Journal of Monetary Economics 150 (2025) 103719

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Discussion of "Singles, Couples, Time-Averaging, and Taxation"

Richard Blundell [a,b,*], Eric French [c,b]

[a] UCL, United Kingdom
[b] IFS, United Kingdom
[c] University of Cambridge, United Kingdom

ARTICLE INFO

Keywords:
Taxation
Labor supply
Retirement
Social security
Life-cycle

ABSTRACT

This discussion comments on "Singles, Couples, Time-Averaging, and Taxation" by Hans Holter, Lars Ljungqvist, Thomas Sargent, and Serhiy Stepanchuk. This work represents an exciting new agenda, combining recent findings in the dynamic micro labor supply and retirement literature with key features of the macro environment by including general equilibrium feedback and a government budget constraint. We point to some areas where the specification could be improved but these do not detract from the key insights in this paper and the significant steps it makes toward building a coherent micro and macro framework for research on labor supply and taxation.

1. Introduction

The paper represents the latest advance in a series of papers that have brought increasing realism when modeling the aggregate effects of taxes. Their paper exemplifies the best of bringing together microfoundations of aggregate behavior with a general equilibrium model to consider the aggregate consequences of that micro behavior. There are many attractive aspects to the approach taken in this paper. Being both microeconomists, we will place more focus on micro aspects of their paper and, at the same time, share some broader thoughts on this area of research.

The life-cycle labor supply setting is at the center of the analysis. Looking across the life-cycle highlights key differences in labor supply. For example, despite strong differences in tax incentives, employment rates are surprising similar for prime age men in US, UK, France, and Germany. Nonetheless, there are strong differences in hours and employment at older and younger ages (Blundell et al., 2011). There are also large cross-country differences in labor supply for women with young children. These differences become accentuated when looking separately at different education groups. This motivates the life-cycle approach to assessing (aggregate) labor supply responses to tax reforms, accounting for fixed costs of work, human capital and families. Precisely the aim of this paper.

The modeling approach makes several advances to better capture the aggregate response to taxes. The consensus view in labor economics (see, e.g., Blundell and MaCurdy, 1999) has been that labor supply is fairly inelastic, at least for on the intensive margin (i.e., the number of hours of work conditional on working) for men. However, there are several reasons to think that labor supply might be more elastic. Holter et al. allow for all of these reasons. First, their model allows for both an intensive and extensive margin of labor supply. Second, they allow for realistic lifecycles that include Social Security and retirement. Third, they allow for households with both husbands and wives. Fourth, they allow for experience capital and persistent productivity shocks. Finally, they allow for general equilibrium responses and government revenue constraints. We discuss the empirical relevance of each of these innovations below.

* Corresponding author at: UCL, United Kingdom.
E-mail address: r.blundell@ucl.ac.uk (R. Blundell).

https://doi.org/10.1016/j.jmoneco.2024.103719
Received 30 October 2024; Accepted 31 October 2024
Available online 8 November 2024

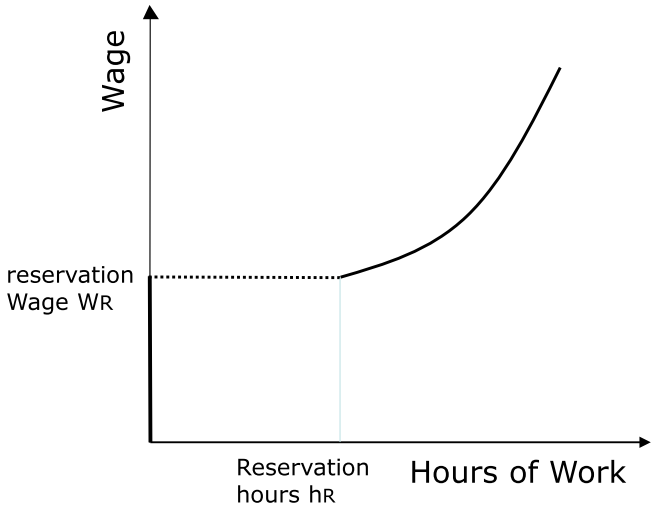

Fig. 1. Labor supply function.

2. The elasticity of labor supply over the life-cycle

The model developed in Holter et al. allows for both an intensive and extensive margin of labor supply. While most studies of labor supply suggest that labor supply is relatively inelastic on the intensive margin, there is some evidence that it is more elastic on the extensive margin, at least for certain groups such as younger individuals, women with young children, and those nearing retirement. By allowing for lifecycles with retirement and households, they are able to model different groups that likely have different labor supply elasticities, including groups with likely high labor supply elasticities.

