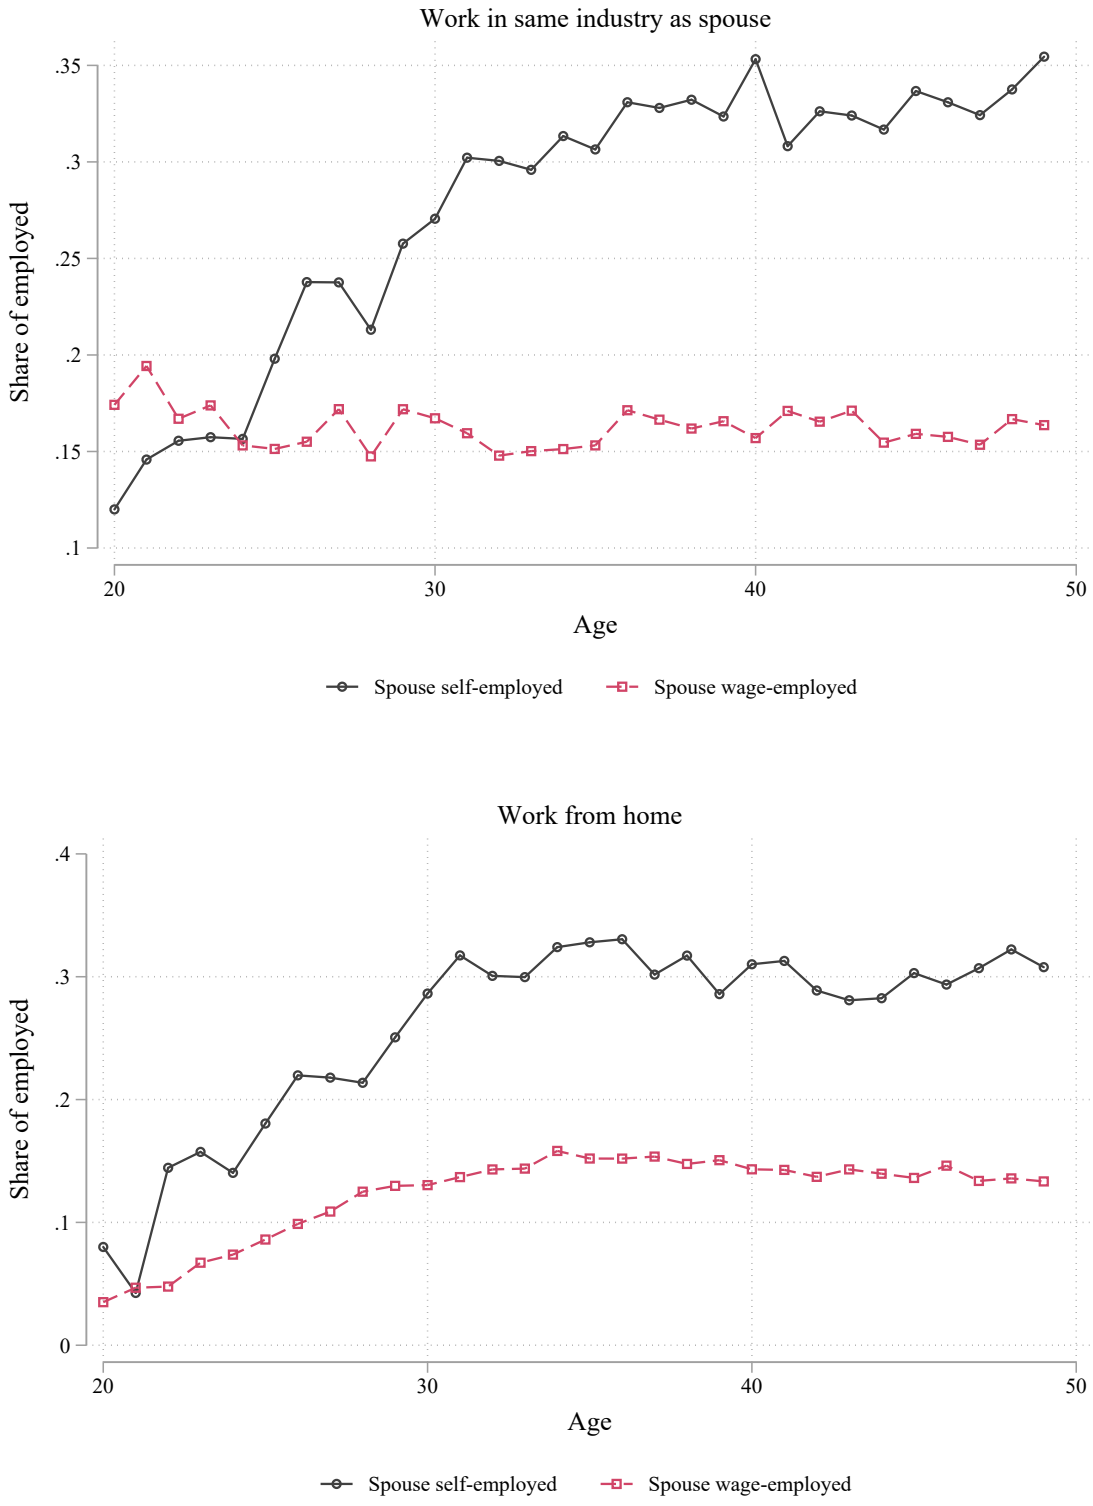

Figure 11: Characteristics of employment by age and self-employment status of spouse for married women. SLID (1996-2010)

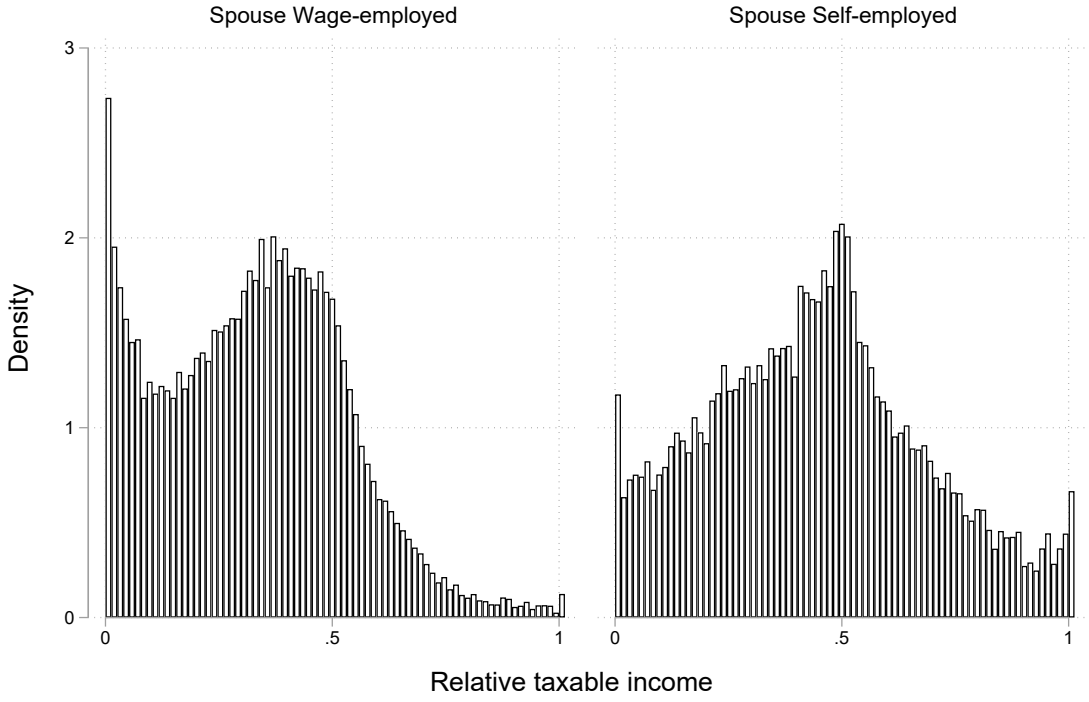

Figure 12: Relative earnings of women aged 25-44 with children by self-employment status of their spouse. Based on SLID/CIS income tax data (1996-2011)