A key finding in economics over the last few decades is the importance of considering both the intensive and extensive margins of labor supply jointly for understanding aggregate labor supply. Consider the utility function for a single (the problem for a couple is similar to the problem for a single, with each member of a couple facing a disutility of work):

$$U^S(c,n,i) = \log(c) - \chi_S^i \frac{(n)^{1+\eta^i}}{1+\eta^i} - F_S^i \cdot 1_{[n>0]}$$

where c is consumption, n is work hours, i is the gender of the individual. The parameters F_S^i and η^i are gender specific to allow for the possibility that female labor supply is more elastic than male labor supply due to differences in preferences. $1_{[n>0]}$ is an indicator equal to 1 when work hours are positive. Note that if $F_S^i = 0$, the marginal disutility of work hours evaluated at 0 h is $\chi_S^i(n) = \chi_S^i(0) = 0$. Thus so long as (after-tax and transfer) wages are positive, the individual will always work positive hours. However, if $F_S^i > 0$, then the individual must pay this fixed (utility) cost if they work any positive number of hours. As a result, it is never optimal to work an infinitesimal number of hours. Thus the optimal decision rule will follow the one in Fig. 1. Below a certain reservation wage level w_R, the individual will not work. Above that reservation wage level, the individual will work a positive amount h_R. The figure highlights the two margins of labor supply: the extensive margin (whether to work at all) and the intensive margin. Both margins will contribute to aggregate labor supply.

While there is a long history of formulating and estimating dynamic models with both intensive and extensive margins labor supply (e.g., Heckman and MaCurdy, 1980 and Blundell et al., 1998), most of these papers assumed no fixed cost of work (see Cogan, 1981 for an early exception). Thus they implicitly assume that an individual on the margin of whether or not to work is an individual on the margin of whether or not to work an infinitesimally small number of hours. Thus, the employment margin is unimportant for aggregate labor supply since the set of individuals at the employment margin only contribute a very small number of hours when they enter employment. Beginning with Gary (1985) and Rogerson (1988), however, macro studies began to bring in an extensive margin of labor supply where individuals were on the margin of whether to work would work significant hours.

Dynamic models that include fixed costs of work are now prominent in the retirement literature (e.g., French, 2005 and references in Blundell et al., 2016). It would be of value to link the current approach back to some of these studies, many of which are carefully matched to the data, and provide measures of both intensive and extensive labor supply elasticities.

As noted by Keane (2011) and Keane and Wasi (2016), labor supply could be more elastic when accounting for human capital accumulation. While wages vary much more than hours over the lifecycle (suggesting that hours are insensitive to wages), when accounting for human capital accumulation, the return to work (wages plus the value of human accumulation) varies much less over the lifecycle, suggesting a larger labor supply elasticity. Below we return to the specification with work experience, persistent shocks and heterogeneity in the determination of wages.

Finally we note that the model in this paper includes family labor supply decisions which are likely important for retirement decisions. There is assortativeness in marriage in the model and there is suggestive evidence that this may have been increasing

Table 1
Delayed retirement credit.

Year of birth	Credit per year
1917–24	3.0%
1925–26	3.5%
1927–28	4.0%
1929–30	4.5%
1931–32	5.0%
1933–34	5.5%
1935–36	6.0%
1937–38	6.5%
1939–40	7.0%
1941–42	7.5%
1943 and later	8.0%

Notes: The midpoint of the sample period the authors use is 1983. Those turning the Normal Retirement Age of 65 in 1983 would then have a delayed retirement credit of 3% per year. See https://www.ssa.gov/oact/quickcalc/early_late.html for details.

in the US, see Blundell et al. (2018, Fig 8). An apparently strong assumption is that of coincident joint retirement but, perhaps surprisingly, this maybe a reasonable approximation. Banks Banks et al. (2015) find that the modal retirement age difference is zero in both the US and the UK even for a subsample of couples where the husband is at least 2 years younger than the wife.

3. Social security and retirement

Holter et al. assume that once an individual turns 65, s/he becomes eligible for Social Security benefits. The individual cannot work and draw benefits at the same time. Furthermore, if the individual forgoes claiming benefits and works an additional year, those benefits are lost. These are not innocuous assumptions and likely leads to an overstatement of the importance of Social Security on labor supply. Since Social Security benefits replace on average 40% of earnings when working, this loss of benefits when working effectively imposes a 40% tax on the labor supply of households, on top of all the other taxes they must pay. It is clear that in the model, Social Security is a powerful labor supply disincentive. The dramatic model-predicted falls in employment at 65 should come as no surprise given Social Security's large implicit tax on work after this age. In reality, however, the work disincentive effect from Social Security is likely less strong.