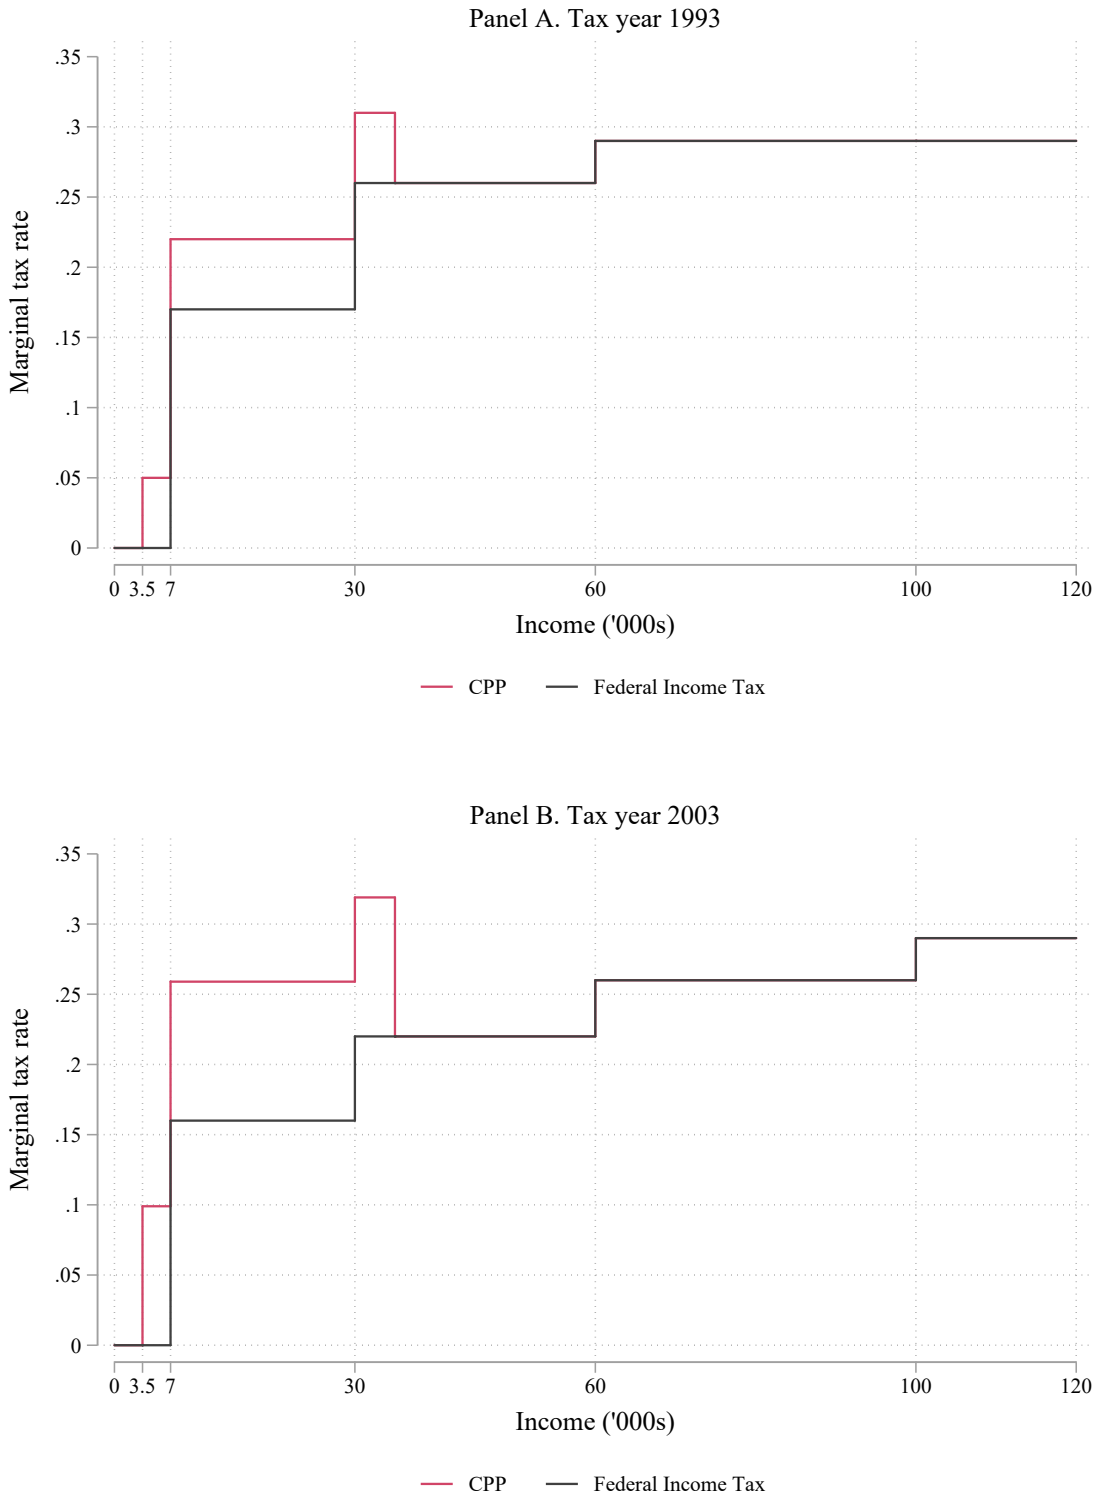

Figure 13: Federal income tax brackets and Canadian Pension Plan contributions for 1993 and 2003; assuming self-employed filer.

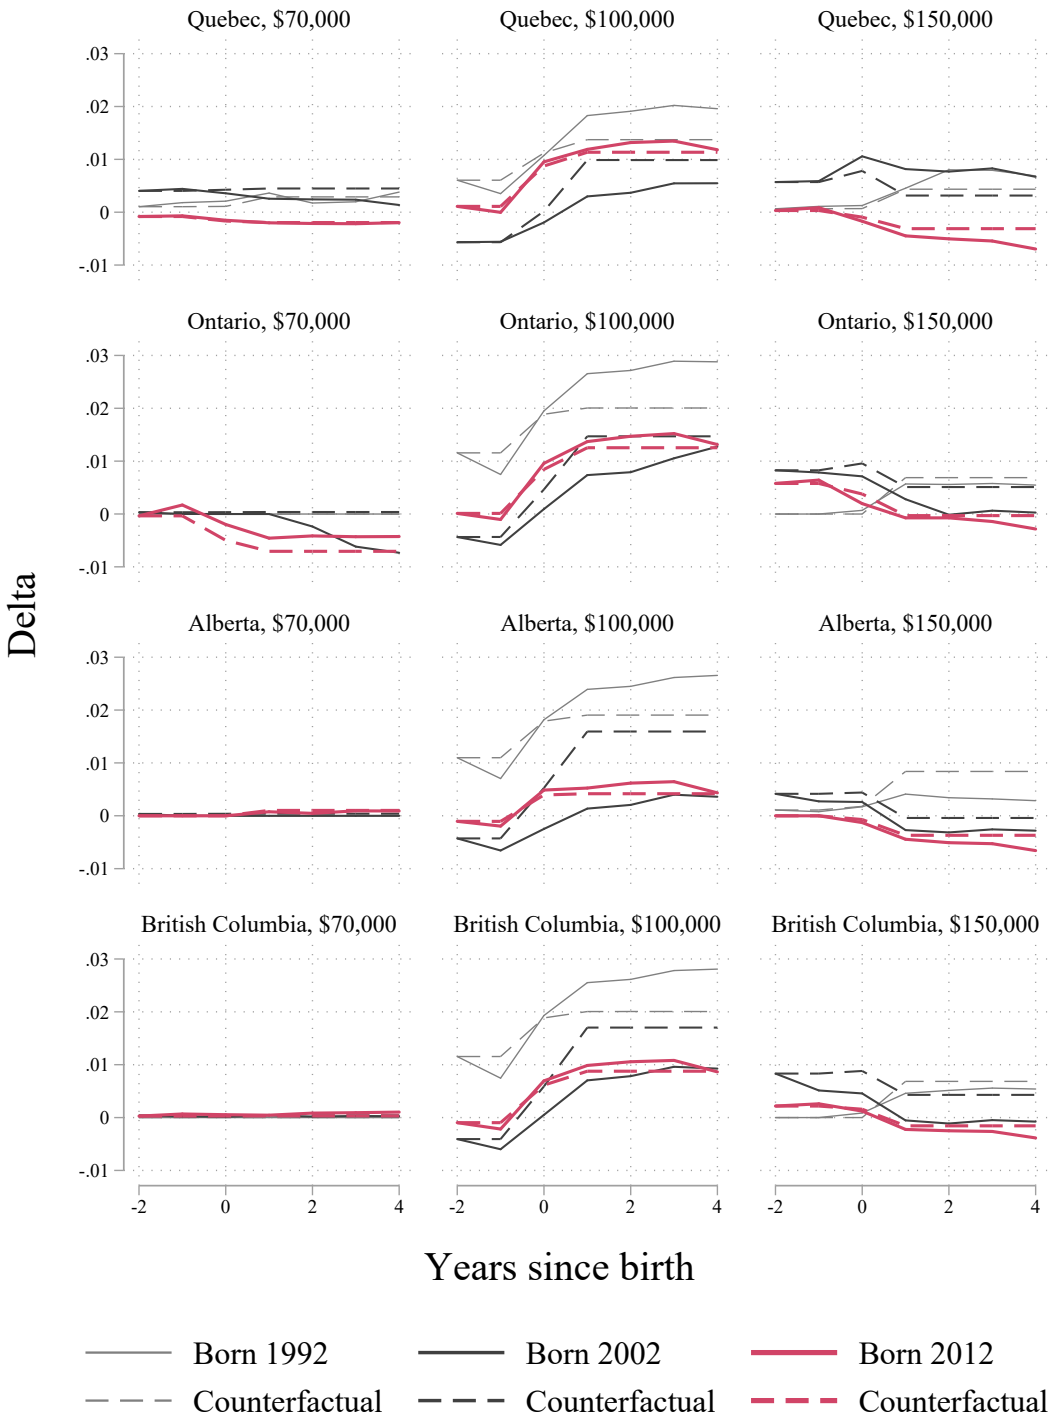

Figure 14: Simulation of delta for a household with a 40:60 income split by province and total household income (in 2000 CAD's)

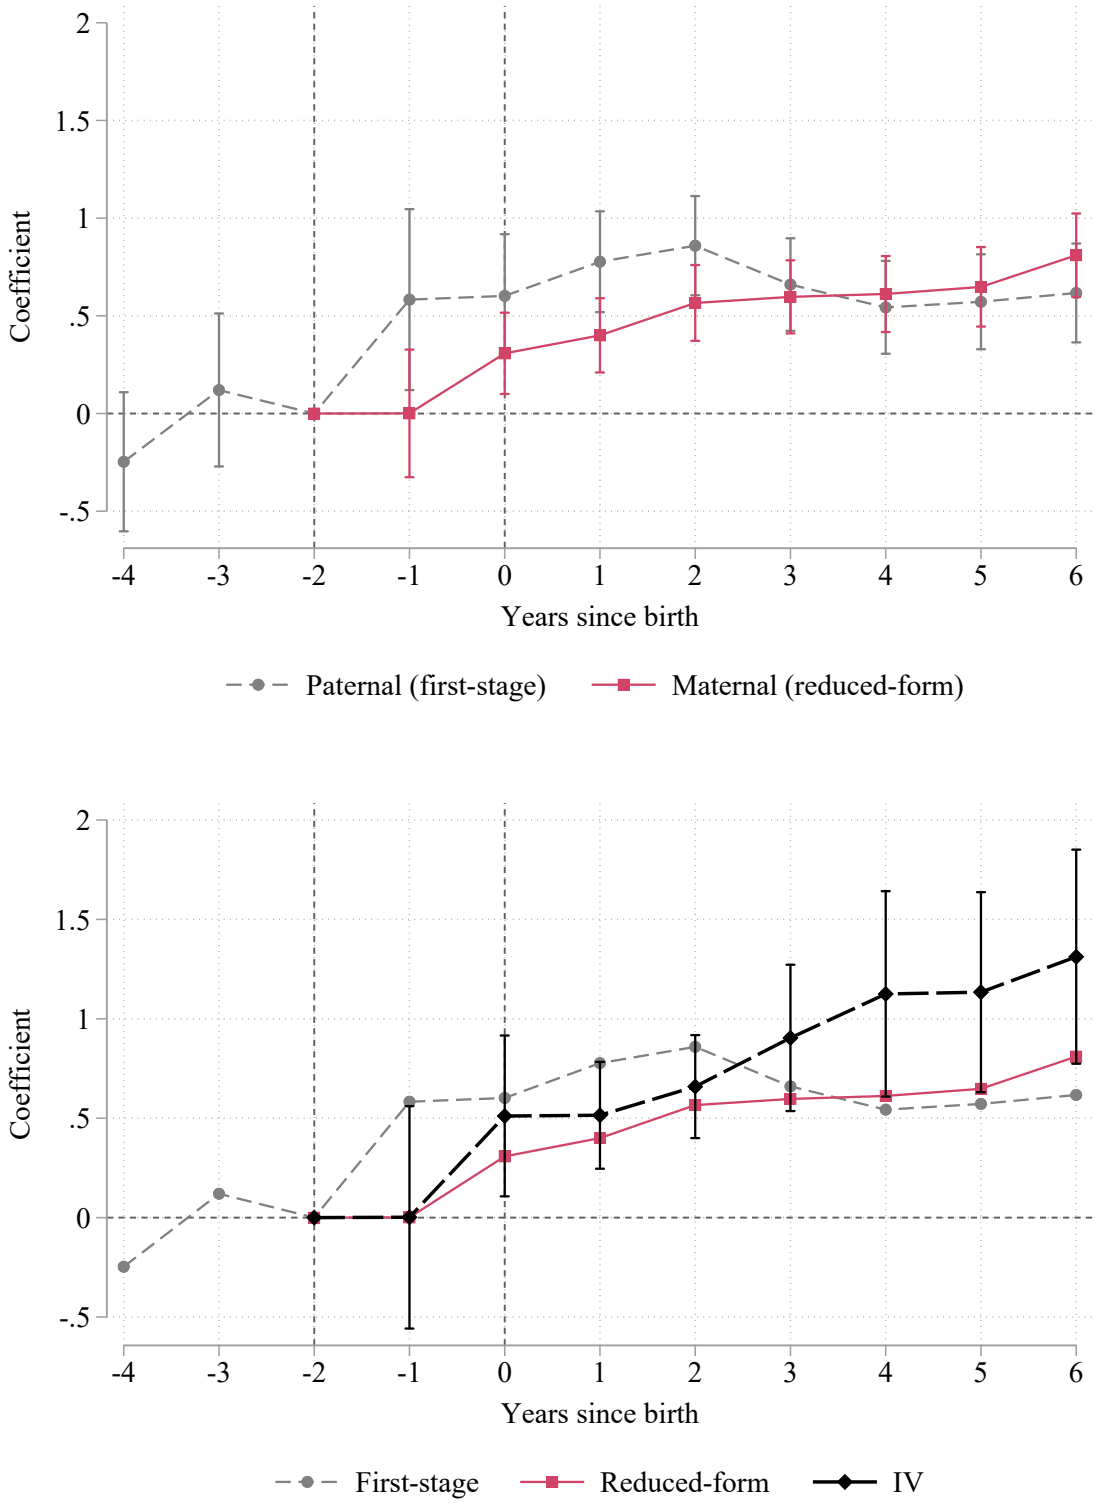

Figure 15: Reduced form results from event-time specific model: with inflation adjustment.