Whereas in the model delaying benefit receipt results in the loss of a year's benefit, in reality delaying benefit receipt leads to higher future benefits. In the US, most individuals are eligible to collect Social Security benefits at the "Early Retirement Age" of 62. For every year prior to the "Normal Retirement Age" (which was 65 for most of the sample period the authors use in the paper and will be 67 for those born 1960 or later) that individuals begin drawing benefits, those benefits are reduced 6.7% for every year of early benefit receipt (which is approximately actuarially fair). After the Normal Retirement Age, delaying benefit receipt leads to higher future benefits through the "delayed retirement credit". Thus delaying benefit receipt raises future benefits, although not necessarily in a way that is actuarially fair. Table 1 shows how the delayed retirement credit depends on year of birth.

To give a sense of the implicit tax from delaying benefit receipt, suppose the delayed retirement credit is 8% per year (which is the current credit), which is actuarially fair. In this case the implicit tax on work is modest since the individual can work and delay benefit receipt in a way that does not impact the discounted value off Social Security benefits. For the cohort born in 1927 or 1928 the delayed retirement credit was 4% per year. A crude back of the envelope calculation is then that the rules replace half of what is actuarially fair. In this sense, the implicit tax from Social Security is about half as big as the model implies. If Social Security benefits replace 40% of earnings from work and the delayed retirement credit replaces half of those lost benefits, then the implicit tax is then 0.5*40% = 20%.

The above calculation likely overstates the implicit tax on work for some households because an individual can work and draw benefits at the same time. While they will potentially face higher implicit tax rates if they draw benefits and work at the same time, for many these taxes will be modest. For those with low earnings, there is no tax penalty for working and claiming benefits at the same time. For higher income individuals, there are tax penalties (see Jones and Li, 2018 for details). First, if household income is sufficiently high, Social Security benefits become taxable. Second, individuals potentially face an "Earnings Test" tax, which can tax away benefits at rates of up to 50% for every additional dollar of earnings gained. But even benefits lost through the Earnings Test lead to future benefits being recalculated upwards.

Here we discuss the Earnings Test in greater detail. The Social Security Earnings Test taxes benefit income for Social Security beneficiaries at a very high rate. If a beneficiary younger than age 70 (or 72 pre-1983 or the Normal Retirement Age post-2000) earns more labor income than a 'test' threshold level of $4920 in 1983 (see Table 2 for other years), benefits are taxed at a 50% rate (33% at or above the Normal Retirement Age) until all benefits have been taxed away. The incentive to draw benefits by age 65 in combination with the Social Security earnings test for Social Security beneficiaries is a major disincentive for work after age 65, but less strong than what is implied by Holter et al. In truth, the "Social Security reform" may have more accurately captured the Social Security rules prevailing during the sample period than in the 'benchmark economy'.

A fuller description of the Social Security benefit rules of the time is described in Myers (1993) a 900 page long book! Obviously, not all of this detail can be built into a single model, and certainly not a model with such richness in so many dimensions such as this. In addition, private pensions and Medicare potentially induce retirement at age 65 (Eric and Jones, 2011). Nevertheless, it is worth noting that the incentives built into the model likely overstate the work disincentives of Social Security.

Table 2
Annual exempt amounts under the retirement earnings test.

Year	Under age 65	Age 65–59	Year	Under age 65	Age 65–69
1975	2520	2520	1988	6120	8400
1976	2760	2760	1989	6480	8880
1977	3000	3000	1990	6840	9360
1978	3240	4000	1991	7080	9720
1979	3480	4500	1992	7440	10,200
1980	3720	5000	1993	7680	10,560
1981	4080	5500	1994	8040	11,160
1982	4440	6000	1995	8160	11,280
1983	4920	6600	1996	8280	12,500
1984	5160	6960	1997	8640	13,500
1985	5400	7320	1998	9120	14,500
1986	5760	7800	1999	9600	15,500
1987	6000	8160			

Notes: All amounts in nominal dollars.

4. The life cycle wage process: experience profiles and idiosyncratic shocks

Log wages, for any individual in any period, are modeled as a function of accumulated experience, unobserved heterogeneity and a productivity shock. These three aspects of wage determination have been found to have a substantive impact on life-cycle responses to tax reform (Blundell et al., 2017). The return to work due to experience gained at work represents the dynamic value of work and will tend to offset the incentive to move in and out of work in response to transitory changes in tax incentives or wage shocks. The shape of the experience profile over the working life will impact the choice of when to retire in response to the implicit tax on work in social security. Individual heterogeneity is important both to account for differential selection of individuals into work and also to ensure that unobserved permanent individual differences are not mistaken as persistent shocks. Finally, the persistence of shocks is important as the higher the persistence the harder it is to smooth these shocks intertemporally. The paper also allows these components to differ by skill and gender as is typically found in the literature.

There are restrictions in the chosen empirical specification of each of these features of the wage process which could affect the simulated life-cycle responses to tax and Social Security incentives. Actual work experience is endogenous in the model as it depends on past labor supply choices and is therefore correlated with past values of persistent shocks and with individual heterogeneity in wages. The paper adopts a 2-step estimator where experience is removed from log wages in an initial step. In general, this will lead to biased parameter estimates of the experience profile or wages and the degree of persistence in the shock process. Typically such a dynamic wage equation is estimated 'within' the estimation of the dynamic structural model to allow labor supply to feedback through experience and wages. This is worth further analysis as the paper estimates an unusually low value for persistence in the shocks to wages. The authors acknowledge this and correctly point out that allowing for heterogeneity and experience likely reduces the estimated persistence in shocks. However, as noted above there still could be a downward bias coming from their two-step approach. A low persistence and biased experience profile matters as it directly impacts the gains to tax smoothing and also impacts the response to tax incentives, both of which are analyzed in the behavioral simulations reported in the paper. The more important the persistent idiosyncratic shocks, the less 'room' for time averaging. Furthermore, persistent shocks create increase precautionary behavior that can dampen labor supply responses.

There are other restrictions on the experience profiles of wages which could usefully be relaxed including an assumption of multiplicative shifts in an otherwise common woman-specific labor market experience profile. Differences by gender are often found to be driven by the extensive margin and part-time work for women with young children which do not enter this model. Finally, there is evidence that the variance of persistent wage shocks varies across the working life with a high variance at the beginning and end of working lives, see the evidence for US men by education, race, and cohort from the US Social Security earnings file matched to CPS in Blundell et al. (2024). One potential implication of a low persistence of shocks is that movements in and out of full-time work at older ages become more likely. If persistence was high and wage shocks were really permanent then these would permanently knock people out of the labor market if they were near retirement age.

Finally, it is also worth making a brief comment on the depreciation of human capital. As Xiaodong et al. (2024) point out, human capital accumulation plus depreciation creates an incentive to cluster hours into a shorter period of the working life. This is an important mechanism in the current paper. There is clear evidence of declining physical health (and more mixed evidence of declining cognitive skills) which impacts labor supply, especially among lower education groups. See the evidence from the US and UK longitudinal health and retirement surveys in Blundell and Britton (2023) which finds that health is a key driver of 'retirement' in US, especially among high school dropouts. Including a health process directly in the model would seem a worthwhile extension.

5. Tax reform simulations with GE effects

As we noted above, the paper shows important and substantive new results on the impact of tax reforms on labor supply. The authors document large labor supply impacts for singles and for couples of moving to a flexible Social Security system.

The heterogeneity in the model by marriage allows new insights of impacts by marriage. For example, in a novel ex-post analysis, the authors examine individual paths of women who draw the poor outcome of staying single, highlighting relative advantage of marriage in the current US tax system. The tax reform analysis shows the efficiency of moving to a flat tax system during working life. Welfare losses at the bottom are somewhat offset by a deduction but this is shown to heavily reduce labor supply gains.

A key issue in deriving the impact of tax reform is the importance of specifying what is done with the revenue raised. The authors highlight the comparison between a tax reform with a lump sum transfer and one without any transfer. This is analogous to an analysis with and without revenue neutrality in public finance, using a lump sum to guarantee neutrality. Thankfully this paper does not assume quasilinear preferences and consequently income/wealth effects from the transfer will be important. They document two mechanisms that differentiate the impact of tax reform with and without lump sum transfers. The first is the GE effect on factor prices and the second is the impact through the government budget constraint. It is hard for a government to spend revenue without impacting labor or capital markets. Consequently, there is a strong interaction between imposing revenue neutrality and the GE effect. The GE impact seems much more important when there is no lump sum transfer and revenue neutrality is not imposed. With the lump sum transfer the GE effects are somewhat attenuated. This seems a key lesson. Of course, it will depend on the reform in question and the method of achieving revenue neutrality but nonetheless an important result.