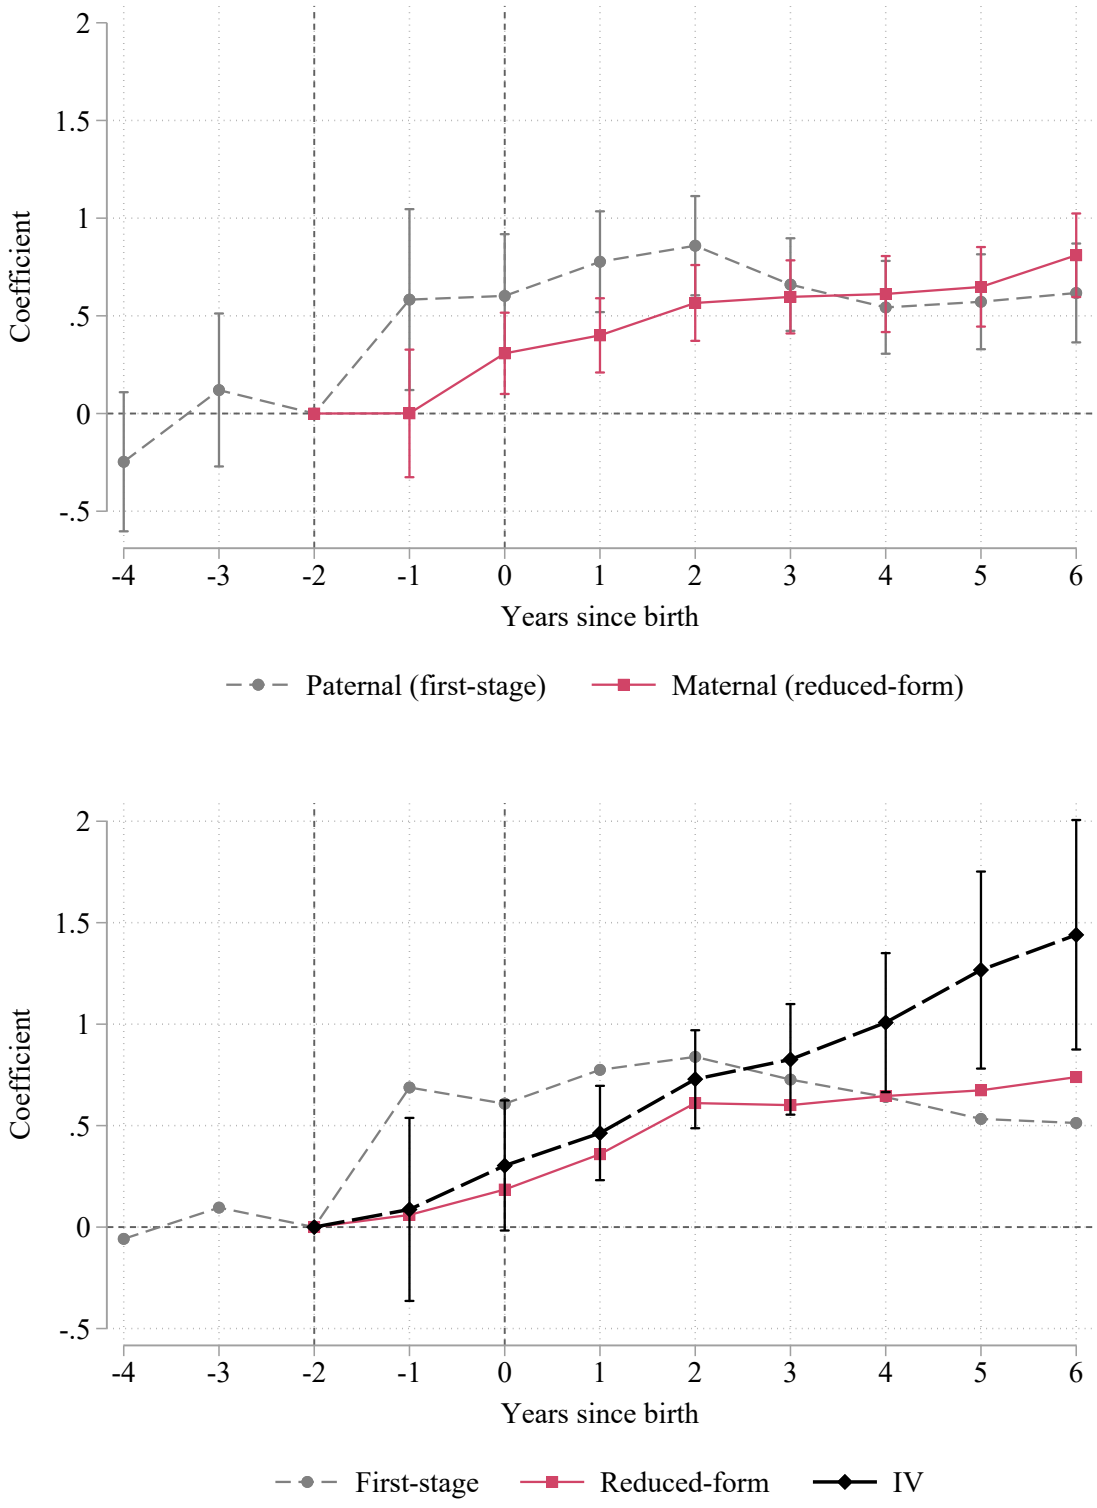

Figure 16: Reduced form results from event-time specific model: without inflation adjustment.

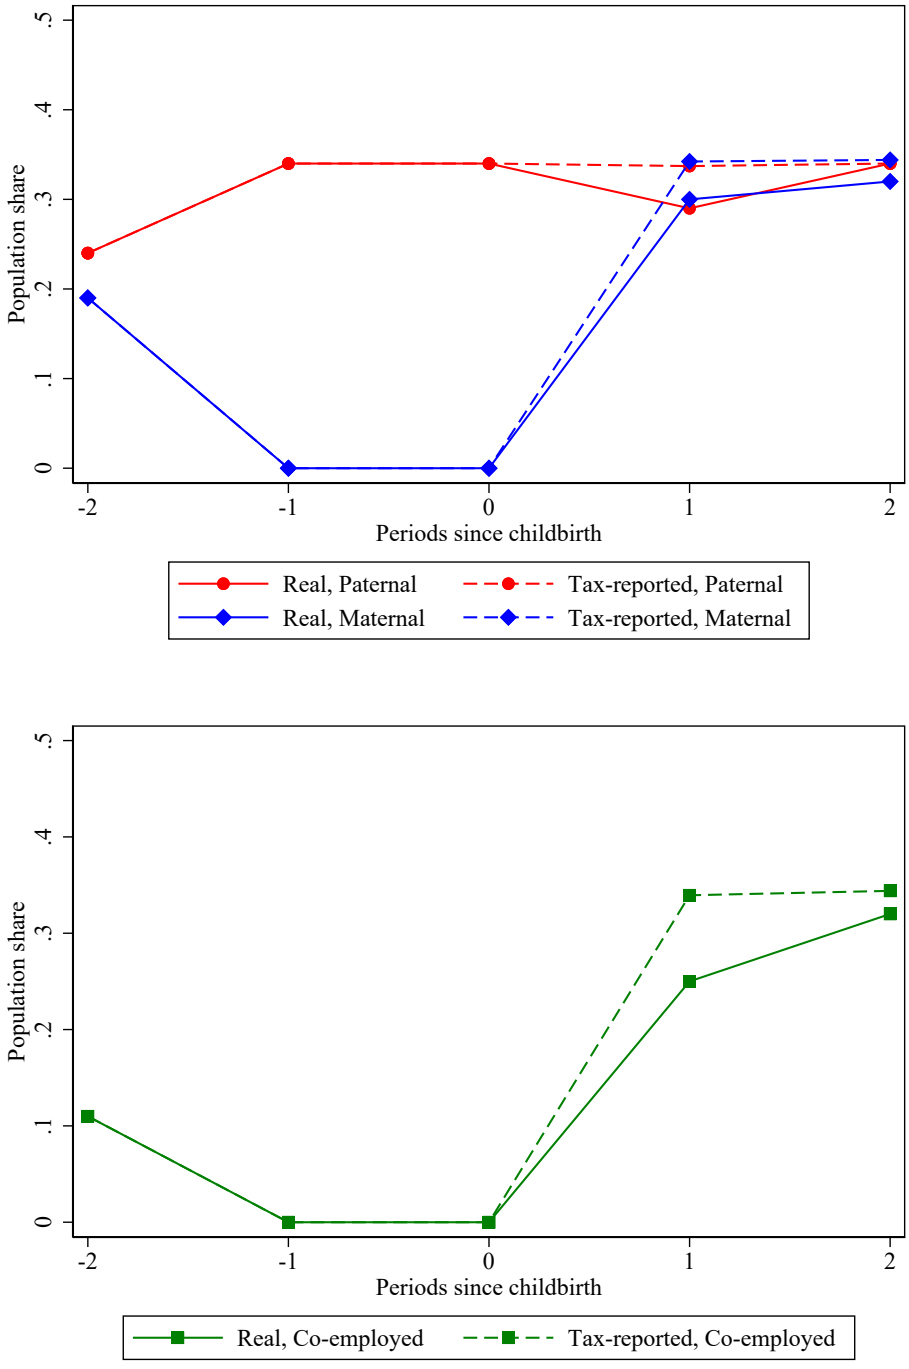

Figure 17: Graphs showing simulated levels of parental self-employment and co-employment in the standard model. Real labour supply denoted by solid line, with tax reported labour supply denoted by dashed line.

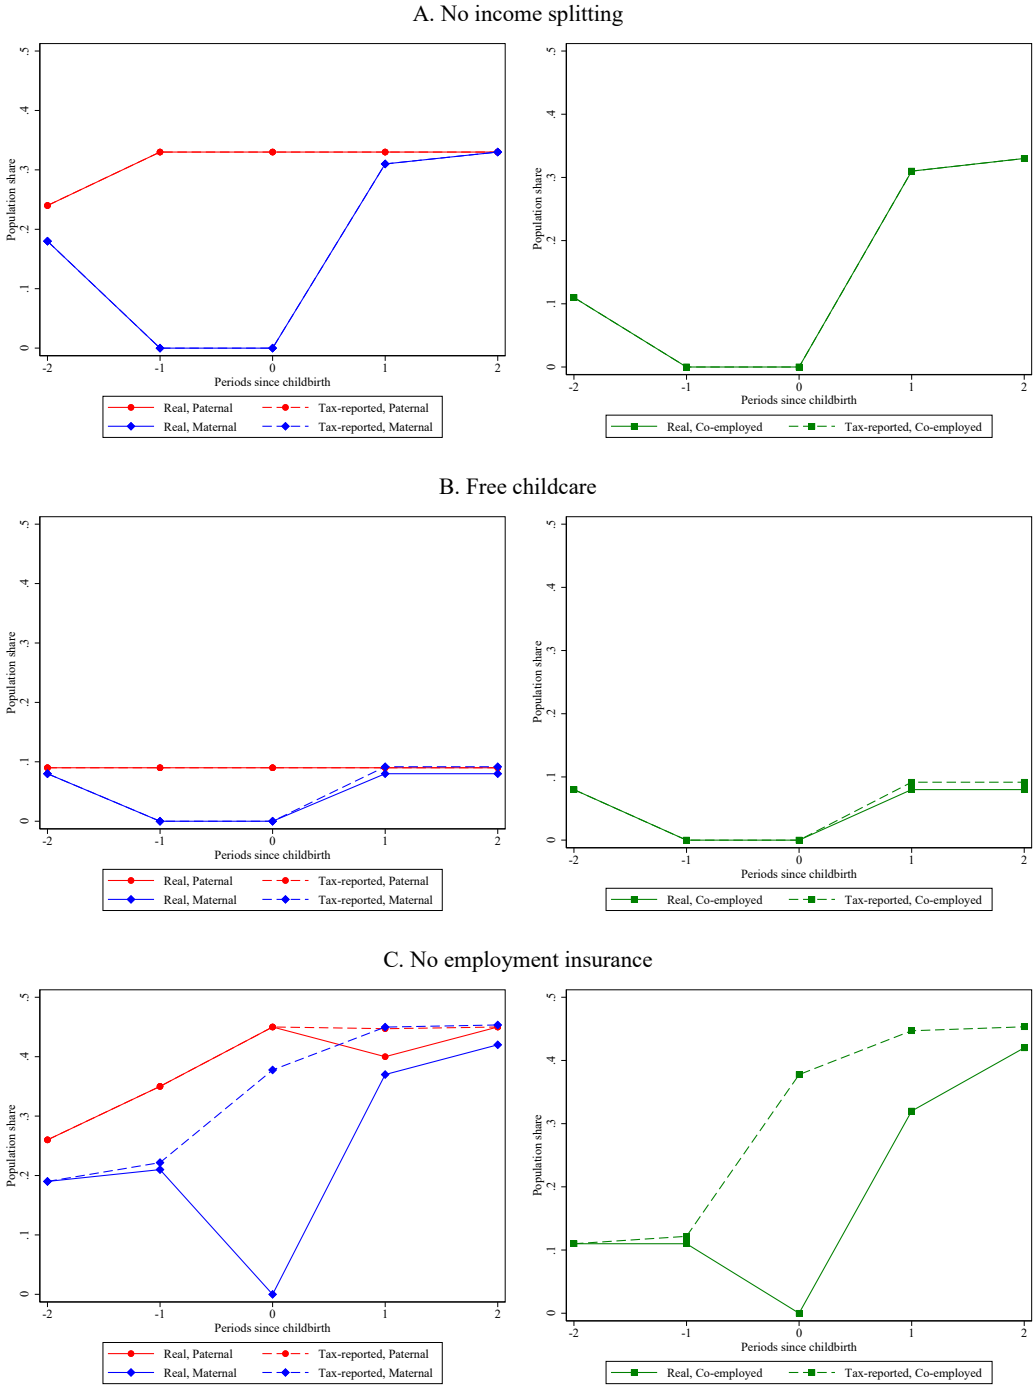

Figure 18: Graphs showing simulated levels of parental self-employment and co-employment under various policy counterfactuals. Real labour supply denoted by solid line, with tax reported labour supply denoted by dashed line.