6. Summing up

The paper develops a life-cycle family labor supply framework with heterogeneity, idiosyncratic risk and experience capital to examine policy substantive policy reforms incorporating general equilibrium and government revenue constraints. It documents sizeable life cycle labor supply responses and tax/social security incentives.

This work represents an exciting new agenda, combining recent findings in the dynamic micro labor supply and retirement literature with key features of the macro environment by including general equilibrium feedbacks and a binding government budget constraint. In our discussion we have pointed to some specific areas where the specification could be improved but these do not detract importance of the key insights in this paper and the significant steps it makes along the path to building a coherent micro and macro framework for research on labor supply and taxation.

Acknowledgments

We gratefully acknowledge funding from the Economic and Social Research Council for the grant for the “Centre for Microeconomic Analysis of Public Policy” at the Institute for Fiscal Studies (RES-544-28-50001) and the Keynes Fund.

Data availability

No data was used for the research described in the article.

References

Banks, James, Blundell, Richard, Rivas, Maria Casanova, 2015. The dynamics of retirement behavior in couples: Reduced-form evidence from England and the US. Institute for Fiscal Studies, mimeo.
Blundell, Richard, Alan, Duncan, Costas, Meghir, 1998. Estimating labor supply responses using tax reforms. Econometrica (1998), 827–861.
Blundell, Richard, Bozio, Antoine, Laroque, Guy, 2011. Labor supply and the extensive margin. Am. Econ. Rev. 101 (3), 482–486.
Blundell, Richard, Britton, Jack, 2023. The impact of health on labour supply near retirement. Monica Costa Dias Eric Fr J. Hum. Resour. 58 (1), 282–334.
Blundell, Richard, Christopher, Bollinger, Charles, Hokayem, Ziliak, James P., 2024. Interpreting cohort profiles of lifecycle earnings volatility. J. Labor Econ. forthcoming.
Blundell, Richard, Eric, French, Gemma, Tetlow, 2016. Retirement incentives and labor supply. In: Handbook of the Economics of Population Aging. Vol. 1, North-Holland, pp. 457–566.
Blundell, Richard, Joyce, Robert, Keiller, Agnes Norris, Ziliak, James P., 2018. Income inequality and the labour market in Britain and the US. J. Public Econ. 162, 48–62.
Blundell, Richard, MaCurdy, Thomas, 1999. Labor supply: a review of alternative approaches. In: Handbook of labor economics. Vol. 3, Elsevier, pp. 1559–1695.
Blundell, Richard, Monica, Costa-Dias, Costas, Meghir, Jonathan, Shaw, 2017. Female labour supply, human capital and welfare reform. Econometrica 84 (5), 1705–1753.
Cogan, John F., 1981. Fixed costs and labor supply. Econometrica 945–963.
Eric, French, Jones, John Bailey, 2011. The effects of health insurance and self-insurance on retirement behavior. Econometrica 79 (3), 693–732.
French, Eric, 2005. The effects of health, wealth, and wages on labour supply and retirement behaviour. Rev. Econ. Stud. 72 (2), 395–427.
Gary, D. Hansen, 1985. Indivisible labor and the business cycle. J. Monet. Econ. 16 (3), 309–327.
Heckman, James J., MaCurdy, Thomas E., 1980. A life cycle model of female labour supply. Rev. Econ. Stud. 47 (1), 47–74.
Jones, John Bailey, Li, Yue, 2018. The effects of collecting income taxes on social security benefits. J. Public Econ. 159, 128–145.
Keane, Michael P., 2011. Labor supply and taxes: A survey. J. Econ. Lit. 49 (4), 961–1075.
Keane, M.P., Wasi, N., 2016. Labour supply: The roles of human capital and the extensive margin. Econ. J. 126, 578–617. http://dx.doi.org/10.1111/ecoj.12362.
Myers, Robert J., 1993. Social Security. University of Pennsylvania Press.
Rogerson, Richard, 1988. Indivisible labor, lotteries and equilibrium. J. Monet. Econ. 21 (1), 3–16.
Xiaodong, Fan, Seshadri, Ananth, Taber, Christopher, 2024. Estimation of a life-cycle model with human capital, labor supply, and retirement. J. Polit. Econ. 132 (1), 48–95.