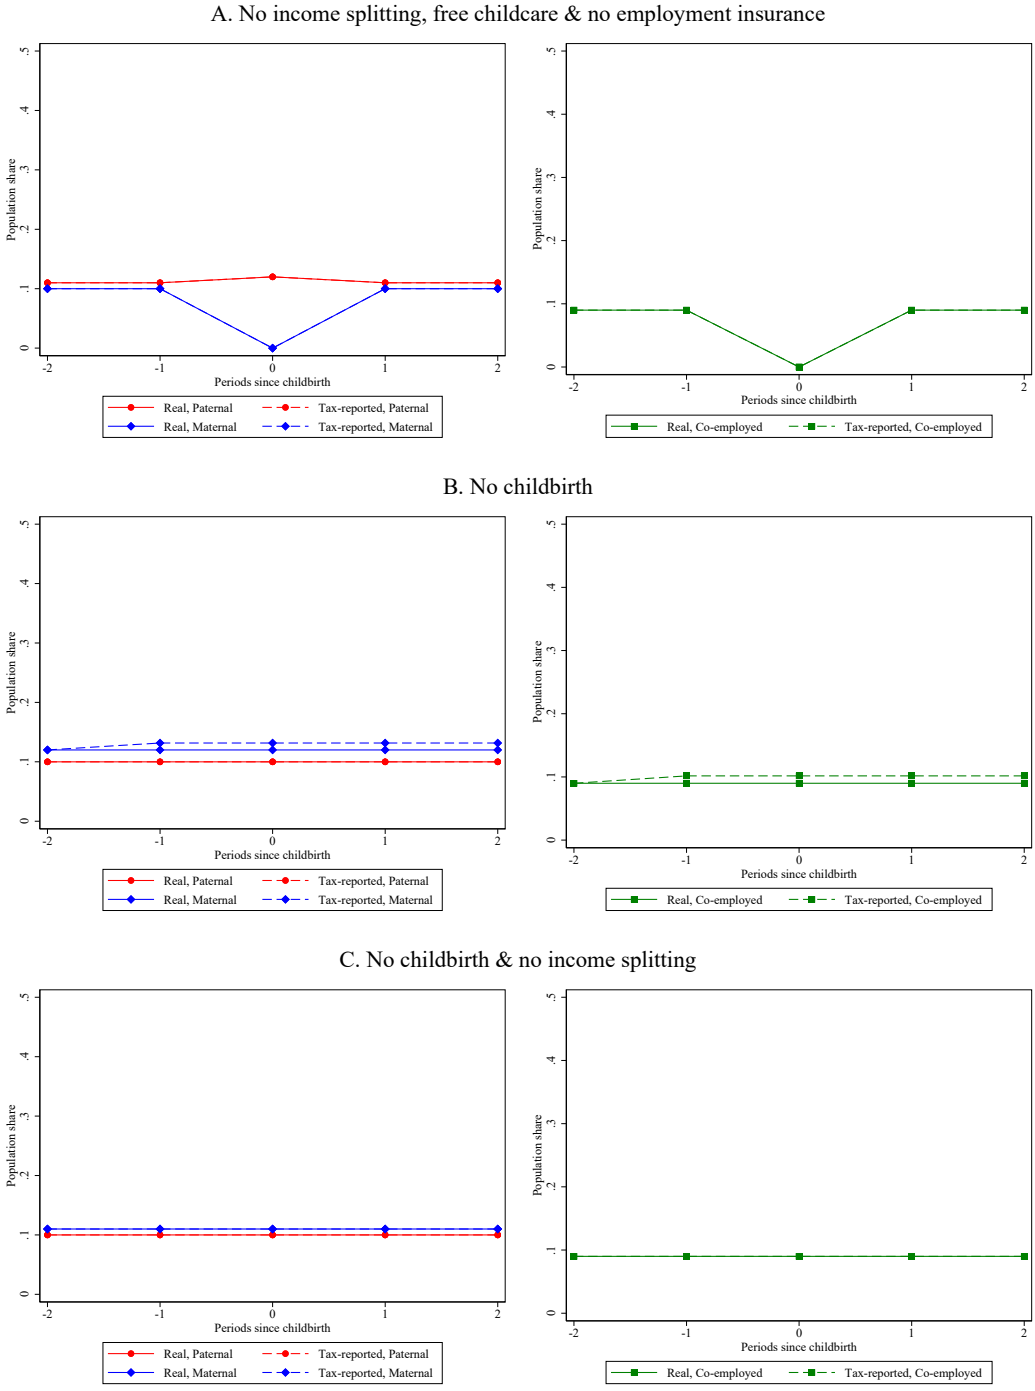

Figure 19: Graphs showing simulated levels of parental self-employment and co-employment under various policy counterfactuals. Real labour supply denoted by solid line, with tax reported labour supply denoted by dashed line.

A. Canada before federal income tax and EI reforms (1990's)

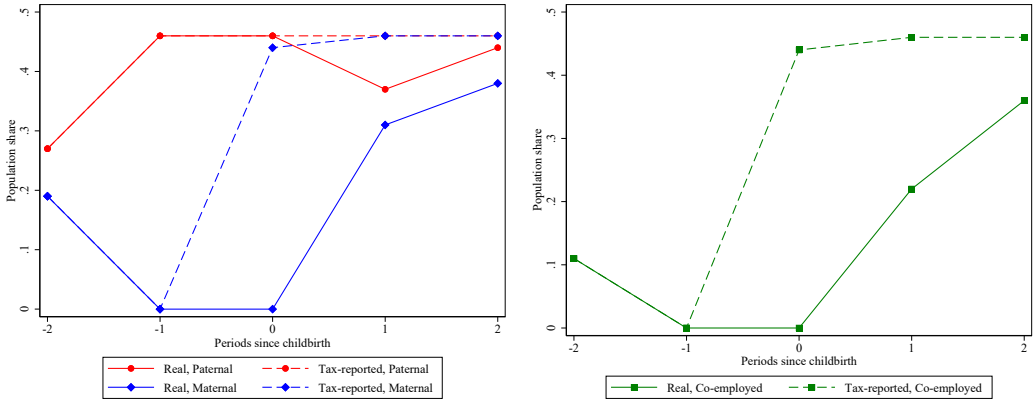

B. Canada after federal income tax reform (2000)

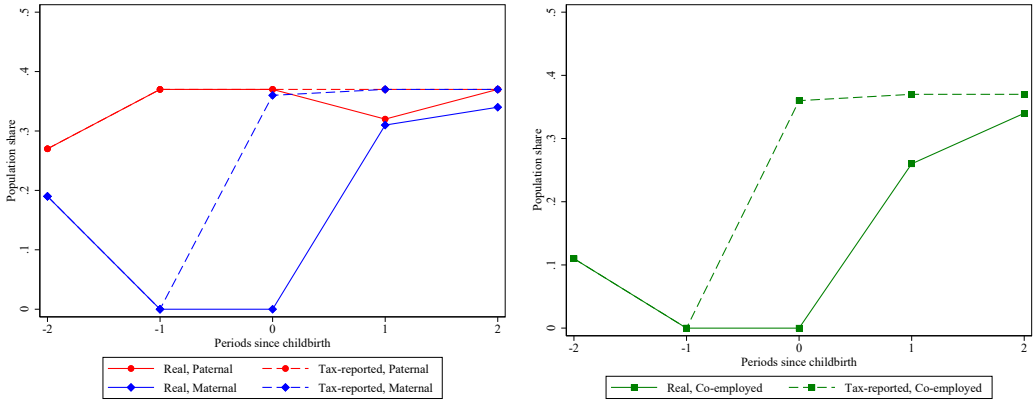

C. Canada after federal EI reform (2001)

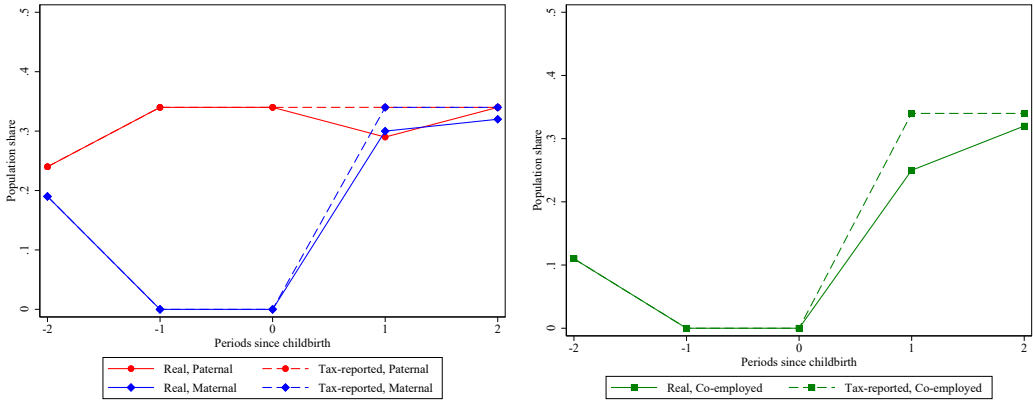

Figure 20: Graphs showing simulated levels of parental self-employment and co-employment before and after the federal income tax reform (2000) and federal EI reform (2001). Real labour supply denoted by solid line, with tax reported labour supply denoted by dashed line.

A. Quebec after childcare reform (1997)

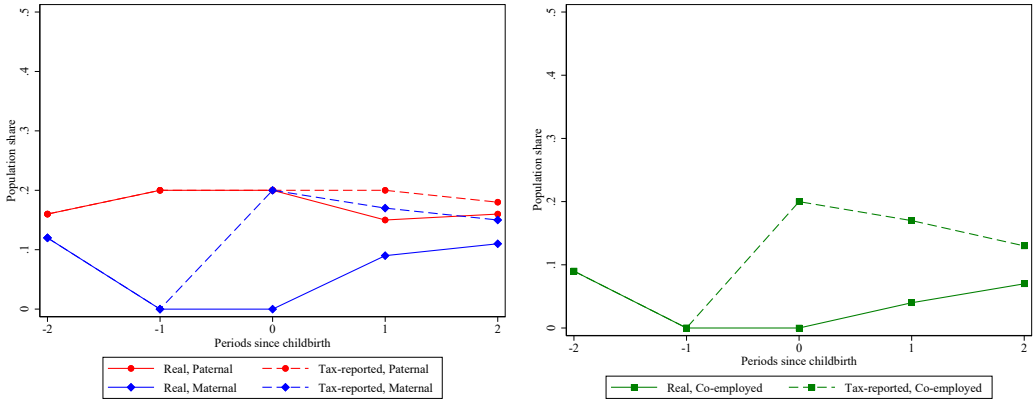

B. Quebec after federal income tax and EI reforms (2001)

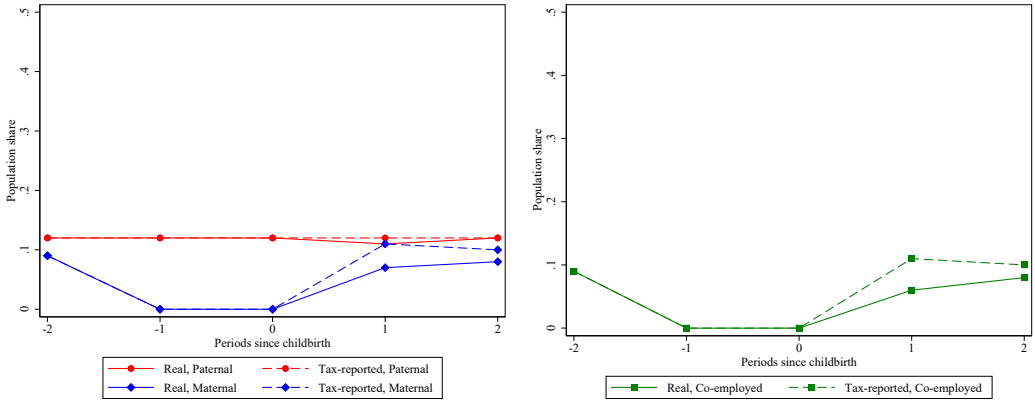

C. Quebec after introduction of QPIP (2006)

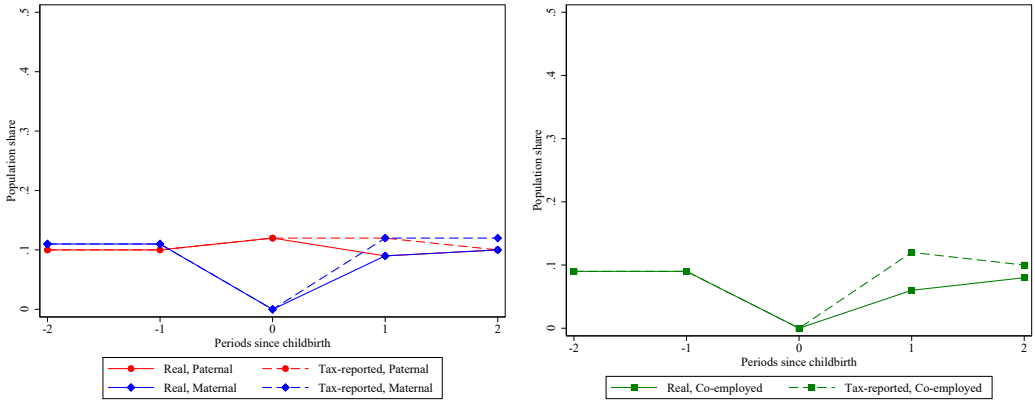

Figure 21: Graphs showing simulated levels of parental self-employment and co-employment for Quebec after various reforms. Real labour supply denoted by solid line, with tax reported labour supply denoted by dashed line.

List of Tables

Table 1: Employment of women by spouse's labour supply and number of children

	Employed	Self-employed (LFS)	Self-employed (TAX)
1 Child	-0.0653***	0.01000***	0.00491*
	(0.00314)	(0.00250)	(0.00276)
2 Children	-0.0837***	0.0181***	0.0147***
	(0.00294)	(0.00235)	(0.00259)
3+ Children	-0.165***	0.0313***	0.0271***
	(0.00359)	(0.00286)	(0.00316)
Spouse self-employed	-0.00648	0.133***	0.150***
	(0.00575)	(0.00459)	(0.00506)
1 Child*Spouse SE	0.0141*	-0.00742	-0.0135*
	(0.00809)	(0.00645)	(0.00712)
2 Children*Spouse SE	0.0274***	0.0162***	0.0151**
	(0.00719)	(0.00574)	(0.00633)
3+ Child*Spouse SE	0.0472***	0.0270***	0.0262***
	(0.00798)	(0.00636)	(0.00702)
Age	Yes	Yes	Yes
Year	Yes	Yes	Yes
Province	Yes	Yes	Yes
N	132262	132262	132262

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2: Job characteristics of women by spouse's labour supply and number of children

	Part-time	Weeks worked	ln(hours/week)	Work home	Same industry	Same occupation
1 Child	0.0662***	-1.170***	-0.143***	0.00430	-0.00199	-0.00967***
	(0.00392)	(0.0998)	(0.00596)	(0.00341)	(0.00364)	(0.00280)
2 Children	0.121***	-1.221***	-0.180***	0.0272***	-0.0229***	-0.0180***
	(0.00369)	(0.0940)	(0.00562)	(0.00321)	(0.00343)	(0.00263)
3+ Children	0.191***	-2.356***	-0.268***	0.0585***	-0.0134***	-0.0185***
	(0.00466)	(0.119)	(0.00709)	(0.00405)	(0.00432)	(0.00332)
Spouse self-employed	0.0383***	-0.364**	-0.00398	0.120***	0.0772***	0.0454***
	(0.00704)	(0.179)	(0.0107)	(0.00612)	(0.00653)	(0.00502)
1 Child, Spouse SE	-0.0193*	0.664***	0.0335**	0.0110	0.0195**	0.0181**
	(0.0100)	(0.255)	(0.0153)	(0.00872)	(0.00930)	(0.00715)
2 Children, Spouse SE	0.00769	1.370***	0.0140	0.0363***	0.0927***	0.0384***
	(0.00889)	(0.226)	(0.0135)	(0.00773)	(0.00825)	(0.00634)
3+ Child, Spouse SE	0.0136	2.253***	0.0148	0.0675***	0.166***	0.0852***
	(0.0101)	(0.257)	(0.0154)	(0.00878)	(0.00937)	(0.00720)
Age	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Province	Yes	Yes	Yes	Yes	Yes	Yes
N	112015	112015	111520	111524	112015	112015

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Coefficients from pooled models, with CPI adjustment

	OLS (1)	OLS (2)	IV (3)	OLS (4)	OLS (5)	IV (6)
α_{-1}	0.0107***	0.00917***	0.0158**	0.00959***	0.00814***	0.0128*
	(0.000250)	(0.000241)	(0.00743)	(0.000257)	(0.000248)	(0.00731)
α_0	0.0151***	0.0106***	-0.000792	0.0145***	0.0101***	-0.00513
	(0.000316)	(0.000296)	(0.00692)	(0.000342)	(0.000322)	(0.00730)
α_1	0.0263***	0.0173***	0.00216	0.0258***	0.0169***	-0.00131
	(0.000436)	(0.000403)	(0.00614)	(0.000472)	(0.000439)	(0.00634)
α_2	0.0336***	0.0201***	-0.00696	0.0329***	0.0195***	-0.00841
	(0.000517)	(0.000468)	(0.00698)	(0.000537)	(0.000490)	(0.00701)
α_3	0.0426***	0.0244***	-0.0249**	0.0419***	0.0238***	-0.0250**
	(0.000618)	(0.000559)	(0.0110)	(0.000641)	(0.000583)	(0.0109)
α_4	0.0516***	0.0288***	-0.0430**	0.0508***	0.0281***	-0.0379**
	(0.000728)	(0.000656)	(0.0172)	(0.000746)	(0.000676)	(0.0163)
α_5	0.0601***	0.0330***	-0.0454**	0.0591***	0.0321***	-0.0352**
	(0.000832)	(0.000747)	(0.0187)	(0.000846)	(0.000763)	(0.0171)
α_6	0.0680***	0.0364***	-0.0683***	0.0669***	0.0354***	-0.0525***
	(0.000956)	(0.000855)	(0.0219)	(0.000968)	(0.000869)	(0.0195)
β_{-1}		0.0570***	-0.197		0.0587***	-0.108
		(0.00308)	(0.282)		(0.00308)	(0.277)
β_0		0.0997***	0.352**		0.101***	0.444***
		(0.00298)	(0.156)		(0.00298)	(0.164)
β_1		0.150***	0.401***		0.150***	0.459***
		(0.00316)	(0.105)		(0.00316)	(0.108)
β_2		0.188***	0.562***		0.189***	0.579***
		(0.00323)	(0.0981)		(0.00323)	(0.0988)
β_3		0.214***	0.794***		0.214***	0.793***
		(0.00324)	(0.130)		(0.00324)	(0.129)
β_4		0.238***	0.990***		0.239***	0.932***
		(0.00329)	(0.181)		(0.00329)	(0.172)
β_5		0.256***	0.995***		0.256***	0.892***
		(0.00333)	(0.178)		(0.00333)	(0.163)
β_6		0.269***	1.165***		0.270***	1.022***
		(0.00336)	(0.189)		(0.00335)	(0.169)
Age (spouse)	No	No	No	Yes	Yes	Yes
ln(Earnings)	No	No	No	Yes	Yes	Yes
Commonlaw	No	No	No	Yes	Yes	Yes

Table 4: Coefficients from pooled models, without CPI adjustment

	OLS (1)	OLS (2)	IV (3)	OLS (4)	OLS (5)	IV (6)
α_{-1}	0.0107***	0.00917***	0.0123**	0.00960***	0.00814***	0.00872
	(0.000250)	(0.000241)	(0.00587)	(0.000257)	(0.000248)	(0.00583)
α_0	0.0151***	0.0106***	0.00715	0.0144***	0.00998***	0.0000627
	(0.000317)	(0.000297)	(0.00592)	(0.000343)	(0.000323)	(0.00626)
α_1	0.0263***	0.0173***	0.00335	0.0257***	0.0169***	-0.00197
	(0.000439)	(0.000407)	(0.00541)	(0.000474)	(0.000442)	(0.00567)
α_2	0.0339***	0.0203***	-0.0138**	0.0330***	0.0195***	-0.0156**
	(0.000522)	(0.000474)	(0.00656)	(0.000542)	(0.000495)	(0.00661)
α_3	0.0429***	0.0246***	-0.0227***	0.0420***	0.0239***	-0.0220***
	(0.000625)	(0.000566)	(0.00852)	(0.000646)	(0.000589)	(0.00834)
α_4	0.0521***	0.0291***	-0.0390***	0.0511***	0.0282***	-0.0321***
	(0.000738)	(0.000666)	(0.0121)	(0.000755)	(0.000685)	(0.0112)
α_5	0.0609***	0.0336***	-0.0656***	0.0597***	0.0325***	-0.0495***
	(0.000841)	(0.000756)	(0.0189)	(0.000854)	(0.000771)	(0.0166)
α_6	0.0689***	0.0371***	-0.0903***	0.0676***	0.0359***	-0.0638***
	(0.000965)	(0.000863)	(0.0242)	(0.000976)	(0.000875)	(0.0202)
β_{-1}		0.0570***	-0.0653		0.0587***	0.0465
		(0.00308)	(0.223)		(0.00308)	(0.221)
β_0		0.0997***	0.171		0.101***	0.327**
		(0.00298)	(0.133)		(0.00298)	(0.141)
β_1		0.150***	0.378***		0.151***	0.470***
		(0.00316)	(0.0920)		(0.00316)	(0.0964)
β_2		0.188***	0.656***		0.189***	0.681***
		(0.00323)	(0.0920)		(0.00323)	(0.0929)
β_3		0.214***	0.763***		0.214***	0.754***
		(0.00324)	(0.100)		(0.00324)	(0.0985)
β_4		0.238***	0.941***		0.239***	0.866***
		(0.00329)	(0.126)		(0.00329)	(0.117)
β_5		0.256***	1.186***		0.256***	1.029***
		(0.00333)	(0.179)		(0.00333)	(0.157)
β_6		0.269***	1.352***		0.270***	1.118***
		(0.00336)	(0.207)		(0.00335)	(0.174)
Age (spouse)	No	No	No	Yes	Yes	Yes
ln(Earnings)	No	No	No	Yes	Yes	Yes
Commonlaw	No	No	No	Yes	Yes	Yes

Appendix

A Simulation parameters

In the following simulation I make the following assumptions:

1. Sample of N=100 households.

2. Wage earnings are assuming to constant through time and uncorrelated within the household. They are drawn from the following distribution:

$$\begin{bmatrix} w^1 \\ w^2 \end{bmatrix} \sim N \left(\begin{bmatrix} 10 \\ 10 \end{bmatrix}, \begin{bmatrix} 1,0 \\ 0,1 \end{bmatrix} \right)$$

3. Self-employment earnings are dependent on wage earnings in the following way,

$$\begin{bmatrix} v^1 \\ v^2 \end{bmatrix} = \begin{bmatrix} w^1 \\ w^2 \end{bmatrix} + \begin{bmatrix} \eta^1 \\ \eta^2 \end{bmatrix}$$

   where the household entrepreneurship shock is correlated across spouses.

$$\begin{bmatrix} \eta^1 \\ \eta^2 \end{bmatrix} \sim N \left(\begin{bmatrix} -0.5 \\ -0.5 \end{bmatrix}, \begin{bmatrix} 1,0.5 \\ 0.5,1 \end{bmatrix} \right)$$

4. No discounting: $\beta = 1$

5. Disutility to work: $\theta^g(h_t^g) = \vartheta^g(h_t^g + s_t^g)$

$$\vartheta^g = U(0, 0.5)$$

6. Cost functions:

$$N_t(l_t) = \nu_t \left((h_1^1 + s_1^1)(h_1^2 + s_1^2) - \nu_s \cdot s_1^1 s_1^2 \right)$$

$$\text{with } \nu_{-2} = \nu_{-1} = \nu_0 = 0, \nu_1 = 5, \nu_2 = 2.5, \nu_s = 0.5$$

$$K(l_t) = \kappa \cdot (s_t^1 + s_t^2 - s_t^1 s_t^2) \cdot (1 - s_{t-1}^1)(1 - s_{t-1}^2)$$

$$\text{with } \kappa = 5$$

7. Tax structure. The marginal tax rate on the first 1.5 units of income is 0, followed by 20% on the next 4.5 units, 30% on the next 4 units, and 40% above 10 units.

8. Policy parameters

$$\psi = 0.55$$